Market Power and Loyalty Redeemable Token Design*

Zhiheng He†
Tsinghua

Kenneth Rogoff‡
Harvard

Yang You§
HKU

October 16, 2024

Abstract

Software and accounting advances have led to a rapid expansion in and proliferation of loyalty tokens, typically bundled as part of product price. Some tokens, such as in the airline industry, already account for tens of billions of dollars and are a major contributor to revenues. An open question is whether, as technology evolves, firms will have a strong incentive to make loyalty tokens tradable, raising regulation issues, including with monetary and banking authorities. This paper argues that for the vast majority of tokens, issuing firms have a strong incentive to make them non-tradable. Our analysis builds on Rogoff and You (2023)'s study of platform currencies to study the dual problem of redeemable tokens, which are vastly more common. The core incentive for token issuance here is that an issuer can earn a higher rate of return on the "float" (tokens issued but not yet used) than its retail customers can, much like a bank. Our main finding is that an issuer earns higher revenue by making tokens non-tradable even though the consumer would be willing to pay a higher price for tradable tokens. We further show that an issuer with stronger market power tends to allow more frequent token redemption, and its revenue is more token-dependent. We test the model's predictions with data on airline mileage and hotel reward programs and document consistent empirical results that align with our theory.

JEL Classification: *G12, G32, G51, M20*

Keywords: *Token issuance, Token-Service Bundling, Redemption, Airline Miles*

*We would like to thank Ke Tang for helpful comments. Yang You, thanks for the funding support of Shenzhen Fintech grant SZRI2023-CRF-03. Zhiheng He, thanks for the funding support of National Natural Science Foundation of China (Grant No. 723B2013). Natalie Kahn, Jinpu Li, Ziwen Sun, and Minghai Xu provide excellent research assistance. All errors are our own.

†Tsinghua University, Email: hezh19@mails.tsinghua.edu.cn.

‡Harvard University, Email: krogoff@harvard.edu.

§HKU Business School, The University of Hong Kong, Email: yangyou@hku.hk.

1 Introduction

Loyalty programs have been a long-lived generation in the evolution of platform assets since the first airline frequent flyer program was launched in 1981. Although it is difficult to compute the total market size, loyalty tokens appear to have reached enormous scale, with many facts speaking to their significance: the value of unredeemed airline miles exceeds $36 billion as of 2021 (based on the airline subset with applicable 10-K reports, see Section 4.2.1); the global market value of the spin-off loyalty management business exceeds $11 billion in 2023;[1] the average American consumer holds 19 (9 active) loyalty program memberships.[2]

We present a simple and tractable model where the token issuer (e.g., a service provider or platform) bundles tokens with unit product sold and decides the quantity of tokens bundled with each sale and whether tokens are allowed to transfer across token holders. When a customer accumulates enough tokens (with the amount being another choice variable), she can redeem them for one unit of the product. In a nutshell, if each sale comes with $1/M$ "free" tokens, consumers face a "buy M get one free" offer. The more tokens bundled, the more the purchase price of one unit/service would rise to reflect the value of the redemption rights. As in Rogoff and You (2023), the core motivation for token issuance is that the issuer has greater outside investment opportunities and so discounts the future more sharply than consumers. Thus a discount wedge drives the issuer's benefits from presale, even at a fair discount for consumers. We generalize our model by allowing the market power of the platform into our economic analysis — product demand decreases as the purchase price rises (e.g., as more tokens are bundled).

Our first key finding is that the issuer's optimal strategy is always to make tokens non-tradable regardless of the market power of tradable tokens, and even if consumers are always willing to pay more for tradable tokens. The intuition lies in the underlying financing mechanism of bundling: it not only bundles tokens with product in the period of sale, but more importantly effectively bundles current and future purchases as consumers need to accumulate tokens over time for a single redemption — we term this as a *time-series bundling*. Bundling enables a de facto cash flow swap that leverages the discount wedge by front loading payments to the issuer. If tokens cannot be traded, consumers value early tokens less relative to later tokens as they need to wait for a longer time to redeem them, but they have to pay the same amount for each purchase. Allowing tradability breaks the *time-series bundling* of token holdings and

[1] See Fortune Business Insights, https://www.fortunebusinessinsights.com/industry-reports/loyalty-management-market-101166.

[2] See the Bond Loyalty Report 2024, https://info.bondbrandloyalty.com/2024-executive-summary.

so eliminates the economic benefits of the cash flow swap. As a result, tradability is still not preferred.

In addition to the aforementioned *time-series bundling*, there are two crucial economic trade-offs of bundling that determine the optimal token issuance design. First, bundling drives consumers to a *forced strategy space* of token holdings, which implicitly requires enough enforcement capacity, otherwise the consumer could give up the whole offer and leave the issuer. Thus, market power enters as a key parameter in the issuer's optimization problem, which in our models is captured by the sensitivity of external (consumer base) losses to unit price increases.[3]

Second, bundling tokens, as compared to selling them separately, has the drawback of *delaying cash flow*, as the issuer can sell tokens only to those with realized demand shocks, rather than to the whole consumer base as with a platform currency. The platform can only fully front-load the cash flow for the consumers who receive consumption shocks in the first period; thus, the delayed cash flow reduces the economic benefits of token issuance.

Making tokens tradable disables the time-series bundling as consumers can then buy loyalty tokens and redeem them without waiting, yet partially compensates the issuer by yielding greater short-term cash flow. As a result, allowing tradability increases token price under any given token bundling issuance regime. Nevertheless, we are able to show the dominance of non-tradability mathematically.

The second key finding is that issuers with greater market power manage to maintain a higher token dependence, defined as the revenue share of token presale. Consequently, a monopoly issuer always favors "buy one and get one free," while micro business may not have the enforcement leeway to benefit from token presales. This implies that industry giants have stronger ability to monetize user bases, further translating their market power into profits through loyalty token issuance.

Our framework allows for several interesting extensions. First, tradability requires stronger commitment power, as issuers can always obtain extra benefits from deviation or suspension after tradability announcements.[4] Second, under perfect competi-

[3]For an extreme example, consumers inevitably accept the monopolist's offer and cannot leave. Therefore, the monopoly issuer suffers no external loss from enforcing bundling. For oligopolies with less market power, however, the potential loss of the consumer base should be taken into account. From another perspective, it is another representation of higher switching costs implemented by large market power (e.g., Klemperer, 1987), linking to the underlying mechanism of customer retention (e.g., Chen, Mandler, and Meyer-Waarden, 2021).

[4]This is because the benefits of tradability are realized once purchases, while deviation rebuilds the time-series building. Put widely, issuers may even be incentivized to interfere redeemability without regulation. This suggests new scenarios for blockchain applications: whereas numerous studies address commitment to token supply using blockchain (e.g., Cong, Li, and Wang, 2022; Chod and Lyandres, 2023; Malinova and Park, 2023), the supply here is deterministic by bundling, but token functionality still needs commitment.

tion, although consumers favor the minimum-cost "buy one and get one free" strategy, the only equilibrium is not to launch any loyalty program, as the only way for issuers to make profit is to deviate. Third, we discuss the potential dominance of tradability: sufficient outside use creates seigniorage and even an over-the-counter market. Their realization relies on tradability of tokens and the embedded permissionless payment by the issuer to the outside business. Traditional issuers fail to utilize outside convenience, and alternatively resort to deep co-operation (e.g. internalized outside use within an alliance) or payment intermediaries (e.g. bundling with credit rewards), while blockchain-based consensus and cryptography technologies may offer similar functionalities with a lower agency cost. Finally, we compare the applicable scope of loyalty tokens to the benchmark platform cash and their complementarity.

After developing the theory, we proceed to empirical analysis based on data from 10-K reports of airlines and hotel groups and their loyalty programs. We find that indeed, companies tend to limit trade in tokens, either by adding fixed and/or proportional fees or restricting transfer flexibility. Consistent with our model, we also find that business scale is a critical determinant of the size of a company's token issuance (normalized by size). In particular, the variation in passenger numbers roughly explains 70% of the standard deviation of token dependence in the airline industry. Using a similar approach in hospitality, we find a one-standard-deviation more guests (0.24 million) is associated with 7.82-percentage-point higher fraction of loyalty liability over operation revenue. Last, we discuss the efforts to seek higher convenience and demand probabilities of tokens by forming aviation alliances. Despite no change in business operation, an aviation alliance would increase the redemption probability, thus increasing willingness to pay for tokens and the issuer's ability to front-load revenue.

Our paper contributes to the large strand of literature on loyalty programs by offering a novel perspective on tokenization, in particular recognizing its value as a financing mechanism. Existing research mostly comes from marketing and operation management, as recently reviewed by Chen, Mandler, and Meyer-Waarden (2021), focusing on how the adoption of loyalty programs influences demand through factors such as customer retention (e.g., Bolton, Kannan, and Bramlett, 2000; Lewis, 2004; Evanschitzky et al., 2012; Kumar and Reinartz, 2016), consumer psychology and behavior (e.g., Sharp and Sharp, 1997; Liu, 2007; Chun and Ovchinnikov, 2019), complementary to user profiling (e.g., Nunes and Drèze, 2006; Rossi, 2018), and evaluation of brand reputation (e.g., Selnes, 1993). The few studies in economics journals primarily on switching costs (Klemperer, 1987, 1995) where loyalty-including economic arrangements can be thought as artificially generating switching costs (Banerjee, 1987), as well as the moral hazard problems from ticketing agencies (Basso, Clements, and

Ross, 2009) where tokens are not directly issued to principal buyers. The fact that loyalty tokens are an asset is recently recognized by Lim, Chun, and Satopää (2024); Chun and Hamilton (2024). Yet surprisingly, (to our knowledge), the fundamental characteristics have not been theoretically modeled.[5]

The rest of the paper is organized as follows. Section 2 sets up the model and draws central implications. Section 3 shows extended discussions. Section 4 presents empirical results and Section 5 concludes.

2 Theory

2.1 Model Setup

We consider a discrete-time infinite-period model, in which a token issuer (e.g. a traditional service provider or digital platform) provides service at unit dollar cost (there is no inflation in the fiat money), and issues tokens representing redemption rights for future product or service. There is a continuum of consumers with unit measure. In every period, each consumer needs a unit of service with probability p. The issuer and consumers have discount rate β^* and β, respectively. Due to greater outside investment opportunities, the issuer values present more, yielding a discount wedge, $\beta^* < \beta$.

Issuance policy: The issuer sells one unit of service and provides $1/M$ unit of loyalty tokens to the consumer. When a consumer accumulates enough fractional tokens to achieve one full unit, her token gives her the right to redeem it for one unit of service.[6] Loyalty tokens are bundled with service delivery and are not sold directly to consumers.

Timeline: In each period, consumers (1) receive tokens (as a reward) for their consumption in the previous period, (2) trade tokens if allowed, (3) receive consumption shocks with probability p, then (4) purchase or redeem for service.

[5] Our paper also adds to the literature on the evolution of redeemable platform currencies and tokenization. Redeemable tokens can be traced back to trading stamps in the 1930s, evolve into the second generation of loyalty tokens (Rogoff and You, 2023), and recently link to modern blockchain technologies (e.g., Sönmeztürk, Ayav, and Erten, 2020), viewed as product tokens in the context of crypto tokens (Cong and Xiao, 2021). Different from the wide literature on token offering in the initial financing stage, we consider the redemption functionality of utility tokens (also as in Chod and Lyandres, 2023, focusing on production competition), and especially uncover the underlying microeconomic mechanism of bundling in the redeemable token issuance.

[6] Here lies an implicit assumption, $M \geq 1$. We extend the entire analysis to the $M < 1$ case in Online Appendix A, i.e., "buy one get N free" strategies. Although our framework allows for its complete exploration, we focus on $M \geq 1$ in the main text for two reasons: (i) a "buy one get N free" strategy imposes significant regulatory pressure and places consumers at extreme risk of exit scam; (ii) as detailed later, the novel feature of loyalty token is the time-series bundling. However, it is naturally absent when $M < 1$ since redemption rights no longer need to be accumulated. Therefore, it is out of our main interest.

2.2 Equilibrium Token Prices

Non-tradable tokens. If tokens are not allowed to be traded, the consumer will enjoy $M+1$ units of goods for free by paying the M transactions, where M is a positive integer. Recall Rogoff and You (2023) shows that the willingness to pay for the i^{th} good delivery can be written as $\left(\frac{\beta p}{1-\beta(1-p)}\right)^i \equiv a^i$. The issuer has pricing power over consumers, and charges the unit price P_M^{NT} to overwhelm the whole willingness to pay of $M+1$ units through M purchases,

$$P_M^{NT} \sum_{i=0}^{M-1} \left(\frac{\beta p}{1-\beta(1-p)}\right)^i = \sum_{i=0}^{M} \left(\frac{\beta p}{1-\beta(1-p)}\right)^i, \tag{1}$$

which yields

$$P_M^{NT} = \frac{1-a^{M+1}}{1-a^M}.$$

Tradable tokens. When tokens can be traded, consumers can sell them to other buyers — that is, a secondary market of loyalty tokens is introduced. Lemma 1 characterizes how tokens are traded among consumers and its market-clearing condition. Through tradability, consumers can avoid holding tokens that are insufficient for redemption. Consequently, the token market clears in the secondary market every period, determining the price $\hat{P}$ by ensuring that consumers are indifferent between holding zero and M tokens (where, again, M is the quantity needed to get a free unit of service).[7]

Lemma 1. ***Token Allocation through Secondary market.***

(i) For any given issuance strategy M, $M \geq 1$, in each period, any user holds M tokens or zero tokens after trading in the secondary market.

(ii) Denote the service price as P_M^T, the secondary market price for unit token $\hat{P}$ is given by

$$\hat{P} = \frac{p+(1-p)a}{M+a} P_M^T. \tag{2}$$

Proof. See Appendix A.1. □

Remark. *The key intuition is that tradability speeds up redemption, as it omits the forced token-accumulating process by introducing the secondary market. Any single token can reliably be redeemed in the next period from its issuance by being combined with tokens from other consumers. In the non-tradable case, however, redemption is not feasible until emerge until M tokens are accumulated.*

[7] Appendix A.1 provides a detailed discussion on the flows of tokens among consumers and shows why they clear at the same price.

For a consumer, one single purchase with fiat money will include unit service plus $1/M$ share of future service that can be sold in the next period. With pricing power over consumers, the issuer could charge a price that fully captures the resale revenue, i.e.,

$$P_M^T = 1 + \beta \hat{P},$$

which solves

$$\hat{P} = \frac{p + (1-p)a}{M}, \text{ and } P_M^T = 1 + \frac{a}{M}.$$

Proposition 1. ***Token price properties.***

Regardless of whether tradability is permitted, service prices would be higher if (i) the demand probabilities p are higher; or (ii) fewer tokens are required for one redemption.[8] *That is,*

$$\frac{\mathrm{d}P_M^x}{\mathrm{d}p} > 0, \; \frac{\mathrm{d}P_M^x}{\mathrm{d}M} < 0, \quad \forall x \in \{T, NT\}.$$

In particular, $\forall M \geq 1$, the tradable tokens are always priced higher, i.e.,

$$P_M^{NT} \leq P_M^T,$$

where the equal sign is obtained if and only if $M \in \{1, \infty\}$ or $p = 0$.

Proof. See Appendix A.2. □

Tradable tokens are always priced higher, as consumers gain excess option value from tradable tokens. The only exception is when $M = 1$ or $p = 0$, in which case consumers naturally accumulate enough tokens to redeem, so there is no trade in the secondary market. At the other extreme, $M = \infty$ is equivalent to not issuing loyalty tokens. Whether tradable or not, higher demand probabilities p increase consumers' willingness to pay for redemption rights, leading to a higher price. On the other hand, smaller M means that more units $(1/M)$ of redemption rights are bundled, and also implies in the non-tradable cases, the rights would be activated earlier.

2.3 Roles of Bundling and Tradability

Compared to a platform currency, the central characteristic of loyalty token issuance is bundling. Now we formally demonstrate the essential role of bundling in three respects:

[8] Note that although M is defined as an integer, the price formulas can be viewed as continuous functions w.r.t. M. The implication naturally holds for integer values. In the following, we allow for the derivatives w.r.t. M without causing confusion.

(1) Forced consumers' strategy space. With bundling, consumers are forced to buy and hold certain number of tokens. To be precise, a consumer holds de facto $k/M < 1$ units of redemption rights after the k_{th} purchase when $1 \leq k < M$, which is always suboptimal under the benchmark issuance — no one would voluntarily buy and hold less than exercisable redemption rights.

(2) Delayed cash flow. Consumers only buy tokens after demand shocks, meaning the issuer misses the front-loading benefits from consumers who do not purchase in the current period. By contrast, in the benchmark platform currency case (where tokens are sold directly by the platform without bundling), the issuer immediately gets front-loaded cash flows (though discounted) from all consumers once tokens are offered. Hence, if the issuer is extremely impatient ($\beta^* \to 0$), bundling sales are suboptimal.

(3) Time-series bundling enables a cash flow swap. A non-tradable issuance strategy effectively bundles future purchases with current purchase, since tokens need to be merged together to become effective, resulting a beneficial "time-series bundling." Put differently, the issuer bundles a series of consumption with ongoing issued tokens in one offer, where the binding force comes from the accumulation of redemption rights. As equation (1) shows, the issuer offers a menu that relates the total price paid to the number of purchases during the accumulation process, and leaves zero consumer surplus. How could the issuer get extra benefits from time-series bundling? The key intuition is: each consumer actually receives different bundled rights in different periods, values them increasingly, but buys at the same price. This effectively enables a cash flow swap that front-loads payments to the issuer.

Figure 1 illustrates how time-series bundling achieves a cash flow swap. Let "after-k right" denote the redemption right that becomes effective after k purchases.[9] Then it is straightforward to prove the consumer's willingness to pay for one "after-k right" is a^{k+1}, whereas the issuer's expected value is a^{*k+1}. In the first purchase, the consumer actually buys $1/M$ shares of "after-$(M-1)$ right." Similarly, the last token for one redemption, $1/M$ shares of "after-0 right," is valued a (a^*) by the consumer (issuer). If the tokens were sold separately (this already requires enforcement), the issuer earns the discount wedge, i.e., the present value of square B plus C as per the issuer (hereafter $B^* + C^*$). On the other hand, the time-series bundling actually sells a token at a constant price, $P_M^{NT} - 1 = \frac{(1-a)a^M}{1-a^M}$, which satisfies

$$\frac{a}{M} \geq \frac{(1-a)a^M}{1-a^M} \geq \frac{a^M}{M}, \quad \forall M \geq 1.$$

[9]Note that the acquired (fragments of) redemption rights remain "inactive" and cannot be used until enough tokens are accumulated.

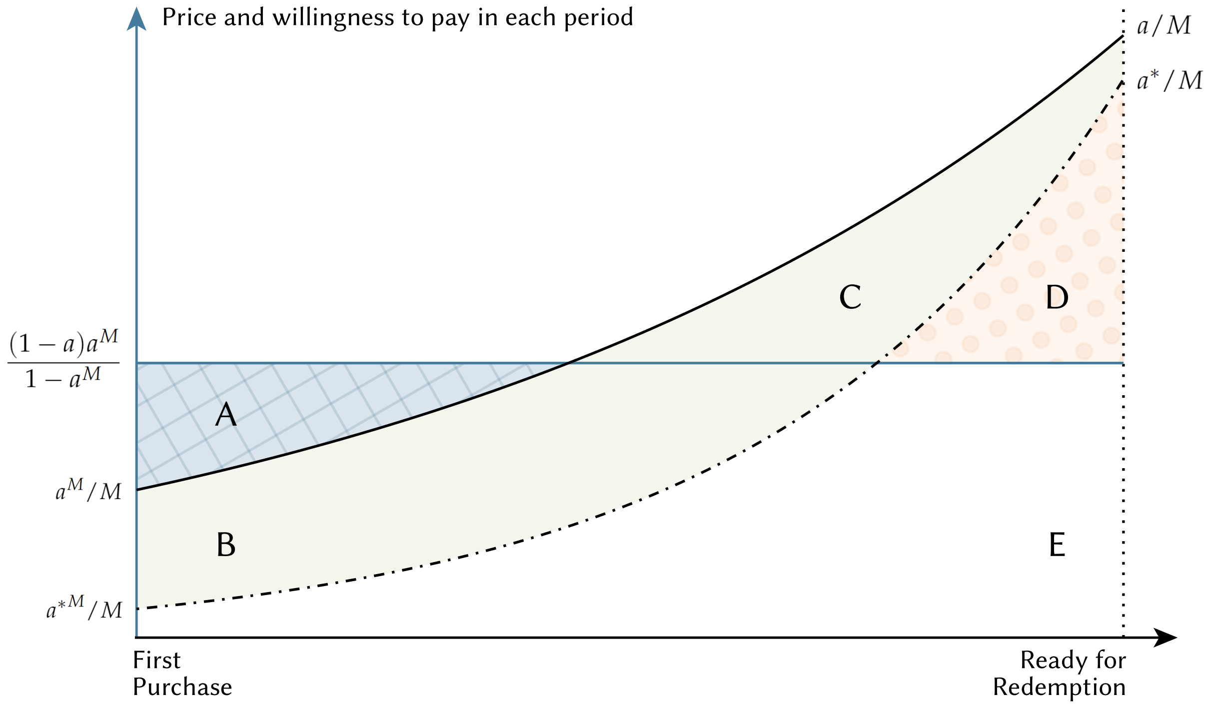

Figure 1. Cash Flow Swap by Time-series Bundling

This figure illustrates how time-series bundling effectively enables a cash flow swap. The solid (dot-dashed) line shows the consumer's willingness-to-pay for (issuer's expected value of) the token in current purchase, starting from the first purchase until the period when enough tokens are accumulated. The blue horizontal line shows the actual token price, $P_M^{NT} - 1$. Squares B, C show the discount wedge. Squares A, D respectively present the issuer's revenue surplus and loss for selling a token in single periods, demonstrating the cash flow swap.

Therefore, consumers overpay in early periods as square A shows, then reduce spending as squares C and D shows. Consumers accept this offer because by (1), this is an offer that makes them just indifferent between buying and not buying, i.e., $A = C + D$. Since the issuer values present more, from its perspective, the present value of A is greater, i.e., $A^* > C^* + D^*$. Consequently, the overall net revenue $(A^* + B^* - D^*)$ is larger than $B^* + C^*$.[10] Therefore, the issuer makes a cash flow swap and yields more front loads than the original discount wedge.

As Figure 2 summarizes, relative to the benchmark, loyalty tokens always take advantage of *time-series bundling* but suffers from *delayed cash flow* by forcing consumers' token holdings. Regarding the *forced strategy space*. It implicitly requires the issuer to have the capacity to implement a bundling issuance, otherwise the consumers may not adopt the issuer and give up the whole offer. As such, we introduce the market power of issuers in later analyses and show how it affects their optimization problem.

[10] Square D does not necessarily exist: when β^* is sufficiently small, $\frac{(1-a)a^M}{1-a^M} \geq a^*/M$. The logic applies similarly, i.e. $A = C \Rightarrow A^* > C^*$, whereas the net revenue becomes $A^* + B^* > B^* + C^*$. The issuer with lower β^* even benefits more from cash flow swap.

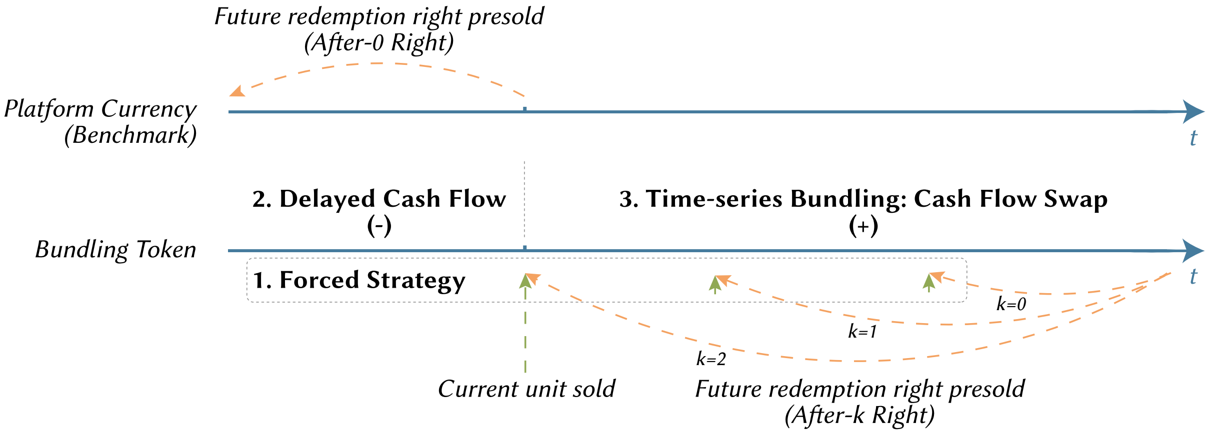

Figure 2. Bundling Token v.s. Benchmark Platform Currency
This figure compares the timeline of issuing bundling tokens (take $M = 3$ as an example) to the benchmark platform currency offering, and illustrates three roles of bundling. The yellow dashed arrows plot the presale of future redemption rights. The green dashed arrows plot current service purchases.

Consider the impact of tradability. First, tradability partially compensates for the delayed cash flow by selling additional reselling rights. From another perspective, the secondary market allows tokens to be held by anyone rather than only who receive demand shocks. On the other hand, tradability undermines the intertemporal correlation of consumer's token holding quantities, thus disables the time-series bundling. In particular, the issuer fails to sell "after-k rights" by adjusting its issuing strategy M, but only "after-0 rights," since tokens are always gathered after trade, as Lemma 1 shows. As a result, there is no room for cash flow swap.

2.4 Issuer's Optimization Problem

So far, we have presumed that consumers accept bundled pricing unconditionally. However, as mentioned in Section 2.3, such an issuance strategy limits their rights to determine token holdings, thus implicitly requiring the issuer to have enough power for implementation. This naturally motivates us to explore market structure, and consider the possibility that consumers may be able to switch within oligopolistic or competitive issuers for similar services. Put differently, if the issuer tries to bundle, it may lose some of its consumer base.

We model this impact as an external penalty related to service pricing in the issuer's optimization problem.[11] In what follows, $c \cdot f(P)$ captures the external penalty, where $f(P)$ is a non-negative function of the service price P. $f'(P) \geq 0$, i.e., higher price

[11]Even if the consumers clearly understand that the higher price is due to the inclusion of more redemption benefits, some are likely limited by budget constraints.

generates more losses.[12] Without loss of generality, we normalize $f(1) = 0$, i.e., when the unit price equals 1 (the service cost), it does not rule out any consumers. $c \geq 0$ is an issuer-specific parameter related to the market power: a lower c indicates a less harmful implementation, corresponding to an issuer with greater market powers, usually associated with large issuers in practice.[13] Monopolist is modeled as the case where $c = 0$, so that consumers have no outside options.

Tin this more general case, where the issuer has less than absolute monopoly power, the issuer's problem is to optimally design the loyalty token issuance strategy (M, x) to maximize the loyalty payoff U_M^x, i.e.,[14]

$$\max_{M \geq 1,\, x \in \{T, NT\}} U_M^x \quad \equiv \max_{M \geq 1,\, x \in \{T, NT\}} Rev_M^x - cf(P_M^x) \tag{3}$$

where Rev_M^x is the aggregate discounted infinite-period revenue when adopting the loyalty program (M, x). It can be further decomposed as $P_M^x Q_M^x$, where Q_M^x is the aggregate discounted fiat sales — although there are always p units of products delivered in one period, some of them are redeemed by tokens. Here equation (3) implicitly restricts c such that the issuer can set the monopoly price in optimal (e.g., $c \leq Q_M^x / f'(P)$, $\forall P \leq P_1^T$), rather than lowering the price to avoid penalty.[15] We prove in Appendix that all the subsequent propositions hold for any $c \geq 0$.

2.5 Monopolist's Optimum and Revenue Characteristics

We first consider the characteristics of revenue Rev_M^x by solving the monopoly case, where the external loss of consumer base is not a concern, i.e., $c = 0$. Then $U_M^x = Rev_M^x$. Consider revenue. The loyalty token issuance (M, x) generates two effects. First, M shapes the share of token-in-advance, or equivalently, the aggregate discounted demand paid by fiat money. Smaller M generates more front loaded revenue, leading to more services paid by redemption in the future. Second, x determines whether reselling options is embedded in the tokens and leaves potential excess earn-

[12]Since this loss is assumed to be directly from the consumers' response to the unit price and is independent of loyalty program designs, $f(\cdot)$ applies to any P_M^x.

[13]From another perspective, it can be considered as an opportunity cost of maintaining the consumer base in an oligopoly with switching costs (e.g., Klemperer, 1987): a large market power (small c) implies less switching cost of consumers from other competitors, which leaves more rooms for price increases.

[14]Since the issuer always faces the same realized demand flow, the present value of the total product cost is constant. We omit the cost term when it is not specified.

[15] For a general $c \geq 0$, the issuer's problem reads $\max_{M \geq 1,\, x \in \{T, NT\}} \max_{P \leq P_M^x} \{Q_M^x P - cf(P)\}$, where note that Q_M^x relates to (M, x) yet is not directly affected by price. That is, if c were sufficiently large, the issuer's optimal pricing might be forced away from the monopoly price. This is consistent with our previous derivation: we presume the issuer has pricing power over consumers to set P_M^x, whereas such pricing power could be destroyed with large c. In an extreme example where $c > Q_M^x / f'(P)$, $\forall P \leq P_1^T$, any price increase would be overall harmful, thus any strategy earns negative payoff relative to the trivial strategy with no token issuance.

ings from charging the reselling option value.

Revenue through non-tradable tokens. On the issuer side, the revenue comes from a series of "buy M and get one free" rounds, with each redemption starting a new round in the next period. Standing at the beginning of each round, the consumer's i^{th} consumption delivers present value (PV) of $\left(\frac{\beta^* p}{1-\beta^*(1-p)}\right)^i P_M^{NT} \equiv a^{*i} P_M^{NT}$. The aggregate revenue reads

$$\begin{aligned}Rev_M^{NT} &= \underbrace{\frac{1}{1-a^{*M+1}}}_{\substack{\text{Discount of}\\\text{All Rounds}}} \underbrace{(p+(1-p)a^*)}_{\substack{\text{Discount of}\\\text{First Purchase}\\\text{in a Round}}} \underbrace{\frac{1-a^{*M}}{1-a^*} P_M^{NT}}_{\substack{\text{PV of}\\\text{M Purchases}}} \equiv Q_M^{NT} \cdot P_M^{NT} \\ &= \frac{p}{1-\beta^*} \frac{1-a^{M+1}}{1-a^M} \frac{1-a^{*M}}{1-a^{*M+1}},\end{aligned} \tag{4}$$

where the aggregate discounted demand $Q_M^{NT} = \frac{p}{1-\beta^*} \frac{1-a^{*M}}{1-a^{*M+1}}$.

Revenue through tradable tokens. First, we solve the number of goods consumers pay with fiat money in each period. Let $Q_{M,t}^T$ represent the fiat sales in period t (i.e., $\sum_{i=0}^{\infty} Q_{M,t}^T = Q_M^T$), and $S_{M,t}^T$ represent the number of outstanding units of redemption rights, i.e. there are $MS_{M,t}^T$ unredeemed tokens in period t. Naturally, $Q_{M,0}^T = p$, $S_{M,0}^T = 0$. With the total demand always equal to p, we obtain $\forall t \geq 0$,

$$p = Q_{M,t}^T + pS_{M,t}^T. \tag{5}$$

On the other hand, the total number of redeemable tokens equals the newly issued tokens plus the unredeemed part from the previous period,

$$S_{M,t+1}^T = \frac{Q_{M,t}^T}{M} + (1-p)S_{M,t}^T. \tag{6}$$

Combining equations (5) and (6) with the initial values, we obtain

$$S_{M,t}^T = \frac{1-\left[1-\frac{p(M+1)}{M}\right]^t}{M+1}, \quad Q_{M,t}^T = \frac{p}{M+1}\left[1-\frac{p(M+1)}{M}\right]^t + \frac{pM}{M+1}.$$

The total revenue for tradable token issuance is

$$Rev_M^T \equiv Q_M^T P_M^T = \sum_{t=0}^{\infty} \beta^{*t} Q_{M,t}^T P_M^T = \frac{p}{1-\beta^*} \frac{M+a}{M+a^*}. \tag{7}$$

Lemma 2. *Revenue and token quantity.*

The aggregate discounted revenue decreases with M regardless of tradability,

$$\frac{\mathrm{d}}{\mathrm{d}M} Rev_M^x < 0, \quad \forall x \in \{T, NT\}; \quad \textit{and} \quad Rev_1^{NT} = Rev_1^T = \frac{p}{1-\beta^*} \frac{1+a}{1+a^*}.$$

Proof. See Appendix A.3. □

Since $Rev_M^x = P_M^x Q_M^x$, the impact of changing M can be decomposed into two opposite parts: smaller M is associated with higher token price, but more redemption rights granted reduce future services paid by fiat money. Lemma 2 ensures that given x fixed, the former takes the dominance.

As a result, without consumer base losses, the monopoly issuer would always prefer $M = 1$, i.e., "buy one and get one free." Comparing the tradable tokens with non-tradable ones. Essentially, they are identical as consumers do not need to trade but always naturally gain sufficient tokens for redemption under "buy one and get one free" policy. This is exactly the limiting case where tradability does not matter.

Tradable vs. non-tradable tokens. In a wider parameter space of (p, M), however, the monotonicity of Rev_M^x w.r.t. M does not reveal the comparison between tradability and non-tradability, which relates to the roles of tradability discussed in Section 2.3.

Figure 3 (a) visualizes the comparative statics of the token price w.r.t. probability p and strategy M. Tradable tokens are always higher priced, since consumers obtain excess option values from selling. The exception appears in $M = 1$ and $p = 0$, where consumers do not benefit from trading. Once the parameters leave these two boundaries, the relative price advantage of tradable tokens increases when p and M become smaller as consumers value tradability more when more redemption rights bundled or they purchase goods more frequently.

Panel (b) shows the revenue comparison. Relative to the non-tradable case, tradability introduces the trade-off driven by two opposite economic forces, *compensating delayed cash flow* by high prices, yet *disabling time-series bundling*. With small M and p, the former effect dominates and $Rev_M^T > Rev_M^{NT}$, as it is far more advantageous to sell more redemption rights to the secondary market than to wait for the next batch of consumers with small probabilities. As M increases, the compensation is reduced because fewer rights are sold, while the time-series bundling is lengthened, making disabling it more costly. As a result, Rev_M^T / Rev_M^{NT} converges to one more quickly than P_M^T / P_M^{NT} in panel (a). Similarly, a greater p leaves delayed cash flow less a concern, diminishing the importance of compensation, and can even make non-tradability more

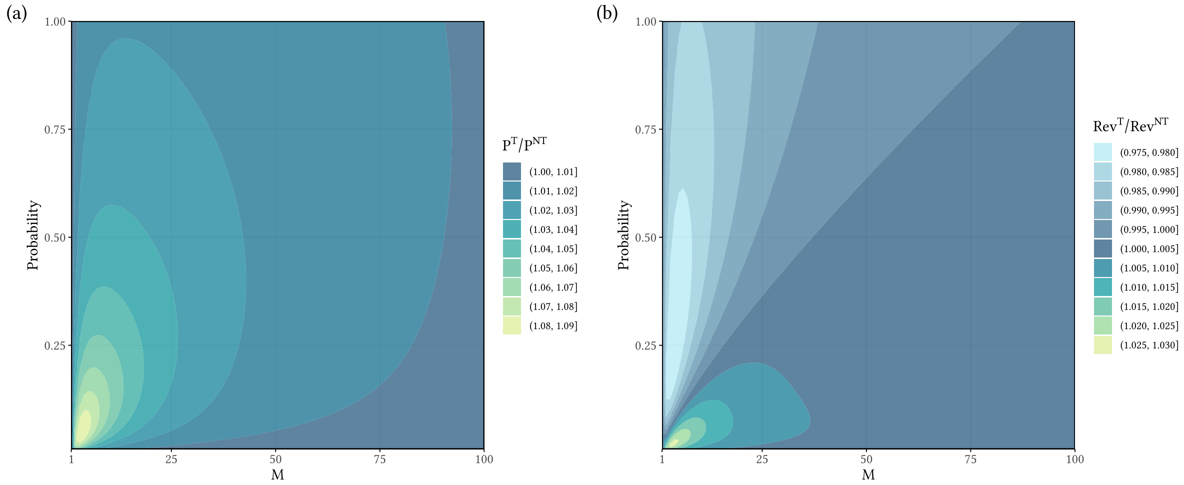

Figure 3. Tradable v.s. Non-Tradable Token Designs

This figure compares tradable and non-tradable tokens over the parameter space (p, M). Panel (a) shows the relative price dominance of tradable tokens, where the coloring scheme visualizes how P^T/P^{NT} changes across the parameter space. Panel (b) shows the revenue of tradable tokens relative to non-tradable tokens, i.e., Rev^T/Rev^{NT}. $\beta = 0.95$ and $\beta^* = 0.9$.

profitable.[16] In general, Panel (b) demonstrates the complex impact of introducing tradability to the aggregate revenue under a given (p, M) and highlights the violation of Proposition 2 in Rogoff and You (2023), which states that the revenue is higher for non-tradable tokens for any given issuance policy.

2.6 Dominance of Non-tradability

Can we compare tradable tokens versus non-tradable token when the issuer makes the optimal choice of M? Of course, if the issuer is a monopolist, the optimal issuance policy is always "buy one get one free" regardless whether tokens are allowed to trade. Thus, we study the tradability in a more general scenario where the issuer chooses optimal M by maximizing (3) with any market power $c \geq 0$.

Recall that Figure 3 (b) shows tradability adds heterogeneous gains or losses of revenue under different strategies M — the issuer needs to find the global optimal payoff across values of M and tradability decisions, after accounting for enforcement losses. Without solving the optimal M analytically, we can prove that any optimal strategy is always to issue non-tradable tokens, regardless of consumption frequency p and market power c.

Proposition 2. ***Dominance of Non-tradability.***

[16] Although we have interpreted the main driving force, the changes of Rev_M^T/Rev_M^{NT} w.r.t. p is non-monotonic, as higher p also slightly shortens the time-series bundling. Further examination of such non-monotonicity detracts from our main focus.

Given any consumption probability $p \in (0,1)$ and market power c, the issuer's optimal strategy is always to issue non-tradable tokens, i.e.,

$$U^{NT}_{M^{NT}_*} \geq U^{T}_{M^{T}_*}, \ \ \forall c \geq 0,$$

where $M^x_* = \arg\max_{M \geq 1} U^x_M$, $x \in \{NT, T\}$. The equal sign holds only if $c = 0$.

Proof. See Appendix A.4. □

The logic of the proof is that the advantage of tradability, higher price, may be achieved (at least partially) through alternative non-tradable approaches, while its complete absence of the time-series bundling is inevitable. To be precise, denote (M, NT) (M non-tradable tokens for one redemption, $M > 1$) as the benchmark strategy. There are two ways to switch to another strategy with a higher unit price: (i) allowing tradability, i.e., (M, T); (ii) shortening the accumulation process, i.e., (m, NT) with $m < M$. Its price also increases as it bundles more shares of redemption right, rather than excess reselling option values. As proved in Appendix A.4, it is always possible to find a strategy (m, NT) that matches the price of (M, T). Once $m > 1$, the token accumulation has not completely vanished — there are always rooms to enable a cash flow swap. Therefore, this strategy brings the same advantage (also external loss) as (M, T), but still benefits from time-series bundling.[17]

The economic intuition is: tradability systematically eliminates time-series bundling by allowing consumers to pool tokens together through the secondary market. The reselling option value (also compensation for delayed cash flow) is essentially the premium of immediately unblocking the token's redeemability. Therefore, the compensation entirely depends on the total sales of redemption rights. When both non-tradable and tradable designs are forced away from $M = 1$ and reduce redemption right sales, non-tradable issuance simultaneously enables a swap as an offset. Put differently, once allowing for tradability, the issuer cannot adjust M to benefit from time-series bundling. Thus, tradability further limits the degree of freedom of issuer in searching the optimal M, thus making the issuer worse off.

An alternative interpretation is that tradability speeds up token redemption, and thus reduces the aggregate discounted fiat sales. Once time-series bundling is in place, tokens are delayed in being activated for redemption, leaving consumers' up-front demands to be met by fiat purchases, they can only pay with tokens when they gave accumulated enough. Then if non-tradable tokens were able to sold at the same high price as tradable tokens, they would generate more front-load cash flow. This is achievable by bundling more redemption rights. In general, both interpretations

[17] Note that in this proof, there is no need for m to be the optimal non-tradable plan M^{NT}_* — In certain scenarios, a non-tradable strategy with lower prices could even be more beneficial.

suggest that time-series bundling is key to the dominance of non-tradability, as it provides a financing tool, increasing front-loading beyond adjusting supply quantities. In addition, the logic and proof in Appendix A.4 are not restricted to parameter choices of c – even if the issuer were not able to set monopoly prices (as footnote 15 mentions), its optimum would be issuing non-tradable tokens.

Link to redeemable platform cash. The dominance of non-tradability echos the core conclusion of redeemable platform cash in Rogoff and You (2023). However, they are achieved through completely different mechanisms, and the non-tradability result correspondingly depends on entirely different intuition. Recall that for the platform currency, the resale market (i.e., allowing tradability) makes it possible to buy tokens in the future, undercutting the current consumers' willingness to hold large quantities. Therefore, competing with the resale market limits the issuer's pricing power in advance. Under a loyalty token design, such friction is systematically eliminated — tokens are bundled with current services, which means consumers have no need nor option to decide the holding quantities from the primary market. The complete control over token supply keeps the resale market from jeopardizing the issuer's power to charge a high price upfront. In fact, the issuer even obtains a higher price, as Proposition 1 shows.

How could non-tradability still take the dominance? We have shown that time-series bundling holds the key. This is also unique to loyalty tokens and non-existent under issuing platform cash, where tokens are always redeemable once held by consumers. Recall Figure 2, to sell one "after-k right" by issuing platform cash, $k > 0$; the issuer has to issue $k + 1$ units at the value of the marginal token, which leaves a positive consumer surplus.

Relative to platform cash, bundling overcomes the competition with the secondary market, while simultaneously creating a cash flow swap if tokens are not permitted to trade. Both tradability and non-tradability are improved by bundling, meanwhile the major cost comes from the delayed cash flow. In Section 3.4, we go further in discussing the adoption choice between issuing loyalty tokens and platform cash.

2.7 Market Power and Token Dependence

This section considers the role of market power in the optimal issuance policy. Issuers with different market power should have heterogeneous capacities to leverage it with loyalty tokens. To evaluate the issuance effectiveness, we introduce an important indicator, token dependence, defined as the revenue share of loyalty token sales under any specific issuance design (M, x), i.e.,

$$\Theta_M^x \equiv \frac{\text{Revenue from Token Sales}}{\text{Total Revenue (Token \& Product Sales)}} = \begin{cases} \frac{a}{M+a}, & x = T; \\ \frac{a^M(1-a)}{1-a^{M+1}}, & x = NT. \end{cases} \tag{8}$$

It is straightforward to show $\mathrm{d}\Theta_M^x/\mathrm{d}M < 0$. Then given the choice x and parameters (p, c), and Lemma 2, a smaller M is equivalent to a higher token dependence, as well as greater revenue. We can compare the issuers' capacities of utilizing the loyalty token financing tool by examining their M under optimal strategy.

Proposition 3. ***Market power and token dependence.***

The issuer with a larger pricing power (smaller c) is able to maintain a higher token dependence (smaller M) no matter whether tradability is allowed. That is,

$$M_*^x(c) \leq M_*^x(c'), \quad \forall 0 \leq c < c',$$

where $M_*^x = \arg\max_{M \geq 1} U_M^x$, $x \in \{NT, T\}$.

Proof. See Appendix A.5. □

Proposition 3 indicates that greater market power is key to maintaining a high token dependence. As Section 2.5 shows, the monopoly issuer's first best is $M = 1$, i.e., "buy one and get one free." However, for oligopolies and even small enterprises, smaller M leads to higher price but incurs consumer base shrinkage. Therefore, $M = 1$ might not be optimal. This logic is straightforward, yet the implication is noteworthy: issuing loyalty tokens serves as a booster for widening the market position gap. Only industry giants are capable of largely monetizing their user base or network traffic via loyalty programs, thus further consolidating their competitiveness. This echoes Banerjee (1987)'s argument, i.e., loyalty program arrangement can effectively generate switching costs for consumers. Then we can infer that not all issuers have this privilege: inherent high switching costs create room for loyalty tokens, further expanding switching costs. Importantly, issuers can issue loyalty tokens to leverage their market power for self-financing meanwhile stabilizing their customer base.

We next use numerical approaches to quantitatively examine M under optimal strategies over the parameter space (p, c), as shown in Figure 4 (a). The simulated optimal strategies, solved by brute force in the strategy space $M \leq 200$, $x \in \{T, NT\}$ for each point (p, c), show to be all non-tradable, confirming Proposition 2. For any given probability p, lower c (larger market power) is associated with a smaller optimal M, equivalently a greater token dependence Θ shown in Panel (b), corroborating Proposition 3. In addition, probability p affects token dependence non-monotonically, as it affects both the total sales and consumers' willingness to pay. With extremely low demand probabilities, consumers do not value future rights. In industries with low demand probabilities (e.g., international airlines), token presale generates significantly advanced cash flow, whereas in frequent-demand industries (e.g., daily retailing), issuing tokens fails to achieve long-term cash flow swap, making presales relatively less crucial relative to fiat sales.

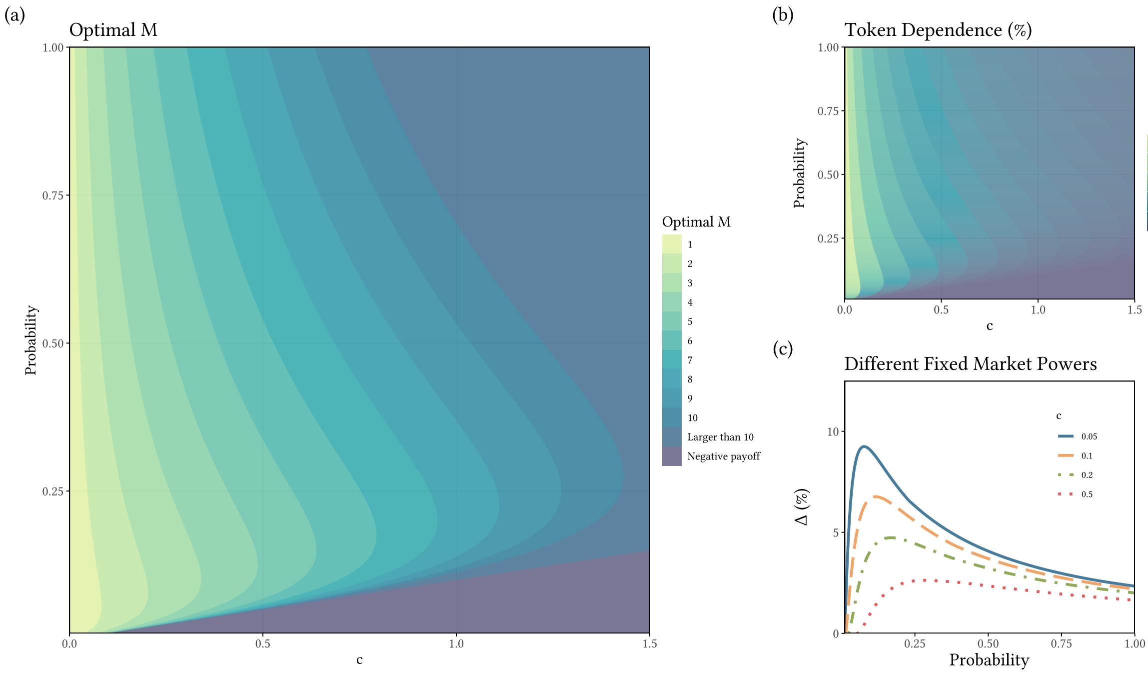

Figure 4. Optimal Issuance and Corresponding Token Dependence

Panel (a) plots the optimal issuance plan M under different probabilities p and market powers c. For each point (p, c), the optimal (discrete) choice of M is solved by brute force ($M \leq 200$, $x \in \{T, NT\}$), confirming the consistent preference for non-tradable tokens in achieving optimal outcomes. 'Negative payoff' refers to cases where all the designs fail to yield positive payoffs, as mentioned in footnote 15. Panel (b) shows the corresponding token dependence Θ^* under the optimal issuance $(M_*^{x^*}, x^*)$. Panel (c) shows how $\Delta = (U^* - Rev_0)/Rev_0$ changes with probability p under different market power c. $\beta = 0.95$, $\beta^* = 0.9$, and $f(P) = P - 1$.

Figure 4 Panel (c) shows an alternative evaluation, $\Delta \equiv (U^* - Rev_0)/Rev_0$, where U^* is the payoff under the optimal issuance plan $(M_*^{x^*}, x^*)$, and Rev_0 denotes the benchmark revenue without token issuance. It is interpreted as the net payoff growth relative to the benchmark of without issuing any tokens. Whereas Θ focuses on revenue generation within an issuance strategy, Δ presents clear theoretical implications on the trade-off of starting a loyalty program.[18] A issuer with greater market power (smaller c) can always generate higher payoff growth by issuing loyalty tokens, re-

[18] We use Θ as the main indicator, as it is closer to practice. As documented in the empirical section, enterprises' accounting statements disclose redemption revenues and unredeemed loyalty liabilities, while the external loss c is mostly unobserved or even non-financial. Online Appendix B provides more discussion on Δ.

gardless of the demand probabilities. Therefore, the ability to maintain a large loyalty scale is a unique advantage of strong market power.

Until now, we have treated the probability p as an exogenous parameter of loyalty token issuer. In practice, a higher p also captures a dimension of market power. Issuers may also strive to change demand probabilities through loyalty programs, for example, by collaborating with other companies to expand the realm of redeemable services or issuing loyalty points with credit card companies to attract more consumers. We further illustrate this point with the airline industry in Section 4.3.

3 Extended Discussion

3.1 Commitment

In the baseline model, we focus on cases where issuers always maintain the same issuance plan. However, in the absence of regulation, issuers may gain by deviating issuance policy, leading to commitment problems.

Proposition 4. ***Unreliable commitment to tradability.***

With naive consumers, unregulated issuers always have an incentive to pretend to allow tradability and switch back to non-tradability in the future.

Proof. See Appendix A.6. □

The key intuition is that the downside of tradability (the failure of time-series bundling) acts across periods, while the upside (compensation from higher prices) is achieved immediately after sales. Therefore, issuers have an incentive to charge a higher price and then prohibit tradability at any time after issuance, i.e., directly impose mandatory intertemporal correlation of token holdings. In other words, issuers always hope that tokens will be redeemed much later after the sale.

It is straightforward to see that with naive consumers and no regulations, the most beneficial strategy is to claim tradability at the beginning and then move back to the overall optimal strategy (M_*^{NT}, NT) when the number of unredeemed tokens reaches a maximum. However, such a strategy is clearly impossible in the long run, since people will realize that the commitment to tradability is unreliable, and in particular, that tradability is suboptimal for issuers. Of course, with clever consumers, such worries would lower consumers' expected payoff of holding tradable tokens. In our model, this leads to an additional discount in the market clearing condition 2.

This suggests a potential upside for the introduction of blockchain and cybersecurity technologies. If issuers were able to announce a credible commitment against deviations in redeemability and tradability, consumers would value the redemption and resale rights without discount. This is particularly crucial for startup issuers with limited reputations. While numerous literature analyzes blockchain embedding for commitment to future token supply,[19] token quantities no longer need to be committed in the context of loyalty tokens. However, the essential goal of the commitment, i.e. token valuation, still needs to be ensured by fixing its functionality.

3.2 Perfect Competition

Assume the issuers offer homogeneous services, and all services are perfect substitues. Then, under perfect competition, $\forall M \geq 1$, an infinite number of issuers offer the issuance strategy with parameter M and earn zero profit. That is, the service price of first M purchases cover the cost of $M+1$ services,

$$P_M^{NT} \sum_{i=0}^{M-1} a^{*i} = \sum_{i=0}^{M} a^{*i} \quad \Rightarrow \quad P_M^{NT} = \frac{1 - a^{*M+1}}{1 - a^{*M}}.$$

Consumers solve the cost-minimization problem, $\min_{M \geq 1} C_M^{NT}$,[20] where

$$C_M^{NT} = \underbrace{\frac{1 - a^M}{1 - a}}_{\substack{\text{PV of First} \\ \text{M Payments}}} \underbrace{P_M^{NT}}_{\text{Unit Price}} \underbrace{\frac{1}{1 - a^{M+1}}}_{\substack{\text{Discount since} \\ \text{First Purchase}}} \underbrace{(p + (1-p)a)}_{\substack{\text{PV of First} \\ \text{Purchase}}} = \frac{p}{1 - \beta} \frac{1 - a^{*M+1}}{1 - a^{*M}} \frac{1 - a^M}{1 - a^{M+1}}. \tag{9}$$

Compare C_M^{NT} with the revenue formula (4). They share the exact same form, except for exchanging positions of β and β^*. Essentially, they represent a pair of duality problems: the issuer's revenue-maximization under complete monopoly and the consumers' cost-minimization under perfect competition ($c = 0$, as explained below). In the monopoly case, consumers do not enjoy any surplus, and the price relates to consumers' discount. Conversely, under perfect competition, issuers do not generate any surplus, causing the price to be determined by the issuers' discount. For the rest three terms, they are discounted values and thus relate to β^* in issuers' problems and β in consumers' problems, respectively.

Remark. *Note that c represents the relative competitiveness in terms of the price sensitivity of the consumer base. The implicit assumption is that there are alternative issuers for these consumers to satisfy their demands. Therefore, a monopoly issuer has complete pricing power $c = 0$, i.e., consumers have no choice but to stay on the platform. Whilst under perfect competition, this term does not exist at all, since the consumer base of any issuer is infinitesimal.*

[19] For example, Chod and Lyandres (2023) also studies the product market competition stage, focusing on the advantages of introducing blockchain technologies and smart contracts to product tokens, showing that commitment to future token sales benefits both giants and new entrants.

[20] Here lies an implicit intuition: the consumer will continuously adopt the specific issuer once purchase, since she has partly pre-consumed its future service through the loyalty tokens.

Or from another perspective, all issuers have no pricing power and maintain prices equal to production costs. Consumers choose service providers randomly, so $c = 0$.

Tradable token price: tokens in the secondary market still clear at the price that satisfies equation (2), while the issuers under perfect competition lose the pricing power to occupy the reselling revenue, and the price equals to the service cost,

$$P_M^T = 1 + \frac{a^*}{M}.$$

Consumers' cost-minimization problem: As the secondary market price ensures that the consumers (sellers and buyers) have the same cost, we only need to solve one of the two, say the sellers. They purchase services through fiat money forever and earn spreads through the secondary market, as the following formula shows, which also appears to be a dual problem of equation (7).

$$C_M^T = \underbrace{\frac{1}{1-a}}_{\substack{\text{PV of All} \\ \text{Payments}}} \underbrace{\left(P_M^T - \beta \hat{P}\right)}_{\substack{\text{Purchase Cost minus} \\ \text{Selling Revenue}}} \underbrace{(p + (1-p)a)}_{\substack{\text{PV of} \\ \text{First Purchase}}} = \frac{p}{1-\beta} \frac{M + a^*}{M + a}. \tag{10}$$

By the nature of dual problems, one would expect that under perfect competition, consumers' cost-minimization would also solve $M = 1$ in optimal, which is proved in the following proposition. However, there is another crucial difference beyond the duality properties — under perfect competition, issuers cannot make profits unless deviate. Similarly to the commitment problems discussed in Section 3.1, the only profitable strategy is to issue loyalty tokens and then prohibit their tradability and even redeemability.[21] As a result, loyalty tokens are not equilibrium under perfect competition. Proposition 5 rationalizes this logic.

Proposition 5. ***Loyalty tokens under perfect competition.***

Under perfect competition, (i) with commitment, consumers prefer $M = 1$ ("buy one, get one free") regardless of whether tradability is permitted,

$$C_1^T = C_1^{NT} = \min_{M \geq 1} \left\{ C_M^{NT}, C_M^T \right\} = \frac{p}{1-\beta} \frac{1 + a^*}{1 + a};$$

(ii) without commitment, the only equilibrium is not to issue tokens.

Proof. See Appendix A.7. □

[21] Although for monopoly issuers, they can also earn excess profits from such deviations, they are less likely to do so because they have enough profit margins and are unlikely to risk their reputation on a one-shot deviation.

Proposition 5 echoes two phenomena in practice: (i) famous loyalty programs tend to be found in oligopolistic and asset-heavy industries, or at least run by industry leaders, rather than retailers under near-perfect competition, e.g., similar micro merchants on Amazon. Successful issuers should not suffer from commitment crises but have revenue creation motives; (ii) in contrast, when merchants, especially retailers, are about to collapse, they tend to offer extra-long time-limited memberships, which are not designed for fulfilling the loyalty rights. This highlights the importance of considering the credibility of loyalty programs.

3.3 Outside Use Cases

Proposition 2 states that without outside use cases, tradability is always dominated by non-tradability. Meanwhile, tradable designs are remarkably common in the new generation of crypto tokens (e.g., Ethereum). This is because these tokens differ in their sources of value, featuring transactional convenience and network effects in digital ecosystems. In the context of loyalty tokens, issuers (largely from classical industries) may not prioritize network effects. Instead, we focus on a unique form of convenience — when there is significant demand for the issuer's tokens outside the loyalty program.

If the tokens can be used for activities outside the issuer's business, e.g., directly using airline miles from airline A to credit coffee purchases from café B, consumers would expect a greater probability for redemption. Realizing this extra demand relies on tradability, whether by trading tokens to a random café customer or transferring from one's airline loyalty account to her own café account. This reselling option value, again as Lemma 1 shows, can be fully charged by the issuer and yield an excess front-loaded cash flow. On the other hand, the extra "cost" is that, each successful outside redemption means an automatic payment from airline A to café B. Then does the issuer always obtain a greater aggregate payoff from outside use? Yes, the logic is that the issuer essentially indirectly finances from the outside companies by issuing permissionless tradable checks. The outside companies are also willing to do that, as their cash flows are not affected.

In the above case, we do not require a deep cooperation between the issuer and outside companies.[22] The permissionless tradability and credit guarantee the execution of outside token redemption, for which the recent blockchain-based consensus and cryp-

[22] Airline alliances are examples of deep cooperation. The alliance becomes a de facto joint issuer. Its business encompasses the routes of alliance members, and importantly, there constitutes no outside use between the members. The joint issuer enjoys a greater demand probability regardless whether allowing token tradability. Section 4.3 provides further discussion. Of course, this advantage require significantly greater cooperative efforts.

tography technologies already enable similar functionality.[23] This kind of outside use also reminds us of a common form of loyalty programs: buy and earn credit card rewards, where banks or digital payment platforms offer a third-party payment channel and validation. Therefore, by introducing new technologies, the issuer may be capable to run a more independent program with significantly lower intermediary costs.

Furthermore, when outside permissionless tradability is widely valid so that not only consumers want to hold tokens, allowing permissionless tradability would create an over-the-counter (OTC) market. Consumers then become OTC brokers who purchase services and replenish token supply, and pay more for outside transaction convenience. This outside convenience creates seigniorage. Here the issuer shares many similarities with stablecoin providers in the cryptocurrency market (e.g. Tether). However, maintaining such price peg through bundling can be extremely costly, especially when outside use far exceeds the fundamental consumption demand.

In general, issuers show no lack of incentives to integrate demands, but subject to technological constraints, they tend to resort to deep co-operation or the introduction of payment intermediaries. New technologies offer the potential to leverage outside convenience through permissionless tradability, but bring with it systemic external risks.

3.4 Loyalty Token versus Platform Cash

Section 2.6 compares the mechanisms by which non-tradability dominates when issuing loyalty tokens and platform cash. Here we further discuss their relationship.

Loyalty tokens and platform cash are viewed as two generations in the history of platform currencies (Rogoff and You, 2023), which are both active in business today and fit technology innovations for the embedding of new features. Given our focus, i.e., discount wedge as the sole source of gains from issuing tokens, they both benefit from service presales. While the key difference is that loyalty tokens bundle service delivery with the token issuance, rather than solving token issuance as an independent decision. In practice, they give examples such as "buy one and get one free" and "token-in-advance," respectively.

Then under what condition does the token issuer prefer loyalty bundling to the "token-in-advance"? First, they are not mutually exclusive — issuers could even adopt them in combination. For example, Starbucks issues stored-value cards as platform cash, and simultaneously runs a loyalty program with redeemable tokens named "stars". Consumers can accumulate stars for each purchase they make, whether paid by platform cash or fiat money. Interestingly, paying for the same order using platform cash accumulates more stars. This sees the two tokens compensating each other's

[23] For example, as introduced in `https://coinmarketcap.com/academy/glossary/permissionless`.

drawbacks. (i) For platform cash, access to future purchase limits the willingness to hold larger quantities at present. The above design creates an incentive to expand current platform cash holdings, where the incentive benefits will only be realized through future consumption and do not add additional cost to the issuer; (ii) For loyalty tokens, the core shortcoming is that it works only after the fiat-money consumption is made. The above design incentives consumers to pay the bundled service by the pre-purchased platform cash.

Second, if the issuer were restricted to adopt only one between loyalty token and platform cash,[24] the crucial consideration would be how high the issuer value front loads. In an extreme example, the impatient issuer ($\beta^* \to 0$) would prefer platform cash, as loyalty tokens have a downside of delayed cash flow. In nonextreme cases, however, the issuer could favor loyalty tokens, as they allow for ongoing issuance with ongoing fiat purchase. It is not hard to see why loyalty tokens are usually more advantageous, as they require additional implementation efforts.

In addition, platform cash may be subject to greater regulatory pressure and higher risk of exit fraud, especially in the initial financing stage, while loyalty tokens are always tied to real products or services (although relatively limited in applicability).

4 Empirical Findings

This section provides empirical evidence for core testable model implications: (i) companies tend to limit token transaction; and (ii) companies with greater market powers are able to maintain higher token dependence. In addition, we discuss practical efforts to obtain higher demand probabilities (in tokens).

4.1 Designs for Blocking Tradability

Our model shows that token tradability is dominated by non-tradability, suggesting token issuers to disable token transactions. We present a narrative analysis to demonstrate this fact. It is worth pointing out that the documented narratives in our analysis may not apply to all past times, could change over time, and may have multiple motivations, whereas our focus lies in discussing their inner logic and implications related to token usage.[25]

As Table 1 shows, none of the airlines or hotels in our sample allow unlimited free transactions. The most direct approach is to ban tradability, labeled "N" in column (1).

[24] This may due to the cost control of developing and maintaining specific token systems, as well as reputation and regulatory considerations.

[25] The narratives collected in this section is based on latest officially-announced terms and conditions by March 2024, and may omit subsidiary specific or related provisions, e.g. exemption policies, regional differences, etc.

However, it may appear overly forceful or uncaring for consumers. Instead, issuers often employ various transfer restrictions to indirectly reach the ban.

Table 1. Transaction Fees and Limitations

Corporation	Tx	$ Cost / 10^3 Units	$ Fee / Tx	Limit	Paper Work	Family Pool	Corporation	Tx	$ Cost / 10^3 Units	$ Fee / Tx	Limit	Paper Work	Family Pool
	(1)	(2)	(3)	(4)	(5)	(6)		(1)	(2)	(3)	(4)	(5)	(6)
Airline							Lufthansa	Y	11.2	0	Y	N	2+5
Alaska	Y	10	25	Y	N	N	Qatar	Y	Y	Y	Y	N	9
American	Y	0	15	Y	N	N	Singapore	W	1	0	Y	N	0+5
Avianca	Y	15	0	Y	N	N	Southwest	Y	10	0	Y	N	N
British	Y	50	0	Y	N	7	Spirit	N					9
Delta	Y	10	30	Y	N	N	Turkish	Y	20	0	N	N	SA
Emirates	Y	15	0	Y	N	8	United	Y	15	30	Y	N	N
AirFrance	Y	Y	Y	Y	N	2+6	VrginAmerica	Y	10	25	Y	N	N
Hawaiian	Y	10	25	Y	N	N	**Hotel**						
VirginAtlantic	Y	25	22	Y	N	10	BestWestern	W	0	0	Y	Y	SA
WestJet	Y	0	50	N	N	N	Choice	N					N
AirCanada	Y	20	0	Y	N	8	Hilton	W	0	0	Y	N	11
Cathay	Y	17	0	Y	N	N	Hyatt	Y	0	0	Y	Y	N
ANA	N					8	IHG	Y	5	0	N	N	N
Etihad	Y	Y	Y	Y	N	9	Marriott	Y	0	0	Y	Y	N
Frontier	Y	25	5.6	Y	N	8	Radisson	W	0	0	N	Y	SA
JetBlue	Y	12.5	0	Y	N	7	Starwood	Y	0	0	Y	Y	N
KoreanAir	N					5	Wyndham	N					N

Notes. This table reports the narrative features of loyalty token transaction. *Tx* clarifies if token transaction is allowed, with "W" indicating that transfers are confined to family groups. The following columns are applicable only when transfer is allowed. *$ Cost / 10^3 Units* documents the transaction fee in US dollars for every 1,000 airline miles or loyalty points (the unit name varies across companies), with "Y" denoting unreported or elaborate fee structures. *$ Fee / Tx* shows fixed processing charge per transaction, with "Y" denoting unreported or elaborate fee structures. *Limit* shows whether there are further transaction limits, e.g. quantity caps and floors, maximum yearly transactions, stepped quantity choices. *Paperwork* conveys if extra processes are required, e.g. form submission and telephonic appointments. *Family Pool* outlines the maximum permissible members in a family pool (or similar group plan) for shared token use. "N" signifies no family pooling, "SA" allows unlimited members with the same address, and * denotes additional restrictions, such as limiting to children accounts. All the contents are collected from official terms and conditions in March 2024.

Fixed and/or proportional fees. The first approach is to add financial costs for transactions, e.g., processing fees and/or proportional transaction fees, as Table 1 columns (2)-(3) show. Fees are usually extremely high relative to the value of the tokens to be transferred. For example, even as one of the most affordable loyalty programs, the `Mileage Plan`™ run by Alaska Airlines only allows consumers to transfer 1,000 to 30,000 miles (i.e., the tokens) in increments of 1,000 miles at a cost of $10.00 per 1,000 miles, plus a $25.00 processing fee per transaction.[26] At their estimated value in US dollars, about 1.5 cents each on average, a 30,000-mile transfer equates to a $450 transaction with a $325 fee — far outweighing the transferor's own use and the transferee's

[26]Documented in the official terms and conditions. Please see: `https://www.alaskaair.com/content/mileage-plan/use-miles/share-gift-miles`.

purchase of points from the airline.[27] Columns (2)-(3) show that such costly transfer plan is a common practice across corporations. As a result, apart from the gray market, token transfers from formal channels are rarely utilized, with only few exceptions, e.g., the transferor decides to withdraw.

Inflexible transaction. It is also common to lower down transfer flexibility, e.g., minimum and maximum amounts and stepped transaction volumes, as labeled by "Y" in column (4). This makes it almost impossible for consumers to achieve optimal holding positions through transaction as in an ideal secondary market. An annual cap on total transaction times or volumes is also common, limiting the market size. Some projects also require additional paperwork (e.g. submitting applications and awaiting approval), labeled as "Y" in column (5). It indicates as if transfer has been opened up for special cases, rather than aimed at providing tradability.

Sharing within a family pool. Interestingly, family pools are often the highlight of loyalty programs — accounts within a pool can share token accumulation. Given that both appear to be conveying redemption rights, why are family sharing welcomed yet transaction prohibited? The explanation is twofold. The key difference is that people do not earn premia from their families. Within a pool, the reselling option is not valued, since no tokens are actually traded but only accumulated and redeemed faster. Therefore, it is equivalent to issue non-tradable tokens to a collective consumer with a larger demand probability. Recall Section 2.7, we show that companies always benefit from increasing p, especially for companies specialized in international routes where p is initially small. We can see from column (6) that international airlines are more keen on family pools and usually set a larger headcount cap. On the other hand, even if there were trading premiums within the family, firms were unable to prohibit them because acquaintances could trade "over-the-counter" based on social connections, just as firms could do little about the gray market. As a result, it is more beneficial to incentivize a collective consumer by adopting the family pool.

4.2 Maintaining High Token Dependence

Our model shows that larger market power helps maintain higher token dependence, which effectively owes to less consumer base loss from token issuance. In this context, we use the consumer counts to proxy market power, and examines the re-

[27] The estimated value refers to *The Point Calculator*. Please see: `https://www.thepointcalculator.com/us/airline/alaska/alaska-miles-value-calculator`. The calculation is based on comparing main cabin fares using both cash and miles, across several dates and destinations.

lationship with token dependence.[28] Section 4.2.1 summarizes our data from 10-K financial reports. Section 4.2.2 examines a sanity check to corroborate our model settings. Section 4.2.3 and 4.2.4 document the main regression results in airline industry and hospitality industry respectively.

4.2.1 Data Summaries and Stylized Facts

We collect data from 10-K financial reports of airlines and hotels.[29] Despite slight variations in indicators disclosed by companies, airlines and hotels typically generate comparable measurements within the industry, respectively. For airlines, we obtain the direct measure of token dependence defined in (8), $\Theta = \frac{\text{Redemption Revenue}}{\text{Operation Revenue}}$.Note that Φ is a share concept relative to total revenue, which reflect how many "rebates" are attached to each unit of sales, and by definition are independent of business scale as the outcome of firm decisions.

As Table 2 summarizes, across our (unbalanced) Company$\times$Year panel sample set, the token dependence of airlines ranges from 0.65% to 13.81% with standard deviation of 3.16%. The size of the consumer base, measured by the annual number of passengers (million), *#passengers*, has mean 77.79 with standard deviation 58.77. They both suggest notable variations across the corporations and periods, similarly for the relative loyalty liability and annual guests of hotels. Additional variables and data of the hotel panel are also summarized, to be introduced later.

Stylized business status. Section 1 has listed remarkable facts on the loyalty business model over the world. Here we focus on the sample set to draw further stylized facts. Aviation is the industry that originated and led the loyalty program since 1981. To show its recent development trend, we form a core balanced panel from 2017 to 2021, including ten major airlines that report loyalty revenues continually.[30] There are totally \$12.46 billion redeemed value, occupying 5.17% of annual operation revenues, and leaving outstanding \$28.59 billion unredeemed in 2017. In 2018 (2019), the indicators reach \$13.66 (\$14.33) billion, 5.32% (5.47%), and \$29.52 (\$30.52), respectively. In 2021 after cut down by the pandemic, despite their aggregate redemption revenue is reduced to \$9.65 billion, it was less affected than the overall business and

[28] Market power might come from other strengths: product segmentation, technological barriers, etc. We aim to choose a dimension that closely relates to our model, whereas the comprehensive measurement of market power is beyond our focus.

[29] Companies are required to disclose the operation status of loyalty programs in ITEM 7, i.e., management's discussion and analysis of financial condition and results of operations. The details for data collection and preprocessing are provided in supplementary materials.

[30] The balanced panel drops the samples with missing data on redemption or operation returns in any of one year from 2017 to 2021, and finally includes Air France, Alaska Airlines, American Airline, Delta Air Lines, Emirates Airline, Hawaiian Airlines, International Airlines Group (IAG), JetBlue, Southwest Airlines, and United Airlines.

Table 2. Data Summary

Variables	N	Mean	S.D.	Min	25%	50%	75%	Max
Panel A: Airline								
RR, Redemption Revenue ($\times 10^9$\$)	69	1.08	1.02	0.05	0.29	0.61	2.18	3.37
OR, Operation Revenue ($\times 10^9$\$)	69	21.45	13.83	0.84	8.43	21.96	31.51	47.01
Token Dependence, RR/OR (%)	69	5.05	3.16	0.65	2.05	5.22	6.50	13.81
# Passengers ($\times 10^6$)	64	77.79	58.77	3.36	33.59	55.08	134.26	215.00
Price per Mile (\$)	69	0.23	0.10	0.15	0.16	0.19	0.27	0.77
Average Flight Distance ($\times 10^3 Miles$)	64	1.35	0.41	0.74	1.08	1.25	1.49	2.68
Panel B: Hotel								
LL, Loyalty Liability ($\times 10^9$\$)	77	1.02	1.53	0.01	0.07	0.23	1.43	6.47
HR, Hotel Revenue ($\times 10^9$\$)	75	4.82	5.17	0.20	1.25	2.43	6.39	20.97
Relative Loyalty Liability, LL/HR (%)	75	15.12	13.20	0.98	5.55	12.02	17.94	59.32
# Guests ($\times 10^6$)	57	0.43	0.24	0.01	0.31	0.45	0.56	1.01
# Properties ($\times 10^3$)	56	4.92	2.53	0.08	3.79	5.29	6.80	9.28
Price per Night (\$)	74	106.67	73.26	22.84	47.38	98.17	131.75	320.38

Notes. This table summarizes the variables of our interest from corporation 10-K reports. The mean, standard deviation, and quintiles are calculated unconditionally from the Corporation×Year panel.

thus occupied 6.13% of annual operation revenues, and kept remarkable outstanding unredeemed values of $36.38 billion — loyalty token in-advance provides a lifeline to some extent in times of external shocks.

Loyalty programs are also widely recognized and adopted by hospitality, especially for massive hotel groups (e.g. Marriott, Hilton and Hyatt) — they are also the earliest ones who start loyalty programs before 2016. In addition, since the pandemic caused a "pause", all hotels experienced a higher backlog of unredeemed loyalty liabilities during 2020-21 (e.g. Marriott in 2020 obtained a relative loyalty liability of 60%). Considering that the loyalty liability actually implies liquidity front loaded to corporations, it somewhat mitigated the impact of the huge turnover decline under external shocks.

4.2.2 Sanity Check of Token Dependence

We start with a sanity check for a deeper understanding of the objective determinants of token dependence — our model features that faster redemption (smaller M) aligns with larger token dependence Θ. Empirical data confirms this implication as the positive correlation shown in Figure 5, where redemption speed is obtained from dividing the annual redemption revenue by the outstanding stock (unredeemed liabilities) by the end of last year.[31] This appears as a general fact that applies to not only comparisons among corporations also pooling the years.

[31] Note that the redemption speed can exceed 1 (100%), since the newly issued tokens can be used within the year, which are not counted in the previous outstanding stock.

In particular, during the pandemic (2020-21), the global airline industry suffered a huge blow. The sudden drop in business led to a natural slowdown in redemption speed. However, the positive correlation holds with an even sharper slope: the token dependence did not decrease dramatically over the period, as the pandemic simultaneously shut down consumption and redemption. This gives us a clearer understanding of token dependence: it is not a size-like concept, and an exogenous consumption shock should not change it unless the shock affects the consumers' choices between using fiat money and tokens.

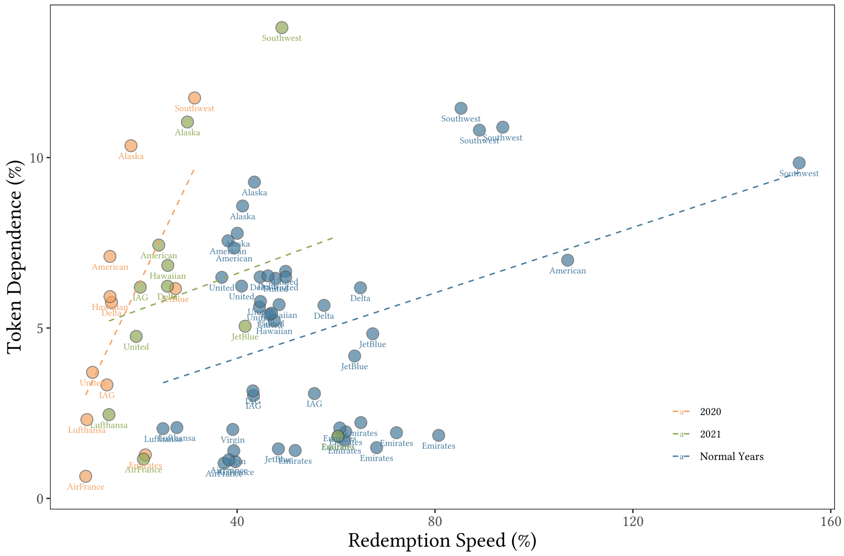

Figure 5. Redemption Speed and Token Dependence

This figure visualizes the sanity check based on the Corporation×Year panel, where the y-axis is the token dependence, and the x-axis presents redemption speed, i.e. the annual redemption revenue divided by outstanding stock in previous yea (%). We color the pandemic periods (2020-21) in different colors and draw the linear fits separately.

4.2.3 Token Dependence of Airlines

We formally test that a strong market power proxied by consumer counts maintains a high token dependence. We estimate OLS regressions with the general specification

$$\Theta_{i,t} = \alpha_0 + \alpha \#Passengers_{i,t} + X'_{i,t}\gamma + b_i + d_t + \epsilon_{i,t}, \tag{11}$$

where the coefficient of interest is α, $X'_{i,t}$ includes potential control variables, b_i and d_t are airline and year fixed effects, respectively.

Table 3 column (1) shows that there is a strong positive unconditional correlation between business scale and token dependence, where the coefficient is 0.016 ($s.e. = 0.007$) — a one-standard-deviation more passengers (58.77 million) is associated with 0.94-percentage-point higher token dependence, explaining about 30% standard deviation of token dependence.[32] Columns (2) and (3) confirm the results with year and airline fixed effects. Notably, with airline fixed effects, R-squared reaches 0.956, suggesting a large share of explanatory power from unobserved time-invariant determinants, including many other dimensions of market power. Column (4) controls other potential correlates of token dependence, including the pandemic-period dummy, average ticket price per mile, and the dummy of whether the airline belongs to an alliance,[33] reflecting heterogeneous implementations of loyalty programs across market segments (cheap or luxury) and alliances. The multivariate regression has a greater explanatory power than column (1), whereas the coefficient of *#Passengers* equals 0.023 ($s.e. = 0.009$), still economically and statistically significant.

There are still concerns in interpreting the above relationship as causal: omitted treatments may be correlated with business scale and also affect the calculation of token dependence. Imagine, an airline experiences a positive productivity shock and gains a larger market share. This positive shock for airline i in year t attracts more passengers (a higher $Passengers_{i,t}$). Meanwhile, the positive shock leads to higher operation revenue, thus mechanically reduces the token dependence $\Theta_{i,t}$ as redemption mainly comes from the loyalty tokens awarded in previous years. As mentioned in Section 2.7, an expansion in market share makes revenue from token redemption relatively less important relative to fiat-money sales, reducing token dependence. On the other hand, a negative productivity shock does not increase the token dependence correspondingly, as the still-existing users' continue their token holdings. Therefore, the coefficient α could be underestimated under the OLS specification (11). It also explains why the coefficient in column (3) is lower than (2) after including airline fixed effects.

We attempt to address these concerns by introducing the average flight distance as an instrument.[34] It reflects the airline's (proactive or reactive) decisions in route operations, i.e. whether short-haul routes with a larger passenger base or longer

[32]We report the standard errors two-way clustered at year and airline×period (normal or pandemic) levels. In particular, the latter clusters the samples from an airline into normal-period and pandemic clusters. The aggregate shock of the pandemic is not the treatment of interest, but brings about observations significantly different from normal periods, of which the difference vastly outweighs other sample variations. If they were treated as correlated as in the conventional cluster formula, it would result in an unnecessarily conservative confidence interval. A more targeted approach is to calculate the causal cluster variance as suggested by Abadie et al. (2023).

[33]Airlines disclose in 10-K reports the revenue passenger mile (RPM), multiplying the number of paying passengers by the distance traveled. We divide the operation revenue (approximately the gross revenue paid by passengers) by RPM to obtain the average ticket price per mile. The alliance dummy equals one if the airline belongs to one of the three major alliances, i.e. Star Alliance, Oneworld and Skyteam.

[34]We use the revenue passenger mile (RPM) divide by aggregate passengers to obtain the average flight distance of a passenger. In precise, denote the number of passengers of flight j as n_j, and the flight distance is m_j. Then the average flight distance reads $RPM\ /\ \#Passengers = (\sum_j n_j m_j)/(\sum_j n_j) \equiv \bar{m}$.

routes.[35] Therefore, the distance relates to the number of passengers. Meanwhile, the distance captures more variations from cross-section, which is less affected by the above concern. Also, the passenger changes explained by the airline's distance decision changes are largely separated from the unexpected passengers from a productivity shock. Therefore, the average flight distance is a potential plausible instrument.

Table 3. Business Scale and Token Dependence of Airlines

Dependent:	Token Dependence, $\Phi^R = \frac{\text{Redemption Revenue}}{\text{Operation Revenue}}$ (%)					
	OLS				2SLS	
Stage:					First	Second
	(1)	(2)	(3)	(4)	(5)	(6)
#Passengers	0.016**	0.023**	0.019*	0.023**		0.037**
	(0.007)	(0.009)	(0.009)	(0.009)		(0.016)
Price per Mile				-20.753***	-29.924	-19.877***
				(5.697)	(50.740)	(5.251)
Alliance Member				-0.795	39.905**	-1.449
				(1.121)	(17.497)	(1.091)
Pandemic				3.597***		
				(0.743)		
Average Flight Distance					-61.636***	
					(11.834)	
Year FE		Yes	Yes		Yes	Yes
Airline FE			Yes			
Observations	64	64	64	64	64	64
R^2	0.081	0.163	0.956	0.532	0.536	
Weak Inst. F-stat					27.1	
Wu-Hausman p-value						0.225

Notes. The dependent and main independent are token dependence (%) and the number of passengers (million), *#Passengers* respectively. Columns (1)-(4) report OLS estimates with year and airline fixed effects and controls including price per mile, alliance member dummy and pandemic dummy. Columns (5)-(6) show the 2SLS specification using average flight distance (10^3 miles) as the instrument. F stats of Kleibergen-Paap test for weak instruments, and p values for Wu-Hausman endogeneity test are reported. Standard-errors two-way clustered at year and airline × periods (normal or pandemic) are in parentheses. ***,**,* indicate statistical significance at the 1%, 5% and 10% respectively.

Table 3 columns (5)-(6) report the 2SLS estimates with instrument variable, the average flight distance. The first-stage estimation displays a significant coefficient with a Kleibergen-Paap F-stat of 27.1, suggesting the distance to be a strong instrument of the number of passengers. As column (6) shows, the coefficient of *#Passengers* (0.037) is much higher than the OLS estimates (around 0.02). This confirms our conjecture, i.e., the omitted determinant problem causes an underestimation in OLS approaches. A one-standard-deviation more passengers is associated with 2.17-percentage-point higher token dependence, roughly explaining 69% standard deviation of the token dependence. Overall, Table 3 indicates that airlines with larger consumer bases are able to maintain higher token dependence.

[35] Longer routes may have other advantages. More importantly, the airlines may not be able to choose freely due to geographical limitations. We do not study the optimal route choice problem, but rather focus on whether the larger number of consumers resulting from a particular route choice has an effect on token dependence.

4.2.4 Token Dependence of Hotels

Similar empirical approach is applied in the hospitality industry. We use the annual number of guests to proxy the business scale of hotels (one counted guest per occupied room), as summarized in Table 2, and use the number of properties as an instrument. The logic is similar (even simpler): it reflects a hotel's capacity and goal to host a larger consumer base. We also control for price per room night and its heterogeneous effect during pandemic. The main difference is the dependent variable. Hotels do not report redemption revenues, possibly due to different disclosure requirements. Alternatively, we use an indirect measure, relative loyalty liability, $RLL = \frac{\text{Loyalty Liability}}{\text{Operation Revenue}}$. While token dependence accounts for realized revenue, *RLL* accounts for the present value of unredeemed loyalty tokens. *RLL* can be used similarly to compare the importance of the loyalty program to the airlines in the cross section.

Table 4 columns (1)-(4) report the OLS and 2SLS estimates. In particular, column (4) suggests that a one-standard-deviation more guests (0.24 million) is associated with 7.82-percentage-point higher relative loyalty liabilities, roughly explaining 59% of its standard deviation. These results are based on pooling regressions (or with year fixed effects). The explanatory power is mainly from the cross-sectional variations of *#Guests*, which captures the comparison of market power. However, a hotel's change in *#Guests* also technically affects *RLL*: a guest's redemption would reduce loyalty liabilities but add to token dependence. Column (5) confirms this conjecture by regressing *RLL* on the change rate of *#Guests* and documenting a negative relationship. With hotel fixed effects in column (6), we also obtain a negative coefficient of *#Guests*, as the variation effectively used in estimation is relatively more from temporal changes within hotels. Despite the unavoidable use of information on time-series changes, the instrument *#Properties* captures the cross-sectional variation more, as its within-group changes are significantly smaller than across-group variations. Then, in columns (7)-(8), we still obtain a positive coefficient of *#Guests* using 2SLS estimation. Therefore, despite challenges from the indirect measurement *RLL*, we still find robust evidence that a larger consumer base helps maintain a higher token dependence.

The relative loyalty liability (*RLL*) provides additional corroboration to our model. It can be seen as a transient state, abstracting away the redemption speed. Then if the redemption or consumption flow is shocked, the *RLL* fails to one-to-one correspond to token dependence. For example, as documented in Section 4.2.1, during the pandemic periods, the airlines suffered from stoppage of routes. Although the high token dependence held up (even increased), it was a combination of simultaneous crashes in total revenue and redemption — there would have been more loyalty liabilities converted to revenue at the regular redemption speed. As a result, we saw a remarkable increase in outstanding loyalty stocks. Similarly, for hotels, token dependence is overestimated

Table 4. Relative Loyalty Program Scale of Hotels

Dependent:	Relative Loyalty Liability, $\frac{\text{Loyalty Liability}}{\text{Hotel Revenue}}$ (%)							
	Impact of Pooling Variations of #Guests				Impact of Longitudinal Change in #Guests			
	OLS		2SLS		OLS		2SLS	
Stage:			First	Second			First	Second
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
#Guests	19.164***	28.152***		32.586***		-21.086		40.623**
	(5.327)	(5.266)		(6.968)		(15.116)		(17.614)
Revenue per Room Night		0.115***	0.003**	0.116**	0.213**	-0.141	0.000	-0.194
		(0.032)	(0.001)	(0.038)	(0.077)	(0.085)	(0.001)	(0.139)
Revenue per Room Night × Pandemic		0.889**	0.004**	0.882**	0.422***	0.327**	-0.001	0.375*
		(0.303)	(0.001)	(0.305)	(0.070)	(0.136)	(0.001)	(0.184)
#Properties			0.083***				0.076***	
			(0.016)				(0.019)	
Guest Growth					-15.925***			
					(3.825)			
Year FE		Yes	Yes	Yes		Yes	Yes	Yes
Hotel FE						Yes	Yes	Yes
Observations	52	52	52	52	50	52	52	52
R^2	0.098	0.696	0.764		0.426	0.937	0.933	
Weak Inst. F-stat				26.8				16.3
Wu-Hausman, p-value				0.264				0.0004

Notes. The dependent and main independent are relative loyalty liability (RLL) (%) and the number of guests (million), #*Guests* respectively. Columns (1) reports OLS estimation of the univariate model. Column (2) controls revenue per room night and its interaction with pandemic dummy, and year fixed effects. Columns (3)-(4) show the 2SLS estimation of the same model as column (2), using the number of properties as an instrument. Column (5) shows the unconditional relationship between *RLL* and the change rate of #*Guests*. Columns (6)-(8) report the OLS and 2SLS estimation of the same model as column (2) with additional hotel fixed effects. F stats of Kleibergen-Paap test for weak instruments, p values for Wu-Hausman endogeneity test are reported. Standard-errors clustered by year and hotel×periods are in parentheses. ***,**,* indicate statistical significance at the 1%, 5% and 10% respectively.

by *RLL* in pandemic years.

Overall, we find evidence in airline and hospitality industries that greater market powers enable corporations to maintain higher token dependence, thus better utilizing loyalty token issuance for liquidity creation. While we make efforts to separate the business scale from other potential determinants, some unobserved factors are also in the scope of market power in our theory. For example, issuers can utilize technology barriers and monopoly a submarket, and thus maintain a large market power. Therefore, representing market power by business scale still economically underestimates its determination of token dependence. Our estimate can be regarded as a lower bound, whereas a more comprehensive analysis should rely on a general metric of market power. In addition, our empirical tests are based on disclosures in 10-K reports, where token revenues are not accounted until the redemption is realized. However, our model emphasizes tokens' crucial roles in liquidity creation before redeemed, which is not fully presented in current related accounting standards.

4.3 Alliance for Higher Redemption Probability

To focus on the central mechanism of loyalty tokens as a finance tool, our model regards demand probability p as an exogenous feature of the located market. In practice, token functionality designs can subtly change consumers' demand probability, which can also be seen as an extension of market power through the loyalty business model.[36] For example, as mentioned in Section 4.1, a family pool effectively creates a "collective consumer," whose demand probability roughly equals the probability that at least one member has a demand. Here, we extend to more designs that could change demand. In particular, in our model, we implicitly assume that people have the same demand probabilities p in fiat money consumption and redemption. Token designs can essentially yield a new demand probability for redemption p^R, in contrast to the demand for (fiat money) consumption p, either by changing convenience, validity period, and redeemable products.

Convenient use. If the tokens were inconvenient to use (e.g. hard to understand, complicated to redeem), people would have a redemption probability p^R lower than p, therefore lower down the valuation of tokens. Then a greater market power would be required to fill this discount loss in implementing the loyalty program. In practice, we see large issuers always applying technologies to their loyalty programs as quick as possible to increase ease of use. For example, as Table 5 shows, airlines and hotel groups and tend to launch mobile apps and enable systems of token management and usage earlier since the iOS App Store launched in July 2008, which becomes apparently standard up to 2020s.

Validity periods. It seems profitable to set a token expiration term, which generates waived tokens and reduces costs for redemption delivery. However, consumers would discount p^R then lower the token valuation. Therefore, this setup is only profitable when most tokens are eventually unredeemed such that the "free lunch" outweigh the drawback of discounted token valuation. To this reason, we see international-focused airlines appear to favor expiration terms (e.g., Avianca, Emirates Airlines, ANA, Singapore Airlines, and Turkish Airlines). Otherwise, airlines commit to never-expiry, labeled by "N" in Table 5. In addition, with possible consumer exits, issuers may adopt a compromise solution, i.e., expire tokens conditional on the holder is inactive for x years, labeled as "xC" in Table 5. This picks up the "free lunch" from inactive consumers without discounting frequent consumers' valuation of tokens.[37]

[36] From the perspective of artificially adding up switching cost (Banerjee, 1987), loyalty program facilitates repeat consumption demands.

[37] This implies that consumer heterogeneity matters in expiration designs. See Sun and Zhang (2019) who develop a model to rationale the trade-off between short and long expiration terms.

Table 5. Redemption Demand Probabilities Narrative

Corporation	Expiry	App Release	Alliance / # Codeshare Partners	Corporation	Expiry	App Release	Alliance / # Codeshare Partners
Airline				Lufthansa	3C	2008-12-08	Star Alliance
Alaska	2C	2010-02-20	Oneworld	Qatar	3C	2012-12-02	Oneworld
American	2C	2010-07-26	Oneworld	Singapore	3	2012-04-06	Star Alliance
Avianca	2	2013-01-18	Star Alliance	Southwest	N	2009-12-18	0
British	3C	2008-07-11	Oneworld	Spirit	1C	2018-11-15	3
Delta	N	2010-09-01	Skyteam	Turkish	3	2017-09-18	Star Alliance
Emirates	3	2014-09-11	15	United	N	2011-10-20	Star Alliance
AirFrance	2C	2010-09-17	Skyteam	VriginAmerica	1.5C	2016-09-23	Oneworld
Hawaiian	N	2010-11-09	5	**Hotel**			
VirginAtlantic	N	2010-01-31	Skyteam	BestWestern	N	2009-12-17	
WestJet	N	2014-05-05	20	Choice	1.5C	2012-04-06	
AirCanada	1.5C	2009-08-20	Star Alliance	Hilton	2C	2013-04-19	
Cathay	1.5C	2009-02-23	Oneworld	Hyatt	2C	2011-10-31	
ANA	3	2010-07-26	Star Alliance	IHG	1C	2010-04-22	
Etihad	1.5C	2016-04-11	24	Marriott	2C	2011-08-05	
Frontier	1C	2015-10-15	3	Radisson	2C	2019-08-08	
JetBlue	N	2012-02-03	7	Starwood	1C	2009-06-05	
KoreanAir	10	2010-07-28	Skyteam	Wyndham	1.5C	2013-12-31	

Notes. The table aggregates key narrative features of loyalty token designs pertinent to the probabilities of token redemption. *Expiry* documents the token's validity period, with "N" indicating never expiration, a numeric value represents years of validity, with "C" indicating expiration date only applies conditional on inactive accounts. *App Release* records the initial launch date of IOS applications for booking and (potentially) redemption affairs. *Alliance / # Codeshare Partners* is only applicable for airlines, which identifies the alliance membership or alternatively lists the number of code-sharing partners. Token accumulation are typically sharable within the alliance or partners. *Expiry* and *Alliance* are collected from official terms and conditions in March 2024.

Alliance. A common way to adjust p^R is to change the redeemable products. For example, many issuers offer "point shops," where tokens can be used to redeem partners' products. In this way, issuers sold tokens with broader uses without expanding their main business (which would be remarkably costly). Here we focus on a prime example in airlines: forming alliances or entering into code-share cooperation, as shown in Table 5. After flying on one airline, consumers can credit tokens in their accounts of partner carriers, or directly interoperate within the alliance. Therefore, consumers have a higher demand probability of redemption p^R than demanding the original airline's flights p, and evaluate higher token prices.[38]

Figure 6 shows the airline networks of three major alliances, i.e., Star Alliance, Oneworld, and Skyteam. We define the "dominant member" of an airport as the alliance member who runs the most routes at that airport within the alliance. Then for each alliance, the map points out all the accessible airports, i.e., covered by the alliance's routes, in different colors and shapes according to their dominant members. Furthermore, we calculate the Herfindahl-Hirschman Index (HHI) of routes in operation at each airport, and summary the distribution of the accessible airports' HHI, plotted in the histogram.

[38] From another perspective, the original airline earns a premium of attracting traffic to its partner.

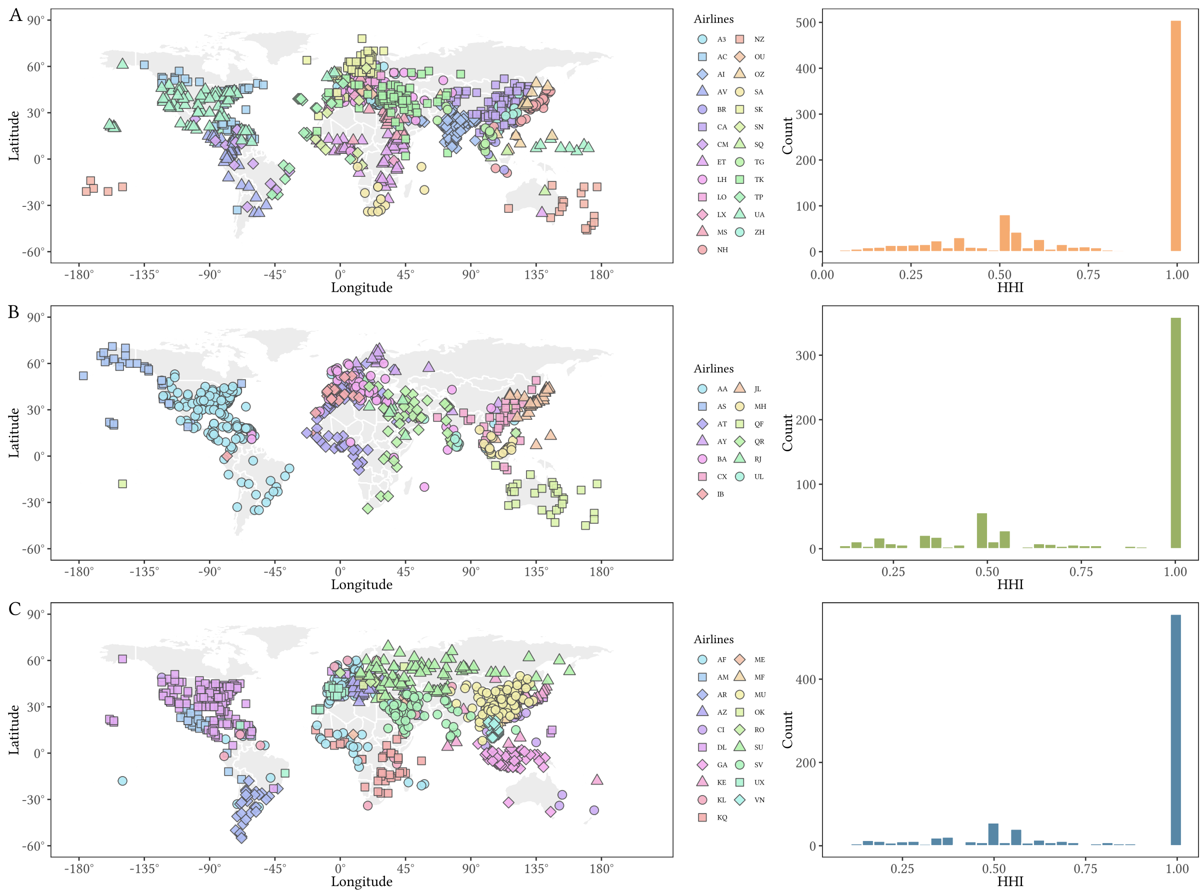

Figure 6. Global Airport Networks of the Alliances

Panels A, B, and C visualize the global airport distribution of Star Alliance, Oneworld and Skyteam, respectively. In each panel, the map depicts the airports covered by the airline alliance's routes. Each airport is shown in different colors and shapes according to the "dominant member" within the alliance who operates the most routes at that airport. The within-alliance Herfindahl-Hirschman Index (HHI) of operating routes at each airport is also calculated and aggregated in the histogram. Data source: OpenFlights Airports Database.

There are two findings. First, each of the three alliances has formed a global airline network to take advantage of global demand for flight out to form their redemption probabilities. Second, the alliance usually does not include two members from the same region to avoid competition within the alliance without effectively increasing p^R. Essentially, the worldwide airport network consists of air hubs and branch airports. As histograms show, a large share of airports are connected by one single member in the alliance, making HHI equal to one. From the airline's perspective, it is better not to cooperate if it does not effectively increase the redemption demand. As documented in Table 5, Southwest Airlines constitutes a typical example: as an airline specializing in short-haul travels, it faces a unique user group and has a monopoly. As a result, building a global network is not a top priority. This logic also applies to those targeting specific user bases, e.g., low-cost and regional airlines. In general, both the alliance and the airline favor partners who are able to increase redemption functionality without adding competition.

5 Conclusion

Our paper studies loyalty token issuance as a financing mechanism and highlights the economic trade-off when token issuance is bundled with consumption. We show that the issuer has an incentive to limit the token secondary market as the tradability eliminates the benefit of time-series bundling, and this loss strictly dominates the welfare gain by selling tokens at a higher price. Loyalty tokens can be viewed as a privileged usage of a financial toolkit to leverage the issuer's market power: more market dominance enables higher token dependence; if all issuers are under perfect competition, loyalty token issuance becomes impossible. These predictions are well supported by analyzing airline mileage and hotel point data.

Our results make twofold implications in regulation: First, regulators need to ensure consumers are aware of token transfer and redemption policy as issuers tend to overstate the "easy-to-use" nature of tokens and induce consumers to pay more than they should pay. Second, if issuers try to issue tradable tokens, they need commitment devices to convince consumers that tokens will be allowed to trade — either token issuance on a public blockchain where consumers can execute permissionless transfers or large-enough punishment enforced by the regulator if the token issuer deviates from the pre-announced token issuance scheme. Outside use cases or payment convenience are necessary for token issuers to be self-disciplined for tradability, meanwhile, outside uses of tokens trigger new challenges of KYC/AML obligations for issuers.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. 2023. "When should you adjust standard errors for clustering?" *The Quarterly Journal of Economics* 138, 1–35.

Banerjee, Abhijit. 1987. *On frequent flyer programs and other loyalty-inducing economic arrangements*. Harvard University.

Basso, Leonardo J, Matthew T Clements, and Thomas W Ross. 2009. "Moral hazard and customer loyalty programs." *American Economic Journal: Microeconomics* 1, 101–123.

Bolton, Ruth N, P KandBramlettMatthewD Kannan, and Matthew D Bramlett. 2000. "Implications of loyalty program membership and service experiences for customer retention and value." *Journal of the academy of marketing science* 28, 95–108.

Chen, Yanyan, Timo Mandler, and Lars Meyer-Waarden. 2021. "Three decades of research on loyalty programs: A literature review and future research agenda." *Journal of Business Research* 124, 179–197.

Chod, Jiri, and Evgeny Lyandres. 2023. "Product market competition with crypto tokens and smart contracts." *Journal of Financial Economics* 149, 73–91.

Chun, So Yeon, and Rebecca W Hamilton. 2024. "Paying with Money or Paying with Points: How Variable Versus Fixed Exchange Rates Influence Loyalty Point Redemption." *Journal of Marketing Research* 00222437241255650.

Chun, So Yeon, and Anton Ovchinnikov. 2019. "Strategic consumers, revenue management, and the design of loyalty programs." *Management Science* 65, 3969–3987.

Cong, Lin William, Ye Li, and Neng Wang. 2022. "Token-based platform finance." *Journal of Financial Economics* 144, 972–991.

Cong, Lin William, and Yizhou Xiao. 2021. "Categories and functions of crypto-tokens." *The Palgrave Handbook of FinTech and Blockchain*, 267–284.

Evanschitzky, Heiner, Balasubramanian Ramaseshan, David M Woisetschläger, 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, 625–638.

Klemperer, Paul. 1987. "Markets with consumer switching costs." *The quarterly journal of economics* 102, 375–394.

Klemperer, Paul. 1995. "Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade." *The review of economic studies* 62, 515–539.

Kumar, Viswanathan, and Werner Reinartz. 2016. "Creating enduring customer value." *Journal of marketing* 80, 36–68.

Lewis, Michael. 2004. "The influence of loyalty programs and short-term promotions on customer retention." *Journal of marketing research* 41, 281–292.

Lim, Freddy, So Yeon Chun, and Ville Satopää. 2024. "Loyalty Currency and Mental Accounting: Do Consumers Treat Points Like Money?" *Manufacturing & Service Operations Management* 26, 1878–1896.

Liu, Yuping. 2007. "The long-term impact of loyalty programs on consumer purchase behavior and loyalty." *Journal of marketing* 71, 19–35.

Malinova, Katya, and Andreas Park. 2023. "Tokenomics: When tokens beat equity." *Management Science* 69, 6568–6583.

Nunes, Joseph C, and Xavier Drèze. 2006. "Your loyalty program is betraying you.." *Harvard Business Review* 84, 124–31.

Rogoff, Kenneth, and Yang You. 2023. "Redeemable platform currencies." *The Review of Economic Studies* 90, 975–1008.

Rossi, Federico. 2018. "Lower price or higher reward? Measuring the effect of consumers' preferences on reward programs." *Management Science* 64, 4451–4470.

Selnes, Fred. 1993. "An examination of the effect of product performance on brand reputation, satisfaction and loyalty." *European Journal of marketing* 27, 19–35.

Sharp, Byron, and Anne Sharp. 1997. "Loyalty programs and their impact on repeat-purchase loyalty patterns." *International journal of Research in Marketing* 14, 473–486.

Sönmeztürk, Osman, Tolga Ayav, and Yusuf M Erten. 2020. "Loyalty program using blockchain." In *2020 IEEE international conference on blockchain (Blockchain)*,, 509–516, IEEE.

Sun, Yacheng, and Dan Zhang. 2019. "A model of customer reward programs with finite expiration terms." *Management Science* 65, 3889–3903.

Appendix

A Omitted Proofs

A.1 Proof of Lemma 1

(i) First, it is obvious that at any time, no one would hold more than M tokens. Otherwise, the user is not able to fully redeem them in one period. Then she could at least buy these excess tokens in the next period at a lower present value.

Consider that the user holds m tokens, $0 \leq m < M$. When $M = 1$, the claim naturally holds, that is, any user holds either zero or $M = 1$ tokens.

When $M \geq 2$, we prove that for any period t_1 and any consumer with m_1 tokens before trading in t_1, and plans to hold m_2 tokens after trading in $(t_1 + 1)$, $0 \leq m_1, m_2 \leq M$, she prefers to hold zero than m tokens after trading in t_1, $0 \leq m < M$. Denote the secondary market price as $\hat{P}$. For any possible (m_1, m_2), the cost of holding m tokens after trading in t_1 reads

$$\underbrace{(m - m_1)\hat{P}}_{\text{trade in } t_1} + \underbrace{p\left[P_M^T + \beta(m_2 - m - 1)\hat{P}\right]}_{\text{trade in } t_2 \text{ if shocked in } t_1} + \underbrace{(1-p)\beta(m_2 - m)\hat{P}}_{\text{trade in } t_2 \text{ otherwise}}, \tag{A.1}$$

where the coefficient of m is $(1 - \beta)\hat{P} > 0$ and note that $m < M$ tokens are not enough for one redemption in t_1. Therefore, a larger m increases the cost of any possible holding strategies. That is, holding $m \in (0, M)$ tokens after trade is strictly dominated by holding zero tokens after trade.

(ii) Note that each user can only receive one token from purchase in one period, combined with the above implication, at the beginning of each period, there are three types of user who hold 0, 1, and M tokens, respectively, denoted u_0, u_1, and u_M. On the other hand, there are only two types, u_0 and u_M after trade. Then the market clearing condition should ensure that for u_1, becoming u_M and u_0 are indifferent, i.e. when purchasing and receiving a new token, it is indifferent to either sell it or buy extra $M - 1$ tokens:

$$\underbrace{P_M^T - \beta\hat{P}}_{\substack{\text{pay for product} \\ \text{and resell token}}} \quad + \quad \underbrace{a(P_M^T - \beta\hat{P})}_{\substack{\text{present value of} \\ \text{next consumption} \\ \text{excluding the value of token}}} \quad = \quad \underbrace{P_M^T + \beta(M-1)\hat{P}.}_{\substack{\text{pay for product and buy token} \\ \text{(and redeem for next consumption)}}} \tag{A.2}$$

This solves a unique price $\hat{P}$,

$$\hat{P} = \frac{p + (1-p)a}{M + a} P_M^T. \tag{A.3}$$

Essentially, $\hat{P}$ simultaneously ensure that for any type, holding M tokens and zero tokens are indifferent:

$$(M-m)\hat{P} = -m\hat{P} + (p+(1-p)a)(P_M^T - \beta\hat{P}), \;\; m \in \{0, 1, M\}. \tag{A.4}$$

Therefore, $\hat{P}$ is the equilibrium secondary market price of one token.

A.2 Proof of Proposition 1

(i) Calculate the derivative directly:

$$\begin{aligned} \frac{\mathrm{d}P_M^{NT}}{\mathrm{d}p} &= \frac{\mathrm{d}P_M^{NT}}{\mathrm{d}a}\frac{\mathrm{d}a}{\mathrm{d}p} = \frac{a^{M-1}(1-a)\left(M - a\frac{1-a^M}{1-a}\right)}{(1-a^M)^2}\frac{\beta(1-\beta)}{(1-\beta+\beta p)^2} \\ &= \frac{a^{M-1}(1-a)\beta(1-\beta)\sum_{i=1}^{M}(1-a^i)}{(1-a^M)^2(1-\beta+\beta p)^2} > 0. \end{aligned} \tag{A.5}$$

$$\frac{\mathrm{d}P_M^T}{\mathrm{d}p} = \frac{1}{M}\frac{\beta(1-\beta)}{(1-\beta+\beta p)^2} > 0. \tag{A.6}$$

(ii)

$$\frac{\mathrm{d}P_M^{NT}}{\mathrm{d}M} = \frac{(1-a)a^M\log(a)}{(1-a^M)^2} < 0; \;\; \frac{\mathrm{d}P_M^T}{\mathrm{d}M} = -\frac{a}{M^2} < 0. \tag{A.7}$$

(iii) $P_M^{NT} < P_M^T$, i.e.,

$$\frac{1-a^{M+1}}{1-a^M} \leq 1 + \frac{a}{M} \underset{Rearrange}{\Longleftrightarrow} \frac{a(1-a^M)}{1-a} \geq Ma^M. \tag{A.8}$$

Note that $a<1$, $\frac{a(1-a^M)}{1-a} = \sum_{i=1}^{M} a^i \geq \sum_{i=1}^{M} a^M = Ma^M$. The equal sign is obtained if and only if $M=1$ or $M \to \infty$.

A.3 Proof of Lemma 2

Treat Rev_M^{NT} as a continuous function of $M \in [1, \infty)$.

$$\frac{\mathrm{d}Rev_M^{NT}}{\mathrm{d}M} = \frac{p\left[(1-a)\log(a)a^M(1-a^{*M})(1-a^{*M+1}) - (1-a^*)\log(a^*)a^{*M}(1-a^M)(1-a^{M+1})\right]}{(1-\beta^*)(1-a^M)^2(1-a^{*M+1})^2}. \tag{A.9}$$

Note that $a, a^* \in (0,1)$, and $\beta^* < \beta \Leftrightarrow a^* < a$. Define a temporary function w.r.t. $a \in (0,1)$, $g(a) = \frac{a^M(1-a)\log(a)}{(1-a^M)(1-a^{M+1})}$. Then $\frac{\mathrm{d}Rev_M^{NT}}{\mathrm{d}M} < 0$ for any $M \geq 1$ is equivalent to $g(a) < g(a^*)$ for any $M \geq 1$, for which the sufficient condition is $g'(a) < 0, \forall a \in (0,1)$.

Further,

$$g'(a) = \frac{a^{M-1}(1-a)\log(a)}{(1-a^M)^3(1-a^{M+1})^2}\left[\frac{M(1-a^{2M+1})}{1-a^M} - \frac{a(1-a^M)}{1-a} + \frac{1-a^{M+1}}{\log(a)}\right] \\ \equiv \frac{a^{M-1}(1-a)\log(a)}{(1-a^M)^3(1-a^{M+1})^2}G(a). \tag{A.10}$$

We now prove $G(a) > 0$.

$$\begin{aligned} G(a) &= \frac{M(1-a^{2M+1})}{1-a^M} - \sum_{i=1}^{M} a^i + \frac{1-a^{M+1}}{\log(a)} \underset{a^i<a}{>} M\left(\frac{1-a^{2M+1}}{1-a^M} - a\right) + \frac{1-a^{M+1}}{\log(a)} \\ &= M\left(\frac{1-a}{1-a^M} + a^{M+1}\right) + \frac{1-a^{M+1}}{\log(a)} \underbrace{\geq}_{\text{F increases in M}} 1 + \frac{1-a^2+a^2\log(a)}{\log(a)} \\ &\underbrace{>}_{\left(\frac{1-a^2+a^2\log(a)}{\log(a)}\right)' = \frac{a\log(a)-1/a}{(\log(a))^2} < 0} 1 + \lim_{a\to 1^-} \frac{1-a^2+a^2\log(a)}{\log(a)} = 0. \end{aligned} \tag{A.11}$$

Therefore, $G(a) > 0$ and recall the formula of $g'(a)$ above, there is only one negative multiplied term, $\log(a)$. Therefore, $g'(a) < 0$ and therefore $\frac{\mathrm{d}Rev_M^{NT}}{\mathrm{d}M} < 0$.

Consider Rev_M^T. Since $a^* < a$,

$$\frac{\mathrm{d}Rev_M^T}{\mathrm{d}M} = \frac{p}{1-\beta^*}\frac{a^*-a}{(M+a^*)^2} < 0. \tag{A.12}$$

A.4 Proof of Proposition 2

Lemma 3. ***Tradable and non-tradable tokens at the same price.***

$\forall M \geq 1$, there exists a unique non-tradable issuance strategy $m \in [1, \infty)$, s.t. $P_m^{NT} = P_M^T$.

Proof.

$$1 + \frac{a}{M} = P_M^T = P_m^{NT} = \frac{1-a^{m+1}}{1-a^m} \quad \Rightarrow \quad M = \frac{a}{1-a}\left(\frac{1}{a^m} - 1\right). \tag{A.13}$$

Note that the R.H.S. increases in m and, when $m = 1$, $\frac{a}{1-a}\left(\frac{1}{a^m} - 1\right) = 1$, $\lim_{m\to\infty} \frac{a}{1-a}\left(\frac{1}{a^m} - 1\right) = \infty$. Therefore, $\forall M \geq 1$, there exists a unit $m \in [1, \infty)$ s.t. $P_m^{NT} = P_M^T$. □

We then show that $\forall M > 1$, $Rev_M^T < Rev_{m(M)}^{NT}$, where $m(M)$ is the unique corresponding m s.t. $P_m^{NT} = P_M^T = P$. Recall (4) and (7), it is equivalent to prove $Q_M^T < Q_m^{NT}$, i.e.,

$$\frac{p}{1-\beta^*}\frac{1-a^{*m}}{1-a^{*m+1}} > \frac{p}{1-\beta^*}\frac{M}{M+a^*} \quad \Leftrightarrow \quad M < \frac{a^*}{1-a^*}\left(\frac{1}{a^{*m}} - 1\right). \tag{A.14}$$

By Lemma 3, we only need to show that

$$\frac{a^*}{1-a^*}\left(\frac{1}{a^{*m}}-1\right) > \frac{a}{1-a}\left(\frac{1}{a^m}-1\right). \tag{A.15}$$

Treat it as a function of $a \in (0,1)$ and note that $m > 1$ by Lemma 3, we obtain

$$\frac{\mathrm{d}}{\mathrm{d}a}\left[\frac{a}{1-a}\left(\frac{1}{a^m}-1\right)\right] = -\frac{(1-a)m-(1-a^m)}{a^m(1-a)^2} < 0. \tag{A.16}$$

Then the above inequality holds since $a^* < a$. This implies that $Rev^T_{M^T_*} < Rev^{NT}_{m(M^T_*)}$, i.e., there is always a non-tradable token-issuance plan that generates more revenue than the optimal tradable plan. Therefore, the overall optimal choice is non-tradable. Note that there is no need to satisfy $m(M) = M^{NT}_*$. Also note that we do not require any restrictions on c – we consider all the possible prices, even lower than the monopoly price. Therefore, this proof applies to the general optimization problem discussed in footnote 15.

A.5 Proof of Proposition 3

Similarly, treat U^x_M as continuous functions of $M \in [1,\infty)$.

Consider U^{NT}_M. When $c = 0$, $M^{NT}_* = 1$ as Lemma 2 shows. Suppose there exists $0 \le c_1 \le c_2$, s.t., $M^{NT}_*(c_1) > M^{NT}_*(c_2)$. Denote them as m_1 and m_2, respectively. According to Proposition 1 and Lemma 2, $P^{NT}_{m_2} > P^{NT}_{m_1}$ and $Rev^{NT}_{m_2} > Rev^{NT}_{m_1}$. Then by the process of optimization,

$$\begin{cases} U^{NT}_{m_2}(c_2) > U^{NT}_{m1}(c_2) \Rightarrow c_2 < \dfrac{Rev^{NT}_{m_2} - Rev^{NT}_{m_1}}{f(P^{NT}_{m_2}) - f(P^{NT}_{m_1})}, \\ U^{NT}_{m_2}(c_1) < U^{NT}_{m1}(c_1) \Rightarrow c_1 > \dfrac{Rev^{NT}_{m_2} - Rev^{NT}_{m_1}}{f(P^{NT}_{m_2}) - f(P^{NT}_{m_1})}. \end{cases} \tag{A.17}$$

Then we obtain $c_1 > c_2$, leading to a contradiction. Therefore, $\forall 0 \le c < c'$, $M^{NT}_*(c) \le M^{NT}_*(c')$.

In the following analysis, we focus on the cases where the optimal plan gains at least zero payoff, which limits the parameter c to satisfy $c \le Q^x_M / f'(P^x_M)$.

Consider U^T_M. The above proof is also applicable for the tradable case. Here we alternatively prove by calculating the exact values of M_*^T:

$$\frac{\mathrm{d}U_M^T}{\mathrm{d}M} = \frac{acf'(P_M^T)}{M^2} - \frac{p(a-a^*)}{(1-\beta^*)(M+a^*)^2} > 0 \Leftrightarrow \left(1+\frac{a^*}{M}\right)^2 > \frac{p(a-a^*)}{(1-\beta^*)acf'(P_M^T)}$$
$$\Leftrightarrow \begin{cases} M < \dfrac{a^*}{\sqrt{\frac{p(a-a^*)}{(1-\beta^*)acf'(P_M^T)}}-1}, & c < \dfrac{p(a-a^*)}{(1-\beta^*)af'(P_M^T)}; \\ M > \dfrac{a^*}{1-\sqrt{\frac{p(a-a^*)}{(1-\beta^*)acf'(P_M^T)}}}, & c \geq \dfrac{p(a-a^*)}{(1-\beta^*)af'(P_M^T)} \quad (M=\infty \text{ when equal}). \end{cases} \tag{A.18}$$

Note that when $c > \frac{p(a-a^*)}{(1-\beta^*)af'(P_M^T)}$,

$$U_1^T = \frac{p}{1-\beta^*}\frac{1+a}{1+a^*} - acf'(P_M^T) < U_\infty^T = \frac{p}{1-\beta^*}, \tag{A.19}$$

then the optimal issuance plan always satisfies $M_*^T = \infty$ when $c > \frac{p(a-a^*)}{(1-\beta^*)af'(P_M^T)}$. When $c < \frac{p(a-a^*)}{(1-\beta^*)af'(P_M^T)}$, $M_*^T = \frac{a^*}{\sqrt{\frac{p(a-a^*)}{(1-\beta^*)acf'(P_M^T)}}-1}$ increases in c. Thus, M_*^T increases in c and tends to infinity when $\frac{p(a-a^*)}{(1-\beta^*)af'(P_M^T)} \to 1$. Combine the two cases, we obtain $\forall 0 \leq c < c'$, $M_*^T(c) \leq M_*^T(c')$.

A.6 Proof of Proposition 4

Consider a simple violation strategy, i.e., the issuer issues tradable token in period 0 and disables tradability from period 1. The advantage is that the price is P_M^T (higher than P_M^{NT}) in period 0 but without tradability in actual. The disadvantage comes from extra loss of consumer base. Since no trade happens, the subsequent token accumulation process is not affected. The excess revenue comparing to the corresponding non-tradable token issuance reads

$$\Delta Rev = p\Delta P - c\Delta P + \beta^* c\Delta f(P) = \left[p - (1-\beta^*)c\frac{\Delta f(P)}{\Delta P}\right]\Delta P, \tag{A.20}$$

where $\Delta P = P_M^T - P_M^{NT} > 0$. As mentioned, the implicit assumption ensures the issuer has a chance to earn non-negative returns, i.e., $c < Q_M^x / f'(P_M^x)$. Similar with the proof process above, $p - (1-\beta^*)c\frac{\Delta f(P)}{\Delta P} > 0$. Then $\Delta Rev > 0$. In fact, there are a series of applicable strategies that allows issuers to earn excess revenue from pretended tradability, as long as the appropriate number of tokens exist that have been bought at a price that includes tradability but have not yet been traded.

A.7 Proof of Proposition 5

(i) Compare (9)-(10) with the revenue formulas (non-tradable and tradable) discussed in Appendix A.3, we note that they have the same forms, respectively, but exchanging the positions of β and β^*, e.g., $Rev_M^T = \frac{p}{1-\beta^*}\frac{M+a}{M+a^*}$, and $C_M^T = \frac{p}{1-\beta}\frac{M+a^*}{M+a}$. Recall the proof procedure in the above section, one can imagine the duality steps to prove the corresponding (opposite) properties of derivatives. However, if we only need to know the optimum, i.e.,

$$C_1^T = C_1^{NT} = \min_{M\geq 1}\left\{C_M^{NT}, C_M^T\right\} = \frac{p}{1-\beta}\frac{1+a^*}{1+a}, \tag{A.21}$$

we can consider the following simpler proofs.

First, we prove that $\forall M > 1$, $C_M^{NT} > \frac{p}{1-\beta}\frac{1+a^*}{1+a}$.

$$C_M^{NT} > \frac{p}{1-\beta}\frac{1+a^*}{1+a} \Leftrightarrow \frac{1-a^M}{1-a^{M+1}}(1+a) > \frac{1-a^{*M}}{1-a^{*M+1}}(1+a^*). \tag{A.22}$$

Note that $a^* < a$, and

$$\begin{aligned}\frac{\partial}{\partial a}\left[\frac{1-a^M}{1-a^{M+1}}(1+a)\right] &= \frac{1-a^2}{a^2(1-a^{M+1})^2}\left[a^2\left(\frac{1-a^{2M}}{1-a^2}\right) - Ma^{M+1}\right] \\ &= \frac{1-a^2}{a^2(1-a^{M+1})^2}\underbrace{\left(\sum_{i=1}^{M} a^{2i} - Ma^{M+1}\right) > 0}_{a^{2i}+a^{2(M+1-i)} > 2a^{M+1}},\end{aligned} \tag{A.23}$$

the above inequality holds. Therefore, $C_M^{NT} > \frac{p}{1-\beta}\frac{1+a^*}{1+a}$.

On the other hand, it is easy to show $\frac{\mathrm{d}C_M^T}{\mathrm{d}M} > 0$. Thus, $C_M^T \geq C_1^T = \frac{p}{1-\beta}\frac{1+a^*}{1+a}$. Also, $C_1^{NT} = \frac{p}{1-\beta}\frac{1+a^*}{1+a}$. Therefore,

$$C_1^T = C_1^{NT} = \min_{M\geq 1}\left\{C_M^{NT}, C_M^T\right\} = \frac{p}{1-\beta}\frac{1+a^*}{1+a}. \tag{A.24}$$

(ii) Now consider the case when there is no commitment. By (i), we have proved that $M = 1$ dominates all the time-consistent token issuance plan (M, x), $M \geq 1$, $x \in \{T, NT\}$. However, without commitment, the consumer would believe the cost equals $C^{\text{Issue}} = [p + (1-p)a]P_1^T + ap/(1-\beta) > p/(1-\beta) = C^{\text{Benchmark}}$.

For Online Publication

October 16, 2024

OA Online Appendix

A Extended Discussion: Buy One - Get N Free

In the main text, we focus on issuance designs $M \geq 1$. Now we consider the relatively rare cases where $M \in (0,1]$, corresponding to "buy one and get N $(= 1/M)$ free." For sake of expression, we let one token represent unit redemption right, and the issuer sells one service bundled with N tokens, $N \in \mathbb{Z}^+$.

Again, we start with the three roles of bundling as shown in Section 2.3. First, the issuer sells unit service bundled with more future rights, thus requiring significantly stronger capacity (market power) for implementation. Second, regarding delayed cash flow, the non-tradable case retains a similar intuition, while the secondary market structure changes significantly: token holders sell "superfluous" tokens to unshocked consumers and never buy themselves, becoming de facto competitive token dealers. Whereas when $M \geq 1$, they either sell or buy to collect "fragmented" rights, so that tokens can clear even without unshocked consumers. As a result, the total token supply becomes crucial to pricing and affects the issuer's cash flow. Third, the core difference is that consumers no longer need to accumulate tokens — they always receive enough tokens for the next redemption upon purchase. That is, there is no time-series bundling even when tokens are non-tradable.

In a word, one may still expect that large market power helps maintain high token dependence, but the dominance of non-tradability may be broken, because the cash flow swap based on time-series bundling is disabled. Now we formally solve the token prices and revenues for the "buy one - get N free" case. We limit our focus on the monopoly case, since it requires strong commitment power to maintain such implementation.

Non-tradable tokens. The monopoly issuer charges the unit price P_N^{NT} to overwhelm the whole willingness to pay for $N+1$ units through one purchase,

$$P_N^{NT} = \sum_{i=0}^{N} \left(\frac{\beta p}{1-\beta(1-p)} \right)^i = \frac{1-a^{N+1}}{1-a}. \tag{OA.1}$$

Similar to (4), the revenue reads

$$\begin{aligned} Rev_N^{NT} &= \underbrace{(p+(1-p)a^*)}_{\substack{\text{PV of First} \\ \text{Purchase}}} \underbrace{P_N^{NT}}_{\text{Unit Price}} \underbrace{\frac{1}{1-a^{*N+1}}}_{\substack{\text{Discount since} \\ \text{First Purchase}}} \equiv D_N^{NT} \cdot P_N^{NT} \\ &= \frac{p}{1-\beta^*} \frac{1-a^{N+1}}{1-a} \frac{1-a^*}{1-a^{*N+1}}. \end{aligned} \tag{OA.2}$$

There is no difficulty in showing $\frac{\mathrm{d}Rev_N^{NT}}{\mathrm{d}N} > 0$, i.e., the issuer always gains greater revenue by bundling more tokens, since it is always able to leave zero consumer surplus. As a result, the monopoly issuer would optimally announce a "buy one to get free forever" strategy, i.e., a lift-time membership at the price $\lim_{N\to\infty} Rev_N^T = \frac{p}{1-\beta^*}\frac{1-a^*}{1-a}$.

This is formally consistent with our main focus, $M \geq 1$: when $M \to 0$, the issuance of loyalty tokens converges to a membership model that yields the largest front load. Proposition 3 still holds, indicating that it requires more market power to implement a "buy one get N free" strategy than "buy M get one free." In practice, regulatory issues seem to be more crucial, as consumers are exposed to significant risks of being scammed by firms promising future consumption and then walking away or going bankrupt. On the other hand, in the extreme case, the membership model effectively makes the consumer a shareholder, raising considerations of capital gains and tax issues, which is beyond our focus.

Tradable tokens. If $N = 1$, it is equivalent to $M = 1$, i.e., there is no trade. When $N > 1$, consumers (after purchase) have "superfluous" tokens for next redemption, thus becoming de facto token dealers. It is easy to show that: (i) selling one token to two people respectively is more beneficial than selling two tokens to one person, because people's willingness to pay for the second token is lower; (ii) there is a trade if and only if there exists two consumers, whose number of token holdings differ by at least 2, otherwise consumers have indifferent valuation on the "superfluous" token. As a result, the secondary market will make tokens as evenly allocated as possible across the entire consumer base. Given that there are $S_{N,t}^T$ outstanding tokens in period t, then each consumer holds $\lfloor S_{N,t}^T \rfloor$ or $\lceil S_{N,t}^T \rceil$ tokens after trade.

Consider the price. Note that token sellers, those who hold "superfluous" are competitive and do not have pricing power. The market clearing price equals the value of the marginal token. In particular, each seller sells $(N - \lceil S_{N,t}^T \rceil)$ tokens at the market-clearing secondary market price,

$$\hat{P}_t = a^{\lfloor S_{N,t}^T \rfloor}(p + (1-p)a). \tag{OA.3}$$

Turn to the primary market. If $\lfloor S_{N,t}^T \rfloor \geq 1$, then there would be no fiat money purchase at time t, i.e., the issuer does not sell tokens. In other periods, the issuer sells $Q_{N,t}^T$ services and can take away all the resellers' surplus. However, it leaves the buyers' surplus, due to the lack of monopoly power of the resellers. In addition, the issuer must maintain a stable service price although $S_{N,t}^T$ could vary across periods.

Therefore, the price satisfies:

$$
\begin{aligned}
P_N^T &= \min_{\{t|S_{N,t}^T\geq 1\}} \left\{ \underbrace{\sum_{i=0}^{\lceil S_{N,t}^T \rceil} a^i}_{\text{self-use for redemption}} \quad + \quad \underbrace{\beta(N-\lceil S_{N,t}^T \rceil)\hat{P}_t}_{\text{reselling option value}} \right\} \\
&= \min_{\{t|S_{N,t}^T\geq 1\}} \left\{ \sum_{i=0}^{\lceil S_{N,t}^T \rceil} a^i + a^{\lceil S_{N,t}^T \rceil}(N-\lceil S_{N,t}^T \rceil) \right\}.
\end{aligned}
\tag{OA.4}
$$

Note that $\forall k \in \mathbb{Z}$, $k \geq 0$,

$$
\left\{ \sum_{i=0}^{k+1} a^i + a^{k+1}[N-(k+1)] \right\} - \left\{ \sum_{i=0}^{k} a^i + a^k(N-k) \right\} = -a^k(N-k)(1-a) < 0.
\tag{OA.5}
$$

Therefore,

$$
P_N^T = \sum_{i=0}^{\max_{t\geq 0}\lceil S_{N,t}^T \rceil} a^i + a^{\max_{t\geq 0}\lceil S_{N,t}^T \rceil}(N - \max_{t\geq 0}\lceil S_{N,t}^T \rceil).
\tag{OA.6}
$$

Consider $\max_{t\geq 0}\lceil S_{N,t}^T \rceil$. Note that

$$
S_{N,t+1} \begin{cases} = S_{N,t}^T - p < S_{N,t}^T, & S_{N,t}^T \geq 1; \\ = [1-p(N+1)]S_t + pN \leq pN, & S_{N,t}^T < 1, \text{ and } 1-p(N+1) < 0; \\ = [1-p(N+1)]S_t + pN < 1-p, & S_{N,t}^T < 1, \text{ and } 1-p(N+1) \geq 0. \end{cases}
\tag{OA.7}
$$

Therefore, $S_{N,t}^T \leq \max\{1, pN\} \Rightarrow \max_{t\geq 0}\lceil S_{N,t}^T \rceil = \lceil pN \rceil$, and

$$
P_N^T = \sum_{i=0}^{\lceil pN \rceil} a^i + a^{\lceil pN \rceil}(N - \lceil pN \rceil).
\tag{OA.8}
$$

Similar to Section 2.5, we solve $\{Q_{N,t}^T\}_{t\geq 0}$ and $\{S_{N,t}^T\}_{t\geq 0}$ jointly. Initially, no one has redemption rights, $S_{N,0}^T = 0$. In each period, only those who do not have tokens choose to purchase by fiat money (note that tokens are allocated evenly), therefore

$$
p = Q_{N,t}^T + p\min\{1, S_{N,t}^T\}.
\tag{OA.9}
$$

The outstanding unredeemed tokens in the next period equals current outstanding minus current redemption plus new issuance, i.e.,

$$
S_{N,t+1}^T = NQ_{N,t}^T + S_{N,t}^T - p\min\{1, S_{N,t}^T\}.
\tag{OA.10}
$$

These two equations can be viewed as generalized versions of (5) and (6), where $S^T_{N,t}$ is always smaller than 1. Then the revenue reads

$$Rev^T_N = P^T_N \sum_{t=0}^{\infty} \beta^{*t} Q^T_{N,t}. \tag{OA.11}$$

Although it is hard to obtain the generic formula of $\{Q^T_{N,t}\}_{t\geq 0}$, we can numerically calculate Rev^T_N based on the recurrence formulas (OA.9) and (OA.10).

Before delving into numerical explorations, we qualitatively discuss how Rev^T_N changes with N. There are two opposite forces that make P^T_N potentially nonmonotonic with N. First, a larger N corresponds to more tokens bundled, yielding a positive force for the service price P^T_N, as the first term in (OA.8), $\sum_{i=0}^{\lceil pN \rceil} a^i$, shows. This is in line with the role of tradability in the "buy M get one free" case, partially compensating for delayed cash flow. Second, a larger N brings a new negative force by changing the value of the marginal token in the secondary market, thus it may lower the service price, as the second term in (OA.8), $a^{\lceil pN \rceil}(N - \lceil pN \rceil)$, shows. This is de facto the other tale of leaving positive buyer surplus. While in the main case, the marginal token always corresponds to the first redemption right, where people trade only to collect the "fragmented" rights, rather than sell "superfluous" rights.

Figure OA1 shows the numerical results of Rev^T_N under different issuance strategy N and consumption probability p, and compares to the non-tradable case, Rev^{NT}_N. In panel A where p is relatively large, it appears that Rev^T_N increasing monotonically with N $(1/M)$. This is consistent with Lemma 2, $\frac{\mathrm{d}Rev^T_M}{\mathrm{d}M} < 0$. This is because the first effect takes the dominance. Meanwhile, non-tradability dominates tradability, since tradability allows competitive resellers in the secondary market, leaving positive consumer surplus.

In panel B, non-tradability is still the dominant strategy. However, Rev^T_N is no longer monotonic, since the negative force, $a^{\lceil pN \rceil}(N - \lceil pN \rceil)$, becomes stronger. With lower p, this term decreases faster with N mathematically. From an economic perspective, a lower p implies a longer interval between two redemptions. Then as the marginal token increases by one, the valuation of tokens suffers from a greater discount.

When p is extremely low as shown in panel C, the dominance of tradability breaks, i.e., $\max_N Rev^T_N > \max_N Rev^{NT}_N$,[39] where the global optimum solves a relatively small N, rather than $N \to \infty$. The reason is twofold. First, when N is small, the negative force is less pronounced as mentioned. In particular, as long as $pN < 1$, the marginal token is always the first token, which means that buyers do not obtain a

[39]Here we focus on monopoly case, similar to Section 2.5. Therefore, $Rev^x_N = U^x_N$, $\forall N \geq 1$, $x \in \{T, NT\}$.

positive surplus. Therefore, the drawback of tradability is disabled. Second, with an extremely low p, consumers largely discount the future willingness to pay. This makes lifetime membership less valuable than immediate transfer rights. As a result, $Rev^T_{N^*} > \lim_{N\to\infty} Rev^{NT}_N$, i.e., it could be worse to sell a membership to a small group than to sell a few tokens to them and let them resell to the others.

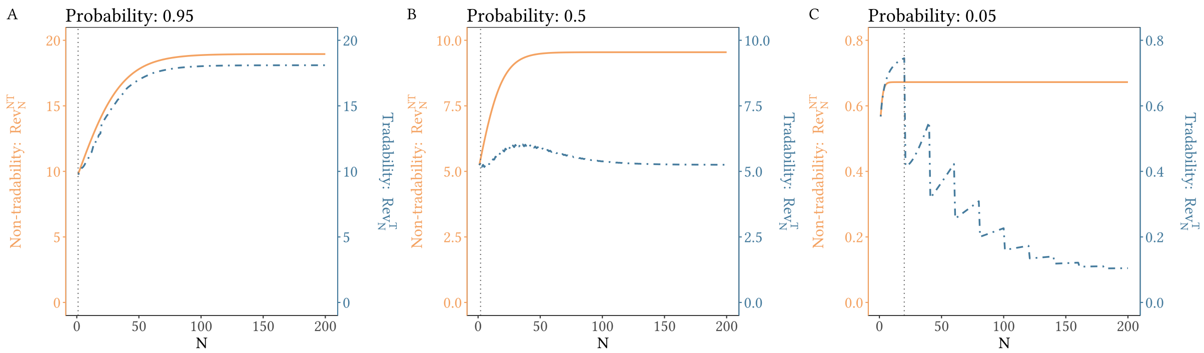

Figure OA1. Revenue Comparison under “Buy One - Get N Free”

This plot compares the revenues of “buy one and get N free tradable / non-tradable tokens”, i.e., Rev^T_N and Rev^{NT}_N, under different N and probability p. In panels A-C, p equals 0.95, 0.5, and 0.05, respectively. In each panel, x-axis is the issuance strategy N, the yellow solid line and its corresponding left y-axis represents revenues under non-tradability cases, Rev^{NT}_N; whereas the blue dashed line and the right y-axis shows the revenues with tradability, Rev^T_N. The gray dotted line plots $N = 1/p$ to illustrate the first step-wise segment. $\beta = 0.95$, $\beta^* = 0.9$.

Similarly, if the issuer is impatient ($\beta^* \to 0$), tradability could always beat non-tradability. In this extreme case, the issuer only earns p purchases in the initial period, and can always charge a higher price from reselling option value when allowing tradability. This counterexample does not apply to our main focus, $M \geq 1$. Although $\forall M > 1$, the issuer takes away the reselling option values, these strategies are all sub-optimal relative to $M = 1$, where no trade happens, and the token price equals the non-tradable one.

In summary, although rare in reality and subject to higher regulatory requirements, the “buy one and get N free” ($M < 1$) case can also be consistently included in our framework. Compared to our main focus, $M \geq 1$, this case naturally disables time-series bundling, and somehow appears to be more like the benchmark case of platform currencies. Whereas the non-tradable issuance converges to a life-time membership model in optimal, tradability yields a non-monotonic service pricing with token bundling. In particular, due to losing the unique advantage from time-series bundling, non-tradable issuance may not take the dominance when immediate purchases are vitally important.

B Alternative Evaluation of Loyalty Token Issuance

In the main text, we mainly use the token dependence Θ to evaluate the effectiveness of loyalty token issuance, which represents the revenue share of presale by loyalty tokens, and corresponds well to empirical proxies. Here we discuss the alternative evaluation Δ in detail, defined as the net payoff growth relative to the benchmark of without issuing any tokens, i.e.,

$$\Delta \equiv (U^* - Rev_0)/Rev_0, \tag{OA.12}$$

where U^* is the payoff under the optimal issuance plan $(M_*^{x^*}, x^*)$, and Rev_0 denotes benchmark revenue (without tokenization).

Figure OA2 explores the variation in optimal loyalty payoffs U^* over the parameter space (p, c), and importantly, examines how much the payoffs are enhanced by adopting loyalty token issuance. Panel (a) visualizes $U^*(p, c)$. The difference in the probability of consumption demand dominates the variation in revenue, which usually reflects the varying characteristics among industries.[40] That is, $Rev_0(p)$ naturally changes with p and in turn increases aggregate revenues.

Figure OA2 (b) examines Δ from a different perspective than Figure 4 (c). Loyalty tokens are most beneficial for relatively small demand probabilities, which partially compensates for the negative impact of the low consumption frequency. However, when the consumption frequency is extremely low, only issuers with absolute power (such as pure monopolies) can profitably implement a loyalty program. In summary, issuers with large market power in industries with relatively low-frequency consumption are most favored for issuing loyalty tokens.

[40]In practice, the unit price may also affect the revenue comparison among industries, which is not our focus. Also, when looking at relative concepts, i.e., Rev^*/Rev_0 and the token dependence Θ, the influence of unit price is eliminated. Thus, for simplicity, we normalize the current unit production cost to 1.

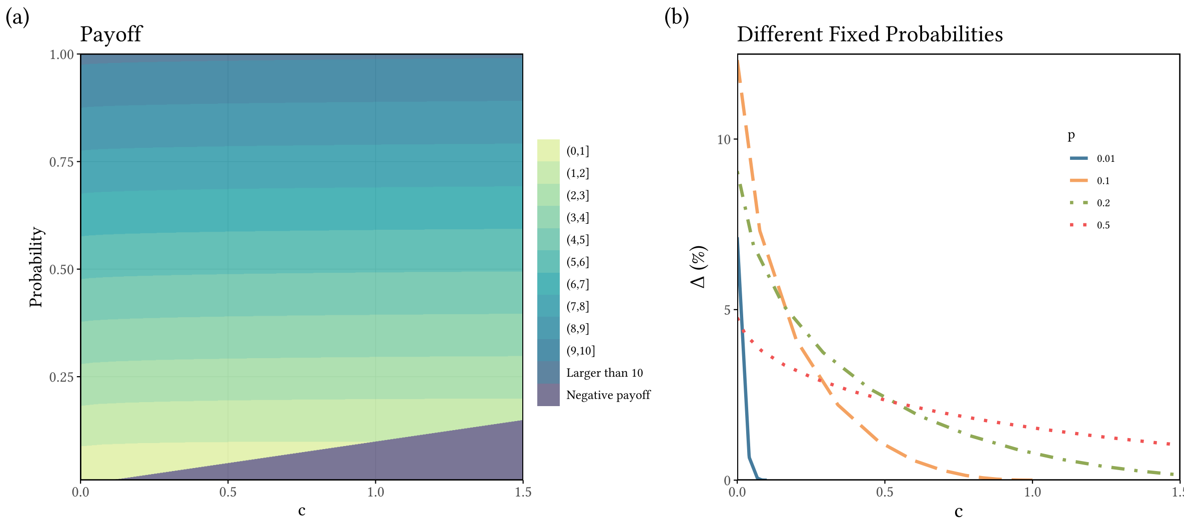

Figure OA2. Payoff Enhanced by Optimal Loyalty Token Designs

Panel (a) plots the aggregate payoff U defined in (3) when adopting the optimal issuance plan for each point (p, c), where the optimal issuance $(M_*^{x^*}, x^*)$ is illustrated in Figure 4. The latter two panels visualizes Δ, the payoff growth rate (accounting for external losses) driven by the optimal issuance, benchmarked against the non-use of any loyalty token offering. In Panel (b), different p are fixed, and the colored curves show how the growth changes with market power. $\beta = 0.95$, $\beta^* = 0.9$, and $f(P) = P - 1$.