



The Optimal Design of Loyalty Programmes

Revenue-based loyalty programmes yield better profits, but consumers don't have to be on the losing end.

Loyalty programmes used to have a simple arithmetic to them. For example, Starbucks customers would get gold status after making a certain number of transactions. Travellers stood to achieve airline elite status after reaching a certain mileage. Sure, the credited flown mileage would vary based on the fare type, but it was still simple enough. Then, from 2014, things got more complicated.

To the dismay of many consumers, loyalty programmes started to impose spending requirements. To qualify as a preferred customer, it wasn't enough anymore to be a regular patron – they had to bring in decent revenues as well. According to Starbucks, shrewd customers had been **maximising** their number of purchases by asking cashiers to ring items individually, slowing down queues in the process. **Throngs** of airline-status seekers successfully hunted for dirt-cheap fares allowing them to meet mileage requirements for next to nothing. In short, it wasn't all that hard to game the system.

We set out to analyse how different loyalty programme designs impact firms and consumers, and whether consumers necessarily lose out when a firm switches to a revenue-based programme or starts imposing dual requirements. What we found is a sweet spot where profit-minded firms can improve

profit without hurting strategic consumers. Our paper, "**Strategic Consumers, Revenue Management, and the Design of Loyalty Programs**", was recently published in *Management Science*.

Many ways to skin a cat

First, our mathematical model showed that revenue-based loyalty programmes do raise firms' profits over traditional quantity-based programmes. However, this increase doesn't come solely from customers forced to spend more. It is also due to lower operation costs, since such programmes thin the ranks of "costly-to-service" customers who multiply bargain purchases to qualify for status. To illustrate the general principle, a customer who makes a \$10 purchase (in a single visit) is more profitable than one who spends the same amount of money over five visits, since the firm incurs some costs for each visit. Even if the firm's revenues stay the same, but costs go down, profits go up.

But what about the customers? Some may benefit from the change, but others may suffer. For example, some consumers may reduce their unnecessary purchases and still qualify for status, but others who otherwise would qualify for status under the old rule might lose the status under the new one. Thus, the negative public reactions could

Visit **INSEAD Knowledge**
<http://knowledge.insead.edu>

be partially justified.

Second, we assessed whether there was a way to design a revenue-based, elite-status programme such that it would increase profit while making no consumer worse off. Economists refer to this as a *Pareto improvement*: one party gets a better deal, but it comes at no one's expense. We showed that such a solution indeed exists. To achieve this, however, the firm might need to sacrifice some of its full potential for profit. In other words, a firm can increase its profit without hurting any consumer, provided it is willing to leave a little money on the table. This optimal scenario is possible if the firm's cost of providing elite status is relatively low (the perks are relatively inexpensive), yet the consumers value that status and are thus willing to pay a little more for the firm's product or service in order to qualify for elite status.

How revenue requirements can best coexist with quantity requirements

In recent years, many elite-status loyalty programmes have added spending requirements on top of the need to reach a certain quantity threshold (e.g. minimum spending and number of legs flown).

Interestingly, our model showed that such dual-requirement programmes do not lead to higher profits compared to programmes strictly based on revenues (e.g. minimum spending alone) – both have the same potential for profit. That is because, in practice, once spending requirements are set to an optimal level, the quantity threshold (e.g. number of purchase transactions, hotel nights, etc.) becomes redundant. However, as dual-requirement programmes do create an extra hurdle, they are bound to trigger negative reactions from consumers.

Our analysis showed that programmes that give consumers a choice between revenue and quantity requirements (either allow consumers to qualify for elite status) can maximise profit just as well as revenue-based programmes. As a bonus, firms can avoid consumer outcry as it is possible to design them in such a way that no consumer will be worse-off. How? The math works if a firm makes it easier for consumers to satisfy the spending requirement than the quantity requirement and thus entices them to take this strategic route. In addition, such programmes can attract loyalty from consumers willing to spend but who wouldn't normally make enough individual purchases (or fly enough miles) in a given year to qualify for elite status.

Key insights

While some consumers are myopic and consider each purchase they make in isolation, the market

Visit INSEAD Knowledge
<http://knowledge.insead.edu>

also comprises a sizeable number of highly strategic consumers, willing to make “odd” decisions in order to reap a future benefit, such as a perk-giving elite status. These consumers may make unnecessary purchases (e.g. **mileage runs**) or pay high prices when lower prices are available (e.g. choose an airfare giving bonus mileage). By switching to a revenue-based loyalty programme, a firm can incentivise consumers to choose a strategic route that naturally results in higher profits.

Contrary to the generally negative public and media reactions, such a switch could lead to win-win outcomes for both consumers and the firm. The best way for a firm to manage the transition is to impose softer rules and allow customers to qualify by meeting either quantity or revenue requirements. This type of design is quite palatable to consumers as it gives them a choice: They can either spend a certain amount or make a certain number of qualifying transactions. If the firm is willing to forgo a little profit, such a programme can be designed to increase firm profits while leaving no customer worse off.

***So Yeon Chun** is an INSEAD Associate Professor of Technology and Operations Management.*

***Anton Ovchinnikov** is a Visiting Professor of Decision Sciences and Technology and Operations Management at INSEAD and a Distinguished Professor of Management Analytics and Scotiabank Scholar of Customer Analytics at Smith School of Business, Queen's University in Canada.*

Found this article useful? **Subscribe** to our weekly newsletter.

Follow INSEAD Knowledge on **Twitter** and **Facebook**.

Find article at

<https://knowledge.insead.edu/operations/the-optimal-design-of-loyalty-programmes-12111>

Download the Knowledge app for free

