

Chapter 22

Frequency Reward Programs

Abstract Frequency reward programs are customer development programs based on the theme, “Buy XXX, get a reward.” “XXX” is usually a required purchase volume, and the reward can be free product, a cash rebate, or even “points” for another company’s reward program. We discuss two ways that reward programs increase sales – points pressure and rewarded behavior – and the empirical evidence for each. We then review the rich economics literature that has endeavored to answer the question, “In a competitive environment, do reward programs increase firm profits?” We review several issues in designing reward programs, including the reward structure, and conclude with a review of reward programs offered by firms including Harrah’s Entertainment and Hilton Hotels.

22.1 Definition and Motivation

Frequency reward programs attempt to increase customer value by rewarding customers proportional to their cumulative purchases, revenues, or profitability. Reward programs are used in a variety of industries such as hotels, car rentals, supermarkets, credit cards, office products, telecom, and gaming casinos. In fact, the airline industry’s frequent flyer programs have in effect created a new currency in the “miles” travelers accumulate. In terms of the simple retention model of customer lifetime value,

$$LTV = \sum_{t=1}^{\infty} \frac{m_t r^{t-1}}{(1 + \delta)^{t-1}} \quad (22.1)$$

frequency reward programs strive to increase both retention rate (r , the customer stays with the firm longer) and the amount purchased (m_t , the customer buys more to accumulate enough “points” to receive a reward).

We differentiate between frequency reward and customer tier programs (Chapter 23). Customer tier programs assign customers to segments or tiers

and deliver different benefits to each tier. Frequency reward programs are a narrower promotional-oriented activity. They focus on the delivery of a single reward – a free flight, an upgrade, a coupon, etc. – based on accumulated points.

Both frequency reward and customer tier programs are often called “loyalty programs.” We do not use that term because loyalty may be a *goal* of these programs, and they do increase purchase frequency. But whether they increase “loyalty”, defined as “a favorable attitude toward a brand resulting in consistent purchase of the brand over time” (Assael 1995, p. 131) is another matter.

22.2 How Frequency Reward Programs Influence Customer Behavior

22.2.1 Mechanisms for Increasing Sales

There are three mechanisms by which the program can increase customer value: points pressure, rewarded behavior, and personalized marketing. Figure 22.1 illustrates.

The points pressure mechanism represents customers increasing their expenditures in order to earn the reward. The attractiveness of the reward obviously increases points pressure. Also, the reward program can create a switching cost for the customer, in that the customer who decides to purchase elsewhere forgoes the opportunity to accumulate points toward the reward (Taylor and Neslin 2005).

The points pressure effect should get stronger as the customer nears the requirements for a reward. First, the reward is subject to less discounting as it looms closer. Second, Kivetz et al. (2006) propose two psychological reasons for points pressure: (1) Goal-gradient hypothesis – this is a behaviorist

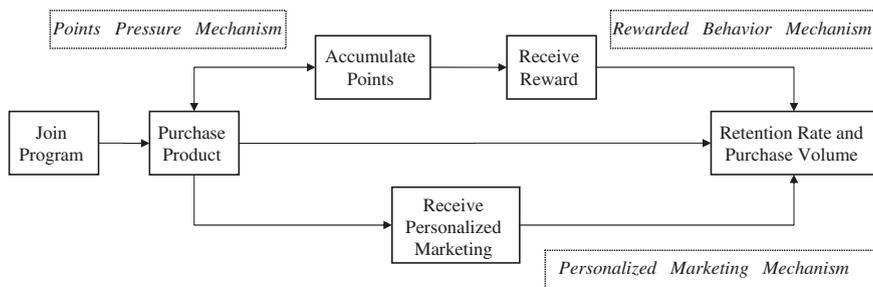


Fig. 22.1 Customer response to frequency reward programs: How reward programs influence retention and purchase volume.

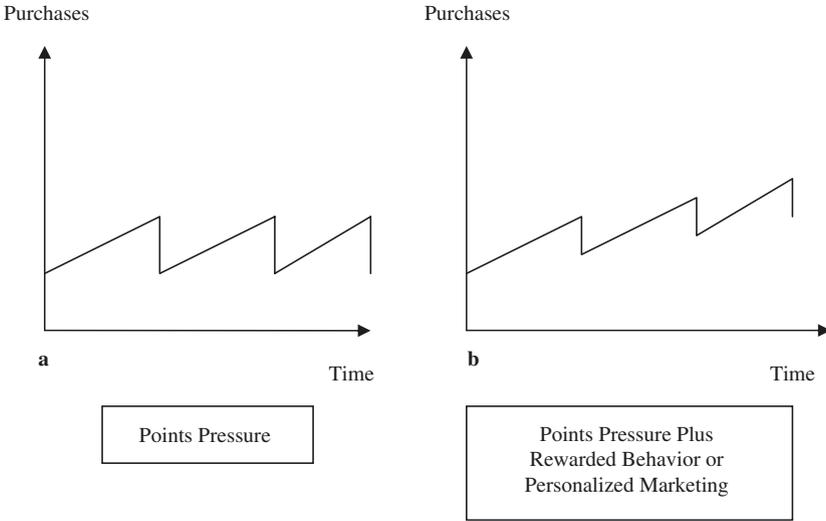


Fig. 22.2 Potential impact of frequency reward programs on purchase rates.

proposition originally due to Hull (1932) where organisms accelerate their efforts as they get closer to a goal. (2) Goal Distance Model – this is based on the psychophysical notion that humans make judgments relative to a benchmark. In the case of reward programs, the benchmark is the amount of effort (e.g., the number of purchases) to reach the goal (earn the reward). Consumers judge their progress by “calculating” the percentage of the distance they have covered toward receiving their reward. The authors hypothesize that motivation to reach the reward is an increasing function of this percentage.

The rewarded behavior mechanism is when customers increase their purchase rate *after* receiving the reward. Behavioral learning theory posits that “rewarded behavior is more likely to persist” (Blattberg and Neslin 1990, p. 22). Taylor and Neslin (2005) note that rewards can increase affect toward the firm, which subsequently translates into higher purchase rates. Whether the rewarded behavior effect is due to behavioral learning or increased affect is important for distinguishing whether rewards truly increase loyalty or just increase purchase inertia (Engel et al. 1995, p. 158).

The third mechanism for increasing retention and purchase volume is through the provision of personalized marketing efforts to members of the frequency reward program. These efforts include individually targeted promotions (for a retail store), cross-selling (for a gaming casino), or personalized customer service (for an airline). These efforts are not rewards *per se* but merely the company making use of what it learns about customer preferences through the customer’s participation in the program.

Figure 22.2 illustrates the impact of the points pressure, rewarded behavior, and personalized marketing mechanisms on customer purchases

(see Taylor and Neslin 2005). Figure 22.2a shows the build up in purchases due to the anticipation of the reward – the points pressure mechanism. Figure 22.2b shows an increase in purchasing above the previous baseline, due either to the rewarded behavior or personalized marketing.

22.2.2 What We Know About How Customers Respond to Reward Programs

22.2.2.1 Laboratory Experiments and Empirical Studies

Lal and Bell (2003) examined programs of the type “Spend \$X over Y weeks and earn a reward”. For example, one program awarded customers with a free ham if the customer spent \$475 over a 6-week period. Lal and Bell grouped customers by baseline sales levels (“Worst,” “Better,” and “Best”). Their findings were important:

- The reward programs increased sales in the period leading up to the reward. This supports a points pressure effect.
- The points pressure effect was strongest among the Worst customers and weakest among the Best customers.
- The programs were generally profitable, due mainly to the increase in sales among the worst customers.
- There was a positive post-redemption impact, strongest among the Worst customer group. The authors had hypothesized that the post-redemption effect would be negative due to consumer stockpiling, so were surprised to find a positive effect. Our interpretation is that this was due to rewarded behavior.

Taylor and Neslin (2005) examine the same type of reward program as Lal and Bell (2003). Customers who purchased \$500 worth of product over the 8-week period leading up to Thanksgiving were awarded a free turkey. Taylor and Neslin detect a points pressure effect in both years the program was run. In Year 1, sales per customer increase \$2.44 per week on a customer baseline sales level of \$37.91, a 6.1% increase, and \$2.61 in Year 2 on a base of \$41.02, a 6.4% increase. The authors calculate profits based on this points pressure effect and find the program to be profitable.

Taylor and Neslin (2005) also measure a rewarded behavior effect. The authors use a switching regression, where the first stage predicts whether a customer will redeem for a reward or not, and the second stage predicts post-redemption sales for redeemers versus non-redeemers. They find an average weekly post-redemption effect of \$14, which on a baseline of \$80, represents an increase of 17.5%. Taylor and Neslin also find that the rewarded behavior

effect is strongest among customers with low current baselines, similar to Lal and Bell (2003).

Kivetz et al. (2006) examine the dynamics of points pressure in a series of field experiments. One involved a coffee shop, where the customer had to purchase 10 coffees in order to receive a free coffee. Another involved a music rating website, where visitors had to rate 51 songs in order to receive a gift certificate. The authors found that both programs induced points pressure. They found that the effect is gradual rather than a step function. This is consistent with the authors' goal-gradient and goal-distance theories. The authors found three additional and intriguing results:

- When the coffee shop offered a two-purchase credit to start off, acceleration toward the goal was faster. This supported the authors' goal-distance theory, because the percentage of the goal achieved was higher the more coffees that were bought.
- After the reward, respondents "reset" their purchase levels to the baseline levels before the promotion. This rules out stockpiling as a reason for the higher pre-reward purchase rates, but does not support rewarded behavior.
- Customers who were closer to their reward program goal were more receptive to promotions. This was investigated using a pencil and paper task, so is exploratory, but the implications for the third mechanism by which reward programs increase sales – targeted promotions – are obvious: Target promotions to customers who are approaching their goals.

Roehm et al. (2002) use laboratory experiments to investigate the impact of reward design on future loyalty, i.e., the rewarded behavior effect. The authors categorize rewards in terms of "cue compatibility" and "tangibility." A cue compatible reward is one that is consistent with what the brand stands for, e.g., a cue compatible reward for a supermarket might be a turkey. Tangibility refers to the directness of the incentive – newsletters and clubs are intangible, whereas price discounts are tangible. The authors hypothesize that an intangible cue-compatible reward would increase loyalty because the customer is encouraged to think about what the brand truly stands for. At the other extreme, a tangible cue-incompatible reward could hurt brand loyalty by encouraging customers to think just about the reward and not about the brand itself.

Roehm et al. create an experimental reward program and choice simulation for *Slice* soft drink. The cue-compatible, intangible reward was admission to a web-site featuring games and puzzles for "a refreshing change of pace" (p. 205). The cue-compatible, tangible incentive was a foam drink insulator. The cue-incompatible, intangible reward was admission to the *Slice* website featuring games and puzzles regarding fitness. The cue-incompatible, tangible reward was gym towels. The results showed that all four incentives were equally enticing in terms of brand choice. However, the impact on loyalty differed by type of reward, as follows:

Reward	Impact on low knowledge subjects	Impact on high knowledge subjects
Cue compatible, intangible	Increase	No change
Cue compatible, tangible	No change	Decrease
Cue incompatible, intangible	No change	No change
Cue incompatible, tangible	No change	Decrease

These results suggested a rewarded behavior effect only if the reward is cue compatible and not too tangible. Also, the impact is only on customers who are less familiar with the brand. Highly tangible rewards can decrease loyalty among customers highly familiar with the product, even if the reward is compatible with positioning.

Roehm et al.'s (2002) results are partially consistent with Lal and Bell (2003) and Taylor and Neslin (2005). In both those studies, the rewarded behavior effect was strongest among light users of the product. However, the rewards were tangible, which would work against finding a rewarded behavior effect. More work is needed to reconcile these findings.

Bolton et al. (2000) studied the impact of a reward program for a European credit card company. The authors collected data for 405 customers. They found that membership in the program did not directly affect customer retention or purchase volume. However, membership in the program decreased negative perceptions resulting from poor service encounters. This may be due to the reward distracting customers from evaluating poor service encounters.

Many reward programs require the consumer to decide whether to enroll in the program. A simple comparison of sales among enrollees versus non-enrollees will therefore be biased because it may be that currently loyal customers are pre-disposed to enroll. Leenheer et al. (2007) control for this in analyzing the impact of reward programs on share-of-wallet (SOW) for Dutch grocery stores. They model SOW as a function of program membership, and program membership as a function of store and household characteristics. The first relationship is estimated using two-stage least squares, since program membership is endogenous; the second relationship is measured using a selection model. The authors find (1) program membership increases SOW by on average 4.1 percentage points, and (2) this translates to a net revenue increase (net of program costs) of €163 per year. This analysis does not determine whether the gain in SOW is due to points pressure or rewarded behavior, but it has a firm footing in claiming a causal impact of reward program membership on firm revenues by controlling for the customer's decision of whether to join the program.

22.2.2.2 Dynamic Structural Models

Frequency reward programs “try to change the customer's choice process from operating in a spot market to operating in a multi-period, contractual relationship market” (Dowling and Uncles 1997). The customer realizes that

purchasing from the company now has ramifications for *future* benefits, so needs to take into account not only the current benefits of purchasing, but future benefits as well. Dynamic structural models are designed to study consumer decision-making when customers take into account both current and future utility, especially when future utility depends on the decisions the customer makes now.

Characteristics of Dynamic Structural Models Relevant to Reward Programs

We first discuss the general structure of dynamic structural models so we do not need to repeat it for each model. The customer’s task is to choose a set of decisions over a time period $\{t, \dots T\}$ to maximize the net present value of their utility:

$$Max_{D_{ik\tau}, k \in C_{i\tau}} \sum_{\tau=t}^T \delta^{\tau-1} \sum_{k \in C_{i\tau}} U_{i\tau}(k) D_{ik\tau} \tag{22.2}$$

where:

- $D_{ikt} = 1$ if customer i chooses decision alternative k at time t .
- C_{it} = Set of decision alternatives available to customer i at period t . These alternatives might include whether or not to purchase from a particular firm; whether to “cash-in” points for a reward, etc.
- δ = Discount factor; the degree to which the customer values future utility. Higher δ means the customer cares more about the future.
- $U_{it}(k)$ = Utility customer i gains for choosing alternative k at time t .

This decision problem can be equivalently expressed using the principle of optimality, which says that the optimal decision for period t is that which maximizes current utility plus the expected value of future utility that results from the decision made in period t . This maximum utility is called the “Value Function”, denoted by V . In particular:

$$V_{it}(S(T)) = Max_{D_{ikt}, k \in C_{it}} \{U_{it}(k) D_{ikt} + E[V_{i,t+1}(S(t+1)) | S(t), D_{ikt}]\} \tag{22.3}$$

where $S(t)$ are “state variables” at time t . State variables influence utility and change over time due either to a stochastic process or to the decision made at time t . A key state variable in frequency reward models is the number of points the customer has accumulated. The future values of state variables such as price are uncertain, so customers may form expectations of these variables (Erdem et al. 2003; Sun et al. 2003).

A Simple Model for Examining Competitive Equilibria

Kopalle and Neslin (2003) investigate competitive equilibria in reward programs. Their model is relatively simple. The customer decides each period

which of two airlines to fly as well as whether to fly at all. The choice set is:

$$C_{it} = \begin{cases} 0 & \text{if customer } i \text{ flies neither airline in period } t \\ 1 & \text{if customer } i \text{ flies airline } ABC \text{ in period } t \\ 2 & \text{if customer } i \text{ flies airline } XYZ \text{ in period } t \end{cases} \quad (22.4)$$

The customer's utility function for each of these choices is as follows:

$$U_{it}(k) = \begin{cases} U_0 & \text{if } k = 0 \\ r_{i1t} - (1 - INV_{i1t})P_{1t} & \text{if } k = 1 \\ r_{i2t} - (1 - INV_{i2t})P_{2t} & \text{if } k = 2 \end{cases} \quad (22.5)$$

where:

U_0 = Utility for not flying, i.e., for the "outside category" (e.g., traveling by car).

r_{ijt} = Customer i 's preference for Airline j in time t . r_{ijt} follows a logistic distribution with parameter a_j .

INV_{ijt} = The number of "points" customer i has applicable to airline j in time t :

$$INV_{ijt} = \begin{cases} 0 & \text{if } D_{ikt} = 1 \quad \text{for } k \neq j \\ 1 - INV_{ij,t-1} & \text{if } D_{ijt} = 1 \end{cases}$$

P_{jt} = The price of airline j in time t .

INV is defined as a zero-one variable. If the customer has no points accumulated for airline j and flies airline j , INV is set equal to 1. If the customer had one point accumulated for airline j and flies airline j , the flight is free (Equation 22.5). If the customer flies another airline or does not fly, INV is set to zero at time t , irrespective of whether the customer had accumulated points as of $t - 1$ or not. That is, customers build up a "credit" by flying a particular airline. They can cash in that credit for a free flight, but if they don't, they lose the credit. This is meant to model points expiration, a characteristic of some reward programs whereby customers lose points they do not cash in.

The state variables are INV_{ijt} and r_{ijt} . The customer is assumed to know the current values of r_{ijt} and the *probability distribution* of r_{ijt} for future periods. This means the customer knows his or her current preferences each airline, but is uncertain about future preferences because future flight requirements and schedules are not known.

Kopalle and Neslin derive insights from their model: (1) The value of the reward (the free trip) is increasing in δ , and hence a customer is more likely to repeat-fly an airline if the discount factor is large. This makes sense – if the customer cares more about the future, the accumulated points or credit will drive current purchase more strongly. (2) The value of the reward is increasing with the price charged by the airline. This is because the reward

is a free flight. This means that expensive tickets are more valuable rewards. This is one reason why upgrades have become so popular. It is also why we often see frequency reward programs implemented by premium brands such as Hilton, Marriott, etc. (3) Reward programs expand the category if U_0 is large. If the outside category is not attractive, the market for Airlines ABC and XYZ will be large whether they have a reward program or not. However, if the outside category is attractive, category sales are currently low, but the program can grow the category.

A Model of an Online Retailer Frequency Reward Program

Lewis (2004) studied the frequency reward program of an online retailer. If the customer accumulated a threshold level of expenditures within a year, he or she would receive a reward (500 miles to be added to the customer's frequent flyer program). The reward is not free goods, but an *indirect* reward that can be applied to another good. We will discuss the question of whether to use direct or indirect rewards in Sect. 22.4.5. The choice set for Lewis' model is:

$$C_{it} = \begin{cases} 0 & \text{if customer } i \text{ does not buy in period } t \\ 1 & \text{if customer } i \text{ buys a small basket } (< \$50) \text{ in period } t \\ 2 & \text{if customer } i \text{ buys a medium basket } (\geq \$50 \text{ and } \leq \$75) \\ & \text{in period } t \\ 3 & \text{if customer } i \text{ buys a large basket } (\geq \$75) \text{ in period } t \end{cases} \quad (22.6)$$

Discretizing the amount bought makes the model easier to estimate. This is not a competitive model but there is an outside good, i.e., the customer may decide not to purchase at all from the website. The utility function is:

$$U_{it}(k) = \begin{cases} \beta_{i0} + \varepsilon_{i0t} & \text{if } k = 0 \\ \beta_{ik} + \beta_{ipk}P_t + \beta_{ick}C_{it} \\ \quad + \sum_{h=1}^4 \beta_{ihk}SH_{iht} + \sum_{g=1}^G \beta_{ig}Rcn_{igt} & \\ \beta_{ifk}FM_{ift} + \beta_{iL}L_{it} + \varepsilon_{ikt} & \text{if } k = 1, 2, 3 \end{cases} \quad (22.7)$$

where:

P_t = Price index for the online retailer at time t .

C_{it} = Indicator variable of whether customer i received a coupon at time t .

SH_{iht} = Indicator variable of whether shipping charge schedule h was available to customer i at time t . The online firm used four different shipping charge schedules that they varied over time.

Rcn_{igt} = Indicator variable of whether customer i was in recency group g at time t . Recency is the time since the previous purchase.

Table 22.1 Frequency reward customer segments (From Lewis 2004)

	Segment 1	Segment 2
Percentage of sample	71%	29%
Mean annual # orders	8.7	14
Mean order size	\$77	\$42
Mean spending	\$672	\$586
Percentage earning reward	15%	1%
β_{iL}	0.75 (p < 0.01)	n.s.

FM_{ift} = Indicator variable of whether customer i was in cumulative purchase group f at time t . The expectation was that larger cumulative purchasing would mean the customer was learning to like the retailer so that the β 's for this variable would be positive.

L_{it} = Value of reward customer i received in time t . The variable was coded as 0 (received no reward), 1 (received 500 frequent flyer miles for accumulating either \$1,000 or \$1,500 in expenditures), or 2 (received 1,000 frequent flyer miles for accumulating \$2,000 in expenditures).

The parameter of most interest is β_{iL} , the utility of the reward. Lewis uses a latent structure formulation to capture customer heterogeneity in the utility parameters, thus yielding segments. Lewis includes an equation similar to the inventory accumulation used by Kopalle and Neslin to keep track of expenditures. If a given purchase will put the customer above a reward threshold, the L_{it} variable is set to 1 for that period. Assuming a positive β_{iL} , this provides an incentive for the customer to purchase.

Uncertainty in the model is incorporated via ε_{ikt} – the customer is uncertain before period t how much utility he or she will derive from decision k . Prices are also assumed uncertain and Lewis assumes customers form expectations for these prices. The state variables are the cumulative level of expenditures (i.e., the points inventory count), recency, frequency, price, couponing, and shipping. Lewis estimated the model and found two segments, as shown in Table 22.1.

The table shows that the majority of the sample derives extra utility from receiving the reward. This coincides with their generally higher level of spending, and the fact that 15% of them did earn a reward. This is in contrast to Segment 2, whose members derive no extra utility from the reward, and in fact, only 1% of them earned it.

The positive coefficient for the reward produces a points pressure effect. This is shown in Fig. 22.3 (based on fig. 2, p. 290 in Lewis' paper). It shows that if a customer has accumulated \$900, the closer he or she gets to the end of the year, the more likely he or she is to purchase. Interestingly, if a customer has accumulated only \$500, there is a *reverse* points pressure effect; the customer is discouraged from purchasing as time runs out. This is a fascinating finding. Remember the customer is working through a dynamic program. At the beginning of the year, achieving \$1,000 in expenditures is a real possibility and this spurs sales. However, toward the end of the year,

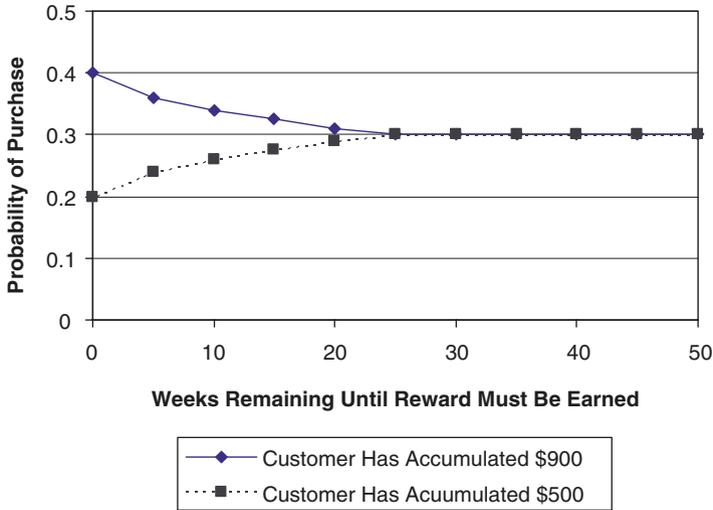


Fig. 22.3 Probability of purchase as time remaining to earn reward decreases (Adapted from Lewis 2004).

if the customer only has accumulated \$500, the possibility is less real and *expected* future utility (V_{t+1} in Equation 22.3) is lower. This suggests one must be careful not to set reward thresholds too high. It might spur sales initially, but if many customers are not be able to make the threshold, they will become discouraged and sales will fall.

Lewis conducted simulations and determined that both the number of purchases and the average customer revenue increases when the reward program is in place versus when it is not, although the gain in revenues is only 2%. Perhaps the reward (500 frequent flyer miles) was not that highly valued, and in fact was not valued at all by 21% of the sample. Also, the threshold level (\$1,000 for the first reward) was might have been set too high.

Modeling both Frequency Reward and Customer Tier Programs

Kopalle et al. (2006) develop a model of an airline’s frequent flyer program. The model is distinct in modeling (1) the decision of whether to redeem miles for rewards, (2) the rewarded behavior effect, and (3) both a frequency reward program and a customer tier program. The choice set is defined as:

$$C_{it} = \begin{cases} 0 & \text{if customer } i \text{ does not fly with this airline in time } t \\ 1 & \text{if customer } i \text{ flies with airline in time } t \text{ but does not cash in} \\ 2 & \text{if customer } i \text{ flies with airline in time } t \text{ and cashes in} \\ & \text{for level 1 reward} \\ 3 & \text{if customer } i \text{ flies with airline in time } t \text{ and cashes in} \\ & \text{for level 2 reward} \end{cases} \tag{22.8}$$

The utility function is:

$$U_{it}(k) = \begin{cases} \varepsilon_{i0t} & \text{for } k = 0 \\ \alpha_{ik} + \sum_{s=1}^S \lambda_{is} E_{ist} + \beta_{i1} \sum_{l=2}^3 D_{ilt-1} & \\ + \beta_{i2} \sum_{l=1}^3 D_{ilt-1} + \varepsilon_{ikt} & \text{for } k = 1, 2, 3 \end{cases} \quad (22.9)$$

where E_{ist} is an indicator variable of whether customer i has accumulated enough miles to qualify for level s of the airline’s customer tier program. If the customer does qualify, he or she automatically is a member of that level until either he or she accumulates more miles and hence qualifies for an even higher tier, or a year has gone by and the customer has not maintained the level of miles needed to remain in the tier.

Note that α_{i2} and α_{i3} reflect the immediate utility of a level 1 reward (upgrade) and a level 2 reward (free flight) respectively. The D_{ilt-1} terms equal 1 if the customer chose alternative 1, 2, or 3 in the previous period (indexed by l); 0 otherwise. Therefore, β_{i1} represents the rewarded behavior effect since it is added to the utility function only if the customer has cashed in for an upgrade or a free flight in the previous period. The β_{i2} term represents state dependence since it is added to the utility function as long as the customer flew with the airline last period, no matter whether it was a regular paid flight, an upgrade, or a free flight.

The state variables are the cumulative number of miles flown and the lagged decision variables. There are two sources of uncertainty. First, as in the other models reviewed in this section, overall utility for each decision is uncertain (the ε ’s). This uncertainty is resolved in the current period, but unknown for future periods. Second is the number of miles the customer will be awarded if he or she flies in some future period $t + x$. Kopalle et al. assume that customers are aware of the average number of miles they fly when they fly with this airline, and factor this in when considering the future.

The authors estimate their model on 200 customers and find two segments. The coefficients for the segments are as follows:

Parameter	Segment 1	Segment 2
Percentage of Sample	6.3%	93.7%
Base utility of flying (α_{i1})	0.62	0.75
Utility of upgrade (α_{i2})	1.56	-5.49
Utility of free flight (α_{i3})	1.77	-9.89
Utility of first customer tier (λ_{i1})	0.32	0.06
Utility of second customer tier (λ_{i2})	0.54	0.14
Utility of third customer tier (λ_{i3})	0.75	0.22
Rewarded behavior effect (β_{i1})	0.45	0.60
State dependence (β_{i2})	0.42	0.38

Segment 1, a distinct minority, has a positive utility for upgrades or free flights. Segment 2 has a negative utility. Since the upgrade and free flights

parameters are a reduced form representations of the short-term gain minus the short-term cost of these rewards, this means that many customers find the benefits of cashing in for an upgrade of a free flight are not worth the hassle. This is consistent with many customers failing to cash in points they earn (Abu-Shalback Zid 2004a). This does not mean that Segment 2 will never cash in. First, there is positive carryover from cash-in, so the customer realizes that after they cash in, they will feel good about it. Second, random variation in ε can make cashing in the highest utility decision. For example, the customer's travel agent might notice the customer had a lot of accumulated miles and offer to arrange the cash-in. This is why, in the actual sample, Segment 1 averaged 1.5 cash-ins per year, while Segment 2 averaged 1.1 cash-ins.

The rewarded behavior effect is positive in both segments. The customer tier program provides significant utility for both segments. The results suggest the following interpretation of the segments. Segment 1 is "loyalty program enthusiasts" They like both the short-term "transaction utility" (Thaler 1985) cashing in for upgrades or free flights, as well as getting the more long-term customer tier reward. Segment 2 is more "customer tier focused." They like special treatment but have no need for upgrades or free flights.¹ Obviously, both frequency reward and customer tier programs have a place in the airline industry.

22.2.2.3 Summary: What We Know About How Frequency Reward Programs Affect Sales

The results from laboratory experiments, surveys, descriptive empirical analyses, and estimated dynamic rational models suggest the following:

- Frequency reward programs do increase sales. The level of increase is high enough to make these programs profitable.
- The points pressure effect is strongly confirmed (Lal and Bell 2003; Neslin and Taylor 2005; Lewis 2004; Kopalle et al. 2006) even to the level of detail that the effect is a gradual increase in sales leading toward the goal rather than a step increase (Kivetz et al. 2006; Lewis 2004; Kopalle et al. 2006).
- The majority of evidence supports a rewarded behavior effect (Lal and Bell 2003; Taylor and Neslin 2005; Kopalle et al. 2006; Bolton et al. 2000).² However, the evidence is not uniform (e.g., Kivetz et al. 2006) and may depend on the nature of the reward and the target group (Roehm et al. 2002).

¹ Note one cannot interpret differences in coefficient magnitudes between segments. Obviously, however, the signs of the coefficients can be interpreted and compared. All coefficients in the table are significantly different than zero.

² See also Leenheer et al. (2007), who find a causal relationship between supermarket frequency reward programs and share-of-wallet.

- Various market segments react differently to reward programs. Both the points pressure and rewarded behavior effects seems to be stronger among light users (Lal and Bell 2003; Taylor and Neslin 2005; Roehm et al. 2002). Dynamic structural models reveal two segments – one very much caught up in immediate rewards; the other segment less interested (Lewis 2004; Kopalle et al. 2006).

22.3 Do Frequency Reward Programs Increase Profits in a Competitive Environment?

An area that has received much attention by economists is whether reward programs produce higher profits in a competitive environment. The answer depends on various characteristics of consumer response and on the types of rewards used. We review the several studies on this issue in this section.

Klemperer (1987a, b) identified two forces that influence the profitability of frequency reward programs: (1) they create long-term monopoly power and hence can increase profits, but (2) short-term competition can decrease profits. The monopoly power stems from the points pressure mechanism, which means that customers with points inventory want to repurchase the product. The short-term competition develops because companies compete more strongly to induce customers to make that first purchase so that they begin to accumulate points. Klemperer (1987a) found that the net result of these two forces depended on the functional form of the demand function.

Beggs and Klemperer (1992) expanded this work by examining the impact of a reward program when customers enter and leave the market. They found that prices and profits uniformly increase in the presence of switching costs induced by reward programs. As enunciated by Klemperer (1995), the firm is faced with the question of either increasing prices to exploit the built-in loyalty of its previous customers who have accumulated points, or decreasing price in the battle for new customers so they can be exploited in the future. Klemperer maintains that current-period exploitation dominates because (1) the future benefits of getting customers is discounted by some factor less than 1, and (2) new customers are less price sensitive in the current period because they realize firms can increase prices in future periods.

One limitation of this work is that the existence of the reward program was exogenously specified; not a decision variable for the firm. Caminal and Matutes (1990) consider offering a reward program to be a strategic decision. The authors found that the profitability of reward programs depended on whether they were of the form of either a “coupon,” promising a price discount but not specifying the future price, or a guaranteed future price, e.g., buy one, get one *free*. The authors found that the coupon reward program increases profits, but guaranteed future price decreases profits. The reason is that the

coupon program allows the firm more latitude for increasing price later on. With a pre-committed price, firms lose latitude and this accentuates price competition. In addition, when firms *choose* between a coupon and a pre-committed price reward program, the result was that the pre-committed price was the equilibrium, implying that reward programs decrease profits.

Kim et al. (2001) examined the rationale for reward programs in the context of heavy users versus light users. The presence of the reward program, its form (efficient versus inefficient), and price were the decision variables. Efficient rewards are inexpensive compared to their value to the customer. For example, a free airline ticket costs the firm very little, but is worth the price of a ticket to the customer. Cash, however, is inefficient in that it is worth the same to the customer and the firm. The authors found that the light user segment can serve as the basis for a reward program. The light users pay higher prices and don't earn rewards. In effect, the light users are subsidizing the heavy users. This equilibrium is profitable unless the light user group is price sensitive, in which case firms compete too strongly for this group, driving down prices to the point that reward programs become unprofitable.

Kim et al. (2004) show how frequency reward programs can be profitable if they are used to manage excess capacity. Many firms face seasonal demand (e.g., hotels). Firms can let customers earn reward points during the high-season, to be redeemed during the off-season. This is an important argument since excess capacity is a chronic problem for so many companies. The strategy hinges on the effectiveness of the reward program, i.e., customers must be willing to earn points during the high season and redeem them during the off season. This may be how reward programs work in industries such as hotels, if rooms are only available as rewards during the off season.

Kopalle and Neslin (2003) study the impact of market expandability and consumer valuation of future benefits on the profitability of reward programs. They assume the consumer follows the decision-making process outlined by Equations 22.4–22.5 and that firms decide whether to use reward programs. The authors find that if the market is not expandable, reward programs are not profitable. The reason is that the power of the reward program is directed at stealing market share from competitors, and competitors react strongly. The competitive impact for a reward program is stronger than that of standard promotions because the reward program locks up customers for more than one period. This creates a higher stakes prisoner's dilemma. When the reward program serves to grow the market, it achieves the benefit of a powerful promotion without precipitating strong competitive response. The result is higher profits.

The authors show that frequent flyer programs appeared in the airline industry in 1981, coincident with the rapid entry of “low-frills” competitors. The authors interpret this as the major airlines using reward programs to enlarge *their* market. They also interpret supermarkets' use of reward programs as a way of protecting their market against mass merchandisers such

as Wal-Mart (see Patton 2002). In both of these cases, the outside competitor is unable or unwilling to respond to the reward program. The low-frills airlines did not have the information system to offer reward programs. Wal-Mart, although it certainly has the information system, would find reward programs inconsistent with its positioning of everyday low pricing and simplicity.

Lal and Bell (2003) propose that reward programs eliminate competition between retailers due to cherry-picking. They assume a Hotelling model where two stores differ in which products are higher priced and which are lower priced. The cherry pickers, located roughly in the middle of the Hotelling line, buy the cheapest items at each store. This creates a transportation cost for the cherry pickers, as they are not always shopping at their closest store. If one store implements a reward program, prices generally increase enough to make the program profitable without turning off the cherry-pickers. However, if both stores implement programs the stores compete away these profits. If in addition there are high and low cost shoppers, and both stores offer frequency reward programs, profits can decrease depending on how much each store discounts its lower-priced items. Overall this presents a mixed view of frequency reward programs. If one store offers them and shoppers are not segmented in terms of shopping costs, the reward program can increase profits. But if both stores implement these programs, and if customers are segmented according to their shopping costs, profits can decrease.

Overall, the theoretical evidence is that reward programs can increase profits if they grow the market, “exploit” a price-insensitive light user group, enable firms to pre-commit to a price discount and not to a specific future price, manage excess capacity, or eliminate cherry picking. Certainly there are examples where these conditions are met. The airline industry may be one. Frequent flyer programs potentially exploit a large segment of light users, and might grow the market. Frequent flyer programs persist today even though the mass influx of new airlines has slackened, and even the lower tier airlines can and do offer reward programs (e.g., Goetzl 2000). However, it is possible that the frequent flyer programs are not growing the market in today’s competitive environment, and in today’s market the infrequent traveler might in fact be very price sensitive, in which case reward programs may be decreasing industry profits.

While empirical evidence has demonstrated that reward programs increase firm profits on a marginal basis (Lal and Bell 2003; Taylor and Neslin 2005; Leenheer et al. 2007), empirical work is needed to see if they increase profits in a competitive industry. Indeed, competitors’ use of reward programs has been found to be one of the most important determinants of a firm’s decision to adopt a reward program (Leenheer and Bijmolt 2003).

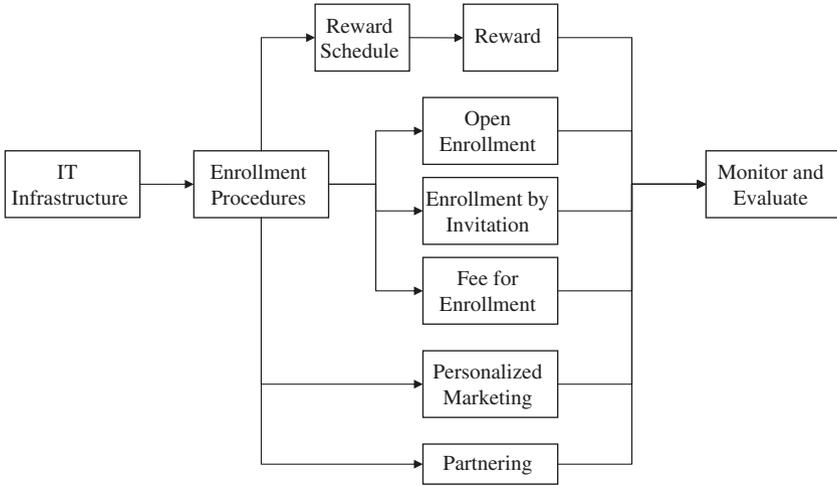


Fig. 22.4 Frequency reward program design considerations.

22.4 Frequency Reward Program Design

22.4.1 Design Decisions

Figure 22.4 depicts the design decisions involved with a frequency reward program. Infrastructure involves the information system for collecting data and awarding rewards. Enrollment procedures, reward schedules, and the reward are fundamental decisions that define the reward program. Partnering has also become a crucial element. Personalized marketing and partnering are potential add-ons. Finally, the program must be monitored and evaluated.

22.4.2 Infrastructure

A reward program must have an infrastructure for collecting customer data, determining which customers get which rewards, and then delivering the rewards. In a small shop this can be done by hand – the proprietor keeps a list of customers and updates their totals whenever the customer makes a purchase. However, when the customer counts reach into the thousands if not millions, major investments are needed in computer information systems.

These investments are not trivial. Even airlines, despite their large information systems, still need to make sure they avoid mishaps such as assigning top passengers to middle seats (Feldman 2002). Information system upgrades range from \$5 to \$10 million for airlines (Feldman 2002). Another company

that one would expect to have ample information technology is American Express. However, the company cited higher costs to run loyalty programs as a cost factor that became a drag on earnings (Lee 2002).

22.4.3 Enrollment Procedures

The three alternatives are: (1) open enrollment, (2) fee for enrollment, and (3) enrollment only by invitation. Most frequency reward programs are open enrollment. Companies are eager to get customers onto their frequency reward program so they can reap the benefits of increased sales, etc. The obvious benefit of charging a fee is that it separates out the best from the worst customers, thus allowing the company to reward its best customers. The cost is the mixed message it gives customers: “We want to reward our best customers, so please pay us \$25 and we’ll reward you!”

Another way to separate out best from worst customers is to make the reward program available only by invitation. Viking Office Products, for example, will not reveal the criteria for being included in their program and does not promote their program (Miller 2001). The benefit of this approach is one can limit the size (and expense) of the program and use complex criteria (previous sales, future potential, etc.) to select members for the program. The downside is the limited opportunity for the reward program to attract new customers.

22.4.4 Reward Schedule

The reward schedule maps customer purchases to rewards. There are two dimensions that define reward schedules: linearity and continuity. Continuity refers to whether the customer earns rewards after each purchase or only after reaching a threshold. Linearity pertains to whether the relationship between purchases and rewards is the same no matter what the level of purchasing. Figure 22.5 illustrates.

The linear continuous schedule is very common. Many credit cards reward customers one frequent flyer mile for each dollar spent. The advantage of this schedule is its simplicity. It doesn’t require the customer to “jump through hoops” to receive a reward, and today’s consumers don’t like to wait (Abu-Shalback Zid 2004b). Another advantage is that it continually rewards behavior, and behavioral learning suggests continual reinforcement is most effective (Rothschild and Gaidis 1981). However, linear continuous rewards do not create points pressure. Consider a credit card program that offers 1 mile per dollar spent. If a competitor has the same reward program, the customer does not lose anything by switching back and forth between the

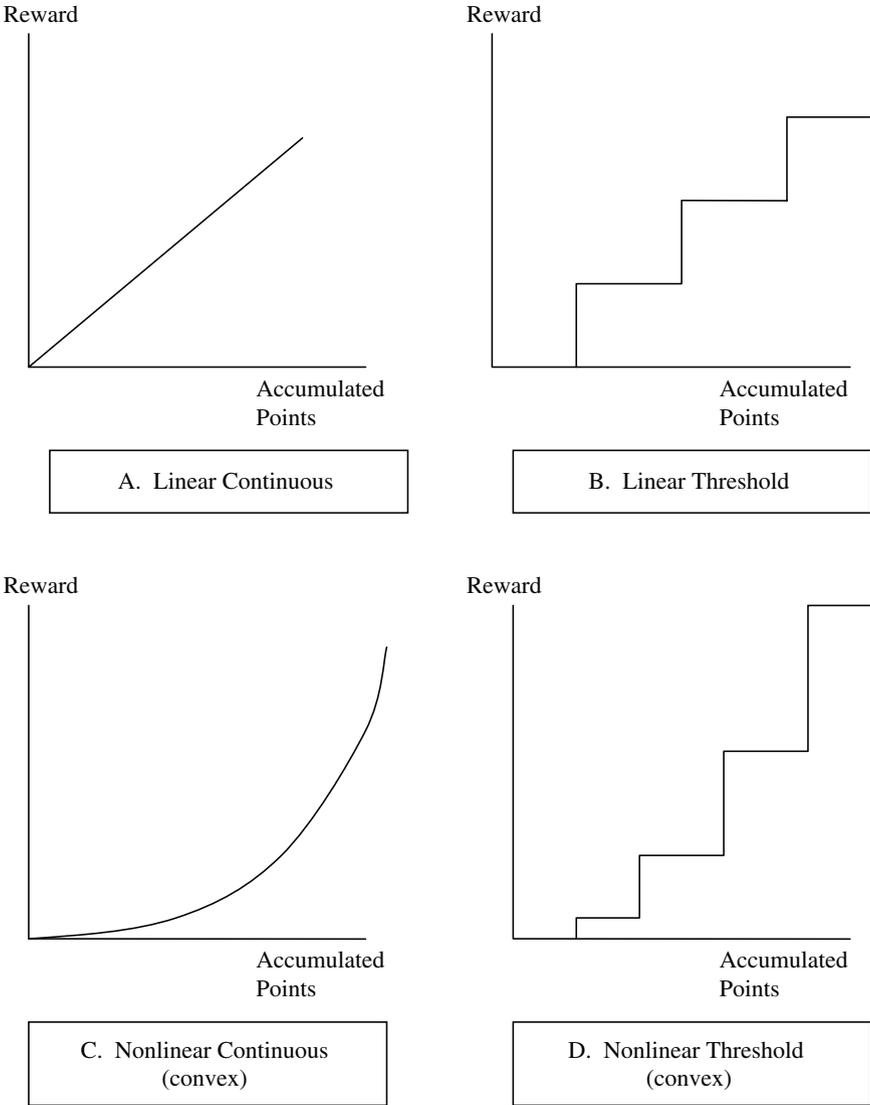


Fig. 22.5 Reward schedules.

credit cards. This nullifies points pressure. This might be addressed by using a continuous but convex reward schedule.

Yi and Jeon (2003) suggest that a linear continuous schedule may be better for low-involvement services. This makes sense as it would appear that high involvement is a necessary ingredient for customers to (implicitly) think through the dynamic rational model (Sect. 22.2.2.2) needed to make decisions under a nonlinear or threshold schedule. In addition, Keh and Lee (2006) find

Table 22.2 Convex (Schedule A) versus linear (Schedule B) frequency reward structures

Points requirement	Schedule A reward	Schedule B reward
25,000	\$100	\$100
50,000	\$200	\$200
75,000	\$400	\$300

that dissatisfied customers prefer constant rewarding, whereas satisfied customers prefer a delayed reward. In summary, despite their disadvantages, linear continuous reward schedules may work best for low involvement services or with dissatisfied customers.

Threshold reward structures set requirement levels at which point the customer receives a reward. For example, 25,000 miles of travel may be required for a free upgrade; a supermarket may reward a free turkey for customers who spend at least \$500 over a 2-month period. The prime advantage of thresholds is they create points pressure. The customer with 24,000 miles on Airline ABC delays an upgrade by switching to Airline XYZ. The disadvantage of thresholds is that behavior is rewarded infrequently, and consumers may focus on the reward rather than on the brand itself. In addition, the design of the exact stepped structure is a challenge. The requirement levels should be set high enough to encourage more frequent purchasing, but not so stringent that customers don't see the reward as a real possibility.

The thresholds can be linear, concave, or convex. For example, Schedules A and B in Table 22.2 both have thresholds, but A is convex while Schedule B is linear. Which program is preferred would depend on customer response, which could be measured using a dynamic rational model. However, it would also depend on competition, so the complete analysis would involve firm's competing for dynamic rational customers using linear versus convex threshold reward schedules.

A related issue is whether points should ever expire. Expiration creates more point pressure (Kopalle and Neslin 2003) and can be used to price discriminate against light users (the infrequent flyer may finally build up enough miles to earn a free flight, but by the time he or she is ready to fly again, the reward has expired). However, expiration dates are more expensive to handle from a data processing perspective, and can cause considerable frustration and "turn-off" for customers.

Another aspect of the reward schedule is how clearly it should be stipulated to the customer. It seems that one would want to make the reward schedule clear. For example, Cunningham (2002), states, "the sheer complexity of points-based schemes can be a turn-off. Many people with loyalty cards do not take advantage of the offers because they do not have the time, or the will, to work out what they are entitled to". Complexity could decrease points pressure because the customer doesn't understand how much he or she needs to accumulate to get a reward. However, there may two potential benefits of complexity: First, the reward serves as a discrimination device,

identifying those who really care about rewards (and will do the homework required to figure out the program). Second, complexity could increase the rewarded behavior effect: (1) The higher effort required to get the reward creates a cognitive dissonance that is resolved in favor of the brand – i.e., the customer asks, “Why did I go through so much trouble for this reward,” and answers, “I must really like this brand” (Dodson et al. 1978). (2) If the customer is unclear of the reward structure, the reward may come as a surprise, surpassing customer expectations, and “delighting” the customer (Rust and Oliver 2000).

22.4.5 The Reward

There are several attributes of the reward that should be considered:

Immediacy: This pertains to the length of time between when the reward is earned and when it is delivered. Smart cards, for example, can let customers know immediately how many points they have accumulated, and the points can be cashed in instantly (Kuchinskas 2000; Lucas 2002; Pepe 2002). The advantage of immediacy is that coupling the reward with the correct behavior creates powerful reinforcement (Rothschild and Gaidis 1981), enhancing the rewarded-behavior effect. However, companies may decide to delay the reward, even if they have the capability to deliver it immediately (Kuchinskas 2000). Decoupling of the reward from the behavior might encourage the customer to focus on the brand experience instead of the reward.

Direct vs. Indirect: Direct rewards are the firm’s product or a similar product. Indirect rewards are a different product or cash. Yi and Jeon (2003) found that for high involvement services (in their case, a hair salon) direct rewards were more effective than indirect rewards, whereas for low involvement services (in their case, a fried chicken restaurant), direct awards provided no advantage over indirect awards. One important form of indirect reward is cash. The advantage of cash is the flexibility it provides to the customer – the cash can be used in any way. Kim et al. (2001) advocate cash because it forces firms to commit to high prices. However, cash is expensive for the company. A cell-phone company can award a new cell-phone that is worth \$200 to the customer because that is how much he or she would have to pay for it at retail. However, the company pays perhaps only \$50. This is obviously less expensive than awarding \$200 cash. In addition, non-cash indirect rewards provide utility in the form of pride or prestige. For example, the salesperson who has earned a unique sports jacket with an insignia can wear that jacket with pride and enjoy fielding questions from peers as to how he or she obtained it (Renk 2002).

Price Discount vs. Pre-Committed Price: As mentioned earlier, Caminal and Matutes (1990) conclude that price discount rewards are profitable, whereas pre-committed prices tie the hands of firms and increase

competition. In their analysis, pre-committed prices were the equilibrium if firms could choose between the two, which suggests price discount rewards should be less prevalent than a pre-committed price. Indeed, the common buy-x-get-one-free nature of many programs would be a prime example of a pre-committed price. However, many retailer reward programs offer discount coupons. The key is to separate the reward from the price of the product. Premiums would do this just as well as a coupon, and are also common rewards.

Monetary Value: A key decision is the monetary value of the reward to the consumer. For example, a credit card offered one point for every dollar spent, and 5,000 points earned a \$50 gift certificate (Polaniecki 2001). That amounts to a 1-cent reward for every dollar spent, a 1% discount. Southwest offered a free flight after 8 round trips (Goetzl 2000). That amounts to a 12.5% discount per flight. A high monetary reward coupled with a reasonably tough reward schedule (e.g., 8 flights for a free flight) can create a lot of points pressure. A higher reward might enhance the rewarded behavior mechanism, since a stronger reward should precipitate more affect toward the firm. However, if the reward is too strong, it becomes the “primary reinforcement” (Rothschild and Gaidis 1981; Rothschild 1987) and reinforces the use of reward programs rather than the use of the product.

Reward Attractiveness: This is similar to the question of monetary value of the reward, but applicable to premiums and other rewards that do not have clear monetary value. The same concern as with monetary rewards exists in that too strong a reward diverts attention from the brand. One has to be careful to anticipate the attractiveness of the reward in advance. Procter & Gamble offered Fisher-Price toys as a reward for frequent purchasers of its diapers. P&G under-estimated consumer response to the program, and months after the end of the program, it still hadn't been able to fulfill all the claims for the toys (Estell 2002a). This problem can be avoided with laboratory pre-testing such as the work by Roehm et al. (2002) and Kivetz and Simonson (2002).

Multiple vs. Single Rewards: The provision of a menu of rewards allows customers to select the reward they like the best, thereby increasing customer utility (Lucas 2002). The main disadvantage of multiple rewards is cost and complexity. For example, there is the question of whether the rewards should be of equal value to the firm or equal value to customers on average, or both. An airline may award a free seat (which costs nothing if it is available) or a travel bag (which has a cost). If too many customers select the travel bag, the reward program could be unprofitable.

Luxury vs. Utilitarian Rewards: Kivetz and Simonson (2002) found that consumers prefer luxurious rewards (facial massage, wine, jewelry, etc.) when program requirements are more stringent. Stringency assuages guilt customers might feel in receiving a luxurious reward. This suggests that if a company uses a convex, threshold reward structure, luxury products should be used for the higher thresholds, while utilitarian products can be used for

the lower thresholds. The complication is customer heterogeneity. Customer A may consider eight flights to be an incredible effort, worthy of a luxury reward. While Customer B may average eight flights a month, so requires little effort. Still, the results suggest that if the program is targeting low-sales customers and sets up a high hurdle, it may be better to use a luxurious reward, since these customers have to work hard to achieve the reward.

Link to Brand Positioning: Roehm et al. (2002) suggest rewards should reinforce product positioning, and that the best way to do it is through intangible rather than tangible rewards. Access to a website with games and contests that reinforce the brand's positioning is better for loyalty than a premium that also reinforced the positioning but is more tangible. The problem with intangible rewards is that they may not create enough points pressure – the prospect of access to a website may not be a strong enough reward to create points pressure (although recall that in Roehm et al.'s experiment, all rewards created the same number of purchases). This means that the gains in retention from the program have to come entirely through the rewarded behavior effect. There are many examples of reward programs that link well to product associations. For example, AT&T Broadband, a television cable company, used premiums such as a Disney watch or a Showtime director's chair (Beeler 2000). Nantucket Nectars, a beverage company, capitalized on their bottle caps that contained clever sayings by using caps as points for the reward program (Estell 2002b).

22.4.6 Personalized Marketing

A major opportunity is to use reward programs to facilitate personalized marketing efforts, especially cross-selling. Cross-selling efforts work best when companies have data on customer purchase history (Knott et al. 2002), and such data are collected via reward programs. As mentioned earlier, Kivetz et al. (2006) found that customers were more receptive to special offers as they got closer to earning a reward. This finding was exploratory, but suggests how frequency reward programs and personalized marketing can be integrated.

It appears that personalized marketing is an under-utilized opportunity for reward programs. Many supermarkets collect purchase data through their frequent shopper programs, but few use those data to personalize marketing campaigns. There are at least three possible reasons for this. First, personalized marketing requires data analysis and supermarkets are not ready to undertake this task (see Leenheer and Bijmolt 2003). Second is privacy. Some shoppers do not want supermarkets scouring their purchase records, looking for "interesting" patterns (Weir 1999). In fact, laws have been proposed that would allow shoppers to forbid supermarkets from *collecting* data on their

buying habits (Weir 1999). Third, supermarkets may be wary of customer dissatisfaction that could result from tailored marketing campaigns. This is because the personalized marketing would take the form of targeted price discounts. Customer A could become jealous if Customer B receives a coupon for a product Customer A wants to buy (Feinberg et al. 2002).

22.4.7 Partnering

An important trend in reward program design is partnering. Partnering takes two forms: “earn” partners and “burn” partners. If members of Company A’s reward program can earn points by purchasing at Company B, Company B is Company A’s earn partner. If Company A’s customers can spend their points by getting discounted or free merchandise at Company B, Company B is Company A’s burn partner. Obviously, Company B could be both an earn and a burn partner for Company A.

The advantage of partnering is that it makes the program more attractive to customers. To the extent that Company A has many earn partners, it is easier for customers to build up their point totals. To the extent that Company A has many burn partners, the rewards are more attractive. Firms have to be careful, however. Sharp and Sharp (1997) examine a program in Australia that involved several retail outlets that served as earn and burn partners for each other. They found that repeat purchasing did not increase at most of the retailers – so many stores were partners that customers did not have to change their purchase habits in order to earn many points.

Another noteworthy partnering arrangement was the AOL AAdvantage Reward Program (Direct Marketing 2000). In this partnership, AOL and American Airlines were earn and burn partners for each other (Regan 2000). AOL allowed its customers to earn miles when they shopped at AOL shopping affiliates, as well as when they flew on American Airlines. The miles could be spent on merchandise or free trips. The program was launched with much fanfare in 2000, but in January 2002, the program discontinued the key feature of allowing customers to earn miles when shopping with AOL affiliates. The burn feature of the program, redeeming miles for merchandise at affiliates, remained, but the earning miles feature has been dropped. What’s more, the financial value of using miles for merchandise has been called to question. One report is that 50,000 AOL AAdvantage points could be used for \$187 in books, but those same miles could be used for a \$500 airplane flight (Drucker 2002).

Despite examples of mixed success, partnering has become a common feature of many reward programs. Many credit cards have programs with multiple burn partners, for example airlines, retailers, and travel agents (Polaniecki 2001).

22.4.8 Monitor and Evaluate

As with any marketing program, it is crucial to monitor and evaluate reward programs. This can be difficult because its desired impact is long-term. Statistical methods (e.g., Sect. 22.2.2.2 and Leenheer et al. 2007) are capable of measuring the points pressure and rewarded behavior mechanisms as well as the overall impact. Another way to evaluate reward programs is through surveys (Bolton et al. 2000), especially if they can be conducted on a before and after basis.

22.5 Frequency Reward Program Examples

22.5.1 Harrah's Entertainment³

Harrah's Entertainment has a highly successful reward program that includes both a frequency reward program and personalized marketing. Harrah's operates 21 gaming casinos across the USA, serving 19 million customers annually. Accordingly, Harrah's needs a large database to store customer information. The data are stored in a Patron Database (PDB) and a Marketing Work Bench (MWB). The PDB houses "raw" data and can be used for simple look-up. The MWB includes data in a form that can be analyzed for personalized marketing efforts. Total size of the databases is 300 GB, updated on a daily basis. The data are at the customer level, and include hotel stay, demographics, customer preferences, event attendance, games played, and marketing offers received and responded to. The data are integrated at the customer level across all Harrah's locations. This is truly a "single view of the customer" (Chapter 25).

Customers accumulate points by using Harrah's facilities, which can be cashed in for merchandise, rooms, food, and cash. This is implemented through the PDB database. Harrah's also uses the MWB database to design various personalized campaigns. These campaigns consist of offers, complimentary passes, etc. The campaigns can be targeted at individuals based on predictive modeling. For example, Harrah's identified consumers with high potential for increased visits based on their residence and gambling "velocity," (literally how fast they play the slot machines). These customers were then mailed offers for cash and food, and their visit frequency increased from 1.1 to 1.4 per month.

The Harrah's reward program is integrated with Harrah's customer tier program. Membership in the different tiers is based on accumulated points thresholds, and each tier has specific privileges such as priority reservations, seating, and check-in.

³ Much of the material in this section is based on Watson and Volonino (2001). See also Watson and Eckerson (2000).

Overall, Harrah's estimates a 72% internal rate of return on their investment in information technology. They report strong increases in visit frequency, profit, and cross-market facility utilization. The Harrah's implementation is a well-publicized "success story" (Swift 2001; Davis 2001; Koller 2001; Rosen 2000; Heun 2000). Two aspects of the program are noteworthy. First, the data are very complete and provide a detailed profile of each customer. Second, the program employs points rewards *and* personalized marketing. The increase in customer profit undoubtedly comes through all three mechanisms in Fig. 22.1: points pressure, rewarded behavior, and personalized marketing.

22.5.2 The UK Supermarket Industry: Nectar Versus Clubcard

The UK supermarket industry has witnessed an intense "tit-for-tat" competition with reward programs. Tesco began the competition in 1995 (Gofton 1998). At the time it was second to Sainsbury's, and wished to retain customers and increase sales. Its "ClubCard" awarded 1 point per British pound spent, and 1 point was worth 1 pence in discounts at Tesco. This equates to a 1% discount. With profit margins of 5–7%, this represented a significant reduction in margin for Tesco (Croft 1995). The question is whether this paid off. Tesco claimed in 1995 to have overtaken Sainsbury's, and that the ClubCard was a major factor (Croft 1995).

There were two issues facing Tesco at this point. First was how Sainsbury's would respond. Indeed Sainsbury's did respond with its Saver Card, a frequency reward program with stepped rewards. The maximum number of points that could be accumulated on one card was 5,000, representing £2,000 purchases, which could be traded in for a discount of £50, a 2.5% savings. However, by 1998 they had also introduced an ongoing Reward Card that rivaled ClubCard (Marketing Week 1998). Second was whether Tesco would be able to utilize the data collected to develop personalized marketing programs that would earn additional profits. This was a crucial step. With a sacrifice of 1% margin on a base of 6% margin, Tesco would have to generate 20% additional revenues in order to break even.⁴

Indeed, Tesco set out to utilize the purchase data collected from its ClubCard members to tailor coupons and other discounts to shoppers. Each quarter ClubCard mailed a communication to members that including points totals, vouchers for redeeming points at Tesco, and targeted coupons. Tesco reportedly creates 100,000 personalized versions of these targeted coupons,

⁴ Let R = base revenues, m = base margin (6%), and δ = incremental revenues. Baseline profits is therefore Rm , and profits under the reward program are $R(m - 0.01) + R\delta(m - 0.01)$. In order to break even, we must have $R(m - 0.01) + R\delta(m - 0.01) > Rm$, or $\delta > (m/(m - 0.01)) - 1 = 0.20$, or 20%.

based on its analyses of purchase data (Gofton 1998). In 1998, Tesco reported a 1/3 increase in its loyal customer base. It also reported coupon redemption rates upwards of 30% (Gofton 1998). All indications are that the personalized marketing aspect of ClubCard has been highly successful, even in competing with Wal-Mart (Rohwedder 2006).

Over time, Tesco also expanded its card to bring in more earn and burn partners. By 2002, earn partners included Alders retail, Beefeater restaurants, H. Samuel jewelry, and several auto and tire centers. Burn partners included an almost bewildering array of recreation and travel opportunities, listed at the Tesco website (www.tesco.co.uk). Another recent burn partner is Air Miles, an air travel rewards program.

While Sainsbury's retained its own reward program during this period, it was not apparently as comprehensive as Tesco's. That changed in late 2002 with the launch of Nectar. Nectar is a reward program representing a consortium of four companies – Sainsbury's, BP gasoline, Debenhams retail stores, and Barclaycard credit card (Kleinman 2002). The earn partners consist of the four partners. The burn partners include Sainsbury's, McDonalds, Blockbuster, and additional companies. In 2002, Nectar claimed 12 million members, compared to 10 million for Tesco (Rogers 2002).

Two fascinating questions emerge: First, how will the programs evolve from this point? One intriguing possibility is that they will become branded products, each consisting of sets of earn and burn partners and various points schemes. Now, the main players in the UK market are Nectar, Air Miles (a frequent flyer program for which Tesco is an earn partner), Tesco's ClubCard, and Boots Advantage. There are other questions for the future. For example, will rewards increase above the 1% of revenues mark that has become more or less the standard?

Second and most importantly, will these programs raise profits for member companies? The sacrifice in margin requires either a substantial gain in revenues or an increase in profits through personalized marketing. Revenues could also increase if Tesco could raise prices. This would be the route to profits suggested by Klemperer and others. Also, heavy users build up more points and therefore have opportunities for burning their points with airlines and other partners that have minimum requirements. So we have the conditions in Kim et al. (2001) where the heavy user segment receives more benefits than the light user segment. As long as the light users are not price sensitive (see Kim and Rossi (1994) for an example where heavier users were more price sensitive than light users), the outcome could be profitable. The problem is that the light users may be price sensitive. First, they may be lower income individuals. Second they may be store-switchers. The important point is that Klemperer and his colleagues' theories are applicable toward thinking about and assessing the profitability of these reward programs.

In addition, the rewards at least for Tesco take the form of cash reductions (i.e., coupons) rather than guaranteed prices. Therefore Caminal and

Matutes' (1990) recommendations also apply. However, the market growth condition of Kopalle and Neslin (2003) would not appear to apply. The supermarket industry in the UK is relatively mature. The other route to profitability is personalized marketing. Tesco's reports suggest that this has been very helpful for them. In any case, the UK supermarket industry provides a fascinating glimpse of the evolution of reward program competition and the emergence of a reward program "industry."

22.5.3 Cingular Rollover Minutes

The wireless telephone industry is concerned about customer churn. Cingular created a reward program that allows customers to "roll over" unused minutes from the previous month into the current month (Thomaselli 2002). This is different from the standard reward program in that points accumulate based on *lower* usage.

Compete, Inc. reported that the program had strong short-term impact reducing churn (Business Wire 2002). However, the impact wore away quickly. Compete collects its data by observing web visit behavior. It measured potential churn as the percentage of a firm's current customers who investigate competitive offerings. Compete found that when Cingular launched its program, potential churn increased by 9–27% among Cingular's competitors, but remained level for Cingular. However, a month or so later, Cingular's potential churn increased by 11% while its competitors' improved. Compete does not have access to actual churn numbers, relying instead on customer search behavior as an indicator. The indication is that Cingular's program caused many competitors' customers to investigate Cingular, while retaining Cingular's current customers. However, things leveled off after the initial advertising campaign.

It is still possible of course that Rollover Minutes paid off. In any case, it is an interesting reward structure that could be used in various subscription services. However, perhaps Rollover Minutes was "merely" a promotion that had a positive short-term effect but little long-term impact. Part of the problem might be the lack of a true points-pressure effect. It seems self-defeating to motivate customers not to use the service now so they can use it more in the future. The points-pressure effect is to encourage *more* use now so customers can get a reward in the future.

22.5.4 Hilton Hotels

The hotel industry is a repeat-purchase business with chronic customer loyalty concerns. This makes it ripe for reward programs. Hilton Hotels offers

Hilton HHonors, open to anyone at no charge (Deighton 2000; Bell et al. 2002). Members earn points by staying at Hilton Hotels, but also through earn partners including various airlines, FTD Florists, and Mrs. Field's Cookies. The points can be used for upgrades and stays at Hilton Hotels, or burned at partners or converted to frequent flyer miles for free flights.

Points are earned at a rate of 10 points per dollar. The dollar value of the points depends on hotel and room availability. For example, 5,000 points could earn 50% off a \$128 room at the Hilton Albuquerque. That equates to a 13% discount off a \$500 investment. In contrast, 25,000 points could earn a free weekend night at a \$239 Hilton Boston Back Bay hotel. That is a 9% discount off a \$2,500 investment.

The program is integrated with Hilton's customer tier program: Blue, Silver, Gold, and Diamond. Membership is based on the number of stays in the most recent year. Benefits differ in terms of points earned per dollar and special upgrade offers. An interesting aspect of the Diamond tier is the lack of specificity in rewards. As reported by Deighton (2000), the strategy was to lower Diamond members' expectations and then "over-deliver" ("delight the customer" (Rust and Oliver 2000)).

Hilton also delivers personalized offers through HHonors (Orr 2000). For example, in any given month, the company has 80–130 potential messages/offers to be included in the customer's statement. The company selects about 14 of these messages based on variables such as previous stays, credit card usage, and demographics.

Hilton had specific motivations in starting HHonors. First was yield management. Hilton often has a pool of unsold rooms. The HHonors program provided the data that produced personalized offers for customers to fill the rooms. Second was to use partners to broaden the rewards it could provide to its customers. Third was to use it as an inducement for hotel owners to become Hilton franchisees – the benefit being access to reward program members who would want to burn points at the franchisee's hotel. Fourth was the more traditional motivation for reward programs – to build relationships with customers. For example, they can cross-sell travel opportunities, and provide better service by knowing customers' preferences for rooms, etc., when the customer checks in.

Its conjoint studies suggested that one in five stays by HHonors was due to membership in the program. These incremental stays are being given away at a steep discount, but travelers spend additional money besides the room when they stay at a hotel. Plus there are potential benefits due to rewarded-behavior and personalized marketing that might increase the visit frequency. Having said this, the company admitted that annual churn among HHonors Diamond members is 40% (Higley 2002). This does not say what churn would have been without the program, but this high number suggests that the industry is very competitive and reward programs do not completely lock in customers.

22.6 Research Needs

Reward programs are a popular form of customer management. There are a number of studies on reward programs but there is much more to be done. There are three basic areas: (1) How do customers respond to reward programs? (2) Are reward programs profitable? and (3) What are the best ways to design reward programs?

Regarding customer response, we have strong evidence that points pressure is a real phenomenon (Sect. 22.2.2), however, more work is needed to pinpoint the existence and impact of the rewarded behavior phenomenon, and the conditions under which it is most prevalent. In short, we need to know whether frequency reward programs, sometimes called loyalty programs, really create loyalty (see Shugan 2005).

We have good theories suggesting the conditions under which reward programs will be profitable in competitive equilibrium. However, these theories consider the points pressure mechanism (customers purchase more to receive the reward) and do not explicitly consider the rewarded behavior and personalized marketing mechanisms. In addition, we need empirical studies to discern the impact of reward programs on profits. The UK supermarket industry reviewed earlier would be a perfect venue.

Finally, there are several aspects of program design that need to be investigated. The overall need is to determine what types of designs work best under what circumstances. When should companies use continuous versus threshold rewards, or convex versus concave versus linear? How instantaneously should rewards be delivered? How transparent should the rules be for earning and burning points? How should earn and burn partners be selected, and what should be the financial arrangement between earn and burn partners? Is the best organizational set-up a company-run program with various earn and burn partners, or a consortium of companies banding together with the program managed by an outsider? Another fascinating issue is the frequent flyer points phenomenon. Frequent flyer points have become perhaps the most prevalent reward used by credit cards and even supermarket programs (see Tesco/Sainsbury example above). Research is needed to understand exactly what it is about this reward that makes it so attractive. A final issue is reward program branding – e.g., Tesco seems to have branded its ClubCard program. The methods to do this, and its benefits, need to be better understood.