XAVIER DRÈZE and JOSEPH C. NUNES*

The authors examine the impact of successfully attaining a goal on future effort directed at attaining the same goal. Using data from a major frequent-flier program, they demonstrate empirically how success contributes to an increase in effort exhibited in consecutive attempts to reach a goal. They replicate the effects in a laboratory study that shows that the impact of success is significant only when the goal is challenging. They also show how progress enhances perceptions of self-efficacy and how successfully completing the task provides an added boost, supporting the notion that self-learning is the principle mechanism driving their results.

Keywords: loyalty programs, rewards, goals, learning, self-efficacy

Recurring Goals and Learning: The Impact of Successful Reward Attainment on Purchase Behavior

Consumers possess voracious appetites for rewards. The enticement of earning a reward and its impact on purchase behavior has led U.S. companies to create more than 2000 loyalty programs (Berman 2006). According to a loyalty census by research firm Colloquy, more consumers are engaging with brands through rewards programs than ever before. In 2009, approximately 1.8 billion memberships were on record in the United States alone, up from 1.3 billion in 2006 (Ferguson and Hlavinka 2009). Furthermore, the loyalty census reports that households actively participate in an average of 6.2 programs, up from 4.7 in 2007. Ferguson and Hlavinka (2009) stress the immense influence of rewards programs on business, noting that rewards have become "the cost of entry for new credit and debit products" (p. 6), the hotel industry "is united in its belief that reward programs continue to offer the best platform for retaining profitable customers" (p. 9), and specialty retailers are fighting Wal-Mart and other discounters through the use of successful reward programs, such as Best Buy's renowned Reward Zone program (p. 11).

In general, research supports a positive connection between loyalty programs built on rewards and customer retention (Lewis 2004; Nunes and Drèze 2006a, b; Taylor and Neslin 2005). However, the efficacy of rewards programs has its critics. Shugan (2005) argues that many programs trade short-term revenues (purchases) for long-term liabilities (promises of future rewards) rather than invest in the customer. Moreover, several empirical studies suggest that loyalty programs generate small effects (Verhoef 2003) or no effect (DeWulf, Odekerken-Schroder, and Iacobucci 2001; Mägi 2003) on purchase behavior. Our research approaches the question of impact with respect to loyalty programs and their rewards differently from previous research. Rather than model the impact of membership in a firm's loyalty program on share of wallet, we examine the impact of successfully attaining a reward in a loyalty program on future effort directed at attaining the same reward in the same program.

Most loyalty programs comprise rewards that can be earned repeatedly, such as free tickets on an airline, free nights at a hotel, or discount certificates at various retailers. We refer to these rewards as "recurring goals" because consumers often work toward a goal (reward) that, once attained, continues to serve as a goal toward which they can work again. It is not clear from previous research what impact successfully earning a reward would have on consumers' motivation to direct future purchases toward the reward-granting firm. Kivetz, Urminsky, and Zheng's (2006) work provides a starting point. They conducted a field study

*Xavier Drèze is Associate Professor of Marketing, Anderson School of Management, University of California, Los Angeles (e-mail: xavier.dreze@anderson.ucla.edu). Joseph C. Nunes is Associate Professor of Marketing, Marshall School of Business, University of Southern California, Los Angeles (e-mail: jnunes@ marshall.usc.edu). Both authors contributed equally and are listed in alphabetical order. Fred Feinberg served as associate editor for this article.

ISSN: 0022-2437 (print), 1547-7193 (electronic)

at a university café in which participating customers were required to make ten coffee purchases to earn a free coffee. For customers who earned the first reward, it took only 2.1 days between the ninth and tenth purchases—the final purchases necessary for the first reward—compared with the 3.1 days it took between the first and second purchases toward the second reward. The authors dub this slowdown in the interpurchase times "postreward resetting."

Kivetz, Urminsky, and Zheng (2006) argue that this deceleration rules out learning as an explanation for the acceleration in purchases (goal gradient) they observed as consumers approached the reward: A deceleration from 2.1 to 3.1 days would imply that whatever was learned had been suddenly forgotten. Figure 1 shows four types of consumer learning. Figure 1, Panel A, portrays a situation in which the goal gradient is the result of procedural learning (e.g., how many purchases are required). In this figure, effort increases monotonically over time; even after the consumer passes the goal, learning continues. Figure 1, Panel B, illustrates the opposite—that is, an absence of learning. As soon as the consumer achieves the goal, effort reverts to its original rate and increases exactly as it had before. Although a postreward reset suggests that the goal gradient effect is not due to procedural learning, it does not preclude other forms of learning. Consistent with prior work, we expect consumers to accelerate purchases as they near a reward, consistent with the goal gradient phenomenon. We state this formally as follows:

H_1: The effort a consumer exerts toward reaching a goal increases as he or she progresses toward the goal.

If progress toward the goal drives effort, when the goal is reached, we might expect a postreward reset and deceleration in purchases after the consumer attains the reward. However, if learning occurs while striving for or when reaching the goal (we provide an in-depth discussion of types of learning in the subsequent paragraph), we would expect this knowledge to affect behavior. We posit that successfully reaching a goal results in a reassessment of the likelihood of success and that this learning manifests as observable increases in effort when reengagement occurs and in successive attempts at reaching recurring goals. Although Kivetz, Urminsky, and Zheng (2006, p. 47) note a "markedly similar acceleration pattern" whether participants were working toward their first or second free coffee, the interpurchase times reported after customers reengaged were slightly lower than those observed the first time (3.1 < 3.2; 2.7 < 2.8), consistent with a partial postreward reset. Kivetz, Urminsky, and Zheng do not test for significance, because the extent of resetting was not their concern. Any boost in effort exhibited toward a recurring goal would undoubtedly be of interest to firms offering loyalty programs. We posit that consumers learn and expect them to exhibit partial postreward resetting rather than full postreward resetting. This leads to the following hypotheses:

H_{2a}: After a consumer reaches a goal, there is a decline in effort when initiating a successive attempt to achieve the same goal (i.e., a postgoal reset).

H_{2b}: The effort exerted initially toward achieving a recurring goal is greater following successful goal attainment than it was previously (i.e., a partial postgoal reset).

Figure 1
TYPES OF LEARNING

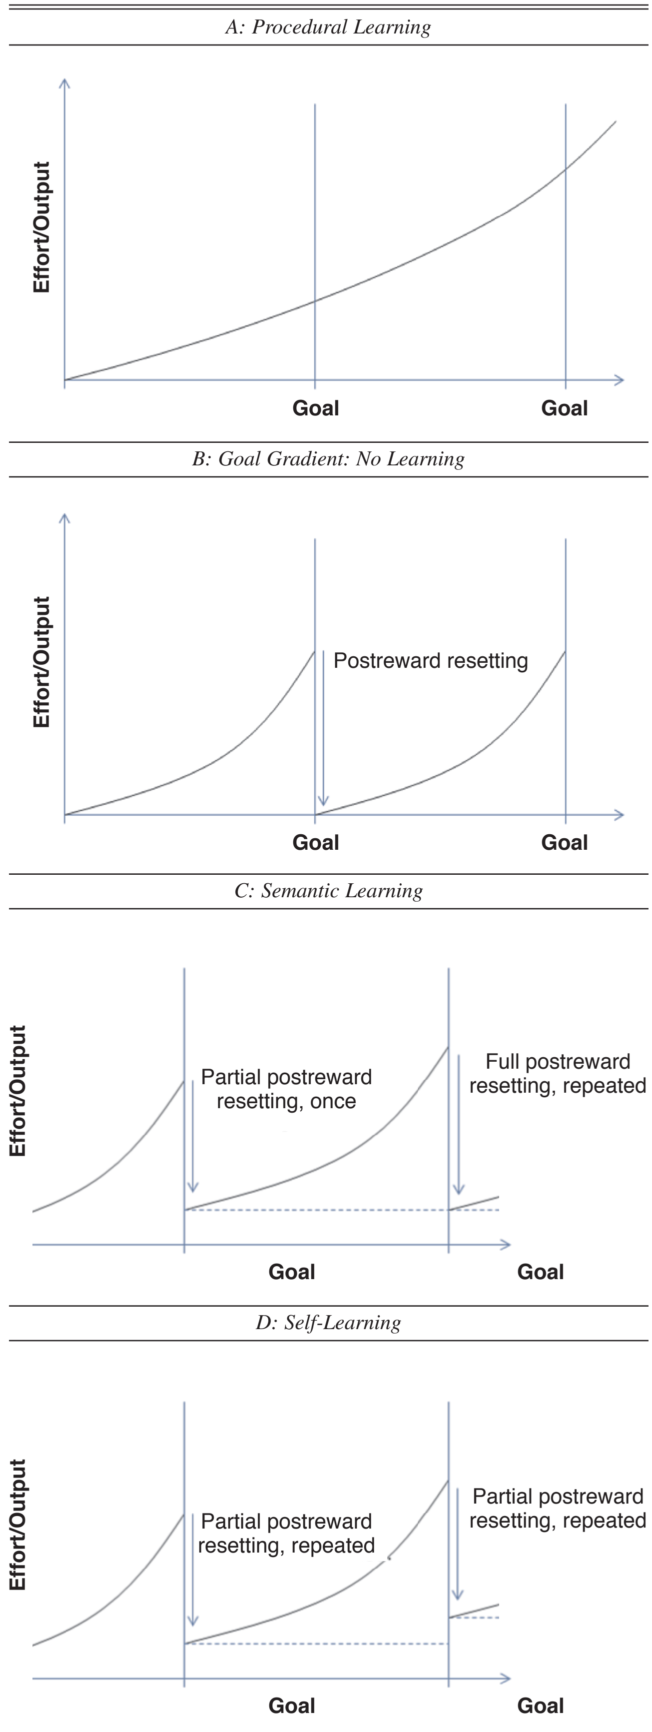

A partial reset could occur for two reasons. First, Figure 1, Panel C, illustrates the case of learning from experience (Hoch and Ha 1986). This occurs when the benefits associated with attaining a goal (reward) are vague until experienced. Consider airlines that provide premier status. Status comes with benefits such as priority boarding and upgrades, which might be difficult to appreciate until experienced. Frequent fliers who understand what to expect when they earn status should be more engaged. Once a flier has experienced these benefits, there is little more to learn, suggesting a partial reset after the first successful reward redemption with subsequent resets returning to this newly elevated level. However, H_{2b} does not limit the increase in effort to occurring only one time. In addition to procedural learning and learning from experience, successful goal attainment can lead to self-learning, resulting in an enhanced sense of proficiency.

Experiencing success can enhance proficiency because it provides clear information regarding how to proceed the next time (procedural learning) while also strengthening beliefs in self-efficacy (self-learning). Self-efficacy refers to a person's perception of how well he or she can execute various courses of actions to deal with prospective situations (Bandura 1982, 1988). Consumers must orchestrate buying behavior in specific ways, scheduling and steering purchases to attain a reward successfully. Previous research has suggested that if a task is difficult enough, such that one success does not guarantee future success, people will reassess their self-efficacy after each success (Dzewaltowski, Noble, and Shaw 1990; Ryan 1970). Therefore, if a task is challenging enough, each successful goal attainment should lead to a reassessment and, in turn, an increase in the base level of effort (depicted in Figure 1, Panel D). Consumers can exhibit effort in many ways. In the current study, increases in effort are exhibited through purchase quantity (reflected by shorter interpurchase times) and store choice (traveling farther away). This leads to the following hypotheses:

H_{3a}: Repeated successes result in recurring partial postreward resets (i.e., each reset occurs at a higher base level of effort).

H_{3b}: For a partial, rather than a full, reset to occur, the goal must be challenging.

Therefore, we expect to observe the impact of success on subsequent effort more than once, which does not rule out learning from experience but supports the incidence of self-learning. Learning from experience occurs when a consumer learns the value of the reward, leading to a one-time reassessment. Self-learning leads to a reassessment of the likelihood of attaining the reward, which results in reoccurring reassessments of self-efficacy. The impact of success through self-learning and perceptions of self-efficacy leads to the next hypothesis:

H_{4a}: Marked progress toward a goal affects a person's assessment of self-efficacy.

H_{4b}: Successful task completion has an additional impact on self-efficacy beyond marked progress alone.

We organize the remainder of this article as follows: We begin by presenting a real-world illustration of partial resetting. Study 1 includes two parts that test H_1, $H_{2a–b}$, and H_{3a}. Using data from a leading international carrier's frequent-flier program, we illustrate partial postreward resetting using elite status as the goal and the rate of miles flown as effort. In Study 2, we test H_1, $H_{2a–b}$, and H_{3b} in a laboratory setting. The results reveal that the impact of success on future effort depends on the consumer perceiving a reward as challenging but not too challenging. In Study 3, we test H_4, documenting the impact of success on perceptions of self-efficacy, the principle mechanism driving our results. We conclude by pointing out some limitations of this work and suggesting avenues for further research.

STUDY 1: EMPIRICAL EVIDENCE OF THE EFFECT OF SUCCESS ON EFFORT

Part 1

We use frequent-flier program data obtained from a major U.S.-based international airline. The airline offers fliers three tiers of status. Tier 3 is reached after flying 25,000 miles in a given year, Tier 2 is reached after flying 50,000 miles, and Tier 1 is reached after flying 100,000 miles. Only miles actually flown qualify for annual status; bonus miles and miles earned through third parties, such as hotels and credit cards, do not count. Each tier entitles members to special privileges and, once earned, is valid until the end of the following year. The airline provides Tier 3 status for life to those who have earned one million miles and Tier 2 status for life to those who have earned two million miles. All earned miles, by flying or other means, such as credit card purchases and hotel stays, count toward reaching the one million and two million marks.

The airline does not offer Tier 1 status for life; this level of status can only be earned by flying 100,000 miles or more in a single year. In light of the structure of this airline's status programs, we restrict our study in Part 1 to a distinct subset of members who have earned two million miles and possessed at least 200,000 miles in their frequent-flier account at the time of the study. For each member, we have 18 months of detailed flight activity; the beginning of the data coincides with the beginning of the 2005 status year. (Status years are offset slightly from calendar years.) These restrictions enable us to accomplish two important things. First, the bank of 200,000 miles suggests that these fliers are unlikely to consider earning 25,000 miles for a free ticket a meaningful goal. Although we realize that restricting our analysis as such does not guarantee that these fliers care only about status, our intent was to select fliers for whom accumulating miles for free flights was of less importance; their primary goal would more likely be earning Tier 1 status. The average flier in our sample has spent two-thirds of his or her lifetime earnings in miles. These fliers use their miles and therefore must value them as a pathway to free flights and other perks. Second, each flier has earned more than two million miles and thus has Tier 2 status for life. They understand the benefits associated with status; any change in behavior associated with reaching Tier 1 status is unlikely to be due to learning from experience. (We tease apart the effects of self-learning from learning from experience more directly in Part 2.) The only meaningful status goal remaining is Tier 1 status. Furthermore, Tier 1 and 2 fliers accrue miles toward rewards and status at the exact same rate, and earning Tier 1 status in 2005 does not facilitate earning Tier 1 status the following year.

We expect the propensity to fly among this subset of program members to increase as they near the Tier 1 status goal (100,000 miles), consistent with the goal gradient effect. In addition, for those who reached Tier 1 in 2005, we expect the propensity to fly at the beginning of the subsequent year to be greater than at the beginning of the current year but lower than immediately before reaching the goal of Tier 1 status (i.e., partial postreward resetting after goal reengagement).

Model. To examine the propensity to fly in a way that is consistent with our hypotheses, we build a random effects proportional hazard rate model. We parameterized the basic Weibull proportional hazard rate model with covariates as follows:

$$H(t, \beta) = \gamma\alpha(\alpha t)^{\gamma-1},$$

$$\text{with } \alpha = e^{X'\beta},$$

where

h = the hazard function,
t = the duration of events that are being modeled (in our case, the interflight time),
γ = the Weibull shape parameter,
α = the Weibull scale parameter,
X = a set of covariates, and
β = the vector of parameters to be estimated.

To account for the way status is earned in loyalty programs, we need to make some adjustments to the simple model portrayed in Figure 1, Panel D. We must account for the fact that status is earned according to calendar years. If a consumer reaches 100,000 miles during year y, he or she gains Tier 1 status for the remainder of the year (y) and for the totality of the following year (y + 1). Miles accrued after earning Tier 1 status during year y do not count toward earning Tier 1 status in year y + 1. Therefore, for travelers who reach 100,000 miles in a given year, we expect effort to drop after reaching this goal and remain low until the beginning of the next year. We state this more formally as follows:

H_5: Effort is reduced after goal completion until the next goal becomes active.

In a traditional buy-ten-get-one-free type of loyalty program, the consumer never fails, because there is always the possibility of eventually reaching ten purchases. However, whenever there is a time limit, success and its impact on self-efficacy depend on reaching the goal in time. Failure could lead a consumer to view him- or herself as less efficacious. Thus, we expect a full reset (or even worse) at the beginning of the new year when a member fails to reach the goal in time. Figure 2 illustrates the predicted effects for customers who are successful in reaching Tier 1, including the hypothesized post–goal attainment dip along with a partial postreward reset at the beginning of a new year. Figure 1, Panel B, reflects the pattern for people who are unsuccessful, substituting the advent of a new year for the first goal.

Our model considers the following: (1) a goal gradient effect leading up to the goal (H_1), (2) a post–goal attainment dip between the time the goal has been reached and the beginning of the next year (H_5), (3) a possible gradient between goal attainment and the beginning of the next year in case the consumer has other goals, and (4) a partial reset when starting the subsequent year ($H_{2a–b}$). We also incorporate random effects coefficients to account for heterogeneity in members' base propensity to fly and individual rates of acceleration. Our full set of covariates is as follows:

Figure 2
PROPENSITY TO FLY

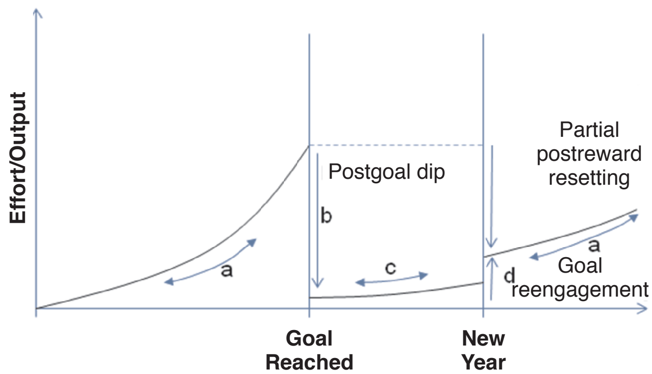

$$X_{iyt}\beta = \beta_0 + \beta_1.NMiles_{iyt} + \beta_2.NMiles_{iyt}.Tier1_i + \beta_3.NPGMiles_{iyt} + \beta_4.Post_{iyt} + \beta_5.Year_{iy} + \beta_6.Year_{iy}Tier1_i + \beta_7.Tier1_i + z_i + zz_i.NMiles,$$

where

$NMiles_{iyt}$ = the number of pregoal miles flown in year y by traveler i before taking flight t,
$NPGMiles_{iyt}$ = the number of postgoal miles flown in year y by traveler i before taking flight t,
$Post_{iyt}$ = an indicator variable that is set to 1 if traveler i has reached Tier 1 status in year y by flight t and 0 if otherwise,
$Year_{iy}$ = an indicator variable that is set to 0 for the first flights taken in the first year of the data set and 1 for the second year, and
$Tier1_i$ = a time-invariant indicator variable that is set to 1 if traveler i reaches Tier 1 status in the first year.

The variables $NMiles_{iyt}$, $NPGMiles_{iyt}$, and $Post_{iyt}$ are reset to 0 at the beginning of the year. The variable $NMiles_{iyt}$ accrues miles for every flight flown until the traveler reaches Tier 1 status, and $NPGMiles_{iyt}$ accrues miles for every flight flown after the traveler reaches Tier 1 status in that year. Thus, together, $NMiles_{iyt}$ + $NPGMiles_{iyt}$ constitute the total miles flown by i up to flight t in year y. To scale the parameters so they are easier to interpret, we divide the number of miles flown by 100,000 (i.e., when reaching Tier 1 status, $NMiles_{iyt}$ is equal to 1). The intercept, β_0, is the base propensity of traveling at the beginning of the year for an average flier who will not reach Tier 1 in Year 1. As such, $\beta_0 + \beta_7$ is the propensity to fly for members who reach Tier 1 in Year 1, and β_2 is the base difference in the goal gradient for fliers who reach Tier 1. In addition, β_3 reflects the gradient post–goal attainment for people who reach Tier 1; this allows for the possibility of other goals. Finally, β_5 is a year intercept that allows for a different propensity to fly in Year 2 versus Year 1.

The random coefficients z_i and zz_i allow for heterogeneity in both a member's base propensity to travel (z_i) and goal gradient (zz_i). A member who flies more may be less likely to increase effort with progress. Thus, we allow the rate of acceleration to be correlated with base propensity and model the random effects as a correlated bivariate normal:

$$(z_i, zz_i) \sim N(0, \Sigma),$$

$$\Sigma = \begin{pmatrix} \sigma_z^2 & \sigma_{z,zz} \\ \sigma_{z,zz} & \sigma_{zz}^2 \end{pmatrix}.$$

Using our model specification, a negative and significant β_1 (NMiles) parameter provides support for H_1. (Increased effort is reflected in shorter interpurchase times.) A positive β_4 (Post) parameter provides support for H_5. Because we rescaled the number of miles flown to be equal to 1 when 100,000 miles are flown, a β_6 (year × Tier 1 interaction) parameter that is significantly smaller than $\beta_1 + \beta_2$ provides support for postgoal resetting (H_{2a}). Furthermore, if this parameter is also significantly greater than 0, it provides support for partial rather than full resetting (H_{2b}). In other words, $\beta_6 = 0$ indicates a full reset, and $0 < \beta_6 < \beta_1 + \beta_2$ indicates a partial reset.

We estimated the model using PROC NLMIXED in SAS. Our sample includes more than 400,000 flights taken by approximately 7000 travelers who, on average, fly slightly more than 85,000 miles per year. We report the results of three model specifications: (1) the random effects model described previously, (2) a model allowing members who have not reached Tier 1 by the last quarter of the year to change the rate at which they fly (i.e., allowing the goal gradient effect become more pronounced as time expires), and (3) a model without random effects. To test the robustness of our model, we also (4) fit a fixed-effect version of the model, (5) substitute the number of trips flown for the number of miles accrued as a marker of progress, and (6) fit the model using only travelers who earned Tier 1 status in Year 1 (and thus possessed status throughout Year 2). The results were substantively the same across all models, leading to the same conclusions regarding our hypotheses, and when comparable, the parameters were of the same sign and magnitude. Therefore, we report results from only the first three models.

Furthermore, to check the face validity of our model, we performed a model-free analysis in which we plotted the average time between flights for different levels of accumulated miles. To control for attrition (i.e., as we increase the number of miles flown, fewer members remain), we examined two specific groups of fliers: those who accumulated between 100,000 and 125,000 miles and those who accumulated between 125,000 and 150,000 miles in Year 0. This ensures that there is no attrition before the 100,000 goal is reached and enables us to compare a group that just makes the goal with one that reaches it more comfortably. Figure 3

Figure 3

MODEL FREE APPROACH

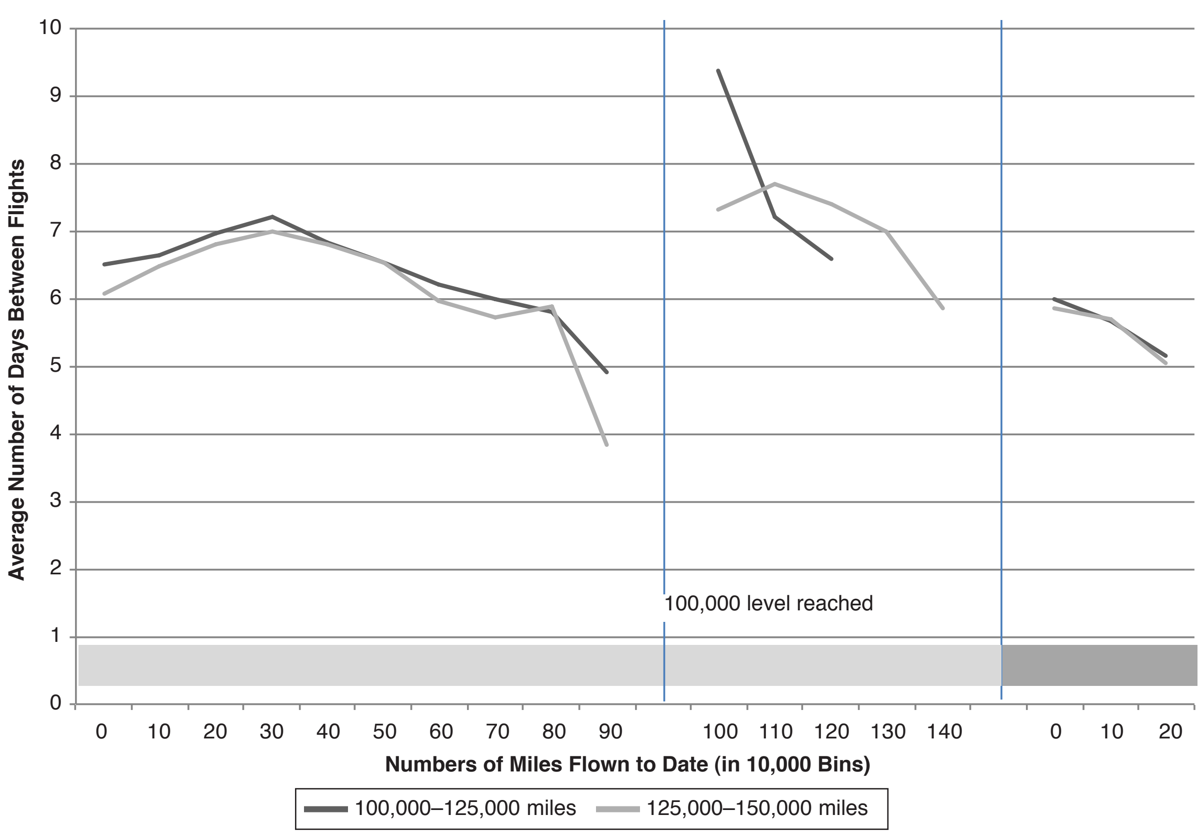

shows the plot of average interflight times as a function of accumulated miles.

This model-free approach shows similar curves or gradients as the 100,000 goal nears for both groups. There is a clear jump when the traveler reaches the goal: Interflight times increase from slightly less than five days, on average, to more than seven (greater than at any point before reaching the 100,000 mark). It also seems that the number of days between flights at the beginning of Year 2 is lower than at the beginning of Year 1. This lends face validity to our principal hypotheses regarding partial postreward resetting.

Results. Table 1 shows the parameter estimates for our model. Negative coefficients indicate a greater propensity to fly, and positive coefficients indicate a lower propensity to fly. For the sake of brevity, we do not report the individual random effects. However, note that the individual-level intercept (z_i) and goal gradient (zz_i) parameters are negatively correlated ($\sigma_{z, zz} = -.0777$, $p < .0001$, implies $\rho_{z, zz} = -.59$), indicating that people with a greater propensity to fly (i.e., those who are more likely to reach the goal) exhibit a smaller goal gradient effect.

We find support for H_1, the goal gradient effect ($\beta_{NMILES} = -.2269$, $p < .0001$). We also find that people who reach Tier 1 in Year 1 not only exhibit a greater base propensity to fly ($\beta_{TIER1} = -.0886$, $p < .0001$) but also exhibit a stronger goal gradient effect ($\beta_{NMILES \leftrightarrow TIER1} = -.1392$, $p < .0001$). At –.3661(–.2269 – .1392), the goal gradient that Tier 1 recipients exhibit is more than 50% larger than for other travelers. In terms of postgoal resetting, we find no year-end effect for those who have not reached Tier 1 status in Year 1 ($\beta_{YEAR} = -.0037$, $p = .74$). This supports the notion of a full reset at the beginning of the year for travelers who did not succeed. In contrast, there is a significant year × Tier 1 interaction ($\beta_{YEAR \times TIER1} = -.0489$, $p < .0001$) such that those who reached Tier 1 status in Year 1 begin Year 2 faster than they began Year 1. This interaction term is smaller than the parameter for the goal gradient ($\beta_{YEAR \times TIER1} = -.0489 < \beta_{NMILES} = -.2269$, $p < .0001$). This indicates that the increased propensity to fly in Year 2 is smaller than the increased propensity to fly observed when reaching the goal in Year 1, thus providing support for a partial postgoal reset ($H_{2a–b}$) for those who succeed. We also compared the reset with the Tier 1 × gradient interaction term to ensure that the reason the reset is partial is not due to a second goal of which we are unaware still being active as the new year begins. We find that the reset is significantly smaller than the interaction term ($\beta_{YEAR \times TIER1} = -.0489 < \beta_{NMILES \times TIER1} = -.1392$, $p < .0001$), providing further evidence of a partial reset.

Moreover, we find support for H_5 ($\beta_{POST} = .2223$, $p < .0001$). This large decrease in propensity to fly validates the notion that Tier 1 status is a goal for these fliers. A test of parameters fails to reject the hypothesis that this drop in propensity to fly is equal to the increase in propensity to fly brought on by the goal gradient ($\beta_{NMILES} = -\beta_{POST}$, $p < .55$). However, there is evidence that this is not the only goal members have in mind. Those who reach Tier 1 have a significantly steeper goal gradient than those who do not ($\beta_{NMILES \times TIER1} = -.1392$). This increase in gradient persists after travelers reach Tier 1 ($\beta_{NPGMILES} = -.1393$, $p < .0001$). A test of parameters fails to reject that these two parameter are equal ($\beta_{NMILES \times TIER1} = \beta_{NPGMILES}$, $p = .49$). This implies that some travelers who reach Tier 1 may have a secondary goal activated after Tier 1 status is attained.

Comparing the estimated parameters for NMILES (H_1), YEAR × TIER1 ($H_{2a, b}$), and POST (H_5) across the three models reported in Table 1 reveals how our hypotheses are supported in a simple model that does not account for heterogeneity in addition to a more complex model in which we attempt to control for heterogeneity and the sense of urgency consumers might feel at the end of the year. Furthermore, if we examine the improvement in log-likelihood from a null model with only an intercept (LL = –1,252,578.78; see the note to Table 1), to a model without random effects (LL = –1,234,497.90), to a model with random effects (LL = –1,222,965.44), we note that approximately 60% of the improvement is due to the fixed parameters representing our model and 40% is due to the random effects that capture heterogeneity in flying and differences in sensitivity to progress (i.e., the goal gradient).

It bears repeating that the aforementioned effects are not due to innate differences in the propensity to fly (i.e., heterogeneity). If it was simply the case that people who fly more often are more likely to reach Tier 1 status in Year 1 and thus are also more likely to fly in Year 2, the effect would be captured by the individual random effects, not an intercept shift in Year 2. Rather, these variables suggest that attaining Tier 1 status in Year 1 leads to partial resetting at the beginning of Year 2; it takes less time between the first and second ticket purchase in Year 2 than in Year 1 for travelers who had successfully reached their goal of Tier 1 status.

Part 2

In Part 2, we examine the behavior of more than 40,000 individual fliers to determine the impact of achieving status on the number of miles flown on the same carrier the subsequent year. More specifically, we examine the number of base miles earned in one year (as in Part 1, only miles actually flown qualify for status) and the likelihood of reaching status the following year for both people who did and those who did not achieve status.

We focus on Tier 3 status. Compared with Tier 3 benefits, fliers might perceive Tier 2 or Tier 1 as more of the same. Fliers are more likely to be upgraded or to board first if they have Tier 1 status, but the nature of the perks is similar. Thus, we do not expect much learning from experience to occur when reaching Tier 1. In contrast, the difference in treatment between Tier 3 fliers and travelers without status is significant. Reaching Tier 3 for the first time opens the door to priority queues, free upgrades, bonus miles, and other perks that make traveling far less onerous. Thus, it seems reasonable to expect the presence of both learning from experience and self-learning. However, learning from experience will only occur the first time the traveler reaches Tier 3 status, whereas self-learning can occur each time Tier 3 status is reached. This enables us to distinguish these two types of learning in this part of Study 1.

If we were to design an experiment to test for the presence of learning from experience and self-learning, we would need two groups of fliers, one that successfully attained status and one that did not, but with random assignment to each group. We could not implement such a field experiment and had to rely on actual travel records. This leads to potential confounds. People who do not reach Tier 3 status

Table 1

STUDY 1, PART 1: PARAMETER ESTIMATES FOR THE INTERFLIGHT HAZARD RATE MODEL

	Base Model			*Base Model with Last Quarter Dummies*			*Base Model with No Random Effects*			
	Estimates	*t-Value*	p-*Value*	*Estimates*	*t-Value*	p-*Value*	*Estimates*	*t-Value*	p-*Value*	*Hypotheses*
Intercept	XXX	153.24	<.0001	XXX	154.90	<.0001	XXX	284.11	<.0001	
NMILES	–.2269	–13.83	<.0001	–.4626	–25.13	<.0001	–.2979	–23.54	<.0001	H_1
NMILESS × TIER1	–.1392	–7.05	<.0001	–.0142	.67	.4998	–.1081	–7.32	<.0001	
NPMILESS	–.1393	–11.79	<.0001	–.3331	–25.84	<.0001	–.4182	–36.53	<.0001	
POST	.2223	35.80	<.0001	.1466	22.70	<.0001	.2690	45.74	<.0001	H_5
YEAR	–.0037	–.33	.7429	–.0346	–3.07	.0021	–.0513	–4.33	<.0001	
YEAR × TIER1	–.0489	–3.96	<.0001	–.0363	–2.94	.0033	–.0370	–2.85	.0044	$H_{2a–b}$
TIER1	–.0886	–5.55	<.0001	–.1116	–6.95	<.0001	–.1011	–12.25	<.0001	
LASTQ				.3464	22.70	<.0001				
LASTQ × NMILES				–.1502	–7.43	<.0001				
Random Coefficients										
$\log(\sigma_z)$	–.8302			–.8282						
$\log(\sigma_{zz})$	–1.1980			–1.1087						
$\sigma_{z,zz}$	–.0777			–.0895						
Weibull shape gamma	1.0220	861.44	<.0001	1.0254	861.12	<.0001	.9362	918.33	<.0001	
Log-likelihood		(1,222,965.44)			(1,222,965.44)			(1,234,497.90)		
N		418,818.00			418,818.00			418,818.00		
BIC		2,446,086.22			2,446,112.11			2,469,112.31		
AIC		2,445,954.88			2,445,958.88			2,469,013.80		

Notes: BIC = Bayesian information criterion, and AIC = Akaike information criterion. We omitted the parameter for the intercept for confidentiality purposes. For comparison purposes, a null model with only an intercept (no random effects) has a log-likelihood of –1,252,578.78. We computed the BIC and AIC coefficients by setting the number of parameters as the number of fixed coefficients and ignoring the number of random coefficients (e.g., for the base model, k = 12).

might have a lower propensity to fly, they might not care about status and prefer splitting their travel across multiple airlines, or they might be less flexible in their travel plans. An established way to address the problem of nonrandom assignment is to use a regression-discontinuity model (Busse, Silva-Risso, and Zettelmeyer 2006; Hahn, Todd, and Van der Klaauw 2001). To do so, we compare travelers who fly slightly less than the 25,000 miles necessary for Tier 3 status with those who travel 25,000 miles or slightly more. We believe that small enough differences in miles flown are more likely due to the vagaries of flight distances and destinations than to a fundamental difference in travel needs. Given that the average flight on the airline we analyzed in Part 1 covers 1248 miles and the minimum number of miles earned for a trip is 500, it is unlikely that someone who travels 24,950 miles in a single year has a basic demand that is fundamentally different from someone who travels 25,050.

We obtained the records of 43,548 travelers who flew between 20,000 and 30,000 miles. The data contain the number of miles flown for each member in 2004 and 2005 and whether they reached Tier 3 status in 2005 and any year previous. We call X_i and Y_i the number of base miles traveler i earned in 2004 and 2005, respectively. Furthermore, let Y_i^- be the number of miles flown in 2005 for a traveler who misses the 25,000-mile goal in 2004 by an infinitesimal margin and Y_i^+ the number of miles flown in 2005 for a traveler who flies exactly 25,000 miles in 2004. If we had enough of such travelers, we could estimate the impact of gaining status as $E(Y^+) - E(Y^-)$. In the absence of a large enough sample of people who barely missed or just made 25,000 miles, a regression-discontinuity model estimates Y^+ and Y^- in the following way: First, we construct an interval of size w around the goal [25,000 – w, 25,000 + w]; second, we estimates the following two linear regressions:

$$Y_i = \alpha^+ + \beta^+(X_i - 25{,}000) + \varepsilon_i,\ \forall X_i\text{: } 25{,}000 \le X_i < 25{,}000 + w, \text{ and}$$

$$Y_i = \alpha^- + \beta^-(X_i - 25{,}000) + \varepsilon_i,\ \forall X_i\text{: } 25{,}000 - w \le X_i < 25{,}000.$$

Then, we compute the impact of reaching status in 2004 on the number of miles flown in 2005 as $\alpha^+ - \alpha^-$. Increasing w leads to more data points being included and helps produce estimates with small standard errors at the expense of reducing the match between the consumers included in the Y^+ and Y^- regressions. If the β^+, β^- corrections do not properly account for changes in propensity to fly as we move further away from 25,000 miles, this may lead to biased estimates. To minimize the chance of biased estimates, we begin with a wide w and estimate the model with a smaller and smaller w until there are too few data points to produce statistically significant estimates. Moreover, we can examine how the parameter changes as the window decreases to ensure that the estimates are stable. We started with a w of 5000 miles and estimated the model in 100-mile decrements. Furthermore, we estimated a single nested model using a dummy variable to indicate whether the traveler had reached 25,000 miles in the focal year. Thus, our model is as follows:

$$Y_i = \alpha^- + \alpha^{\pm}I(X_i) - \beta^-|25{,}000 - X_i| + \beta^{\pm}|25{,}000 - X_i|I(X_i) + \varepsilon_i,$$

and

$$\forall X_i\text{: } 25{,}000 - w \le X_i < 25{,}000 + w,$$

where $I(X_i) = 1$ if $X_i \ge 25{,}000$ and 0 if otherwise.

In this model, $\alpha^{\pm}$ provides the impact of reaching 25,000 miles in 2004 on the number of miles flown in 2005 directly. In addition to this linear model of miles flown in 2005, we estimated a logistic regression that models the likelihood of earning status in 2005 as a function of the number of miles flown in 2004. We report the results of both models.

We estimated each model in two steps. First, to get the cleanest test of whether successful goal attainment leads to an increase in effort in successive attempts, we examine only travelers who never reached the status level before 2004. A positive and significant estimate for $\alpha^{\pm}$ would provide support for the positive impact of success. However, it could be indicative of either learning from experience or self-learning (or both). To examine self-learning in the absence of learning from experience, we performed the same analysis for travelers who had attained status at some point before 2004. Learning from experience should have occurred the first time they attained status (before 2004), and any effect of successfully earning status should reflect self-learning. For these travelers, we add the number of times they had reached status in the past (NSTATUS) and its square (NSTATUS2) as covariates. We expect NSTATUS to be significant and positive; however, given that self-learning should only be expected to increase a person's subjective likelihood of success up to a point, we expect NSTATUS2 to be significant and negative.

Table 2 shows parameter estimates and associated statistics. The coefficients for $\alpha^{\pm}$ are positive and significant in all four models. Furthermore, coefficients for NSTATUS are positive and significant, and coefficients for NSTATUS2 are negative and significant in both models that include fliers who attained status repeatedly. None of the β^- or $\beta^{\pm}$ parameters are significant, indicating that the number of miles flown in 2005 is not sensitive to deviations from 25,000 miles. This is a good indication that the interval size (w) is not too wide. Figure 4 illustrates the effect of reaching status on subsequent behavior; here, we display the probability of reaching status in 2005 as a function of the number of miles accrued in 2004.

Examining program members who never earned status before 2004 (solid line), we observe that the likelihood of earning status in 2005 (reaching 25,000 miles) increases by 50% for those who barely attained the same goal in 2004 compared with those who barely missed it (25,000 miles). The effect is just as pronounced for members who reached status in the past (dashed line, aggregated across all repeat earners). In the case of first-timers (solid line), the jump is due to a combination of learning from experience and self-learning. For those who had status in the past (dashed line), learning from experience is unlikely, but there is evidence of self-learning.

In Part 2, we cannot account for fliers who choose to switch carriers as soon as or shortly after they reach 25,000 miles. They may prefer to work toward earning status on another airline rather than enjoy the benefits of status on the studied airline. Although we do not doubt these fliers exist, we find it unlikely that they drive our effects. Nevertheless, Studies 2 and 3 control for this behavior.

Table 2

STUDY 1, PART 2: PARAMETER ESTIMATES FOR REGRESSION DISCONTINUITY MODELS

	Miles Flown				*P (Reach 25,000 Miles)*			
	First-Timers		*Repeaters*[a]		*First-Timers*		*Repeaters*[a]	
Most Effective	*Estimate*	*t-Statistic*	*Estimate*	*t-Statistic*	*Estimate*	χ^2	*Estimate*	χ^2
α^-	XXX	21.02**	XXX	11.37**	XXX	51.42**	XXX	67.09**
$\alpha^\pm$	1183.5	1.99*	2689.7	2.41*	.71	4.73*	.52	5.67*
β^-	–.82	–1.38	.48	.39	.0003	.45	–.00006	.07
$\beta^\pm$	–.13	–.20	–1.08	–.84	–.0004	.85	8.1E-07	.00
NSTATUS			1322.2	7.95**			.22	53.81**
NSTATUS2			–57.7	–3.43**			–.01	9.37**
w	1300		1200		900		1300	
N	15,845		8478		9956		9269	
Y^-	11,682		13,979		9.2%		17.4%	
Y^+	12,865		16,668		17.1%		26.0%	

*Significant at $p = .05$.
**Significant at $p = .01$.
[a]For a traveler who has reached status once before.
Notes: We omitted the parameter for the intercept for confidentiality purposes.

Figure 4

STUDY 1, PART 2: EMPIRICAL PROBABILITY OF ATTAINING STATUS FOR FLIERS WHO HAVE AND HAVE NOT SUCCEEDED PREVIOUSLY

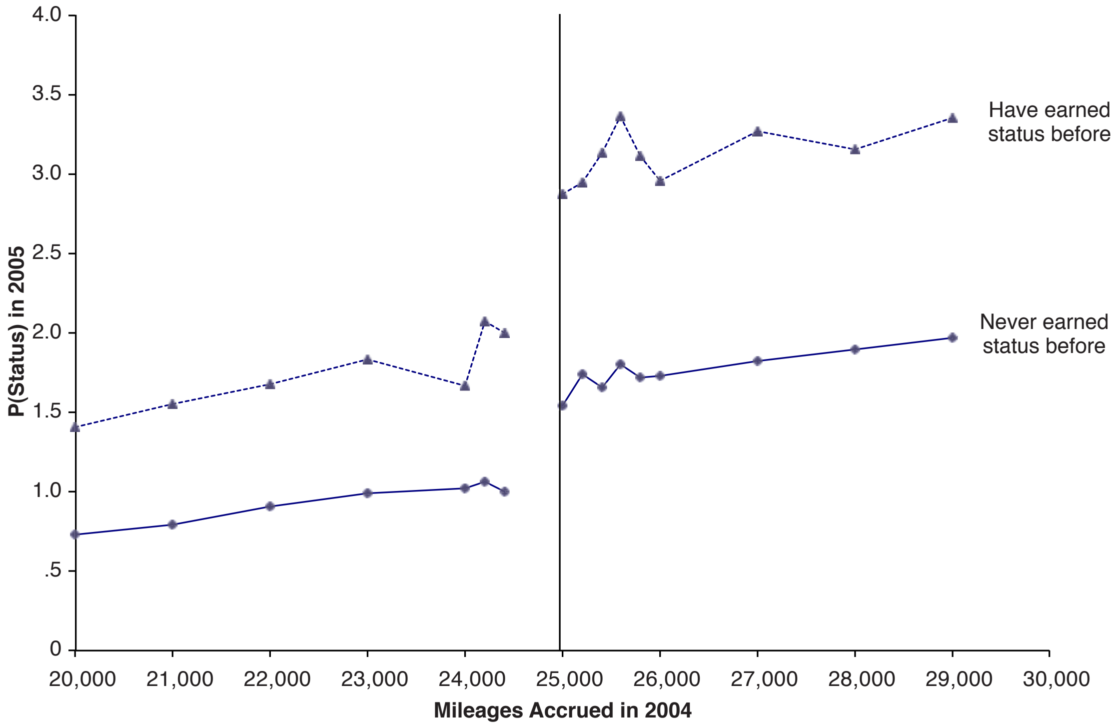

Notes: Probabilities are expressed relative to the probability of reaching status for a traveler who had never reached status before 2004 and barely missed reaching 25,000 miles in 2004.

Discussion

In Study 1, we use real-world data to demonstrate empirically how success contributes to an increase in effort exhibited in successive attempts to reach the same goal. "Effort" refers to consumers' orchestrating their buying behavior in specific ways. In Part 1, those who attained Tier 1 status exhibited a greater propensity to fly on the focal carrier, as exhibited through partial resetting: After earning Tier 1 status, the time between flights slowed but not to their original levels. In Part 2, we observe the positive impact of successfully attaining status and its frequency (more than once vs. once) on the likelihood of earning Tier 3 status in the next period.

It is crucial to note that all successes are not equal. Self-efficacy can be instilled and strengthened by experiencing success, but only when success requires perseverance. Successes that come too easily are unlikely to lead people to update their sense of self-efficacy (Wood and Bandura 1989). Therefore, a program that issues small rewards frequently is unlikely to lead to self-learning. At the opposite extreme are goals that are too lofty. Creating larger rewards with greater purchase requirements should lead to more overall effort—but only up to a point. Studies have shown that motivation dissipates as the likelihood of task completion diminishes beyond the realm of possibility (Garland 1984; Mento, Cartledge, and Locke 1980). For firms using reward programs, the effort required and size of the reward can be scaled up or down, dictating the number of redemption opportunities available and the purchase activity required. A program's divisibility, as we call it, determines how easy or difficult each reward is to earn (e.g., 250 points leading to a $10 certificate vs. 1000 points leading to a $40 certificate). We predict that decreasing divisibility will serve to reinforce perceptions of self-efficacy to the extent that the person believes he or she can successfully complete the task.

STUDY 2: THE IMPACT OF DIVISIBILITY ON EFFORT

In Study 2, we test H_{3b} in addition to revisiting H_1, $H_{2a–b}$, and H_{3a} by exploring consumer reactions to differing degrees of difficulty in attaining a reward. The results illustrate how self-learning diminishes with too much divisibility (rewards attained too easily) and motivation diminishes with too little divisibility (rewards perceived as too lofty or out of reach).

Method

Respondents. Participants were 300 undergraduate business students who completed the survey voluntarily. Of these, 7 failed to answer the questions, resulting in 293 usable surveys.

Stimuli and design. We used a 3 (divisibility: reward every $100, $500, or $1,000 spent) × 4 (accumulated spending: $25, $425, $525, or $925) between-subjects, full factorial design. Participants were asked to imagine that they had moved to a new city and there were two large grocery stores near their house that carried essentially the same assortment of goods. The stores were no different in terms of price, layout, cleanliness, and service. However, respondents were told that the two grocers differed significantly in their frequent-shopper programs. Store B offers a 10% cash refund on every dollar spent at the store. For example, if a shopper spends $50, he or she receive a $5 refund immediately and thus paid only $45. Store A offers cash back at various increments or reward levels that vary according to the degree of divisibility ($100, $500, or $1,000). Note that the cumulative rewards the stores offer are identical (10% of cumulative purchases); only the frequency of disbursement differs.

The first question asked respondents which store they believed they would shop at more frequently. The study next described their progress toward the goal at Store A. It explained that "one store is near the freeway, which you often take to work, while the other is closer to a main surface street, which you take to and from work instead of the freeway when you leave during rush hour. While they are equal distances from your house, they are in opposite directions. As a result, you have shopped at both stores repeatedly during the past six months, as convenience has often been a factor." We manipulated accumulated spending by stating the following: "At Store A, you have currently accumulated $25 [$425, $525, or $925] in purchases. (Remember, when you reach an accumulated spending level of $100 [$500, $1,000] they will refund $10 [$50, $100].)" Then, the scenario indicated that it was the weekend (rush hour would not be a factor) and asked respondents to imagine that they were heading out to buy $50 in groceries for a party they were planning that evening. They were asked to indicate the store they would choose to shop at, A or B.

We expect the likelihood of going to Store A to increase as the program becomes less divisible and the reward level increases from every $100 to every $500. However, we expect the likelihood of going to Store A to decrease as the program becomes even less divisible; when the reward level increases from $500 to $1,000, the goal begins to seem unattainable to many. In addition, in accordance with the goal gradient effect (H_1), when the only reward available is at $1,000, consumers will become increasingly inclined to go to Store A as their accumulated spending increases from $25 to $925.

With regard to the impact of success on effort, for travelers who earn $50 for every $500 spent, we expect the likelihood of going to Store A to be greater when accumulated spending is $525 (one success) than when it is $25 (no successes), even though the size of the reward and the effort required is identical in both cases (H_{2b}). Similarly, we expect the likelihood of going to Store A to be greater for those who have accumulated $925 in spending than for those who have accumulated $425, despite both groups being $75 away from earning a $50 reward. Again, this is because those who have spent $925 have experienced success (H_{3a}). Consistent with H_{3b}, we do not expect an impact of success for recurring goals within the $100 reward conditions, because the task is not challenging enough to prompt self-learning.

Results

Only 4 of the 293 respondents (1.4%) indicated that they would regularly visit Store A, suggesting that Store B was vastly preferred in the absence of any accumulated spending. Figure 5 depicts the reported probability of going to

Figure 5

STUDY 2: IMPACT OF DIVISIBILITY ON STORE CHOICE

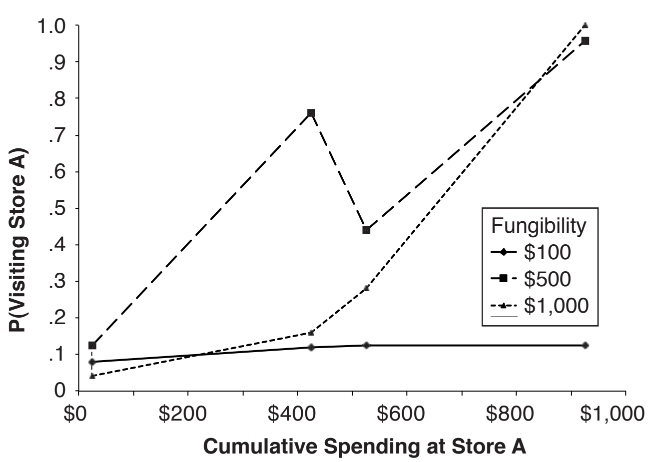

Store A for each of the three reward levels ($100, $500, and $1,000) and each of the four levels of accumulated spending ($25, $425, $525, and $925). As the graph demonstrates, the likelihood of going to Store A is affected by both divisibility—the level of reward redemptions—and accumulated spending.

As we predicted, and in support of H_{3b}, effort increased with a decrease in divisibility. The average likelihood of going to Store A was greater in the $500 condition (57%), when shoppers were paid $50 for every $500 in spending, than in the $100 condition (11%), when shoppers were awarded $10 for every $100 spent ($\chi^2 = 15.2$, $p < .01$). In addition, the likelihood of going to Store A decreased when divisibility decreased too much. The likelihood of going to Store A was lower (36%) when shoppers were offered $100 for every $1,000 spent than when they were awarded $50 for every $500 spent ($\chi^2 = 8.6$, $p < .01$). In other words, the $500 reward level induced a greater likelihood of visiting Store A than either the $100 or $1,000 reward level.

As we also expected, the likelihood of going to Store A increased monotonically in the $1,000 reward condition, consistent with the goal gradient effect and in support of H_1. Tests of proportion revealed that those who spent $525 were more likely to choose Store A than those who spent only $25 ($PA_{525} = 28\%$ vs. $PA_{25} = 4\%$; $\chi^2 = 3.9$, $p < .05$), and those who spent $925 were more likely to go to Store A than those who spent $525 ($PA_{925} = 100\%$ vs. $PA_{525} = 28\%$; $\chi^2 = 10.2$, $p < .01$). Although the likelihood of choosing Store A in the $425 condition (16%) falls halfway between those who spent $25 (4%) and those who spent $525 (28%), neither difference is statistically significant. Nevertheless, the overall pattern of results supports the notion that effort increases as participants near their goal.

Moreover, as we expected, those in the $500 condition who accumulated $525 in spending were more likely to choose Store A than those who had accumulated $25 in spending (44% vs. 12.5%, respectively; $\chi^2 = 5.35$, $p < .05$), in support of H_{2b}. If success was not motivating, we would expect shoppers in both conditions who were $475 away from earning a $50 reward to be equally inclined to go to Store A. Likewise, we would expect the percentages to be equal for $425 and $925 in accumulated spending (each $75 away from earning $50). This is not what occurred. In line with H_{3a}, those with $925 in accumulated spending were marginally more likely to choose Store A than those with $425 (96% vs. 76%, respectively; $\chi^2 = 3.11$, $p < .10$). These results conflict with notions of expected value and discounting: The likelihood of visiting Store A varied significantly despite shoppers being equidistant from the same size reward. It seems that loyalty programs benefit from people taking into consideration whether they have successfully reached the goal in the past when deciding what to do in the future. We also tested whether the likelihood of going to Store A in the $100 reward condition varied with accumulated spending. An analysis of variance using the CATMOD procedure in SAS indicated no difference across any of the four levels ($p = .95$). Consistent with H_{3b}, for past successes to matter, goals must be challenging.

Discussion

Study 2 shows the importance of divisibility—the number of rewards and the level at which they are awarded—on self-learning. Although the overall amount a shopper would receive remained constant for anyone who accumulates $1,000 in spending, our study indicates two effects: If the program is too divisible (but not perfectly, as in Store B), success does not teach the consumer anything. In contrast, if the program is not divisible enough, it may offer great incentives, but only for those who are already committed to the program; thus, they will be demotivating to those who deem the reward level as too lofty. An adequate balance must be struck.

The key point this study makes salient is the impact of success on effort. Goal attainment, which is a function of divisibility, leads people to reassess their aptitude, which affects future effort toward recurring goals. Accordingly, respondents who supposedly confronted a challenging goal and succeeded were more likely to choose Store A. We should emphasize the relative size of the effect: At the $525 level, respondents in the $500 condition were more likely to go to Store A than respondents in the $1,000 condition, despite having less accumulated spending toward the next reward ($25 vs. $525, respectively) and a smaller reward ($50 vs. $100). It seems that success served as a powerful motivator in Study 2. Admittedly, Study 2 does not isolate the mechanism. In Study 3, we examine the impact of goal attainment on self-learning and self-efficacy more directly.

STUDY 3: THE IMPACT OF GOAL ATTAINMENT ON SELF-PERCEPTIONS

We designed Study 3 to gauge how task completion and the successful attainment of a reward affect perceptions of self-efficacy and, thus, self-learning. Our goal was to design a task in which success would provide meaningful information about participants' ability to succeed, but one in which we could control success in individual trials so that all participants experienced the same sequence of successes and failures (to ensure comparability across conditions). We also needed participants to be oblivious to our manipulation of the outcomes; therefore, the study was presented as a game in which the goal was to determine whether the experimenter was lying by judging her facial expressions.

Method

Respondents were 70 undergraduate business students who participated in this along with several other studies for course credit. Respondents were brought into the lab individually under the premise that they were participating in a study on their ability to interpret facial expressions. It was explained that in business, just as in poker, people make credibility judgments on the basis of subtle facial cues. Each respondent was shown the eight of diamonds from a standard deck of cards. The experimenter explained that a card would be drawn randomly from the deck. Regardless of what card was drawn, the experimenter would say the card was "greater" than the eight of diamonds. This implied the card drawn was either a nine, ten, jack, queen, king, or ace. The respondent's task was to determine whether the experimenter was lying (the card was actually a two, three, four, five, six, or seven). Because there is an equal number of card values less than and greater than eight, on average their success rate should equal 50%, unless that person was proficient at detecting a bluff. They were told expert poker players can identify when a person is bluffing on average two-thirds of the time, or in 66% of trials.

The design was a single factor experiment with two conditions: 10 trials or 30 trials (cards drawn). Those in the 10-trial (30-trial) condition were told that if they exceeded 66% correct in the task, or were successful in at least 7 of 10 trials (21 of 30), they would receive a $4 ($12) reward. In other words, the number of trials required to complete the task varied, not goal difficulty. There was no reason for respondents to believe that any more proficiency was required in either condition. In addition, we held the exact sequence of successes and failures constant across conditions (we describe how subsequently). Thus, any change in self-efficacy would be the result of successfully completing the task and earning the reward.

Before commencing, respondents rated their ability to determine whether someone is bluffing and their ability to read facial cues on an 11-point scale (–5 = "much worse than average," and 5 = "much better than average"). Then, respondents rated their subjective likelihood of scoring as well as professional poker players (66% or better) on a 7-point scale (1 = "not at all likely," and 7 = "extremely likely"). Both groups answered the same questions after completing 10 trials (a success in the 10-trial condition and one-third progress in the 30-trial condition). Those in the 10-trial condition were invited to repeat the task (a recurring goal), and both groups completed the same questionnaire after 20 trials (two successes and two-thirds progress, respectively).

During the study, the experimenter used a screen to hide the cards drawn from the deck, which included a hidden pocket containing a preordained sequence of cards. This ensured that every participant got seven of their first ten responses correct and seven of their second ten responses correct in an identical way—the same cards in the same order. Objectively, all respondents experienced the exact same sequence of success and failure; However, those in the 10-trial condition completed the task successfully once and then twice, while those in the 30-trial condition were only one-third and two-thirds of the way toward attaining their goal when asked to assess their individual aptitude. Respondents did not know the experiment was rigged and believed that they were learning something about their lie-detection abilities; no one during the debriefing suggested any other reason for the study.

We expected an increasing number of successful trials to affect assessments of self-efficacy for all participants, consistent with H_{4a}. More important, in line with H_{4b}, we predicted that those in the 10-trial condition who successfully reached their goals (7 of 10 and 14 of 20 correct) and received the rewards would believe they were learning more about their face-reading capabilities and thus experience a greater change in their perceptions of self-efficacy.

Results

Three participants in the 10-trial condition did not want to perform the task a second time; thus, we have 32 responses for the post-20-trial evaluation measures in that condition. Because the measures used different scales (–5 to +5 for "bluffing" and "facial" and 1 to 7 for "success"), we standardized the data before starting the analysis. We collapsed the three measures given the high degree of agreement between them ($\alpha = .83$). To ensure that the results were not driven solely by a reassessment of the likelihood of success but rather a reassessment of ability, we also ran the analyses using only the two ability measures ($\alpha = .84$). The results were essentially identical; therefore, we present the analysis of the composite of all three measures.

To test for effective randomization, we compared the premeasurement (self-assessments before beginning the task) across conditions. We did not find any significant difference ($M_{10,\ pre} = -.49$ vs. $M_{30,\ pre} = -.61$; F = .44, n.s.). Our within-subject design enabled us to normalize each participant's score by subtracting each respondent's pretask rating (their self-assessment before receiving any feedback) from later measurements (after 10 and after 20 trials) before making any comparisons across conditions. Figure 6 shows the average changes in ratings.

An analysis of variance revealed a change in perceptions of self-efficacy that varied across conditions ($M_{10} = .70$ vs. $M_{30} = .43$; F = 13.68, $p < .01$), a significant change in these perceptions as respondents completed more trials (M_{pre} vs. M_{+10} vs. M_{+20}; F = 60.26, $p < .01$), and a significant interaction between the two (F = 3.74, $p < .05$). Individual contrasts reveal that there is a significant improvement in perceived ability after ten trials in both conditions ($M_{10,\ pre} = .0$ vs. $M_{10,\ +10} = .92$, $p < .01$; $M_{30,\ pre} = .0$ vs. $M_{30,\ +10} = .59$, $p < .01$). This supports H_{4a}: Progress toward the goal affects a person's assessment of self-efficacy. As predicted, the interaction reveals that this increase is larger for those who viewed 70% correct after ten trials as completing the task successfully than for those who viewed ten trials as one-third progress toward the overall goal ($M_{10,\ +10} = .92$ vs. $M_{30,\ +10} = .59$, $p < .01$). This result supports H_{4b}: Successfully reaching the goal has an additional impact.

The gap in self-perceptions widens further after 20 trials, or between two successful task completions and two-thirds completion ($M_{10,\ +20} = 1.17$ vs. $M_{30,\ +20} = .69$, $p < .01$). In addition, the changes in ratings from 10 to 20 trials is marginally significant in the 10-trial condition ($M_{10,\ +10} = .92$ vs. $M_{10,\ +20} = 1.17$, $p = .056$) but not in the 30-trial condi-

Figure 6

STUDY 3: HOW PERCEPTIONS OF APTITUDE CHANGE WITH SUCCESS

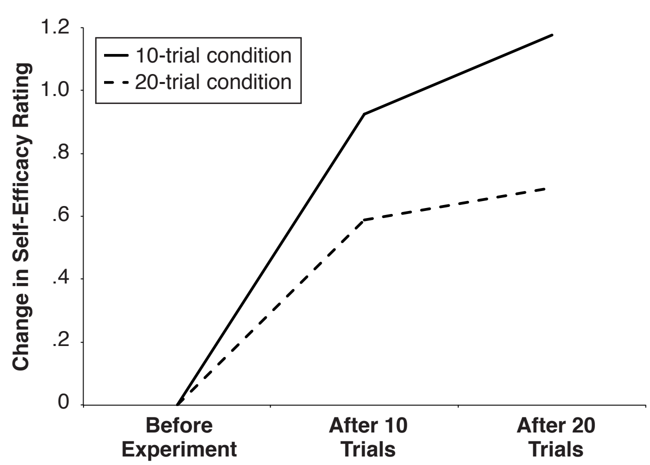

tion ($M_{30, +10}$ = .59 vs. $M_{30, +20}$ = .69, p = .42). Finally, it is worth noting that the change for those in the 30-trial condition after 20 trials was less than the change for those in the 10-trial condition after only 10 trials, and this difference was marginally significant ($M_{30, +20}$ = .69 vs. $M_{10, +10}$ = .92, p = .068).

Discussion

In Study 3, we examined the impact of success as it pertains to self-learning. The only difference across conditions was whether participants believed that the task was completed or still in process. We find that despite identical sequences of successes and failures on individual trials, successfully reaching a goal and earning the reward differentially affects perceptions of self-efficacy. Successfully completing a task composed of 10 trials (7 correct) did more to boost both the participants' perceptions of their own face-reading skills and the likelihood of success in the future than a task in which completing 10 trials meant that they had completed one-third of the task successfully. Successfully completing a task composed of 10 trials did even more than successfully completing 20 trials in the 30-trial task. We find unequivocal support for H_{4a} and H_{4b} and our hypothesis that successful goal attainment leads to self-learning and a change in self-efficacy. In turn, this increases a person's perceived ability of reaching the same goal again.

GENERAL DISCUSSION AND CONCLUSION

"If I did it once, I can do it again" is a common mantra for those who have attained success. However, there is no indication that people will try as hard or harder the next time. This research shows that not only can the successful do it again, but they will frequently work harder to help ensure that success happens again. The findings illustrate how success in a recurring goal framework enables consumers to learn something about themselves and thus leads them to amplify their effort in successive endeavors toward the same reward.

We argue that the increase in effort brought on by successfully reaching a goal and earning a reward is partly due to self-learning. People believe that success is a reflection of their ability to coordinate their efforts, overcome obstacles, and do whatever is necessary to succeed again. While we recognize that self-learning is likely only part of the story, its contribution is significant and worth noting. In Study 1, we used real-world frequent-flier program data to show how success fosters reengagement; successful frequent fliers begin the next year flying more frequently. In addition, reaching the goal of earning status affects a flier's likelihood of success in subsequent attempts at earning status. In Study 2, we show that increasing divisibility (from $1,000 to $500), which reframes a larger task as several smaller tasks and disburses a reward in smaller increments, can boost effort and consumers' willingness to reengage. Conversely, the study showed that too much divisibility ($100) was demotivating. Therefore, loyalty programs that offer people multiple redemption opportunities must balance the attractiveness of a reward with an appropriate level of difficulty in attaining success. If the only reward offered was a television for one million points, this may prove too difficult and thus unattractive for most consumers. Conversely, offering magazine subscriptions for 100 miles may be uninspiring. Consider American Express highlighting that rewards points can be used at 160 stores for more than two million products. While this might encourage consumers to carry an American Express card, our work indicates that loyalty may be better accomplished with more moderate divisibility, such that it is not so easy to cash out rewards points, but when someone does, he or she feels successful.

Finally, in Study 3, we controlled the process such that participants experienced the exact same sequence of success and failure on individual trials: Their hit and miss rate remained constant. However, when we framed getting 70% correct in a sequence of 10 as the successful completion of a task accompanied by a reward, respondents elevated their perceptions of self-efficacy, showing that success indeed leads to learning—that is, learning about one's abilities.

All three studies showed that more than one success matters. Thus, not only does goal attainment result in increased effort the second time around, but successive successes further elevate effort. Although we show how successful goal attainment and the accompanying reward serves to increase consumers' efforts in subsequent undertakings, it is important to emphasize that this increase can occur after more than one or two successes.

This research is not without its limitations. We suggest that goals or reward levels should be challenging enough to foster self-learning, but the precise level of difficulty will vary across situations. We do not offer a methodology for determining this level but rather make managers aware of how changes in divisibility (i.e., adding or deleting reward levels) might affect their programs. We also did not examine the impact of success on people's emotional response, such as the feelings of satisfaction and pleasure it should elicit (Isen 1987). It is possible that affect regulation, a process in which the positive affect associated with the reward is sought after, moderates persistence. However, it is not clear whether positive feelings would encourage people to try again or whether people would attempt to protect the positive state associated with success by not trying again. Further research could examine the determinants of these opposing responses.

In Study 3, we show how successfully reaching a goal and earning a reward results in enhanced self-learning and perceptions of self-efficacy, though outside the context of loyalty programs. We used a context pretested to ensure that our respondents had little previous experience and were uncertain about their abilities. We should point out that judging facial expressions is not entirely dissimilar to consolidating purchases in that perceived self-efficacy does not involve the subskills necessary to achieve the goal (e.g., observing subtle facial cues, detecting a change in mannerisms, interpreting gestures) but whether the person feels efficacious. In other words, determining whether someone is lying and purchasing successive flights to accumulate miles include both unpredictable and ambiguous elements.

From a practical perspective, all possible successes may not be entirely within the control of the firm. For example, frequent-flier miles, by far the most ubiquitous alternative currency, are becoming interchangeable with other alternative currencies; some can now be redeemed at numerous second-party vendors. This could result in outside rewards that qualify as successes. Earning 25,000 miles for a free ticket may no longer be the quintessential goal. The firm's goals might be diluted by such changes in the marketplace. Other firms, such as Southwest Airlines, are strategic in limiting the uses of their alternative currency by issuing credits in proprietary, nondivisible increments (flight segments). Fliers on Southwest quickly learn whether consolidating eight flights is something they can or cannot do repeatedly. However, the airline's less divisible currency (segments as opposed to miles) limits its flexibility as a reward mechanism, which has led Southwest partners to begin issuing credits in partial segments. The firm's objective should be to identify the right amount of divisibility—that which engages high-value customers and enables them to reach their goals, thus keeping them coming back for more.

REFERENCES

Bandura, Albert (1982), "Self-Efficacy Mechanism in Human Agency," *American Psychologist*, 37 (2), 122–47.

——— (1988), "Perceived Self-Efficacy: Exercise of Control Through Self-Belief," in *Cognitive Perspectives on Emotion and Motivation*, J.P. Dauwalder, M. Perrez, and N.H. Frida, eds. Dordrecht, the Netherlands: Kluwer Academic, 37–61.

Berman, Barry (2006), "Developing an Effective Customer Loyalty Program," *California Management Review*, 49 (1), 123–48.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer (2006), "$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotion," *The American Economic Review*, 96 (4), 1253–70.

DeWulf , Kristof, Gaby Odekerken-Schroder, and Dawn Iacobucci (2001), "Investments in Consumer Relationships: A Cross-Country and Cross-Industry Exploration," *Journal of Marketing*, 65 (October), 33–50.

Dzewaltowski, David A., John M. Noble, and Jeff M. Shaw (1990), "Physical Activity Participation: Social Cognitive Theory Versus the Theories of Reasoned Action and Planned Behavior," *Journal of Sport and Exercise Psychology*, 12 (4), 388–405.

Ferguson, Rick and Kelly Hlavinka (2009), "The Big Sort: The 2009 Colloquy Loyalty Marketing Census," *COLLOQUYtalk*, (April), 1–20.

Garland, Howard (1984), "A Cognitive Mediation Theory of Task Goals and Human Performance," *Motivation and Emotion*, 9 (4), 345–67.

Hahn, Jinyoung, Petra Todd, and Wilbert Van der Klaauw (2001), "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69 (1), 201–209.

Hoch, Stephen J. and Young-Won Ha (1986), "Consumer Learning: Advertising and the Ambiguity of Product Experience," *Journal of Consumer Research*, 13 (2), 221–33.

Isen, Alice (1987), "Positive Affect, Cognitive Processes, and Social Behavior," in *Advances in Experimental Social Psychology*, L. Berkowitz, ed. New York: Academic Press, 203–253.

Kivetz, Ran, Oleg Urminsky, and Yuhuang Zheng (2006), "The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention," *Journal of Marketing Research*, 43 (February), 39–58.

Lewis, Michael (2004), "The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention," *Journal of Marketing Research*, 41 (August), 281–92.

Mägi, Anne W. (2003), "Share of Wallet in Retailing: the Effects of Customer Satisfaction Loyalty Cards and Shopper Characteristics," *Journal of Retailing*, 79 (2), 97–106.

Mento, A.J., N.D. Cartledge, and Edwin A. Locke (1980), "Maryland vs. Michigan vs. Minnesota: Another Look at the Relationship of Expectancy and Goal Difficulty to Task Performance: 1966–1984," *Organizational Behavior and Human Decision Processes*, 25 (3), 419–40.

Nunes, Joseph C. and Xavier Drèze (2006a), "The Endowed Progress Effect: How Artificial Advancement Increases Effort," *Journal of Consumer Research,* 32 (4), 504–512.

——— and ——— (2006b), "Your Loyalty Program Is Betraying You," *Harvard Business Review*, (April), 1–9.

Ryan, T.A. (1970), *Intentional Behavior*. New York: Ronald Press.

Shugan, Steven M. (2005), "Brand Loyalty Programs: Are They Shams?" *Marketing Science*, 24 (2), 185–93.

Taylor, Gail Ayala and Scott A. Neslin (2005), "The Current and Future Sales Impact of a Retail Frequency Reward Program," *Journal of Retailing*, 81 (4) 293–305.

Verhoef, Peter C. (2003), "Understanding the Effect of Customer Relationship Management Efforts on Customer Retention and Customer Share Development," *Journal of Marketin*g, 67 (October), 30–45.

Wood, Robert and Albert Bandura (1989), "Social Cognitive Theory of Organizational Management," *Academy of Management Review*, 14 (3), 361–84.