Market Lett (2010) 21:149–162
DOI 10.1007/s11002-009-9089-y

Two-stage lotteries and the value of unresolved uncertainty

Ido Erev · Ernan Haruvy

Published online: 15 September 2009

Abstract We study two-stage lotteries wherein, in the first stage, the consumer may be awarded a lottery ticket for some choices and, in the second stage, some of the lottery tickets win prizes. In a series of Internet experiments, we examine store choice and repeat visits in response to immediate and two-stage lottery incentives. We show that a delayed resolution of uncertainty may dramatically increase the desired response over immediate resolution. However, this effect is due to an interaction of several potentially contradictory effects. When prizes are frequent, the effectiveness of the two-stage lottery is expected to decline.

Keywords Reinforcement learning · Lottery · Secondary reinforcements · Anticipation · Multiple contacts

1 Introduction

Consider the example of a consumer facing a choice between two retailers that are comparable to one another on product selection and prices. One way to increase store traffic in such a competitive environment is through the use of prize promotions. If a consumer is to receive a prize for visiting a store, which prize would be more effective—an immediate prize or a prize drawing at a later date? Extant research indicates that most consumers prefer immediate prizes. However, the preference by consumers for the more immediate prize does not necessarily make it more effective in increasing the frequency of visits to the promoting store. In this

I. Erev
Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel
e-mail: erev@tx.technion.ac.il

E. Haruvy (✉)
School of Management, University of Texas at Dallas, P. O. Box 830688, Richardson, TX 75083, USA
e-mail: eharuvy@utdallas.edu

work, we examine the possibility that a delayed prize drawing will be more effective at increasing store traffic.

The effectiveness of a prize promotion is typically measured by frequency of response. As such, basic concepts from *operant conditioning* (Skinner 1953) are useful. According to Nord and Peter (1980), increasing the frequency of desired behavior crucially depends on three important aspects. First is the *power* of *positive reinforcement*. A positive reinforcement is defined as a reward that serves to increase the frequency of a given behavior (Nord and Peter 1980; Peter and Nord 1982; Rothschild and Gaidis 1981). Its *power* (Rothschild and Gaidis 1981; Peter and Nord 1982) is loosely defined as the effectiveness of the reward stemming from its desirability. Second is the *timing of reinforcement*. *Immediate reinforcement* follows immediately after the desired behavior. *Delayed reinforcement* takes place after time has passed following the desired behavior. Peter and Nord (1982) note that "most research has shown that immediate reinforcement is more effective than delayed reinforcement". Early empirical works that established the effectiveness of immediate reinforcers include Mowrer and Ullman (1945), Kamin (1959), and Walters and Demkow (1963).[1] More recently, the preference for immediacy has been shown in the context of mental accounting (e.g., Thaler 1981; Thaler and Sheffrin 1981).[2]

The third important aspect is the *reinforcement frequency*—how often a subject receives a positive reinforcer. Nord and Peter (1980) note that the *variable ratio schedule*, where the reinforcer is awarded probabilistically rather than deterministically (e.g., every second time or third time), is the most effective in producing persistently high rates of behavior. Accordingly, one focus of this article is on the probability of receiving the reinforcer.

The prize itself is a positive reinforcer, but it may be received by the consumer too infrequently to generate a substantial effect over time. Here, we propose adding *lottery tickets* for a future prize drawing as an alternative or complementary form of reinforcement. Lottery tickets are defined as any object which communicates a specific probability of winning a future prize. As secondary reinforcers, lottery tickets are clearly weaker reinforcers than the ultimate prize itself (Rothschild and Gaidis 1981).[3] However, to the extent they serve as reinforcers, they can be made as frequent as needed. A prize promotion that uses *lottery tickets* as reinforcers can be thought of as a *two-stage lottery*, where, in the first stage, the reinforcement schedule pertains to the award of the lottery ticket itself and, in the second stage, a drawing takes place to determine the final prize. In this case, the reinforcement frequency for a given instance of the desired response is determined by multiplying the probability of receiving the lottery ticket by the probability of the ticket winning a prize.

[1] In Walters and Demkow (1963), an experimenter forbids a child to play with a popular toy and then leave the child alone in the room. A loud buzzer sound followed any attempt to play with the toy. It was shown that when the sound followed immediately, it was much more effective in discouraging such attempts than the same unpleasant sound with some delay.

[2] As with every body of research, there are notable works that suggest that preference for immediacy may be violated (e.g., Ronen 1971; Hirsch 1978; Loewenstein 1987), but these do not directly apply in the present setting.

[3] Primary reinforcers have intrinsic utility, while secondary reinforcers have no such utility and must be converted. They are thus generally considered less powerful reinforcers (Rothschild and Gaidis 1981).

The current work's main contribution is demonstrating that lotteries can be conceptualized as reinforcers and, therefore, unresolved uncertainty may actually improve the effectiveness of reinforcement. As long as consumers ascribe a nonzero probability to the event that the lottery ticket will result in a final prize, then receiving the lottery ticket may itself provide a degree of reinforcement. Our goal is to explore whether the standard result regarding timing—i.e., that immediate reinforcement is superior to delayed reinforcement—will hold in a lottery context. The timing of the uncertainty resolution and probability of winning determine the frequency and valence of reinforcement.

We approach the prize promotion as a problem of mechanism design. Specifically, we break the prize promotion decision into two elements—timing and frequency—fix other important elements—particularly the prize budget—and abstract from other design aspects. Based on insights from the extant literature, we propose that delaying the resolution of uncertainty of winning a large prize and instead offering lottery tickets towards winning that prize can improve the frequency of reinforcement and therefore result in increased effectiveness of the promotion, while maintaining the desirable effects of small probability events. We find in three sets of field implementations of these recommendations that the results support the proposed mechanism in the specific settings we looked at.

The context of Studies 1 and 3 is a repeated choice among online music stores. One of the two stores provides a possible prize (a prize promotion) and the other does not. The expected consumer surplus from the two stores is the same. We want to see how manipulating the prize distribution, particularly the introduction of lottery tickets to the promotion mix, can increase the frequency of choice of the store with the prize promotion. The main difference between study 1 and study 3 involves the number of instances of the final prize. In study 1, we examine a rare large prize, whereas in study 3 we examine a frequent small prize. We keep the expected value of the promotion (the probability of winning × the size of the prize) the same across the two studies.

The context of study 2 is maximizing visits to a recruitment website that was in use for experimental studies at the university. In contrast to Studies 1 and 3, the only choice subjects get is to visit the site on a given day or not, and the dependent variable of interest is the frequency with which subjects visited the recruiting site.

The results from Studies 1 and 2 show that the two-stage lottery design can dramatically increase the effectiveness of promotions relative to immediate resolution when the large prize is rare. Study 3 demonstrates that the lottery procedure does not increase effectiveness when the final prize is frequent to begin with.

2 Study 1

The objective of study 1 is to examine whether it can be beneficial to the success of a prize promotion to introduce lottery tickets for a prize instead of awarding the prize directly and immediately.

The context of the investigation is choice among virtual stores on the Internet. We examine four conditions detailed by Table 1. First, we examine a condition where we

Table 1 The experimental conditions of study 1

Condition	Probability of winning an immediate prize	Probability of getting a lottery ticket (p_1)	Probability that the lottery ticket results in winning a final prize (p_2)
1	1/100	0	NA
2	0	1/3	3/100
3	0	2/3	3/200
4	0	1	1/100

provide an infrequent but immediate prize (condition 1) and compare that to conditions 2–4, where we increase the frequency of reinforcement through lottery tickets to be followed by a later drawing.

We vary the frequency of the lottery reinforcement between conditions 2–4 in order to check the robustness and sensitivity of the effect of the lottery reinforcement to different reinforcement frequencies. Conditions 2–4 differ on the probability of getting a lottery ticket (p_1) as well as the probability that a lottery ticket results in winning the prize (p_2). However, we keep the prize the same at 100 tokens ($10) in all four conditions and fix the prize budget so that the expected value of the prize is 1 token ($0.10) in all four conditions. This means that the unconditional probability of winning the prize (probability of receiving a lottery ticket×probability that a ticket wins) remains the same (1/100) across conditions in this study.

The probabilities in the four different conditions translate to different frequencies of reinforcement. If one accounts for the lottery ticket as a positive reinforcer, then the expected frequency of reinforcement per 100 instances of the desired response is 1 in condition 1, 33 in condition 2, 66 in condition 3, and 100 in condition 4. Hence, condition 1, where the prize is immediate, is the lowest frequency reinforcement condition, and condition 4, where the lottery ticket is received with every desired response, is the highest frequency reinforcement condition. In the operant conditioning literature (Nord and Peter 1980; Peter and Nord 1982), condition 4 in study 1 would be known as a continuous reinforcement schedule since subjects receive a lottery ticket for every visit.

2.1 Method

2.1.1 Subjects

Recruitment posters around campus and announcements in classes invited students to participate in an Internet experiment. Upon responding to this ad, the responder was given a web address that corresponded to one of four experimental conditions. One hundred sixty-three subjects participated in the four conditions, where 44, 42, 36, and 41 subjects participated in conditions 1, 2, 3, and 4, respectively.

2.1.2 Procedure

The web site provided the instructions for the relevant condition. It explained that the experimental task was to purchase 40 CDs over a period of 2 weeks, with at least 1 h between purchases. The participants received an initial budget of 500 tokens and

were told that their final payoff would be determined by the number of tokens left in their account at the end of the experiment. The tokens would be converted to money with a conversion rate of 10 tokens to $1. In addition, one in ten participants would be randomly chosen to receive one of the CDs they selected.

The CDs could be purchased in either of two virtual stores. The price of each CD was 9 or 10 tokens depending on the store. The more expensive store had a prize of 100 tokens with probability 0.01. Specifically, the instructions stated: "If you visit Store A, you can win 100 tokens. Each visit has a 1 in 100 chance of winning. The chance of winning the prize in each visit is independent of other visits". Thus, the participants were asked to choose between a cheaper store and a more expensive store that offered a lottery. The objective attractiveness of the two stores (in term of expected cost and quality of the CDs) was identical and known to the participants.

2.1.3 Procedure for shopping trips

Upon any given shopping trip, the participants were instructed to click on whichever store they chose to shop at. Upon entering a store, the participants were given a choice between two CDs. Pictures of the CDs were shown. The CDs displayed were randomly determined from a database of 70 available CDs, with equal probability for any CD to be displayed. The same database of CDs was used for both stores and in all conditions. The CDs in this database were carefully pre-tested with a random sample of the subject population for recognition and desirability. The initial sample for pre-testing came out of Amazon top recommendations. Following pre-testing, we picked an equal mix of older well-established artists and newer top of the charts artists and relatively balanced mix of genres. Artists in the data of 70 CDs ranged from Alicia Keys, Shakira, Beyonce, and Eminem to Madonna, Kenny G., Bob Marley, Pink Floyd, Prince, Celine Dion, Shania Twain, Janet Jackson, and others. Since subjects have different tastes and music backgrounds, not all subjects recognized all artists, but all artists came up high on average on recognition and likability.

2.1.4 Design

Four experimental prize conditions were compared. The prize (100 tokens) and the probability of ultimately winning it (p_w=1/100) were the same in all four conditions. Table 1 summarizes the four conditions.

The four conditions differed with regard to the timing of the resolution of the uncertainty. In condition 1, the immediate condition, the uncertainty was resolved immediately after the choice. In other words, there is only one stage—stage 1. Subjects in condition 1 each received a lottery ticket on every visit to the promoting store, with a 1/100 chance to win 100 tokens, and they learned whether they won or not upon entering the store (immediate resolution). That is, subjects in the immediate condition got a chance for a prize with every visit. Subjects in conditions 2–4 experienced delayed resolution. Only a portion of purchases received a lottery ticket when they entered the promoting store (depending upon p_1), and then those that received a ticket learned at the end of the experiment whether or not they won. Subjects in condition 2 received a lottery ticket with only 1/3 probability (so two thirds of the subjects entering the store in this condition did not receive a lottery ticket). Those who received a ticket then had a 3/100 probability of winning 100

tokens, with the outcome determined at the end of the experiment. Conditions 3 and 4 differed from condition 2 in that the probability of winning the ticket was 2/3 and 1 in conditions 3 and 4, respectively, and the probability of winning the prize with each ticket was 3/200 and 1/100 in conditions 3 and 4, respectively.

2.2 Results

Figure 1 presents the proportion of choices of the promoted store over the 40 trials in each of the four conditions. The results show that all four promotion procedures were effective: Although the expected value from visiting the two stores is identical (cost of 9), the observed choice rates of the promoted store were significantly higher than 50% in all four conditions. The exact rates and t test statistics are 67% ($t[28]=2.22$, $p<.01$), 86% ($t[41]=8.10$, $p<.01$), 88% ($t[35]=8.76$, $p<.01$), and 81% ($t[40]=5.93$, $p<.01$) in conditions 1 (immediate condition), 2, 3, and 4 respectively.

Note that condition 1—the immediate condition—led to the lowest choice rate of the promoted store. The other conditions led to higher choice rate of the promoted store than the immediate condition ($t[84]=2.54$, $p<.01$; $t[78]=2.73$, $p<.01$; and $t[83]=1.73$, $p<.10$; for conditions 2, 3, and 4, respectively).

Examination of the three conditions with delayed resolution reveals similar proportions of choices of the promoted store (86% in conditions 2, 88% in condition 3, and 81% in condition 4). The difference between the three conditions is insignificant ($F[2,116]=0.62$, $p>.50$).

Over time, the learning curves in Figure 1 appear relatively flat. The attractiveness of the promoted store slightly increases in condition 3 and slightly decreases in the other conditions, but the difference between the slopes of the linear curves of conditions 2, 3, and 4 is insignificant ($F[3,472]=0.86$, $p>.40$).

2.3 Discussion

We find that a delayed chance at winning a prize is better than an immediate chance. What could explain this apparent preference for delayed rewards? The answer may

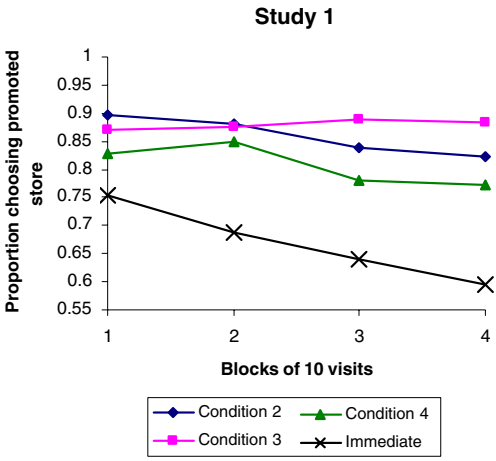

Fig. 1 Proportion of visits to the promoted store in study 1 as a function of condition and time

be that the lottery tickets themselves serve as reinforcers (albeit secondary, Rothschild and Gaidis 1981). If so, then the very low reinforcement frequency of the immediate reward condition (1/100) is inferior to conditions 2–4 (the Delayed reward conditions) where subjects are receiving frequent reinforcement in the form of lottery tickets, which carry with them the chance of a later prize.

Another likely interpretation for the difference between the immediate and delayed conditions is that subjects in the immediate condition lost (did not receive a reward) almost every time, which is discouraging.[4] In the delayed conditions, subjects did not know during the experiment whether the tickets they received (i.e., positive reinforcement) had won, so they kept on gathering chances to win. The similarity of the three conditions with delayed resolution suggests that as long as reinforcement is received with a sufficiently high probability in the course of the experiment (p_1 between 1/3 and 1), the prize promotion will be effective.

Hence, in terms of the effects discussed in Section 1, the effect of frequency dominates the effect of timing in study 1. The frequency of reinforcement appears most critical in this set of conditions. In that context, the lottery tickets and delayed resolution permitted us to increase the frequency of reinforcement without altering the expected cost of the promotion. Though the introduction of the lottery tickets resulted in a delayed resolution of uncertainty, the delay was not detrimental in this case.

Note that the effect of frequency did not appear to improve the success of the promotion between conditions 2, 3, and 4, despite the increased reinforcement frequency. This may be due to a diminishing incremental value of reinforcement frequency once frequency achieves a moderate level. In other words, increasing frequency from 1 in 100 to 1 in 3 is more consequential than increasing it from 1/3 to 2/3 or from 2/3 to 1/1.

3 Study 2

The main goal of study 2 is to explore the applicability and robustness of the results of study 1 in a different setting. Whereas study 1 focused on participants' choices among two online music stores, study 2 focuses on the tendency to enter a promoted web site.

Another goal of study 2 is to examine how prize choice might influence consumer preferences. The goal is particularly important in light of the assertion early in Section 1 that consumers would typically prefer an immediate prize.

3.1 Method

3.1.1 Subjects

Recruitment posters around campus invited subjects to visit a web site used for the purpose of recruitment for experiments. Subjects were informed that the experiments

[4] Rothschild and Gaidis (1981) noted that the lack of reward can be perceived as negative reinforcement (punishment): "Intermittent reinforcement…may be seen as punishment by the consumer" (p. 72).

advertised on the site would not be announced on any other forum. Instead of the web site's address, the posters only listed an email address. Upon sending an email, a subject was given a web address that corresponded to one of two conditions—immediate and delayed. A separate group of subjects was recruited into a third condition—the survey condition—intended to examine preferences over the immediate and delayed conditions. A total of 171 subjects participated in this study. Of these subjects, 53 were in the immediate condition, 64 were in the delayed condition, and 54 were in the survey condition. The subjects were all students and predominantly undergraduates.

3.1.2 Design

As explained above, there were two recruitment conditions—immediate and delayed—and one survey condition. Under the immediate condition, subjects were told instantly of the outcome of the lottery. In the delayed condition, with probability $p_1 = 1/3$, subjects received a message that they had won a lottery ticket that would give them the prize money of \$100 at the end of the month with probability of $p_2 = 3/250$. In the survey condition, a different group of subjects was asked to indicate a preference over the two prize promotions of immediate and delayed.

3.1.3 Procedure for recruitment conditions

Upon logging in, subjects were told that experiments would run in the coming weeks and were asked to log in daily to check for recruitment announcements. They were further told that as an added incentive, they would be eligible to enter a drawing for \$100 once a day (and no more than once a day). Each visit resulted in a chance of 1/250 to win \$100. Chances were independent of each other and other subjects. The experiment lasted 72 days.

In both conditions, subjects saw the following message upon visiting the site: "Welcome to the Experiment Registration Page. What do you need to do to get paid? Occasionally researchers will post studies on this web page. You will need to log in to the studies page as often as you can to find out about studies. To encourage you to check this web site often, we are offering prize money. Every day you log in, you may win (a lottery ticket for a drawing for a prize of \$100) \$100. One in 250 logins will win \$100. Winnings are independent of each other, so if you log in today and win, this will not affect you chance of winning again tomorrow. To log in, you simply need to provide your name and email".

Upon entering their name and email in the appropriate textbox, subjects observed the response, which was "sorry, you did not win this time" or, depending on the condition, either a lottery ticket ("Congratulations. You won a lottery ticket for the \$100 prize! One in 83 tickets will win \$100") or an immediate prize announcement ("Congratulations. You won the \$100 prize!").

3.1.4 Procedure for survey condition

To examine preferences over immediate and delayed, a separate group of 54 undergraduate students was asked to indicate preferences over two scenarios.

Subjects were told that announcements for experiments would be posted on two websites, website A and website B, and that each website would reward visits with a chance to win $100 according to a scenario corresponding to one of the above conditions. Website A corresponded to the delayed condition and website B corresponded to the immediate condition. A detailed description of the prize distribution and odds in each condition was provided to the subjects. Subjects were asked to choose between A and B and then to provide an estimate for the number of days per month they would visit the web site they chose. Unlike all the other conditions described in this paper, this task was a survey and not a paid experiment.

3.2 Results

Figure 2 presents the main experimental results: The proportion of entries as a function of time in the two conditions. Over the 72 days, the average participant entered 6.81 times (SD=8.33) in the immediate condition and 11.64 (SD=15.17) in the delayed condition. The difference between the two conditions is significant ($t[115]=-2.07$, $p<.05$). The advantage of the delayed procedure appears to be robust over time.

Survival analysis reveals that the stark difference between the conditions occurred between 20 and 50 days into the subjects' visiting duration. In that duration interval, subjects in delayed condition were more than 50% as likely to survive. For example, 37.5% of the delayed condition subjects survived past 20 days. In contrast, only 24.5% of the immediate condition subjects survived past 20 days.

In addition to examining the frequency of visits to the web site, we asked an independent group of survey respondents to indicate which of the two prize promotions they preferred. When survey participants were asked to choose between the two designs, 33 out of 54 respondents expressed a preference for the immediate design. Subjects were also asked to provide an estimate for the number of days they would visit their chosen web site in a month. The estimates were 23 and 18 days for immediate and delayed, respectively, but the difference was not significant.

3.3 Discussion

The immediate benefit of the above experiment was in the improvement of our recruitment methods. With identical effort and cost, we are now able to improve our reach to potential experimental subjects by over 71%. The more frequent reinforce-

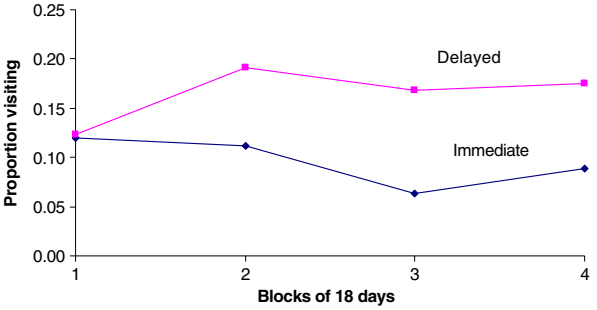

Fig. 2 Proportion of visits to the web site in study 2 as a function of condition and time

ments result in more repetitions of the desired response, even though the expected monetary reward remains the same.

Given the overwhelming success of the delayed prize promotion in generating more visits to the web site than the immediate prize promotion, it is important to ask whether this success is due to the inherent liking the subjects had for the delayed promotion or due to the more frequent reinforcement structure. The apparent preference for the immediate prize promotion indicated by the survey responses suggests that if subjects were allowed to choose their promotion, they would prefer the immediate prize promotion. However, when faced with the two reinforcement schedules, the delayed promotion is more successful from the point of view of the promoter. This suggests that reaction from experience fundamentally differs from reaction to verbal description of the problem.

4 Study 3

Study 1 showed that frequent lottery tickets for a delayed prize could be very beneficial to the success of a prize promotion relative to an immediate but *infrequent* immediate prize. In study 3, we increase the reinforcement frequency for the subjects in the immediate condition. To maintain equivalence between conditions in study 3, we increase the second stage probability of winning for the delayed condition relative to study 1.

The working hypothesis so far has been that the benefit from delayed resolution of uncertainty lies in the ability to make reinforcement *more* frequent through lotteries for the delayed prize. When the immediate prize is rare, this increase in frequency through lotteries is clearly a beneficial endeavor because it helps ensure desired behavior will not decrease over time. While there are obvious disadvantages to delay, stemming from preference for immediate rewards, they are outweighed by the benefit of frequency. However, when the immediate prize is already sufficiently frequent to sustain desired behavior, then there is no benefit to delaying the prize. That is, the beneficial effect of delay may not hold if the prize is frequent to begin with. The goal of study 2 is to examine whether the main result of study 1 holds for a prize that is frequent.

To do this, we examine a condition where we provide a small but frequent immediate prize and compare that to a condition where we make the prize delayed through lottery tickets (as in study 1). In comparison to study 1, we varied both the prize and the probability of winning it. However, we fixed the prize budget so that the expected value of the prize is 1 in all conditions of both Studies 1 and 2. From an economics perspective, expected value (to both the consumer and the store) is the key determinant in decision making, and it was deemed critical to keep this constant in order to avoid undesirable confounds from an economics viewpoint. Given the constraint of fixed expected value, the unconditional probability of winning the prize follows from the prize itself.

4.1 Method

4.1.1 Subjects

Subjects were university students who were recruited through posters and announcements in classes. Upon responding, each responder was given a web address that

corresponded to one of two experimental conditions. Sixty-three subjects participated in two conditions, refereed to as immediate and delayed. Thirty-one and 32 subjects participated in conditions immediate and delayed, respectively.

4.1.2 Procedure

We use the same procedure as in study 1—shopping for CDs in virtual stores—but with frequent prizes that are realized with probability 5/6. Briefly, subjects are instructed to purchase 40 CDs. They are told they will have a chance to receive one of the CDs they purchase and that their earnings will be reduced by the price of the CDs. Upon any given shopping trip, the participants were instructed to click on whichever store they chose to shop at. One store had a low price and the other had a higher price as well as a prize drawing.

4.1.3 Design

We study two conditions, immediate and delayed. Table 2 shows the details of the two conditions.

The expected value of the promotion was 1 in both conditions and identical to the four conditions of study 1. To eliminate any confound with the probability of the lottery, the probability of receiving the lottery (p_1) ticket is fixed at 1. Unlike study 1, where the unconditional probability of winning the prize was 1/100, here it is 5/6. In condition immediate, shoppers receive a prize of 1.2 in an average of five out of six visits to the promoted store. In condition delayed, shoppers receive a lottery ticket with every purchase at the promoted store, and these lottery tickets have a probability of 5/6 of winning a 1.2 token prize, where tickets are resolved independently.

4.2 Results

Table 3 presents the proportion of choices in the promoted store in study 3 and the similar conditions (same p_1) in study 1.

The exact rates are 62.5% (SD=40.1%) in the immediate condition, and 63.5% (SD=38.7%) in the delayed condition. A comparison of study 1 and study 3 shows that the increase in the frequency of the immediate prize (the change from study 1 to study 3) eliminated the positive effect from the delayed resolution of uncertainty. That is, delayed resolution helped the promoting store in study 1 (low probability), but not in study 3 (high probability).

Table 2 The experimental conditions of study 3

Condition	Probability of winning an immediate prize	Probability of getting a lottery ticket (p_1)	Probability that the lottery ticket results in winning a final prize (p_2)
Immediate	5/6	0	NA
Delayed	0	1	5/6

Table 3 Comparison of the choice rate of the promoted store in study 1 and study 3

Condition	Study 1	Study 3	*p* value
Immediate	0.67	0.64	.630
Delayed	0.81	0.63	.048

4.3 Discussion

We believe that the results support the assertion that the observed behavior reflects a joint effect of frequency and timing. The two effects lead to contradicting predictions in study 3. Frequency implies a positive effect for delayed resolution, and timing implies a negative effect. Thus, if one interprets the finding of study 1 as evidence in favor of the frequency effect, the similar rates observed in the two conditions of study 3 could be explained with the assertion that the two effects cancel each other out (alternatively, neither effect might be present).

Figure 3 presents the observed learning curves. The figure depicts the proportion of choices of the promoted store in four blocks of ten visits in both conditions of study 3.

In comparing study 1 and study 3, the key construct is the power of the reinforcer. As discussed in Section 1, this is loosely translated as the attractiveness or desirability of the reward (Rothschild and Gaidis 1981; Peter and Nord 1982). The most attractive reward in all store choice conditions (Studies 1 and 3) is the immediate reward of 100 tokens. It is also the least frequent and therefore not very effective. The next in desirability is a lottery ticket for the chance to win 100 tokens (study 1, conditions 2–4). While less desirable, this lottery ticket is more frequent and therefore more effective. In study 3, the potential prize is 1.2 tokens. This is not very desirable but it is frequent. Looking at the immediate conditions in Studies 1 and 3, we see that the two are quite comparable, so there is an offsetting tradeoff between power and frequency. The frequency of the immediate prize in study 3 is sufficient in achieving the goal of sustaining the desired behavior. So, the benefit from increasing the frequency through lotteries does not outweigh the cost of reducing power in study 3. This, we think, is why the effect of the lottery is not positive in this study.

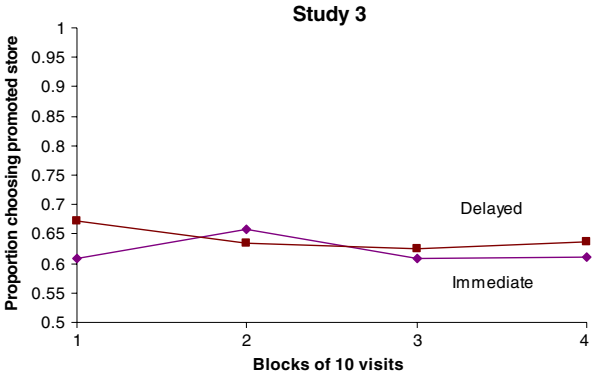

Fig. 3 Proportion of visits to the promoted store in study 3 as a function of condition and time

5 Discussion and implications

Prize promotions have received a great deal of attention in the academic literature. A common wisdom hitherto has been that immediate prizes are more effective than delayed prizes. In this work, we challenged this wisdom with a demonstration of the effectiveness of lottery tickets for delayed prizes. The reason that the lottery tickets were effective is that they increased the frequency of reinforcement. That is, both frequency and timing need to be considered when designing a prize promotion.

While promotion effectiveness improved with frequency of reinforcements through lottery tickets when the immediate prize was rare, we did not observe additional benefits to frequency when moving from a frequent lottery to a somewhat more frequent lottery. Nor did we observe a positive effect when moving from a frequent immediate prize to a more frequent lottery. That is, the main effect is strongest when moving from a rare immediate prize to a frequent lottery.

We presented three Internet studies that evaluated the effect of the timing of uncertainty resolution. The results demonstrate that when the probability of winning the prize is low, delayed resolution may be more effective than immediate resolution of uncertainty.

Study 1 gave participants a choice between two Internet stores selling CDs. One of the stores used a prize promotion with attractive but infrequent prizes. The results suggest that the effectiveness of the prize promotion is highest with a two-stage prize implementation.

In study 2, participants were asked to visit a website as often as possible and were given monetary prizes as incentive to do so. Unlike study 1, where the focus was on store choice, study 2 focused on frequency of visits to the web site. The purpose of the study was to examine whether the advantage of the two-stage promotion could be replicated in this different task. It showed that two-stage promotions could significantly increase the frequency of visits to a promoted web site.

Study 3, in an Internet store choice environment, sought to establish boundary conditions when prizes can be frequently awarded. It demonstrated that the benefit of the two-stage procedure disappeared when immediate prizes were already frequent.

The results of all three studies suggest that both frequency and timing contribute to the positive effect of delayed resolution of uncertainty. Under this interpretation, the effectiveness of the two-stage procedure is maximized when it provides frequent—but not too frequent—rewards. This interpretation explains the observation that in the context of lotteries, consumers behave “as if” they “prefer” delayed resolution of uncertainty. However, it is important to stress that two-stage lottery is not “preferred” over immediate resolution in the usual sense of the word. A questionnaire in the second part of study 2 showed that subjects, when asked what they “preferred”, actually stated that they preferred immediate resolution. Instead of preference, the effectiveness of the two-stage mechanism hinges on the frequency of reinforcement and the impact on choice that such higher frequency entails. When the prize is rare, replacing the rare immediate prize by frequent lottery tickets increases the frequency of reinforcement and therefore increases the likelihood of the reinforced action. This explanation also makes it clear why study 3 did not show the same effect observed by study 1 and study 2. In study 3, the prize was frequent, so the reinforcement was already frequent to begin with. Adding lottery tickets is not going to increase the frequency of reinforcement, and

therefore, from a reinforcement perspective, it is not expected to have an effect. It is therefore important to stress that two-stage lotteries did not appear to have an effect beyond their reinforcement value.

References

Hirsch, M. (1978). Disaggregated probabilistic accounting information: the effect of sequential events on expected value maximization decisions. *Journal of Accounting Research, 1978*, 254–269.

Kamin, L. J. (1959). The delay of punishment gradient. *Journal of Comparative & Physiological Psychology, 52*, 434–437.

Loewenstein, G. (1987). Anticipation and the value of delayed consumption. *Economic Journal, 97*, 666–684.

Mowrer, O. H., & Ullman, A. D. (1945). Time as a determinant of integrative learning. *Psychological Review, 52*, 61–90.

Nord, W. R., & Peter, J. P. (1980). A behavior modification perspective on marketing. *Journal of Marketing, 44*(2), 36–47.

Peter, J. P., & Nord, W. R. (1982). A clarification and extension of operant conditioning principles in marketing. *Journal of Marketing, 46*(3), 102–107.

Ronen, J. (1971). Some effects of sequential aggregation in accounting on decision making. *Journal of Accounting Research, 9*(2), 307–332.

Rothschild, M. L., & Gaidis, W. C. (1981). Behavioral learning theory: its relevance to marketing and promotions. *Journal of Marketing, 45*, 70–78.

Skinner, B. F. (1953). *Science and human behavior*. New York: Macmillan.

Thaler, R. H. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters, 8*(3), 201–207.

Thaler, R. H., & Sheffrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy, 89*, 392–406.

Walters, R. H., & Demkow, L. (1963). Timing of punishment as a determinant of response inhibition. *Child Development, 34*, 207–214.