Journal of Behavioral Decision Making
J. Behav. Dec. Making, (2008)
Published online in Wiley InterScience
(www.interscience.wiley.com) **DOI**: 10.1002/bdm.593

Future-Biased Search: The Quest for the Ideal

SUZANNE B. SHU*
Southern Methodist University, Texas, USA

ABSTRACT

Decision-makers with ideal candidates already in mind often extend search beyond optimal endpoints when searching for the best option among a sequential list of alternatives. Extended search is investigated here using three laboratory experiments; individuals in these tasks exhibit future-bias, delaying choice beyond normative benchmarks. Searchers' behavior is consistent with setting high thresholds based on a focal ideal outcome without full attention to its probability or the value of second-best alternatives; the behavior is partially debiased by manipulating which outcomes are in the searchers' focal set. Documenting future-bias in sequential search tasks offers new insights for understanding self-control and intertemporal choice by providing a situation in which thresholds may be set too high and myopic behavior does not prevail. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS sequential search; optimal stopping; focalism; intertemporal choice

INTRODUCTION

In an article titled 'An endless odyssey: Looking for Mr. Perfect in an imperfect world', journalist Barbara Buchholz describes recent trends of well-educated, successful women who are forever looking for the perfect mate (Buchholz, 2003). In interviews, the women both describe their theoretical ideal husbands and the men they have already dated who they judged to be acceptable but less-than-perfect candidates. Although they do not know when or even if their ideal candidate will appear, the possibility of his existence motivates each woman's odyssey and causes her to reject all other alternatives.

This struggle to find the best candidate or best match from a sequential string of potential alternatives happens in many search problems, whether it is finding a mate, selecting an applicant for a job, choosing which version of a new technology to purchase, deciding when to take a vacation day, or even choosing the right day to open a special bottle of wine. The decision-maker in these tasks often has an ideal candidate in mind against which other candidates are measured; the challenge is whether to select a currently available option or to continue searching for a better alternative with the risk of later realizing that the best alternative has already been forgone.

* Correspondence to: Suzanne B. Shu, Southern Methodist University, Cox School of Business, PO Box 750333, Dallas, TX 75275, USA. E-mail: sshu@smu.edu

Sequential search models that grapple with similar problems have been studied in operations, economics, and marketing, from both a theoretical and an experimental perspective (Gilbert & Mosteller, 1966; Kamien & Schwartz, 1972; Weitzman, 1979). The candidate search problem, often called the Secretary Problem (or, appropriately, the Marriage Problem or Dowry Problem), is perhaps the most well-examined of these problems (Ferguson, 1989; Samuels, 1991). However, in contrast to the single women in Buchholz's article who are on an 'endless odyssey', experimental laboratory work on search has regularly documented that decision-makers tend to end search too soon when compared to the predictions of the optimal models (Beardon, Rapoport, & Murphy, 2006a, 2006b; Seale & Rapoport, 1997, 2000; Zwick, Rapoport, Lo, & Muthukrishnan, 2003).[1]

The goal of the present study is to investigate whether the type of search behavior documented anecdotally, in which individuals search 'too long' for a perfect match, can be observed in an experimental setting with clear incentives and full information. If extended search can be observed in an experimental setting, is the 'searching for perfect' motivation described by journalists and advice writers a valid explanation for the observed behavior? And, when searching for a perfect candidate, are individuals properly accounting for the value of second-best alternatives and/or the probability that the perfect candidate will arrive? This paper begins by describing common characteristics of search tasks in which individuals anecdotally focus on a single best outcome, similar to the daters searching for a perfect mate. An optimal timing model is derived to determine the policy that maximizes payoff for a task of this type, and a laboratory experiment is conducted to test whether searchers stop earlier or later than the optimal benchmark. The results indicate that decision-makers do, in fact, continue searching beyond the stopping points predicted by the optimal model. Using these empirical results, in combination with a review of relevant psychological processes, I develop two predictions: that searchers overestimate the probability of a desired outcome, and that they underestimate the value of a second-best outcome. These predictions are then tested in Experiment 2. Experiment 3 more directly tests the underlying psychological behavior by manipulating the salience of certain desirable outcomes, and offers some evidence that it may be possible to debias searchers by encouraging them to focus on a wider range of outcomes. The results of the three experiments suggest that individuals in these tasks behave consistent with a strategy in which their search centers on obtaining the focal outcome, without full attention to its probability or the value of second-best alternatives. Finding and understanding this type of delayed choice behavior in sequential search tasks offers important new insights in our understanding of self-control and inter-temporal choice by providing us a situation in which the more typically observed myopic behavior does not prevail.

FUTURE-BIASED SEARCH

Consider, as an example of a sequential search with a focal ideal outcome, the problem of a traveler deciding on which occasion to use a free ticket provided by an airline. If the free ticket is unlikely to be replaced once exercised, the individual may desire to ensure that the flight on which it is used is the 'best possible match', preferably the most expensive of any flights to be taken within a specific timeframe. This is the individual's desired focal outcome. On each possible travel occasion the traveler must decide whether this occasion is the best match or whether a future trip may be even better. For trips on which the free ticket is not used, the traveler will need to buy a ticket out of pocket. Note that there are several important characteristics of such a search task. First, the free ticket cannot be retroactively applied to a past trip. Second, the traveler has some *a priori* knowledge of the variation in ticket prices; he can easily distinguish between a reasonable price and a

[1] An exception to this finding is when searchers do not adapt their strategy to incorporate higher exogenous search costs and search too long relative to an optimal model that does properly account for the increased costs, as documented in Zwick et al. (2003).

DOI: 10.1002/bdm

high price for a given trip. Third, after the free ticket has been exercised, he will continue to experience other travel occasions that might have been better candidates. Finally, in situations where the traveler decides to delay using the free ticket, he will incur a cost accordant to the price of the current trip. Examples similar to our traveler include not only the single daters seeking a mate, but also individuals choosing when to use a gift certificate to a restaurant they might not otherwise visit, or even wine drinkers deciding on which occasion to open a rare or special bottle of wine.

Consistent with the real world features of these examples, the following assumptions are made for the types of tasks explored in the context of extended search:

(1) Previously seen opportunities cannot be recalled. An option can be selected at the time it appears, but once the individual has rejected that opportunity and moved on to the next one, it is no longer available to be chosen.
(2) Full or partial knowledge of the distribution of outcomes is available. The individual has some sense of what qualifies as the 'best match' when it is observed, in an absolute sense.[2]
(3) Observation of other possible alternatives may continue even after the choice is made.
(4) The individual's payoff is directly related to both the value of the option chosen and the cost of options not pursued.

These assumptions are applicable not just to some of the examples already noted, such as finding a mate or opening a bottle of wine, but also to many other everyday search tasks. The inability to recall a previously seen opportunity is especially pertinent for decisions with a temporal component; opportunities that are available today (a one-day sale, a sunny afternoon) cannot always be recreated at the decision-maker's future whim. When it comes to knowledge of outcome distributions, individuals in many common search tasks have at least partial, if not full, knowledge of the possible range of outcomes. This knowledge may be exogenously provided (e.g., online product comparisons or historical price distributions) or accumulated by the searcher through personal experience. For example, although an employer hiring a secretary may not know the exact distribution of quality in her universe of candidates, she does have some idea of what a high-quality secretary is capable of relative to a low-quality secretary, and can assess individual candidates accordingly. As searchers acquire more knowledge or experience with candidate outcomes, they are more likely to assess those outcomes according to their absolute location within the broader distribution, rather than relying on relative rank.

The third assumption, continued observation of other options after search ends, also occurs regularly for common tasks; almost every shopper has had the experience of making a major purchase but still checking the weekly advertisements to see if a better price has been announced (a contributing factor to so-called buyers' remorse). While these four assumptions can apply to trivial, repeated search tasks, they also apply to searches that are more rare, done as few as half-a-dozen times (opening expensive wines) or even one or two times in a lifetime (finding a mate). Since searchers in these more rare environments have so few trials in which to learn an optimal strategy, any bias in their search process is likely to persist.

The particular experimental paradigm used in this paper is built around the scenario of an individual deciding at what point to cash in a free flight coupon while purchasing a series of airline tickets. All four of the assumptions outlined above are maintained in the experimental task. In addition, there is assumed to be no temporal discounting due to the short duration of the task, and no risk aversion given the small dollar amounts earned by the participants. This specific scenario was chosen because it is a situation familiar to the

[2]Note that this differs from many sequential search tasks like the Classical Secretary Problem, in which the individual only knows an outcome's relative rank rather than where it falls on the full distribution of outcomes. In cases where full information is provided in Secretary Problem tasks, participants appear to use threshold-based stopping rules (Rapoport & Tversky, 1970; Lee, O'Connor, & Welsh, 2004).

DOI: 10.1002/bdm

individuals participating in the experiment and each decision point has clear financial implications. Since the distribution of outcomes and all other parameters are known to the participant in advance, dynamic programming can be used to derive an optimal solution to this task, which serves as a benchmark for the observed behavior in the laboratory. Solving for optimal behavior yields a series of thresholds for the ticket prices at which the free trip should be used in each period; the size of the threshold decreases as the periods progress, according to the expected value of the remaining periods. Derivation of this solution is provided in the Appendix along with a numerical example for a given distribution of ticket prices.

Given the optimal solution for the task, the goal of Experiment 1 is to examine individuals' search behavior in an environment where a 'best match' is difficult to achieve—in other words, in a world where the most desired outcome is of a relatively low probability compared to other acceptable outcomes. When faced with a large majority of low-valued outcomes and a small minority of high-valued outcomes, will they choose the first high-valued outcome they see? Or will they hold out and accept only 'the best'? And how will this search behavior compare to the optimal model's recommendations? The distribution of outcomes selected for Experiment 1 allows us to answer these questions.

EXPERIMENT 1

Method

A computer-based task fitting the four assumptions previously described was tested with 28 undergraduates at a Midwestern university. Experimental participants are told that they are required to buy a series of 15 weekly airplane tickets from a pre-defined distribution of ticket prices, which is provided in advance. The 15-week task is completed three times in separate trials. Each participant begins with an endowment of funds ($9000) from which to buy the weekly tickets for all three trials. As encouragement to spend as little as possible, participants receive a two-part final payment consisting of a $1 base payment and an additional $.20 for every $100 left in the endowment at the end of the experiment.[3] Actual payments ranged from $1 to $5 with an average of $3. Each participant has access during each 15-week period to one reward coupon that can be used for a single free flight and two discount coupons that can each be used for 30% off. After completing the three separate 15-week trials, participants filled out a questionnaire about their experience. A copy of the experimental instructions is provided in Attachment 1.

The distribution of ticket prices consists of six possible outcomes; three outcomes are low priced and high in probability [$100, 55%; $200, 20%; $300, 19%], while the other three outcomes are high priced and low in probability [$400, 3%; $500, 2%; $1000, 1%]. Note that the most highly desirable ticket price for using the free flight coupon, the $1000 ticket, is of low enough probability that it does not occur for every participant. This distribution allows us to observe whether participants will accept any high value, low probability outcome when it appears, or whether they will focus only on the single most desirable outcome. Other aspects of the task mirror the search environments described earlier in this paper: only one option can be chosen out of the series of alternatives, alternatives will continue to be observed even after the choice is made, and the searcher has enough information about the possible outcomes that he or she carries some idea of a 'best match' in mind during the search process.

There are four main dependent measures for this study: the total amount spent per trial, the value of the ticket on which the free trip is used, the timing of the free trip usage, and finally the participant's reported happiness with their use of the free trip on each trial. The first three measures (total cost, ticket value, and timing) can be compared to the benchmark of the normative optimal timing model to determine if individuals

[3]The endowment was set high enough to allow most participants to end with positive balances. Participants were not penalized for overspending their endowment; only two participants did so, at respective ending balances of −$30 and −$250.

DOI: 10.1002/bdm

Table 1. Optimal model predictions for the experimental frequent flier task used in Experiment 1

		Week														
Cost	State	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
$100	1–2	P	P	P	P	P	P	P	P	P	P	P	P	D	D	F
	1–1	P	P	P	P	P	P	P	P	P	P	P	P	P	D	F
	1–0	P	P	P	P	P	P	P	P	P	P	P	P	P	P	F
	0–2	P	P	P	P	P	P	P	P	P	P	P	P	P	P	D
	0–1	P	P	P	P	P	P	P	P	P	P	P	P	P	P	D
$200	1–2	P	P	P	P	P	P	P	P	D	D	D	D	D	F	F
	1–1	P	P	P	P	P	P	P	P	P	P	P	D	D	F	F
	1–0	P	P	P	P	P	P	P	P	P	P	P	P	P	F	F
	0–2	P	P	P	P	P	P	P	P	P	P	P	D	D	D	D
	0–1	P	P	P	P	P	P	P	P	P	P	P	P	P	D	D
$300	1–2	D	D	D	D	D	D	D	D	D	D	F	F	F	F	F
	1–1	P	P	P	D	D	D	D	D	D	D	F	F	F	F	F
	1–0	P	P	P	P	P	P	P	P	P	P	F	F	F	F	F
	0–2	P	P	P	D	D	D	D	D	D	D	D	D	D	D	D
	0–1	P	P	P	P	P	P	P	P	P	P	D	D	D	D	D
$400	1–2	D	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	1–1	D	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	1–0	P	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	0–2	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D
	0–1	P	D	D	D	D	D	D	D	D	D	D	D	D	D	D
$500	1–2	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	1–1	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	1–0	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	0–2	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D
	0–1	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D
$1000	1–2	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	1–1	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	1–0	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F
	0–2	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D
	0–1	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D

For each possible ticket price (e.g., $100), current state of coupon availability, and week number (1–15), the table indicates whether to purchase the ticket at full price (P), use the free trip (F), or use a discount coupon (D). The second column is a state variable indicating how many free trip and/or discount coupons remain available; for example, 1–2 indicates an availability state of one free trip and two discount coupons.

are systematically deviating from the optimal selection; see Table 1 for a solution to the optimal behavior for this particular task. The timing and ticket value measures for the two discount coupons per trial can also be analyzed relative to the optimal model. The happiness measure, which is each participant's response to the question 'how happy are you about your use of the free trip coupon' on a 0 (not happy) to 10 (very happy) scale, is collected through the post-trial questionnaire. Since it is possible that individuals receive some utility from the search process itself and are thus better off by delaying choice, this happiness measure provides some indication of whether individuals who search longer are consistently more or less happy with their choices than those who have not delayed.

In addition to the overall comparisons of the empirical results to the optimal model for the dependent measures of amount spent, ticket value, and usage timing, a set of specific testable 'usage rules' can be defined as outputs of the optimal model. These usage rules can be defined as:

DOI: 10.1002/bdm

1. Never end a trial without having used the free trip coupon.
2. Never end a trial without having used both discount coupons.
3. If a $1000 ticket appears at any time, use the free trip coupon if it is still available.
4. If a $500 ticket appears at any time, use the free trip coupon if it is still available.
5. If a $400 ticket appears in week 2 or any time thereafter, use the free trip coupon if it is still available.
6. If a $300 ticket appears in week 11 or any time thereafter, use the free trip coupon if it is still available.
7. Do not use the free trip coupon on any tickets under $400 before week 11.

These seven usage rules are consistent with the optimal solution to the task presented in Table 1.

Results

Results for the four main dependent measures, along with the predictions from the optimal model, are provided in Table 2.[4] Compared to the normative benchmark from the risk-neutral optimal timing model, subjects spend, on average, more per trial empirically ($2289) than the optimal model predicts ($2250; $t(83) = 1.66$, $p = .05$). As might be expected from the higher than optimal total spending, the average value of the ticket on which the free trip is used is significantly lower for the empirical results ($373) than the optimal benchmark ($415; paired t-test significant at $t(83) = -2.11$, $p = .02$).

However, less than optimal spending simply indicates that individuals are not behaving according to the optimal model; these deviations from optimal may be random rather than consistently earlier or later than optimal. Analysis of the timing of the free trip usage will tell us whether the bias in the usage is systematically in one direction or another. Analyzing results from all trials, we see that the timing of the actual use (11.5) is indeed later than the timing of the optimal use (10.0; paired t-test significant at $t(83) = 3.44$, $p < .001$). This can also be seen clearly by categorizing usage behavior relative to the optimal; dividing participants' trials into those where the free trip coupons were used earlier, the same as, or later than optimal shows that although the majority of free coupons were used optimally (49 of 84 trials), there were still far more trials in which coupons were used later (28) rather than earlier (7) than optimal (see Figure 1). At the participant level, 20 of the 28 participants had at least one trial with later than optimal free coupon use. Is this delay due to a search for a 'perfect' outcome? Looking at specific instances of individuals faced with high value outcomes, there are 48 events across all trials for which the optimal model calls for using the free trip on a $300, $400, or $500 ticket in the first 12 periods; of those 48 events, participants delayed using the free trip in 23 cases (48%). Compare this to 8 events in which a $1000 ticket occurs in the first 12 periods; these participants used their free trip in all 8 cases, consistent with the optimal model. Thus, individuals do seem to be delaying their free trip use beyond the optimal choice by focusing their search on the $1000 ticket and forgoing other high value outcomes along the way.

This behavior can also be analyzed relative to the 'usage rules' defined earlier. The first rule that the free trip coupon should always be used before the trial ends, would appear a difficult one to violate since a

Table 2. Average results for all 15-week trials from Experiment 1 for the four main dependent variables

	Cost (total dollars spent)	Value of ticket for free coupon use	When free coupon used	Happiness with free use
Emprical	$2289 (38.5)	$373 (31.4)	11.5 (0.44)	5.65 (0.3)
Optimal	$2250 (39.9)	$415 (25.0)	10.0 (0.46)	n.a.

Both optimal model predictions and empirical results are presented. Standard errors are shown in parentheses.

[4]Optimal model predictions are calculated per participant based on the actual sequences observed during the experiment.

DOI: 10.1002/bdm

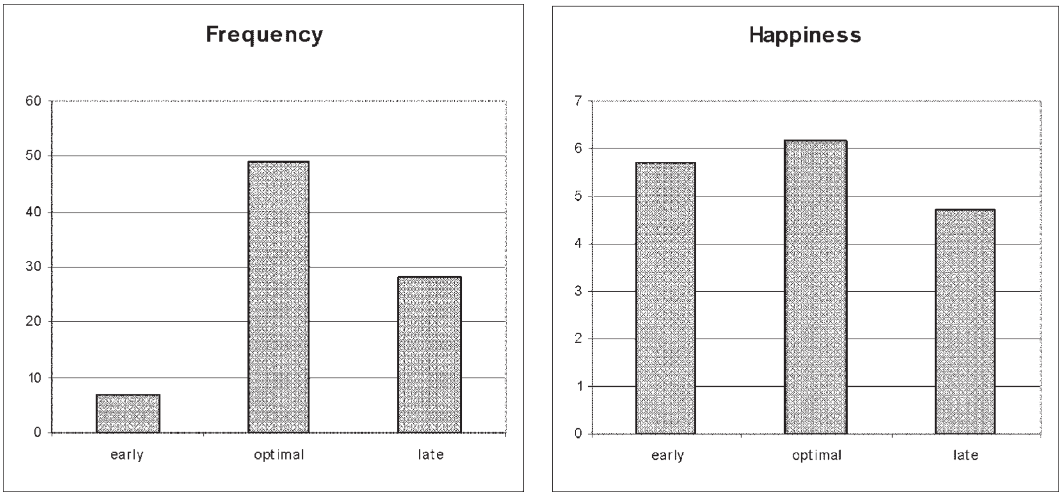

Figure 1. Comparison of actual free trip usage to optimal model predictions in Experiment 1; usage was coded as either earlier than, same as, or later than optimal. The number of trials in each category (figure on left) shows that late usage is more prevalent than early usage. Average reported happiness for each category (figure on right) shows that individuals who used the free ticket later than optimal were also less happy about their usage, suggesting that they do not receive extra utility from delay

participant can always use the free trip in the final period, week 15. However, of 84 completed trials, there were four instances in which participants ended the trial without using the free trip (these individuals typically used a discount coupon in week 15 instead). Usage rule two was also violated, with 8 trials out of 84 in which not all of the discount coupons were utilized. Rules three through six are the most important ones relative to the behavior of interest, since violations to these rules indicate non-optimal delay in using the free trip. Not surprisingly, rule three had no violations, since participants used the free ticket in all 11 situations where they were still holding it when a \$1000 ticket occurred. Rules four through six were more easily violated; there were eight violations in 24 cases where the free ticket could be used on a \$500 ticket, four violations in nine cases with a \$400 ticket, and 21 violations in 37 cases with a \$300 ticket.[5] Finally, rule seven measures a different type of behavior: non-optimal early use of the free ticket coupon on a low value ticket. In only 4 of the 180 applicable situations did this behavior occur.

If individuals are receiving some satisfaction from the search process independent of the actual outcome they receive, then extended searching may be a utility-maximizing behavior. Some previous work on future-bias has suggested that individuals delay due to increased utility from anticipation or savoring (Loewenstein, 1987). On the other hand, recent work on maximizing versus satisfying choice strategies finds that individuals who consistently ‘search for the best’ often feel worse about the (objectively better) outcomes they receive than individuals who do less exhaustive searching (Iyengar, Wells, & Schwartz, 2006). For this experiment, the per-trial happiness measure collected in the post-task questionnaire provides some confirmation that participants are not receiving additional utility from delaying their selections. A comparison of the happiness measures for participants’ trials divided into the categories of early, optimal, and late usage demonstrates that those trials in the late usage category led to participants being significantly less happy that those with trials in the optimal usage category ($\mu_{\text{late}} = 4.7$ vs. $\mu_{\text{optimal}} = 6.2$, $t(75) = 1.98$, $p = .02$),

[5]The number of cases given here for each ticket size are only the ones in which the participant was still holding the free ticket coupon; \$300 and \$400 tickets occurred in higher frequency than reported in these specific counts.

DOI: 10.1002/bdm

while participants with trials in the early usage category are not significantly less happy than with optimal use ($\mu_{\text{early}} = 5.7$ vs. $\mu_{\text{optimal}} = 6.2$, n.s. at $t(12) = 0.55$, $p = .30$). Thus, individuals who delay usage by searching for the most expensive ticket are not more satisfied with their results; instead, late usage leads to lower happiness than that reported both by participants who behaved according to the normative predictions of the optimal model and by those who used their free trip too early.

One additional question for happiness is whether it is sensitive to the value of the ticket on which the free trip coupon was used. In other words, was a participant who used the coupon on a $1000 ticket significantly happier about their usage in that trial than one who used it only on a $200 ticket? A regression of individuals' per trial happiness measures found a significant effect of free trip value on happiness ($t(83) = 7.28$, $p < .001$); in other words, participants who used their free coupon on higher-valued tickets are indeed happier than those who did not. Additional regressions that include timing of free trip use, differences in timing compared to optimal, and deviation of optimal trial cost show that these additional variables have no significant effect on happiness. Happiness therefore appears to be highly dependent on outcome, and individuals in this task do not appear to be receiving any extra utility benefit by extending the search process.

The usage of the two 30%-off discount coupons provided for each trial can also be analyzed relative to the optimal model. Participant behavior for these coupons is interesting because we might expect participants who delay use of the free trip coupon in early periods to substitute use of one of the discount coupons instead. If this is the case, then empirical usage of the discount coupons may happen sooner than the optimal prediction. Instead, results for the difference between optimal usage time and actual usage time for the discount coupons suggest that participants are also delaying use of these coupons, consistent with the results from the free trip coupon (paired *t*-tests for empirical vs. optimal usage for both coupons significant at $p < .001$). Table 3 shows this usage difference per condition for both discount coupons 1 and 2, along with the differences in optimal and empirical ticket value for each coupon. Figure 2 shows the cumulative probability of use at each period for both coupons; empirical usage can be seen to lag optimal use in nearly all 15 periods. As both the table and figure indicate, both discount coupons were used significantly later than the optimal model suggests. This later than optimal use does not have a financial impact for discount coupon 1 since the value of the ticket on which the coupon is used is not significantly different from optimal. However, for discount coupon 2, later than optimal use does lead to the coupon being used on a significantly less valuable ticket ($268 vs. $319, $t(83) = 3.43$, $p < .001$). Closer analysis of the data shows that the second discount coupon went completely unused in 10% of the trials; in other words, participants in these trials were still holding an unused discount coupon after the 15 periods ended.

Finally, to check for learning within the three trials that each participant completed, analysis was done of the difference between optimal and empirical results across trials to determine whether individuals were improving relative to the optimal benchmark in their later trials. No significant difference was found for these measures for total cost, value of free trips, or timing of use between the three trials. This is likely a result of the small number of trials each participant experienced; completing a much larger number of trials (e.g., more than 20) may lead to different results. Note however that the small number of trials is similar to many real-life sequential searches that are only repeated a small number of times (e.g., finding a mate).

Table 3. Timing usage and ticket value average results for both discount coupons used in Experiment 1, including both optimal model predictions and empirical results

	When coupon used	Optimal use	Value of ticket for coupon use	Optimal value
Discount coupon 1	5.6 (0.43)	4.2 (0.31)	$304 (8.8)	$302 (2.4)
Discount coupon 2	10.9 (0.41)	9.1 (0.36)	$268 (16.6)	$319 (15.9)

Standard errors are shown in parentheses.

DOI: 10.1002/bdm

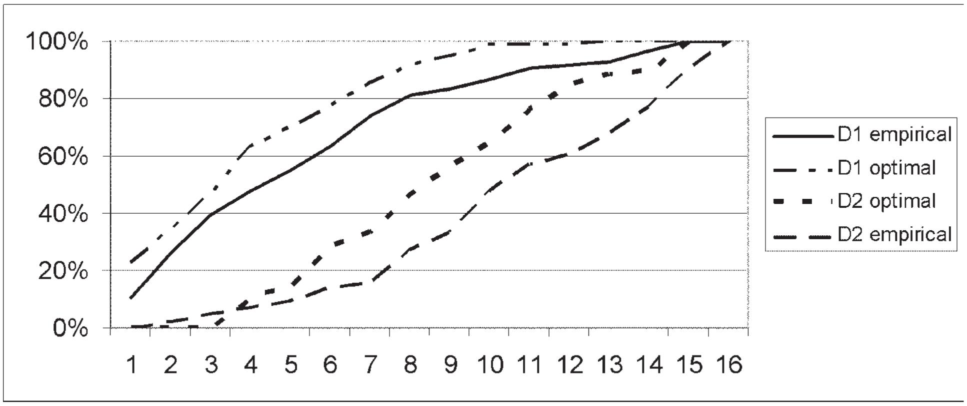

Figure 2. Cumulative probability per week of having used each of the two discount coupons in Experiment 1, including both the optimal model prediction and observed behavior from participants for each discount coupon

The overall results of Experiment 1 suggest that there is a main effect of delay, in which individuals put off using their free trip longer than an optimal model would predict. Rather than accept the first high value outcome that they encounter (e.g., a $500 ticket), participants appear to be overly focused with finding a ‘perfect match’, the single most expensive $1000 ticket, regardless of its low probability. They then forgo early opportunities to use the free trip on other high value tickets, hold the free trip too long, and are finally forced to ‘waste’ the free trip on an inexpensive flight toward the end of the 15-week trial. The result of the delay is that they pay more per trial than the optimal model would predict and the average value of the free trip is lower than with the optimal model, consistent with the data.

The delayed choice behavior described in Experiment 1 is an example of future-biased choice.[6] The main hallmark of future-bias is that individuals systematically make the choice to postpone use of a valuable option when normative considerations may suggest otherwise. Since the individual is primarily forward-looking, the behavior is labeled as ‘future-bias’, which contrasts it from the myopic ‘present-bias’ described for many other intertemporal choice situations (Ainslie, 2001; Loewenstein & Elster, 1992).

THE ROLE OF FOCALISM

The results of Experiment 1 suggest that participants seem to be searching for an ideal but rather improbable outcome (the $1000 plane ticket), bypassing other good alternative outcomes that an optimal model would suggest selecting. In a sense, they act as if they have a rule-based approach to their choices with a specific per-period threshold below which they will not choose. The use of thresholds as a choice rule in sequential search problems is well supported in the empirical findings for many Secretary Problem tasks, although participants in most of those tasks end search earlier than optimal (Seale & Rapoport, 1997, 2000 and Zwick et al., 2003). Note that such a threshold-based rule is also consistent with the predictions of the optimal model. However, there is a significant difference in the empirical threshold use in this paper’s results as compared to threshold use in other search tasks; where thresholds in many Secretary Problem experiments are set *too low* by participants, individuals in this task appear to set their thresholds *higher* than the optimal predictions, resulting in later selection rather than earlier.

[6]Similar behavior has also been labeled as hyperopia (Kivetz & Simonson, 2002) or reverse time inconsistency (Loewenstein, 1987).

DOI: 10.1002/bdm

Why might the threshold error be in the opposite direction in these tasks than in other search tasks? First note that this task has some important differences from other experimentally tested sequential search tasks. The two assumptions that most differentiate this task from previous efforts are full knowledge of the distribution and observation of other outcomes after the decision is made. These two differences may contribute to the higher thresholds observed in Experiment 1. Consider, first, how distribution knowledge differs from the use of relative ranks in other search tasks. With only knowledge of the relative rank of each sequential outcome, an individual who observes an early outcome with rank one must evaluate whether this outcome is truly the best option or whether its absolute rank will turn out to be much lower. The salience of the high relative rank may cause him to lean toward believing that it is the best option, in which case he will end search immediately. In contrast, an individual who knows the full distribution knows for certain that an outcome is or is not the best among all possible outcomes; his dilemma is whether or not an even better option from within the distribution will later occur. In this case, the salience that this outcome is not the most extreme outcome within the distribution pushes him toward continuing the search. Continued observation of other outcomes may also encourage extended search as compared to tasks where searchers immediately stop seeing additional candidates. Actively continuing to observe new candidates may heighten the anticipated regret that comes from choosing a poor option early and learning how much better of an option was later available, while search tasks that end immediately may moderate the regret by showing the other outcomes only after the searcher no longer has active control over the task.

Search thresholds in these tasks may also be set too high because of specific features of the item and/or due to the psychological effects that come from thinking about the likelihood of achieving a ‘best match’. The threshold is likely to be set high as a conservative measure when an item’s usage is highly constrained. Setting a high threshold becomes a form of self-control to ensure that a rare item is not used indiscriminately, and thus helps guard against later regret or guilt from having used it inappropriately.[7] This will be most true for singleton items versus items that are one of many (or easily replaced); consider, for example, how usage may be more heavily constrained for the last bottle of wine in a case than for the first bottle. It will also be most true for infrequently purchased items that are tightly constrained by mental accounting and other self-control rules (Heath & Soll, 1996; Thaler, 1985, 1999).

Once a threshold is set, certain psychological effects of thinking about that preferred outcome may exaggerate its overly conservative nature. The tendency for individuals to put undue weight on a focal outcome has been explored in its effects on probability estimates, affective forecasting, and other judgment tasks (Koehler, 1994; Rottenstreich & Kivetz, 2006; Wilson, Wheatley, Meyers, Gilbert, & Axsom, 2000). Both optimism bias and the planning fallacy are related to the tendency to overemphasize a desirable focal outcome (Kahneman & Lovallo, 1993; Kahneman & Tversky, 1979). Tatarkiewicz (1976) wrote, ‘in anticipating a coming event we have it alone in mind, and make no provision for other occurrences’. The term focalism refers to the tendency of individuals to not think much about the consequences of other possible events when focused on a single preferred event or outcome. When focalism occurs as part of an agent’s explanation for some outcome or imagination of a specific scenario, it can lead to greater confidence in the belief that such an outcome will occur (see Koehler, 1991 for a review). It can also cause the probability of that outcome to be overestimated relative to the probability of alternate outcomes. Finally, focalism also contributes to a durability bias, in which individuals underpredict their ability to adapt to a non-preferred outcome; in other words, their affective forecasts for life after the event assume that they will be more strongly affected for a longer period of time than they actually are (Wilson et al., 2000). As a result, ‘second-best’ outcomes are heavily undervalued relative to the focal outcome.

[7]Note that this use of high thresholds differs substantially from the typical thresholds that we see in self-control problems, which are often too low. Low thresholds allow individuals the opportunity to break their own self control rules (e.g., to find loopholes for giving in to temptations, as per Cheema & Soman, 2006). High thresholds, however, result in individuals never allowing themselves to indulge, even when deserved (see, e.g., Kivetz & Simonson, 2002).

DOI: 10.1002/bdm

The multiple effects of focalism have implications for the search behavior investigated here. First, focalism's effect on durability bias suggests that the searcher will underestimate the satisfaction he would feel when the chosen outcome is merely suitably good (rather than perfect). Second, focalism may cause individuals to be overly confident about their own likelihood of finding a best match. Rather than accurately considering that a perfect candidate may never arrive, the individual continues to delay under the belief that the desired option is yet to come. Together, the effects of setting a high threshold and the associated focal thinking about a perfect match lead to two predictions for the types of tasks described in the previous section:

Prediction 1: Because participants undervalue less-than-best outcomes, observed choice behavior will be independent of the value of the second best outcome (i.e., whether the second best outcome has a relatively high or relatively low value).

Prediction 2: Because participants are overconfident about the occurrence of the target outcome, observed choice behavior will exhibit greater levels of delay relative to the optimal model when the most focal outcome is relatively rare.

To test whether these predictions hold in the task used in Experiment 1, a new laboratory study was run in which probability of the focal outcome and value of the second best outcome are manipulated.

EXPERIMENT 2

Method

To test whether focalism helps explain delays in choice beyond optimal levels, Predictions 1 and 2 were tested using a computer-based study with 85 undergraduates at a Midwestern university. Participants were paid a flat \$3 for their participation. As in Experiment 1, each participant is required to buy a series of 15 weekly airplane tickets from a predefined distribution of ticket prices. He or she also has access to one reward coupon that can be used for a single free flight during the 15-week period. The explicit distribution of airfares was provided to subjects in advance. Subjects were told that their goal was to spend as little as possible from a predefined endowment of money for travel (\$9000).[8] After completing three 15-week trials, participants filled out a questionnaire regarding how happy they were with their usage of the free trips and whether they felt that they had used their free trip at the appropriate time.

To test whether searchers both underestimate the effects of alternate good outcomes and are overconfident about the occurrence of the focal outcome, one of the airfare distributions manipulates the value of the second best outcome while another distribution manipulates the probability of the best outcome, all relative to the basic distribution. For the alternate outcome manipulation, the second best outcome is either a low value (\$500) or a higher value (\$750) relative to the focal outcome (\$1000). Probabilities of the outcomes are held constant between these two distributions. (The distributions for all three conditions are provided in Table 4.) For the outcome probability manipulation, the distribution of airfares was manipulated such that some subjects were very likely to face a \$1000 ticket during the 15-week trial (5% probability per week), while others were relatively unlikely to experience such a ticket (1% probability per week). Note that the expensive, high probability condition is stochastically dominated by the low probability condition such that those participants in the high probability condition should almost always spend more on their trials.

To simplify comparisons between the conditions, participants were tested using a yoked design; each participant was matched with participants from the other two conditions who saw a sequence of airfares

[8]As in experiment 1, participants were not penalized for overspending their endowment; only 8 out of 85 participants did so.

DOI: 10.1002/bdm

Table 4. Distributions of airfares for each of the three conditions in the experimental frequent flier task used in Experiment 2

Condition 1	Condition 2	Condition 3
1% chance of $1000	1% chance of $1000	
2% chance of $500	2% chance of $750	
2% chance of $400	2% chance of $400	5% chance of $1000
20% chance of $300	20% chance of $300	20% chance of $300
20% chance of $200	20% chance of $200	20% chance of $200
55% chance of $100	55% chance of $100	55% chance of $100

Participants saw the distribution only for his/her assigned condition.

matched to his own. In this way, a direct comparison can be made of the same sequence across all three conditions. If individuals use a threshold-based rule of saving the free trip for the focal outcome (here, the $1000 ticket), regardless of probability or alternate outcomes, then they may forgo early opportunities to use the free trip on lesser tickets, hold the free trip too long, and finally be forced to 'waste' the free trip on an inexpensive flight toward the end of the 15-week trial. Such behavior will lead to the apparent future-biased search observed in Experiment 1. In addition, the design of the three conditions allows testing of the two predictions defined above. The first prediction is tested by comparing behavior in Conditions 1 and 2; according to Prediction 1, individuals in the condition with the higher value second-best outcome ($750) should behave no differently than those in the condition with the lower value second-best outcome ($500). The second prediction is tested by comparing Conditions 1 and 3; participants facing a less favorable (more expensive) distribution will use their free trips in a more timely manner because their focal outcome will arise more often. Thus, they will delay less, use their free trips on higher value tickets, and be closer to the predictions of the optimal model. I also predict that these participants in Condition 3 will be happier about their free trip use as a result of being able to use it against the 'best match' outcome.

To see this last prediction more clearly, consider a scenario involving Conditions 1 and 3. An occurrence of a $400 ticket in Condition 1 will appear as a $1000 ticket in the matched sequence for Condition 3. The optimal timing model developed earlier recommends use of the free trip in both situations. However, the participant in Condition 1 may decide to not use the free trip at this point, preferring instead to wait for the unlikely occurrence of a $1000 ticket, whereas the Condition 3 participant will use his free trip immediately. The Condition 1 participant then finds himself holding an unused free trip in the final weeks of his trial, leading to later usage, usage on a less valuable ticket, and less satisfaction with his experience. Because the sequences are matched across conditions, any such difference in usage timing suggests that participants are sensitive to the *existence* of the focal (most expensive) ticket but not sufficiently sensitive to the actual probability of that event versus other events. An important implication of this behavior is that individuals may actually be happier with their timing decisions under stochastically dominated distributions, since distributions with a higher probability of extreme events (here, Condition 3's higher probability of a $1000 ticket) provide more opportunities to actually use the free trip than distributions with an overall lower average ticket price.[9]

As with Experiment 1, there are four main dependent measures for this study: the total amount spent per trial, the value of the ticket on which the free trip is used, the timing of the free trip usage, and finally the participant's reported happiness with their use of the free trip. We can again use the normative optimal timing model as a benchmark, solved for each of the three distributions.

[9]Whether individuals would realize this implication *ex ante*, and whether they would consciously choose such a stochastically dominated distribution, remains unexplored in this paper.

DOI: 10.1002/bdm

Table 5. Average results for all 15-week trials for the four main dependent variables, for each of the three conditions in Experiment 2

Condition		Dollars spent	Value of ticket for free coupon use	When free coupon used	Happiness with free use
1 (base)	Empirical	\$2380 (43.5)	\$289 (20.5)	11.4 (0.51)	5.4 (0.34)
	Optimal	\$2275 (43.0)	\$379 (16.5)	9.8 (0.3)	
2 (high alt.)	Empirical	\$2362 (43.0)	\$363 (29.9)	11.3 (0.52)	6.0 (0.48)
	Optimal	\$2313 (42.7)	\$401 (22.3)	11.2 (0.37)	
3 (high prob.)	Empirical	\$2388 (48.6)	\$528 (36.7)	10.7 (0.39)	6.7 (0.47)
	Optimal	\$2333 (38.5)	\$569 (33.4)	10.9 (0.40)	

Both optimal model predictions and empirical results are presented. Standard errors are shown in parentheses.

Results

Overall results

Results for three of the main dependent measures, along with the predictions from the optimal model, are provided in Table 5. Since each participant completed three 15-period trials, results are averaged per individual before being compared between conditions. Compared to the normative benchmark from the risk-neutral optimal timing model, subjects spend, on average, more per trial empirically (\$2377) than the optimal model predicts (\$2307; $t(84) = 5.59$, $p < .001$). This difference between empirical and optimal is also significant within each condition: \$2380 versus \$2275 for Condition 1 ($t(28) = 4.25$, $p < .001$); \$2362 versus \$2313 for Condition 2 ($t(28) = 2.39$, $p = .02$); \$2388 versus \$2333 for Condition 3 ($t(28) = 3.00$, $p = .005$). As might be expected from the higher than optimal total spending, the average value of the ticket on which the free trip is used is significantly lower for the empirical results (\$393) than the optimal benchmark (\$450; paired *t*-test significant at $t(84) = -5.22$, $p < .001$). This difference between empirical and optimal ticket value is also significant within each condition: \$289 versus \$379 for Condition 1 ($t(28) = -4.54$, $p < .001$); \$363 versus \$401 for Condition 2 ($t(28) = -1.99$, $p = .06$); and \$528 versus \$569 for Condition 3 ($t(28) = -2.58$, $p = .02$).

However, non-optimal spending does not necessarily imply that there is a consistent bias in timing of use; participants could be using the free trips too early just as well as too late. To determine whether future-bias is really a problem, we need to look at when participants are using their free trips versus what the optimal model predicts. Combining all conditions into a single analysis, we see that the timing of the actual use (11.1) is later than the timing of the optimal use (10.6; paired *t*-test significant at $t(84) = 1.93$, $p = .06$). Thus, individuals do seem to be holding their free trips for later use than optimal. Is this delayed usage a result of an overemphasis on the most extreme \$1000 outcome? To determine if this kind of focal thinking is occurring, we can see how the manipulations of alternate outcomes in Condition 2 and probabilities in Condition 3 affect the behavior.

Results for Conditions 1 and 2

The results for Conditions 1 and 2 suggest that participants in both conditions have the same strategy for using their free trip, consistent with the idea that they are focused on the ‘best match’ outcome (the \$1000 ticket) and not attending to the value of the second-best outcomes (\$500 or \$750). There is no significant difference between the conditions in how much subjects actually paid per trial (\$2380 in Condition 1, \$2362 in Condition 2; paired *t*-test of 1 vs. 2 is n.s. at $t(28) = .65$, $p = .52$). When the difference is taken between these empirical values and each condition’s optimal spending level, we find that participants in Condition 1 are farther from optimal than those in Condition 2; deviation from optimal cost is −\$104 in Condition 1 and −\$49 in Condition 2 ($t(28) = -1.99$, significant at $p = .056$).

A comparison of the timing of use between the two conditions also finds no significant difference ($\mu_1 = 11.4$ vs. $\mu_2 = 11.3$ is n.s. at $t(28) = .20$, $p = .84$). These absolute differences between the conditions are

DOI: 10.1002/bdm

not significant, but there is a difference in timing behavior between the conditions when compared to the optimal benchmark. The deviation in timing of free trip use is -1.6 in Condition 1 and -0.02 in Condition 2 (paired t-test shows 1 vs. 2 significant at $t(28) = -2.58$, $p = .016$). This suggests that participants are not behaving any differently between the two conditions, but since the normative model recommends later use in Condition 2, participants in this condition end up being closer to the optimal benchmark. In other words, although the optimal model adjusts the recommended search strategy when the value of the second best outcome changes, these individuals do not.

Since the results suggest that participants in both conditions are using the same rule of thumb (waiting for the \$1000 ticket) to guide their decisions, the effect of that rule on their happiness is roughly equivalent. Participants in Condition 2 were not significantly happier than those in Condition 1 ($\mu_1 = 5.4$ vs. $\mu_2 = 6.0$ n.s. at $t(28) = -1.06$, $p = .29$). Thus, on the three main dependent variables of total cost, timing, and happiness, participants in these two conditions perform the same, as predicted. These results suggest that individuals are so focused on thinking about the most extreme outcome that they neglect to take into consideration the value of alternate outcomes which may also be appealing, consistent with Prediction 1. However, focal thinking may also cause the individual to overestimate the probability of the optimal occasion ever occurring. To explore this aspect of the behavior, Condition 3 manipulates the probability of the \$1000 outcome.

Results for Conditions 1 and 3

Results for the conditions in which the probability of the optimal outcome was manipulated suggest that participants are more susceptible to future-bias and delayed choice when optimal outcomes are relatively rare. As expected, Condition 1 participants who faced a low probability for the most extreme outcome delayed selection longer than those in Condition 3 since they encountered outcomes that met their usage criterion less often. When the absolute deviation between the empirical and optimal benchmark is compared across conditions, we find that the deviation in timing of free trip use is significant between Conditions 1 and 3 (deviation from optimal timing is -1.6 in Condition 1 and .16 in Condition 3; a paired t-test shows 1 vs. 3 significant at $t(28) = -3.31$, $p = .002$). Participants are also significantly happier with their free trips in Condition 3 as compared to Condition 1 ($\mu_1 = 5.4$ vs. $\mu_3 = 6.7$, paired t-tests significant at $t(28) = -2.75$, $p = .01$).

Although the expected value of trial cost is significantly lower for Condition 1, there was no significant difference between the conditions in how much subjects actually paid per trial (\$2380 in Condition 1 and \$2388 in Condition 3; paired t-tests of 1 vs. 3 are n.s. at $t(28) = -1.04$, $p = .30$).[10] There is a difference in the value of the ticket on which the free trip is used between the conditions; the deviation from optimal for the value of the ticket is \$90 in Condition 1 and \$40 in Condition 3 (a paired t-test is significant at $t(28) = 1.92$, $p = .06$).

Another way to see whether individuals are paying appropriate attention to the actual distribution of alternatives is to calculate the likelihood that participants successfully use their free trip on a ticket of \$400 or more, the top 5% of the price distribution. A graph of these likelihoods for all three conditions is provided in Figure 3. Note that, for all conditions, the likelihood of actual use is lower than what the optimal model suggests within these ticket prices. However, this difference is largest for Condition 1, where the focal \$1000 tickets are less likely to occur, consistent with the earlier analysis of deviation from optimal use in this condition. Individuals in Condition 1 appear to be so focused on the \$1000 ticket that they are less likely to use the free trip on other high value tickets in the \$400-and-over range.

We can further consider how actual timing deviates from the optimal benchmark for the two conditions by coding each participant's usage decision as either earlier than, exactly the same, or later than the optimal recommendation. A graph of the frequency of each type of deviation, shown in Figure 4, demonstrates that

[10]It would have been nice to find that participants in the condition 3 actually spend *less* per trial than those in condition 1, since those in condition 3 are able to perform better relative to optimal and use their free trip coupon on more expensive airfares. However, the higher expected value of condition 3 works against this outcome, so the actual average cost per trial is essentially flat across the two conditions.

DOI: 10.1002/bdm

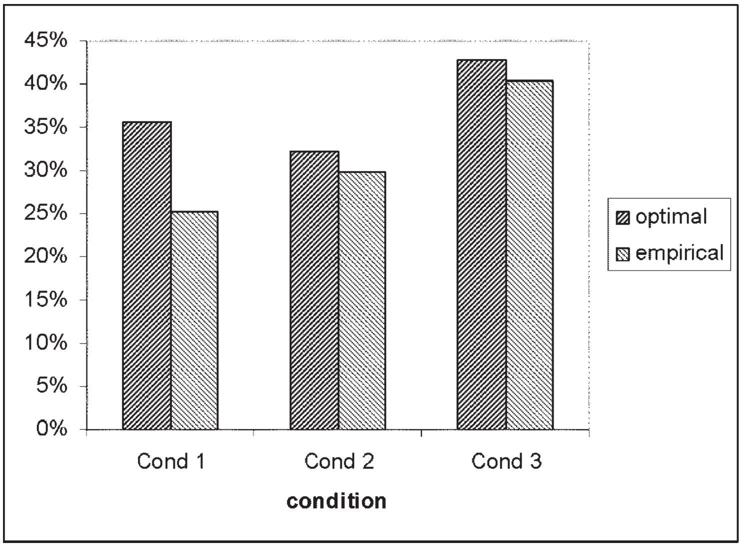

Figure 3. Likelihood that a free trip will be used on a ticket from the top 5% of the distribution (i.e., tickets above $300) for Conditions 1, 2, and 3 of Experiment 2

late usage is much more frequent in Condition 1 and that 'just right' usage is more frequent in Condition 3, consistent with predictions. This delay in timing is also evident when the data are evaluated on a period-by-period level; Figure 5 shows the cumulative probability that the free trip has been used at each of the 15 periods for the three conditions, respectively. As evident on these graphs, empirical usage lags optimal usage for all periods in Condition 1 but closely mirrors optimal usage in Condition 3.

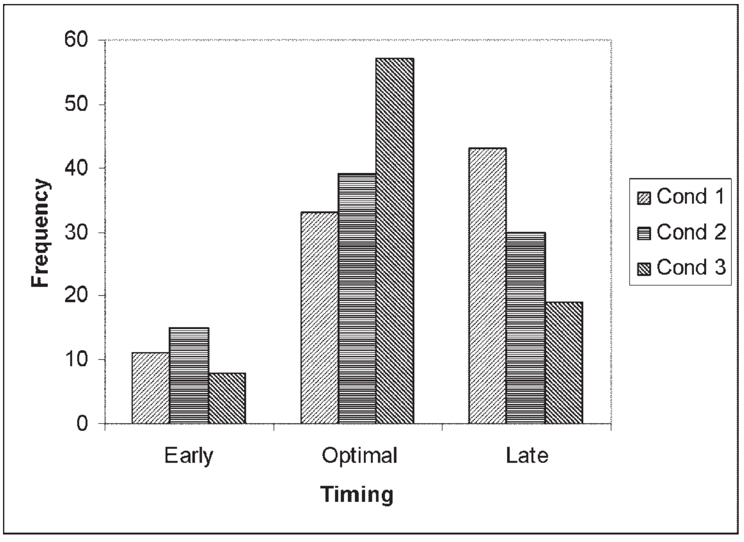

Figure 4. Comparison of actual free trip usage to optimal model predictions in Experiment 2; usage was coded as either earlier than, same as, or later than optimal. Participants in Condition 1 were more likely to be later than optimal, while those in Condition 3 were more consistent with the optimal predictions

DOI: 10.1002/bdm

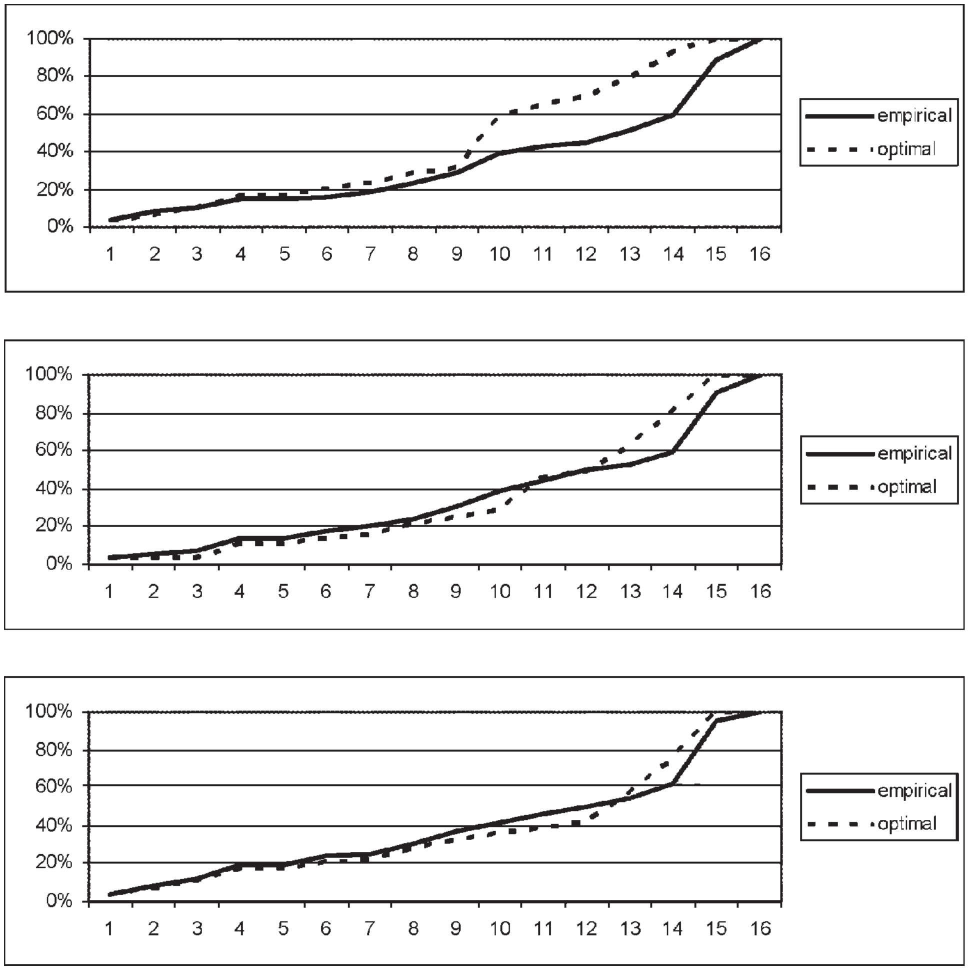

Figure 5. Cumulative probability of having used free trip, per week, for Condition 1 (upper figure), Condition 2 (middle figure), and Condition 3 (lower figure). The optimal model prediction is shown as a dashed line, while observed behavior from participants is shown as a solid line

Questionnaire results

As with Experiment 1, participants were asked to complete a questionnaire after finishing the three trials. When asked whether they felt that they had used their free trip at the appropriate time, too early, or too late, those in Conditions 1 and 2 were significantly more likely to report using it later than appropriate, while those in Condition 3 report using it at the right time ($\mu_1 = 4.9$, $\mu_2 = 4.2$, and $\mu_3 = 5.4$; a repeated measures ANOVA shows all conditions significant at $F(2,82) = 2.83$, $p = .06$). Note that this self-reported timing appropriateness is consistent with the actual timing results reported above, suggesting that individuals are accurately realizing that they are holding the free trip longer than they should. Consistent with this awareness of waiting too long, participants in Conditions 1 and 2 are significantly less happy with their free trip use than participants in Condition 3 ($\mu_1 = 5.4$, $\mu_2 = 6.0$, and $\mu_3 = 6.7$, Condition 3 significantly different at $t(84) = -2.8$, $p = .007$). As with Experiment 1, delaying usage, which occurs more often in Conditions 1 and 2 than in Condition 3, does not provide greater happiness (or presumably greater utility) than using the free trip in an optimal way.

DOI: 10.1002/bdm

Discussion

The goal of this laboratory experiment was to further understand the psychology that contributes to future-bias in a sequential search problem. As with Experiment 1, distributions of outcomes were used that had a single, low probability extreme outcome (the $1000 ticket). Observed behavior of individuals facing these distributions was consistent with a search strategy in which this highly focal outcome acts as a 'best match' selection threshold. By using distributions that varied in the probability of the focal outcome and the value of the second best alternative, we were able to verify that the predictions of focalism held for this task. The level of future-bias was unaffected by providing a higher value second-best outcome (Prediction 1). A distribution with a lower probability of the focal outcome led to more future-bias than a stochastically dominated distribution with a higher probability of the focal expensive outcome (Prediction 2). Interestingly, participants experiencing this stochastically dominated distribution of outcomes were happier with their experience and closer to the optimal predictions, probably because the goal of obtaining the focal outcome was easier to achieve.

Although the future-biased behavior observed in Experiment 2 is consistent with the two predictions, and thus provides some support for the argument that focalism is contributing to delayed choice, a more direct test of how focalism affects sequential search is needed. The primary goal of Experiment 3 is to more directly investigate the role of focalism by testing whether a direct manipulation of the focal outcome has an effect on the future-biased behavior evidenced in Experiments 1 and 2. To do this, the experiment asks participants to answer questions about the likelihood of certain ticket outcomes appearing during their trials. The outcomes referenced in these likelihood questions are manipulated between conditions to be either the most extreme outcome or a group of three outcomes (the most extreme and its two closest alternatives). By asking about outcomes other than the 'best' match, the alternative outcomes become included in the decision-maker's focal set. If successful, the previously documented future-biased behavior should be debiased and participants' choices should be closer to the optimal model. This leads to a new prediction:

> *Prediction* 3: Participants who are asked to think about outcomes other than the most extreme outcome will expand their focal set and display less future-biased choice than those participants who think only about the most extreme outcome.

Experiment 3 also provides an opportunity to test some alternative theories regarding the observed behavior. One explanation for making non-optimal choices is that participants simply do not understand the probability distribution of outcomes, perhaps believing that extreme outcomes are far more probable than they actually are. In this experiment, participants' beliefs can be directly assessed by analyzing their responses to the likelihood questions used to manipulate which outcomes are in the focal set. This experiment also provides an opportunity to look for any effects of local features of the sequences—in other words, does a participant who sees a series of outcomes '100, 200, 400' react differently than a participant who sees a series '300, 200, 400'? Finally, since regret may also be playing a role in future-bias, new questions were added to the post-experiment questionnaire to explore how much regret participants are experiencing for certain decisions they made within the experimental task.

EXPERIMENT 3

Method

The basic task employed in this experiment is the same as for Experiment 2: a frequent flier paradigm in which participants search for the best occasion on which to use a free trip. The distribution of ticket prices was the same as used in Condition 1 of Experiment 2. Ninety-six undergraduates at a Southern business school completed the experiment on computer. To manipulate which outcomes were most focal for participants (thus

DOI: 10.1002/bdm

directly testing the focalism explanation) and to collect information about participants' expectations regarding the probability of certain outcomes, a new set of questions was inserted between the instructions and the start of the three trials. The questions asked participants to assess both the likelihood (described in probabilistic terms) and the expected frequency of either the most extreme outcome ($1000 ticket) or a larger set of three outcomes (the $400, $500, and $1000 tickets). Participants were asked to answer these questions from memory and were not able to look directly at the outcome distribution from the instructions while they answered the questions. Both probability and frequency measures were taken based on research suggesting that participants find it easier to understand frequencies than probabilities (Gigerenzer, 1996; Slovic, Monahan, & MacGregor, 2000); any confusion by participants about the actual likelihood of certain outcomes should be picked up by one or both of these measures.

To check for local sequence effects, and to be able to ask specific questions that directly addressed real choices made by the participants, a subset of six pre-generated sequences were used rather than the unlimited set of randomly generated sequences used in Experiments 1 and 2. The sequences were first randomly generated, and four sequences were then partly modified to be able to test local pattern effects. More specifically, the modified patterns consisted of two pairs of sequences in which the first four outcomes were matched, but the second outcome was varied, within each pair. Patterns A and B were adjusted to begin with the outcome sets [$100, $400, $100, $100] and [$100, $300, $100, $100], respectively. Patterns C and D were adjusted to begin with the outcome sets [$200, $100, $200, $400] and [$200, $300, $200, $400], respectively. If individuals make assumptions about trends in local sequences, then the different ticket values appearing in period two may affect later choices on high-value tickets; for example, an individual observing pattern D may believe that the entire trial will include high cost tickets and be more likely to delay use than an individual who observes pattern C. Each participant saw only three of six sequences, in random order. After completing all three trials, participants answered questions about their experienced regret regarding actual choices made during the three trials. The questions focused on the amount of regret they felt from using a free trip too early (e.g., using it in an early period and experiencing a $1000 outcome later in the sequence) and/or too late (e.g., bypassing an opportunity to use it on a $400 outcome early in the sequence).

Results

As with Experiments 1 and 2, the overall results indicate that participants are delaying choice past the recommendations of the optimal model. Compared to the normative benchmarks, subjects spend more per trial ($2365 vs. $2263, $t(95) = 9.91$, $p < .001$), use the free trip on less valuable tickets ($286 vs. $383, $t(95) = 9.11$, $p < .001$), and use it later in the trial (11.0 vs. 5.8, $t(95) = 12.62$, $p < .001$).

These results become even more interesting when separated out for the two focalism conditions, for which some participants were asked likelihood questions about only the most extreme outcome ($1000) and others were asked about a larger set of outcomes ($400, $500, and $1000). If focalism is a contributor to the observed future-bias behavior, then we should find that increasing focus on the larger set affects which outcomes are seen as focal and debiases the behavior. The main dependent variables per condition are shown in Table 6. Results for both conditions continue to show future-bias, with participants in each condition significantly spending more per trial, using the trip on less valuable tickets, and using it later than the optimal model predicts. However, the results also show that, as predicted, subjects in the condition with the larger focal set are less extreme in their results that those in the condition where only the most extreme outcome was made focal. Comparing the two conditions, participants who assessed likelihood for the larger focal set spent less per trial ($2344 vs. $2386, $t(47) = 2.44$, $p = .009$), used the free trip on more valuable tickets ($318 vs. $255, $t(47) = 4.24$, $p < .001$), and used it earlier in the trial (8.8 vs. 13.2, $t(47) = 8.07$, $p < .001$) than those who assessed likelihood for only the most extreme outcome. Thus, although increasing the number of focal outcomes does not entirely debias the behavior, it does appear to reduce it relative to the group for whom only

DOI: 10.1002/bdm

Table 6. Average results for all 15-week trials for the main dependent variables, for each of the two focalism conditions in Experiment 3

Condition	Cost (total dollars spent)	Value of ticket for free coupon use	When free coupon used
1 ($1000 focal)	$2386 (34.1)	$255 (16.8)	13.2 (0.27)
2 ($\geq$$400 focal)	$2344 (37.9)	$318 (14.6)	8.8 (0.45)
Optimal	$2263 (44.5)	$383 (5.7)	5.8 (0.31)

Both optimal model predictions and empirical results are presented. Standard errors are shown in parentheses.

the $1000 outcome is focal. Further analysis of each of the six sequences used in this experiment show that these significant differences per condition are also obtained within each individual sequence.

One benefit of the focalism manipulation is that it also allowed us to collect likelihood estimates for various outcomes. Analysis of these likelihood estimates provides a measure of whether participants actually understand the probability distribution they have seen in the experiment. If subjects have understood the probability distribution given to them in the instructions then there should be a significant difference in their answers between the two conditions; specifically, those asked about the larger set of outcomes should indicate higher likelihoods than those asked only about the most extreme outcome. Responses to the probabilistic likelihood question were collected on an 8-point scale with 1 being 'absolutely impossible', 4 being '50/50 chance', and 8 being 'completely certain'. The average response to this question was significantly higher for individuals in the condition with the larger outcome set (2.73 vs. 2.29, $t(47) = 2.42$, $p = .01$). For the question regarding frequency, participants indicated how often they expected one of the listed ticket prices to appear during the 45 weeks included in the three trials. Results are shown in Table 7; again, subjects in the condition with the larger outcome set expect those outcomes to occur more regularly than subjects in the other condition. Compared to the normative answers, participants in both conditions appear to be fairly accurate in their frequency predictions. These results suggest that the experimental subjects are both sensitive to differences in outcome probabilities and reasonably accurate about the actual frequency of outcomes, so future-bias is not simply an effect of not understanding the likelihood of the outcomes.

An investigation of local sequence effects is possible through direct comparisons between some of the outcome sequences used in this experiment. For the first two matched sequence pairs, patterns A and B, the question was whether seeing a $400 ticket versus a $300 ticket in period two would affect their probability of using the free trip on later high value tickets; in other words, does the effect of forgoing use on an early $400 ticket change their later usage behavior? A comparison of the free trip usage for these two patterns finds no difference in when their free trip was actually used (13.2 vs. 13.6, $t(94) = .42$, $p = .34$); subjects in both groups delay use until the end of the trial. For the other two matched sequences (patterns C and D), we are interested in differences in usage on the $400 ticket in period four after bypassing either a $100 or a $300 ticket in the second period. (Note that using the free trip coupon on the $400 ticket is the optimal use for both sequences.) Of the 48 participants who see each sequence, 16 who observe the $100 outcome use the free trip on period four's $400, and 10 who observe the $300 outcome make the same decision. This difference is not significant as tested by a logistic regression on the period four usage decision ($z = -1.54$, $p = .12$). Interestingly, a closer inspection does reveal differences in the period four usage for these two sequences according to which condition participants were in; including condition in the logistic regression shows that participants for whom the $400 was made focal were significantly more likely to use the free trip in that period ($z = -4.0$, $p < .001$), consistent with the earlier results per focalism condition. Thus, local features of the observed sequences do not appear to be influencing the probability that the free trip gets used optimally, but the focalism manipulation does have a significant effect on their decisions.

After completing the experimental task, a final set of questions was given to participants to ask them about any regret they may have experienced due to certain decisions made during the experiment. Regret was measured on 5-point scale (1 = very high regret, 5 = no regret). The regret questions differed in accordance

DOI: 10.1002/bdm

Table 7. Experiment 3 results for pre-task survey in which participants were asked to predict the frequency of certain outcomes occurring within 45 weeks of observations

	Predicted frequency of outcome occurrence						
Outcomes	0	1–3	4–6	7–10	11–15	16–25	26+
\$1000 only	10	29	3	3	1	2	0
\$400, \$500, or \$1000	0	25	9	8	5	0	1

Condition 1 participants predicted the frequency of only the \$1000 airfare (1% probability) while Condition 2 participants predicted the frequency of the top three airfares combined (5% probability).

with the actual decisions made by participants but can be categorized into two main types: regret about using the free trip coupon early in a trial and experiencing a higher value ticket later (\$1000 or \$500), or regret about bypassing a high value ticket (\$400 or \$500) early in a trial and being stuck using the free trip on a lower value ticket later. Sixty-nine participants answered a question about early use and 80 participants answered a question about late use (some participants answered both types of questions if they experienced both types of outcomes within the three trials). Analyzing regret levels according to type, participants appear to express higher levels of regret from too early use than too late use (2.05 for early, 2.75 for late, $t(144) = 3.88$, $p < .001$). Thus, participants do seem to be reporting higher regret from early use.

Experiment 3's results provide additional evidence for a focalism-based explanation for the future-biased behavior observed in the search task from Experiment 1. By manipulating which outcomes were most likely to be part of the searcher's focal set (namely, whether only the most extreme \$1000 ticket or the larger group of three high value tickets were salient), behavior was influenced such that participants with the larger focal set achieved outcomes closer to those prescribed by the optimal model. Increasing the focal set thus appears to partially debias the delay behavior observed in the earlier experiments. It is also important to note that, from the participants' perspective, the manipulation to increase the focal set is a subtle one; rather than providing explicit instructions to think about more than the single most extreme outcome, the experiment simply asked about the frequency of the various outcomes. This suggests that efforts to debias the focal thinking that leads to future-bias need not be heavy-handed, so long as they successfully draw the searcher's attention to the larger set of alternative outcomes.

In addition to directly testing the role of focalism, Experiment 3 simultaneously provided the opportunity to test some alternate explanations for the behavior observed in Experiments 1 and 2. First, the perceived likelihood and frequency information collected from participants at the start of the experiment provide some evidence that participants do understand the probability distribution of the outcomes, and that they are reasonably accurate in their predictions. Second, by using pairs of patterns with similar sequences of outcomes, we were able to test whether local sequence properties were affecting the behavior. Any effects from local sequence patterns appear to be insignificant, at least for the set of patterns tested here. Finally, a set of post-task questions were used to gather some information about how regret may be related to decisions to delay; results from these questions demonstrate that participants' experienced regret from 'too early' usage is higher than regret from 'too late' usage. If decision-makers accurately predict that they will feel more regret about early use, then the observed future-bias in this task may be a positive way of avoiding that regret.

GENERAL DISCUSSION

The goal of this paper has been to examine how certain types of search tasks, specifically those with low probability focal outcomes, can lead to systematically non-optimal behavior in which search is continued beyond the benchmark provided by a normative model. This propensity to delay choice (and continue search) beyond the optimal is an example of future-biased choice. By using a task in which searchers knew the full

DOI: 10.1002/bdm

distribution of possible outcomes, knew that search would continue even after a selection was chosen, and had in mind a preferred but relatively improbable 'best' outcome, we were able to document evidence of future-biased behavior. One possible psychological explanation for such behavior, focalism, was supported by showing that manipulation of both the probability of the focal 'best' outcome and the value of the second-best alternative had little effect on the actual search behavior. In all three conditions of Experiment 2, searchers behaved as if the focal outcome was likely to be encountered, and they delayed choice until that focal outcome appeared. More specifically, their behavior is consistent with a heuristic based on setting a threshold for selection, where that threshold is set too high (and thus choice is delayed). Experiment 3 provided additional support for the focalism explanation by demonstrating that forcing searchers to think more directly about second-best outcomes reduced the level of delayed choice. Manipulating the outcomes that are included in the focal set thus appears to partially (but not completely) debias the searchers by lowering their selection threshold.

Based on the evidence of future-biased choice in the task explored in this paper, it is interesting to consider what other types of search tasks are also likely to lead to future-bias. The primary driver for future-bias explored here has been the effects of focalism: thinking about the best possible match causes searchers to reject lesser alternatives. However, there may also be psychological effects due to anticipated regret which cause the decision-maker to delay. First, anticipated regret may be a strong driver of which outcome is chosen as the focal outcome. In the experimental task used in this paper, observed behavior was consistent with the most expensive outcome being focal; selection of this outcome as focal could be based on a concern that regret levels will be highest if the free trip is already used when this outcome occurs, as compared to the regret from still holding the free trip at the end of the trial. The results for regret levels for different types of outcomes in Experiment 3 support this explanation. Second, the nature of this task, in which search continues even after the choice is made, is one that allows for much greater possibility of regret than an environment in which search ends immediately. This paper has only superficially investigated the aspects of regret inherent in this search environment; additional work deserves to be done that more closely examines the role of regret as a determinant of future-bias. For example, while Experiment 3 demonstrated that experienced regret is higher for early use than for late use, it is unclear whether anticipated regret will follow the same pattern. If regret is a contributor to the future-biased choices observed in this task, then we may also find evidence of future-bias in tasks without full information about the distribution of outcomes. For example, a more traditional Secretary Problem search task, adapted so that observation of other alternatives continues after the selection is made, would provide an interesting environment in which to examine this problem.

In addition to considering the search environments that lead to future-bias, it may also be that certain types of items are more likely to have their use delayed than others. Based on the psychological drivers of future-bias, the behavior is most associated with items whose use is constricted by strict self-control rules and associated high thresholds. Thus, items with a high indulgence or luxury value could be as susceptible to future-bias as the single-use or limited items explored here, due to the fact that consumers institute strict self-control rules to restrict replacement of highly indulgent items (Thaler, 1985; Wertenbroch, 1998). Anecdotal examples like the saved special bottle of wine or future exotic vacation would fit this model. Indeed, the Wall Street Journal has instituted an annual event known as 'Open That Bottle Night' to encourage individuals saving special wine bottles to finally open and consume them. Early evidence from studies on consumption of these types of restricted luxury items suggests that individuals do indeed attempt to match the usage of special items with tightly defined special occasions, and that the thresholds for what constitutes such an occasion are often set too high, leading to future-bias (Shu, 2006). Just as with the high thresholds set by participants in the search task investigated here, an extreme focus on the focal 'best' outcome is one of the underlying psychological bases for the behavior.

While learning from experience was not evident among participants in these three experiments, it should not be assumed that these individuals cannot learn to make more optimal choices. It may be that repeating this experimental task a larger number of times would allow participants time to learn a more optimal strategy and

DOI: 10.1002/bdm

begin to correct their bias naturally. However, the small number of trials used in these tasks was intentionally designed to mirror many real-life search tasks that individuals might face only a few times during a lifetime, such as the wine drinkers with special bottles in the previous paragraph. It should also be noted that learning in such real-life search tasks is further complicated by the asymmetric nature of the outcomes. A future-biased searcher with unlimited timeframes who delays because of a belief that a better option will arise can continue to search indefinitely without being proven wrong. In contrast, an individual who resists future-bias and accepts an early outcome will at best be satisfied with that outcome, and at worst experience significant regret from having chosen an outcome too early. In other words, searchers may more frequently experience disappointment from choosing too early than from choosing too late, thus impeding their ability to learn.

The implications for the future-bias observed in this search task are that although individuals often exhibit a desire to end search early or experience an outcome immediately, there may sometimes be environments in which the opposite behavior predominates. Understanding such examples of future-bias provides us with additional insight into the boundaries of self-control and temporal myopia by investigating situations for which the typical predictions of consuming too early do not hold. Where previous research has demonstrated the ways in which individuals fail at self-control efforts, cases of future-bias seem to suggest that there are situations in which individuals are overly successful at restraining themselves. When future-bias happens for items as mundane as the free airline trip coupons studied here, and also for more highly controlled but highly indulgent items like bottles of fine wine, then individuals are procrastinating their own enjoyment. Such behavior is not new; the Roman poet Horace addressed the problem of delaying pleasure in 23 BC, and his prescriptive advice for avoiding future-bias remains true today: '*carpe diem, quam minimum credula postero*'.[11] In seeking the best option from a sequential search in which future alternatives will still be observed, decision-makers may be wise to follow his advice.

APPENDIX

Derivation of optimal policy

Consider an individual who faces a sequence of n events of varying magnitude, in which the probability distribution of outcomes is known in advance. The individual's goal is to select the event with the highest value at the time the event occurs. This could be selecting the best (most positive) option from a series, or exercising an option to avoid the worst (most negative) option. The model developed here will be for the latter situation of exercising an option to avoid a cost, although the general model applies equally well to other versions of the problem. More specifically, we will consider the problem of an individual holding frequent flier miles who must determine the optimal time to cash in those miles on a free trip when faced with a sequence of airfares from a known distribution.

We will begin with the assumption that the individual holds enough frequent flier miles to be allowed a single free trip (this assumption can be modified later). In addition, he faces a sequence of n occasions in which he must purchase an airline ticket. He knows only the cost of the current ticket for certain, and he also knows the distribution of possible prices for future tickets. This distribution of $j = 1$ to m discrete possible prices is represented by $(x_j\ p_j)$, where p_j is the probability of occurrence for each price x_j; $X = \sum_{j=1}^{m} p_j x_j$ is then the expected value of a ticket in any one period. The cost of the current ticket in the sequence is defined as c_i, $i = 1$ to n, where c_i is drawn from the distribution $(x_j\ p_j)$. The decision-maker must thus make the decision of whether to purchase the current ticket at its known posted price c_i and save his frequent flier miles

[11]'Seize the day, put no trust in tomorrow' Horace, *Odes* (I, xi).

DOI: 10.1002/bdm

for a future trip, or use the frequent flier miles at this time and have to pay full prices for all future trips. The individual's goal is to minimize the expected cost.

To simplify terminology, define a state-dependent variable $\Psi(f, i)$ to represent the expected cost in period i of all future periods dependent on whether or not the free trip is still available when entering that period. Defining f as 1 if the free trip is still available, and 0 otherwise, $\Psi(0, 3)$ would thus represent the expected cost of all periods from 3 to n when the free trip has already been used. Since the free trip will always be used in the final period if it has not been used previously, $\Psi(1, n) = 0$ by definition. If however, the free trip has already been used before reaching the final period, then $\Psi(0, n) = X = \sum_{j=1}^{m} p_j x_j$ (the expected value of the final ticket). Similarly, the expected value of all remaining weeks once the free trip is used can be written as $\Psi(0, i) = (n - i + 1) \times X$.

Start by considering the decision in the next-to-last period, $i = n$- 1, when the free trip has not yet been used. The individual faces a known cost of c_{n-1} in this period. If he chooses to use the free trip in this period, then he will pay 0 now but will have to pay the full price, c_n, in the final period, where c_n has the expected value of the overall distribution of airfares, $\Psi(0, n) = X$. If he chooses to save the free trip for the next period, then he will pay c_{n-1} for certain now and $\Psi(1, n) = 0$ in the next period. His decision takes the form:

$$\begin{aligned}&\text{Use now : cost} = 0 + \Psi(0,\ n) = 0 + X\\&\text{Use later : cost} = c_{n-1} + \Psi(1,\ n) = c_{n-1} + 0\end{aligned}$$

Thus, he will use the free trip now if $c_{n-1} > \Psi(0, n)$. The expected cost of period $n-1$, before c_{n-1} is known for certain, is given by:

$$\begin{aligned}\Psi(1, n-1) &= \mathrm{pr}(c_{n-1} < \Psi(0,\ n)) * (\mathrm{EV}(c_{n-1}|\mathrm{pass}) + \Psi(1,\ n)) + \mathrm{pr}(c_{n-1} > \Psi(0,\ n)) * (\Psi(0,\ n))\\&= \mathrm{pr}(c_{n-1} < \Psi(0,\ n)) * ((X|x_j < \Psi(0,\ n)) + \Psi(1,\ n)) + \mathrm{pr}(c_{n-1} > \Psi(0,\ n)) * (\Psi(0,\ n))\end{aligned}$$

where $(X|x_j < \Psi(0, n))$ is the expected value of the subset of the distribution for which $x_j < \Psi(0, n)$. In other words,

$$(X|x_j < \Psi(0,\ n)) = \frac{\sum_{j=1}^{m} p_j x_j}{\sum_{j=1}^{m} p_j} \quad \forall x_j < \Psi(0,\ n)$$

Now consider the decision in the preceding period, $i = n-2$. If the free trip is used in this period, then the total remaining costs will be 0 (for this period) plus $\Psi(0, n-1)$ (for the remaining two periods). If the free trip is not used, then the total remaining costs will be c_{n-2} (for this period) plus $\Psi(1, n-1)$, which is defined above. Thus, the free trip will be used in this period if:

$$0 + \Psi(0,\ n-1) < c_{n-2} + \Psi(1,\ n-1)$$

In other words, the free trip will be used if the current cost is greater than the threshold defined by the difference between the expected value of future periods with and without the trip. We can generalize the decision in any period i as being between $0 + \Psi(0, i + 1)$ if the free trip is used now, and $c_i + \Psi(1, i + 1)$ if it is saved for a future period. Defining the per period usage threshold as T_i, where

$$T_i = \Psi(0,\ i+1) - \Psi(1,\ i+1)$$

we can then state that the free trip will be used when c_i is greater than T_i. This allows us to define the expected value of the remaining periods from i to n, when the free trip is still valid and c_i is not yet known, as

$$\Psi(1,\ i) = \mathrm{pr}(c_i < T_i) * ((X|x_j < T_i) + \Psi(1,\ i+1)) + \mathrm{pr}(c_i > T_i) * (\Psi(0,\ i+1))$$

DOI: 10.1002/bdm

Table 8. Optimal model predictions for the sample distribution of the Appendix [$100, 0.30; $300, 0.40; $400, 0.25; $700, 0.05]

Period	1	2	3	4	5	6	7	8	9	10
$100	Save	Save	Save	Save	Save	Save	Save	Save	Save	Use
$300	Save	Save	Save	Save	Save	Save	Save	Save	Use	Use
$400	Save	Save	Save	Save	Save	Use	Use	Use	Use	Use
$700	Use	Use	Use	Use	Use	Use	Use	Use	Use	Use

Table entries show the results of the use free trip vs. save free trip (buy ticket) decision at each period/ticket value combination, conditioned on the free trip having not been previously used.

Solving this at every period provides a suggested usage pattern for every week, for every possible ticket price c_i given the current state of $\Psi(f,\ i)$.

For a discrete set of possible outcomes, the solution then works as a series of thresholds. The threshold at which the free trip should be used becomes lower as the number of remaining periods in the task decreases. Consider a simple example with 10 periods ($n = 10$), with possible ticket prices and associated probabilities of $100 (30%), $300 (40%), $400 (25%) and $700 (5%). The expected value for any one ticket is then $X = \$285$. Table 8 shows the optimal solution for each ticket price for each period, indicating whether the free trip should be used or saved, conditional on not yet having been used. In periods 1 through 5, the free trip should only be used if the ticket costs $700; in periods 6 through 8, it should be used for tickets of $400 or $700; in period 9, used for tickets of $300 and higher; and, finally, used for any ticket price in period 10.

ATTACHMENT 1: SAMPLE INSTRUCTIONS FOR EXPERIMENTAL FREQUENT FLIER TASK (EXPERIMENT 1)

Frequent flier game

Imagine that you have decided that you will fly from Chicago to Denver, Colorado every weekend to visit a close friend. You will make this trip for 15 consecutive weeks. Luckily, there is a new airline called Mountain Express that flies between Chicago and Denver regularly. However, this new airline has some unusual policies. For example, you are not able to book a flight more than 1 week in advance. Their policy is that Mountain Express will post the weekly round-trip fare every Monday, and that is when you can buy your ticket for the coming weekend. Airfares can change from week to week, but Mountain Express has promised that fares will vary according to a specific distribution:

1%	Chance of	$1000
2%	Chance of	$500
3%	Chance of	$400
19%	Chance of	$300
20%	Chance of	$200
55%	Chance of	$100

Since you are part of their frequent flier program, Mountain Express has rewarded you with several coupons. The first coupon allows you one free trip; in other words, instead of paying full fare for one of your trips, you can take that trip for no cost. They have also given you two discount coupons, which will allow you get 30% off the posted price for that week. All three of these coupons will expire in 15 weeks.

The computer game you will be playing today is to simulate purchasing the airplane tickets for the 15 weeks that you will be flying to Denver. In each week, you will see Mountain Express' posted airfare for that week. You will then need to decide whether you will buy the ticket at full fare, use your free trip, or use

DOI: 10.1002/bdm

one of your discount coupons. The game will be repeated three times (labeled as trial 1, trial 2 and trial 3). In each of the three trials, you will need to buy 15 weekly tickets, and you will have the three coupons to use within those 15 weeks.

Your goal is to spend as little money on these trips as possible. You will begin the game with an 'endowment' of $9000. The amount that you spend on weekly airfare will be subtracted from this endowment. At the end of each trial and the end of the game, you will be told how much is left in your endowment.

After the game is over, you will complete a survey about your experience.

ACKNOWLEDGEMENTS

I appreciate the comments of three anonymous reviewers, and of seminar participants at the Behavioral Decision Research in Management Conference, the University of Chicago, INSEAD, Harvard, and UCLA. I am also grateful to my dissertation advisors, Richard Thaler, George Wu, Yuval Rottenstreich, and France Leclerc, for their patience and guidance on this project.

REFERENCES

Ainslie, G. (2001). *Breakdown of will*. Cambridge: Cambridge University Press.

Beardon, J. N., Rapoport, A., & Murphy, R. O. (2006a). Sequential selection and assignment with relative ranks: An experimental study. *Journal of Behavioral Decision Making*, *19*, 229–250.

Beardon, J. N., Rapoport, A., & Murphy, R. O. (2006b). Sequential observation and selection with rank-dependent payoffs: An experimental test of alternative decision rules. *Management Science*, *52*, 1437–1449.

Buchholz, B. (2003). An endless odyssey: Looking for Mr. Perfect in an imperfect world. *Chicago Tribune*, August 6, C1.

Cheema, A., & Soman, D. (2006). Malleable mental accounting: The effect of flexibility on the justification of attractive spending and consumption decisions. *Journal of Consumer Psychology*, *16*, 33–44.

Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, *4*, 282–296.

Gigerenzer, G. (1996). The psychology of good judgment: Frequency formats and simple algorithms. *Journal of Medical Decision Making*, *16*, 273–280.

Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, *61*, 35–73.

Heath, C., & Soll, J. B. (1996). Mental budgeting and consumer decisions. *Journal of Consumer Research*, *23*, 40–52.

Iyengar, S., Wells, R., & Schwartz, B. (2006). Doing better but feeling worse: Looking for the "best" job undermines satisfaction. *Psychological Science*, *17*, 143–150.

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, *39*, 17–31.

Kahneman, D., & Tversky, A. (1979). Intuitive prediction: Biases and corrective procedures. *TIMS Studies in Management Science*, *12*, 313–327.

Kamien, M. I., & Schwartz, N. L. (1972). Some economic consequences of anticipating technical advance. *Western Economic Journal*, *10*, 123–138.

Kivetz, R., & Simonson, I. (2002). Self-Control for the righteous: Toward a theory of precommitment to indulgence. *Journal of Consumer Research*, *29*, 199–217.

Koehler, D. J. (1991). Explanation, imagination, and confidence in judgment. *Psychological Bulletin*, *110*, 499–519.

Koehler, D. J. (1994). Hypothesis generation and confidence in judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 461–469.

Lee, D. L., O'Connor, T. A., & Welsh, M. B. (2004). Decision-making on the full-information secretary problem. In K. Forbus , D. Gentner , & T. Regier (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

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

Loewenstein, G., & Elster, J. (1992). *Choice over time*. New York: Russell Sage Foundation.

DOI: 10.1002/bdm

Rapoport, A., & Tversky, A. (1970). Choice behavior in an optimal stopping task. *Organizational Behavior and Human Performance*, *5*, 105–120.
Rottenstreich, Y., & Kivetz, R. (2006). On decision making without likelihood judgment. *Organizational Behavior and Human Decision Processes*, *101*, 74–88.
Samuels, S. M. (1991). Secretary problems. In B. K. Gosh , & P. K. Sen (Eds.), *Handbook of sequential analysis*. New York: Marcel Dekker.
Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the 'secretary problem'. *Organizational Behavior and Human Decision Processes*, *69*, 221–236.
Seale, D. A., & Rapoport, A. (2000). Optimal stopping behavior with relative ranks: The secretary problem with unknown population size. *Journal of Behavioral Decision Making*, *13*, 391–411.
Shu, S. B. (2006). Occasion matching for indulgences. SMU working paper.
Slovic, P., Monahan, J., & MacGregor, D. G. (2000). Violence risk assessment and risk communication: The effects of using actual cases, providing instruction, and employing probability versus frequency formats. *Law and Human Behavior*, *24*, 271–296.
Tatarkiewicz, W. (1976). *Analysis of happiness*. Warszawa, Poland: Polish Scientific Publishers.
Thaler, R. H. (1985). Mental accounting and consumer choice. *Marketing Science*, *4*, 199–214.
Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, *12*, 183–206.
Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, *47*, 641–654.
Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing Science*, *17*, 317–337.
Wilson, T. D., Wheatley, T., Meyers, J. M., Gilbert, D. T., & Axsom, D. (2000). Focalism: A source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, *78*, 821–836.
Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*, 503–519.

Author's biography:

Suzanne B. Shu (PhD) is an Assistant Professor of Marketing at Southern Methodist University. Her research interests include intertemporal choice and judgments, self control, and marketing of financial products.

Author's address:

Suzanne B. Shu, Southern Methodist University, Cox School of Business, PO Box 750333, Dallas, TX 75275, USA.

DOI: 10.1002/bdm