

The Effect of Feedback and Financial Reward on Human Performance Solving ‘Secretary’ Problems

James Campbell (james.campbell@psychology.adelaide.edu.au)

School of Psychology, University of Adelaide
South Australia, 5005, AUSTRALIA

Michael D. Lee (mdlee@uci.edu)

Department of Cognitive Sciences, University of California, Irvine
Irvine, CA, 92697-5100

Abstract

The secretary problem is a recreational mathematics problem, suited to laboratory experimentation, that nevertheless is representative of a class of real world sequential decision-making tasks. In the version of the problem we consider, an observer is presented with a sequence of values from a known distribution, and is required to choose the maximum value. The difficulties are that a value can only be chosen at the time it is presented, that the last value in the sequence is a forced choice if none is chosen earlier, and that any value that is not the maximum is scored as completely wrong. Previous research has found large individual differences in people’s ability to behave according to the known optimal solution process. In addition, there is some evidence that, at least under some conditions, these differences are stable, in the sense that there are no significant learning effects. We examine the effect of financial reward and of feedback on people’s performance over a series of 120 five-point problems, in a 2×3 between-subjects design. Our main finding is that people perform very similarly in all six experimental conditions, and there is no evidence people learn to perform better in any condition.

Introduction

Many real world decision-making problems are sequential in nature. A series of choices is made available over time, and it is often efficient (and sometimes even necessary) to make a selection without waiting to be presented with all of the alternatives. In this paper, we use a recreational mathematics problem known as ‘secretary problems’ (see Ferguson 1989 for a historical overview) to study human decision-making on a sequential optimal stopping problem in a controlled laboratory setting.

In secretary problems, an observer is presented with a sequence of possible choices, and must decide whether to accept or reject each possibility in turn. The number of choices in the complete sequence is fixed and known, and only the current alternative in the sequence is presented to the observer. The goal is for the to observer choose the best possibility in the sequence, under a 0-1 loss function (i.e., if they choose the best alternative their decision is correct, but any other choice is regarded as completely incorrect). In the original formulation of the secretary problem, the rank of the




Figure 1: A sample secretary problem of length 5, with the sequence of values shown by circles, and the optimal threshold shown as a solid line.

current alternative, relative to those already seen, is presented. In the full information version, sometimes known as the ‘Cayley’ problem, that we study here, observers are presented with a numeric score drawn independently from a known distribution for each alternative.

Solving Secretary Problems

Gilbert and Mosteller (1966) provide a thorough and useful overview of early mathematical analysis of several versions of the secretary problem. For the full information version we study, the optimal decision process requires choosing the first value that is the maximum value observed in the sequence thus far *and* exceeds a threshold level for its position in the sequence. Gilbert and Mosteller (1966, Tables 7 and 8) detail these optimal thresholds and the associated probabilities of making a correct decision.

For the five-point problems we study, where the values are two decimal point numbers uniformly dis-

tributed on the interval $[0, 100]$, the optimal thresholds are 82.46, 77.58, 68.99, and 50.00 for the first four positions respectively. Figure 1 provides a graphical example of a five-point problem, showing a sequence of values and the optimal thresholds for each position in the sequence. Following the optimal decision process for these problems the expectation is that about 64% of problems will be solved correctly.

The existence of a known optimal solution process distinguishes secretary problems from other difficult combinatorial optimization problems, such as visually presented Traveling Salesperson Problems (TSPs), that have recently been studied in the context of human problem solving abilities (e.g., MacGregor & Ormerod 1996; Vickers, Butavicius, Lee, & Medvedev 2001). In particular, it permits the measurement of human performance in terms of adhering to the optimal process (which can always be achieved), rather than achieving the optimal outcome (which cannot, and so constitutes an inherently noisy measure of performance).

Individual Differences and Learning

In this context, Lee (2006) observed that, over a total of 147 participants, each completing one of two different sets of 40 problems, there was evidence of individual differences, but no evidence of learning. In other words, the proportion of times the optimal solution process was followed differed between participants, but did not appear to change as the same participant answered additional problems.

Burns, Lee and Vickers (in press) seized on this suggestion of stable individual differences, and explored the relationship between performance on secretary problems and standard psychometric measures of cognitive abilities. Within a standard Cattell-Horn-Carroll framework of intelligence, these authors demonstrated that performance on the Secretary Problem loaded on fluid intelligence (Gf), with performance on the problem also being shown to load approximately 0.4 on a general ability factor, (g). Interestingly, this g-loading was comparable to that of the Digit Symbol task from the Wechsler Adult Intelligence Scale. It was tentatively concluded by Burns et al. (in press) that performance on the Secretary Problem might be able to be used as a measure of cognitive ability, but that further investigation was necessary. In particular, they noted that the possibility people's performance might improve (or, more generally, change) over repeated trials required further investigation.

There has been little additional relevant research considering the possibility of learning over repeated trials in secretary problems. Seale and Rapoport (2000) were inconclusive as to whether learning effects were present in rank order versions of the Secretary Problem. Bearden, Murphy and Rapoport (2005) reported very small learning effects for an extended

'multi-attribute' secretary problem, but never explicitly tested the rival hypothesis that there was no learning, which would seem a more parsimonious explanation for the raw data they display.

Perhaps most importantly, the experiments in which Lee (2006) found stable individual differences did not provide any feedback to participants regarding the quality of their decisions, did not provide any financial reward or other motivation for good decision-making, and involved only relatively small problem sets. In this study, we undertake a more thorough investigation of learning effects, by manipulating both the type of feedback that is provided, and whether or not financial reward is given, and by using a much larger set of 120 five-point problems.

Feedback

The general use of the feedback available after making decisions is essential for adaptation and survival, but it seems likely that different types of feedback will have different influences on decision-making (e.g., Einhorn & Hogarth, 1981; Jacoby, Troutman, Mazursky & Kuss 1984). One prominent suggestion (e.g., Wofford & Goodwin, 1990) is that people's decisions will tend not to change when they are given positive feedback, but will tend to change when given negative feedback.

For example, Rimm, Roesch, Perry and Peebles (1971) investigated the role of non-informative and blank feedback administered randomly along with positive and negative outcome feedback in a sequential decision making task. Their results, and subsequent re-analysis by Spence (1972), suggested that when people were given non-informative feedback after a decision, they were likely to make similar decisions in subsequent problems. It appeared that decision makers were interpreting the feedback as indicating correctness. Levine, Leitenberg and Richter (1964) suggest this sort of behaviour is generalized from experience, as everyday decisions that are correct are often not followed by feedback, whereas incorrect everyday decisions are often followed by immediate feedback.

It seems obvious that feedback with more information regarding the decision process and outcome will generally be more effective in improving performance. One useful distinction is provided by Jacoby et al. (1984), who adapted the notions of 'outcome feedback' and 'cognitive feedback' from Social Judgement Theory. Outcome feedback is made up of information regarding the accuracy of a response, whereas cognitive feedback involves information underlying the how and why of this accuracy. Under this view, cognitive feedback has a higher information value than outcome feedback, because it augments the predictive value of indicating decision accuracy with the explanatory value of allowing the decision-maker to understand the quality of their decision. Jacoby et al. (1984) reported that feedback with both explanatory and predictive value

is most effective for promoting high levels of performance in decision-making tasks. Furthermore, they suggested that good decision makers were very effective at ignoring outcome feedback, when it contained neither explanatory nor predictive value.

Financial Reward

Camerer and Hogarth (1999) reported that “a search of the American Economic Review from 1970-1997 did not turn up a single published experimental study in which subjects were not paid according to performance” (pp. 31). Many psychological experiments, however, do not use financial rewards, with some studies questioning their capability to eliciting high performance (e.g., Hertwig & Ortmann, 2001).

There is some evidence, however, supporting the value of offering financial rewards in psychological experimentation. For example, in a decision-making study, Parco, Rapoport and Stein (2002) reported “when learning is possible, monetary payments may bring the decisions closer to the predictions of the normative models.” (pp. 296). Camerer and Hogarth’s (1999) meta-analysis concluded that financial reward can improve performance under some circumstances, particularly in judgment and decision tasks.

Bonner, Hastie, Sprinkle and Young’s (2000) review found that for judgment and choice tasks, the most effective incentive scheme was a quota payment schedule, with individuals receiving a flat rate irrespective of performance until a certain target level of performance (quota) is reached. Once this quota is achieved, the individual receives a bonus. In one of the few studies incorporating feedback and incentive, Bucklin, McGee and Dickinson (2003) successfully used a piece-rate scheme of payment, paying a pre-defined amount of money for each correct response to increase performance. These authors concluded that feedback amplified the positive effect of financial reward on performance. In other words, financial reward, when combined with feedback in the form of the percentage of correct answers, produced higher performance from the decision makers when compared to a combination of either reward and no feedback, or base-pay and feedback.

Experiment

Our experiment involves six conditions in a 2×3 between-subjects design. Financial reward is either provided or not provided, and there are three types of feedback: no feedback, outcome feedback, and fullfeedback.

Participants

The financial reward groups included 12, 12 and 13 participants for the full, outcome and no feedback conditions, respectively. The no financial reward groups

included 14, 12 and 12 participants for the full, outcome and no feedback conditions, respectively. All participants were drawn from the population of University of Adelaide first year students, and all of the groups had broadly similar age and gender distributions.

Method

Basic Procedure Each participant completed the same 120 five-point problems, which were divided into 12 blocks of ten. Each participant attempted the 12 blocks in the same order, however, the ten problems within each block were randomized. For each of the problems, participants were sequentially presented with five numbers ranging from 0.00 to 100.00, with the task being that the maximum value be selected. When a participant chose a number, they rated their confidence from ‘definitely wrong’ to ‘definitely correct’ on a nine-point confidence scale.

Feedback Manipulation Having made a choice, participants receiving no feedback were presented with the next problem, and so were not informed as to whether they made the correct choice. The overall score tally was not displayed to these participants.

Participants receiving outcome or full feedback were shown a simple ‘correct’ or ‘wrong’ message after they had made their selection, indicating whether their choice was the maximum in the sequence.

Participants receiving full feedback condition were additionally shown graphically all five numbers in the problem as a bar graph, annotated with the actual numbers in their digit form, together with arrows highlighting their choice and the maximal number in the sequence. In both of the feedback conditions, an overall score tally was displayed.

Financial Reward Manipulation Participants in the no financial incentive conditions were asked to “try their best to obtain as many correct answers as possible” with no extrinsic reward. Participants in the financial incentive conditions, were informed that there was a monetary reward for high performance. The incentives followed a quota-piece rate scheme. A \$5 reward is paid to the participants in the financial incentive group regardless of their performance with an additional \$5 reward being paid after every 12 correct responses once the participant has answered 40% of the problems correctly, with a ceiling imposed on the payments after 80% of responses had been correctly answered, such that the maximum a participant could earn was \$30.

Results

We consider the results from two perspectives. First we examine the central question of learning: whether there is evidence of people improving over repeated trials, and how their change in performance depends on the experimental manipulations. To anticipate the




Figure 2: The average proportion of times participants in each experimental condition followed the optimal decision rule for the first 40, second 40, and third 40 problems. One standard error is shown.

results, we find no evidence of learning in any of the six experimental conditions. Accordingly, we also consider the results in terms of overall performance across all of the trials. Here we find some interesting dependencies on the experimental manipulations.

Learning Figure 2 shows the average proportion of times participants in each experimental condition followed the optimal decision rule for the first 40, second 40, and third 40 problems they completed. We make two main observations. The first is that it appears learning did not occur across the 120 problems, since there is no evident improvement in performance. The second is that financial incentive may have had an effect on performance, since those conditions with reward tend to outperform those that did not.

To test these possibilities, we used a recently developed Minimum Description Length (MDL) clustering technique that is well suited to making inferences about the similarities and differences between learning curves (Navarro & Lee, 2005). The technique involves defining statistical models for the data generating process, and then partitioning them using the Normalized Maximum Likelihood criterion (Risannen, 2001). We consider a range of models that make different meaningful assumptions about the relationships between the learning curves for the six experimental conditions. For each of the models, we find the number of bits used by the Maximum Likelihood code, which is essentially a measure of goodness-of-fit, and the number of bits used by the Normalized Maximum Likelihood code, which is essentially a measure of

Table 1: Four partitioning models for the six learning curves, and their MDL evaluation. The partition model is shown by the bracketing of two-character strings, with the first giving the feedback condition (F=full, O=outcomes, N=none) and the second giving the financial reward (N=no reward, R=reward).

Model	Fit	Comp	Tot
(FN,ON,NN,FR,OR,NR)	366	164	530
(FN,ON,FR,OR,NR) (NN)	320	273	593
(FN,ON,NN) (FR,OR,NR)	322	290	612
(FN)(ON)(NN)(FR)(OR)(NR)	259	682	941

the complexity of the model associated with the code. Summing these two numbers gives the total number of bits used by the Normalized Maximum Likelihood code, which is an overall measure of the likelihood of the model that balances goodness-of-fit and complexity. The lower the total number of bits, the lower the description length, and the more likely the statistical model.

The fit, complexity and total bit counts for four of the competing models we considered are detailed in Table 1. It is clear that the most likely model of the data is one that simply assumes all six learning curves belong to the one partition, using only 530 total bits of information. Thus the data supports the null result that there is no difference between each of the respective experimental groups. The second most likely model we found, needing 593 total bits, assumes that all the curves belong in one partition except the curve representing the no feedback and no financial reward condition. The third most likely model we found, needing 612 bits, assumes that there are two partitions, one with all the learning curves of those receiving no financial reward, and one with all the curves of those receiving financial reward. While these models are less likely than the null result, they are much more likely than the saturated model, shown for comparison, which assumes each of the learning curves belongs to its own partition, and needs 941 bits.

Our conclusion from this analysis is that the best justified inference we can make is that none of the performance curves differ from one another in any major way. Despite some interesting and interpretable suggestive trends, it appears neither feedback nor financial reward change the performance curves significantly.

Although not reported in detail here, an extremely similar pattern of results is obtained by looking at the change in mean confidence over trials, or by looking at every ‘yes’ or ‘no’ decision in solving the problems, rather than just the final choices. There is no evidence of systematic change over trials, nor of significant differences between the experimental conditions.




Figure 3: Average proportion of decisions following the optimal decision rule for all six conditions. One standard error is shown.

Overall Performance Figure 3 shows the average proportion of times participants in each experimental condition followed the optimal decision rule over all problems. To make statistical decisions about the possible effects of the experimental manipulations, we again generate different statistical models making different assumptions about how the experimental manipulations affect the dependent variable, and then used (approximate) Bayesian model selection to choose between them.

The first model, *constant*, assumes that neither feedback nor reward had any effect on performance, and so all the means are captured by the same single parameter. The *reward* model assumes that only the presence or absence of financial reward affect performance. The *feedback* model assumes that only the type of feedback affects performance. The *one-way* model assumes that both manipulations affect performance in an independent way. This model assumes not only that feedback has an effect on performance, but that financial rewards also effects performance, and that this effect is constant for all feedback types. The *suggested* model assumes that there is a dependency between the manipulations such that the presence of financial reward affects full and no feedback conditions in a constant fashion, but has no effect on outcome feedback conditions. The *full* model assumes that both manipulations affect performance and that they interact in a completely unrestrained way. This saturated model has six parameters: one for each of the six experimental groups.

Figure 4 shows the log maximum likelihood fits (up to an irrelevant additive constant), using a Gaussian likelihood function defined by the means and variances




Figure 4: Best fits of six alternative models to the proportion of correct decisions..

Table 2: Bayesian analysis for six statistical models of how experimental manipulations effects the proportion of correct decisions.

Model	Fit	Comp	BIC	BF
Constant	38.81	1	40.60	> 100
Incentive	11.73	2	15.31	37.38
Feedback	31.30	3	36.68	> 100
One-Way	5.16	4	12.32	8.39
Suggested	0.90	4	8.07	1
Full	0	6	10.75	3.82

in Figure 3, for each of the six models. From these fits, and the parametric complexity of the models, the Bayesian Information Criterion (BIC) can be calculated. The BIC values allow Bayes factors between each pair of models to be estimated, quantifying how much more likely one model is than another (see Kass & Raftery 1995).

The results of this analysis are detailed in Table 2. The most likely model is the suggested model, with the full model being 3.82 less likely, the one-way model being 8.39 times less likely. The remaining models are far less likely. Accordingly, we conclude there is evidence for an interesting interaction between the feedback and financial reward manipulations at the level of overall performance. In particular, it seems people receiving full feedback or no feedback perform better when given financial reward, but the same is not true for people given intermediate outcome feedback.

Conclusions

Our main finding is that, regardless of feedback or financial reward, people’s ability to follow the optimal

decision process did not improve over the course of 120 problems. One possible reason for this is the uncertain nature of the feedback itself, given the probabilistic relationship between following the optimal decision process and actually choosing the maximum value.

An additional intriguing possibility relates to the numerical format used to present the values in the current (and most previous) experiments. Learning the optimal decision process relies on tuning a series of threshold values, and it might be that four digit decimal numbers are not representations cognitively amenable to continuous small adjustment. It would be interesting to repeat essentially the same experiment, trying different formats (or even modalities) for presenting the stimuli, such as lines of different lengths, or tones of different pitch, and consider whether learning thresholds becomes a more natural cognitive process.

Finally, we note that our results have implications for the suggestion of Burns et al. (in press) that Secretary Problems could be used as measures of cognitive ability. The lack of learning or practice effects is a highly desirable property in this context. In addition, by comparing the first and second halves of the problem sets, we found a medium-to-high test-retest reliability, and the preservation of individual rankings. While much work remains to be done, it may be that this intuitive problem-solving task provides a useful window onto some aspects of intelligence.

References

- Bearden, J. N., Murphy, R. O., & Rapoport, A. (2005). A multi-attribute extension of the secretary problem: Theory and experiments. *Journal of Mathematical Psychology*, *49*, 410–422.
- Bonner, S., Hastie, R., Sprinkle, G., & Young, S. (2000). A review of the effects of financial incentives on performance in laboratory tasks: Implications for management accounting. *Journal of Management Accounting Research*, *12*, 19–64.
- Bucklin, B., McGee, H., & Dickinson, A. (2003). The effects of individual monetary incentives with and without feedback. *Journal of Organizational Behaviour Management*, *23*, 65–94.
- Burns, N. R., Lee, M. D., & Vickers, D. (in press). Individual differences in problem solving and intelligence. *Journal of Problem Solving*.
- Camerer, C., & Hogarth, R. (1999). The effects of financial incentives in experiments: A review of capital-labour-production framework. *Journal of Risk and Uncertainty*, *19*, 7–42.
- Einhorn, H., & Hogarth, R. (1981). Behavioural decision theory: Processes of judgement and choice. *Annual Review of Psychology*, *32*, 53–88.
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science* *4*(3), 282–296.
- Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *American Statistical Association Journal* *61*, 35–73.
- Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioural and Brain Sciences*, *24*(3), 383–451.
- Jacoby, J., Troutman, T., Mazursky, D., & Kuss, A. (1984). When feedback is ignored: Disutility of outcome feedback. *Journal of Applied Psychology*, *69*(3), 531–545.
- Lee, M. D. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*(3), 1–26.
- Levine, M., Leitenberg, H., & Richter, M. (1964). The blank trials law: The equivalence of positive reinforcement and nonreinforcement. *Psychological Review*, *71*, 94–104.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association* *90*(430), 773–795.
- MacGregor, J. N., & Ormerod, T. C. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics* *58*, 527–539.
- Navarro, D. J., & Lee, M. D. (2005). An application of minimum description length clustering to partitioning learning curves. *Proceedings of the 2005 IEEE International Symposium on Information Theory*, pp. 587–591. Piscataway, NJ: IEEE.
- Parco, J., Rapoport, A., & Stein, W. (2002). Effects of financial incentives on the breakdown of mutual trust. *Psychological Science*, *13*(3), 292–297.
- Rimm, D., Roesch, R., Perry, R., & Peebles, C. (1971). Effects of blank versus non-informative feedback and “right” and “wrong” on response repetition in paired-associate learning. *Journal of Experimental Psychology*, *88*(1), 26–30.
- Spence, J. (1972). Effects of blank versus non-informative feedback and right and wrong on response repetition in paired-associate learning. *Journal of Experimental Psychology*, *94*(2), 146–148.
- 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*(3), 221–236.
- Vickers, D., Butavicius, M. A., Lee, M. D., & Medvedev, A. (2001). Human performance on visually presented travelling salesperson problems. *Perception*, *32*(7), 871–886.
- Wofford, J., & Goodwin, V. (1990). Effects of feedback on cognitive processing and choice of decision style. *Journal of Applied Psychology*, *75*(6), 603–612.