Optimality and Adaptivity in Quota and Level Cutoff Strategies

Joel Huber and Fred M. Feinberg

Joel Huber, the Fuqua School of Business, Duke University, Durham, NC, 27708; Fred Feinberg, University of Michigan Business School, 701 Tappan Street, Ann Arbor, MI, 48109. Correspondence regarding this manuscript, including requests for data allowing theoretical model comparison in tabular format, should be addressed to the second author by e-mail to fred.feinberg@umich.edu, or by surface mail to the above address. The authors are listed reverse-alphabetically; both contributed equally to the research. The authors wish to thank Bob Meyer, Wes Hutchinson, Markus Christen and Steve Hoeffler for their suggestions on an earlier draft of this manuscript.

Optimality and Adaptivity in Quota and Level Cutoff Strategies

ABSTRACT

A cutoff strategy is one that reduces search cost by limiting the number of alternatives to be evaluated, possibly resulting in a less satisfactory alternative being chosen. We empirically examine the effectiveness of two types of cutoff strategies: a *Level* strategy, which examines all options surpassing a particular minimal value of the screening attribute, and a *Quota* strategy, which specifies the number of examined alternatives as ranked by the screening attribute. A behavioral simulation with explicit search and quality measures assesses the extent to which respondents approach optimality in the use of these strategies. Respondents are found to be more effective using the Quota strategy: they achieve greater profitability and require less time to enact the appropriate cutoff. While these differences between Level and Quota methods are significant in themselves, a more intriguing finding is that with either cutoff method respondents adapt correctly, albeit suboptimally, when the parametric environment of the decision problem is altered. We show that this correct adjustment occurs in response to changes in three environmental variables: the quality of the screening attribute, the cost of evaluation, and the value of the Default available if none of the alternatives generated are acceptable. Such 'bottom-up' adaptivity occurs despite a pronounced lack of conscious understanding of the comparative statics of the cutoff problem.

1 Introduction

As individuals and firms consider choices in the marketplace, cutoff opportunities abound. These can reflect prices they will not pay, suppliers they would not patronize, and brands they do not consider. Cutoffs offer the benefit of rendering search more manageable, but entail the penalty of possibly screening out a superior alternative. This fundamental trade-off between search difficulty and outcome quality has been recognized in the economics of information literature since Stigler's (1961) classic article, blossoming after that time into a major theme in economics research and finding varied application in theoretical and empirical models of choice (Tversky 1972, Weitzman 1979, Shugan 1980, Gensch 1987, Hauser and Wernerfelt 1990, Roberts and Lattin 1991).

The use of cutoffs as a simplification tool has also been recognized in descriptive models of choice, typically in the form of a conjunctive rule whereby an alternative is rejected if it contains any unacceptable attribute values (Wright 1975, Lussier and Olshavsky 1979, Payne 1982, Nedungadi 1990, Russo and Leclerc 1994). Extensions of this research have shown that such cutoffs depend on consumers' processing costs (Grether and Wilde 1984), on attribute utility (Klein and Bither 1987) or range (Payne, Bettman and Johnson 1988), and on the covariation among attributes (Bettman, Johnson, Luce and Payne 1993, Huber and Klein 1993).

These studies have demonstrated that humans adapt to cutoffs in ways that are, generally speaking, appropriate. However, it has been very difficult to assess the effectiveness (profitability) of cutoffs relative to any absolute standard, especially those involving (perhaps gradual) convergence toward a theoretically-optimal solution. Such questions are difficult to answer definitively, as invoking a cutoff involves a trade-off between search cost and the quality of the ultimate choice, two variables that are not easily represented on the same scale. That is, consumers can rate or rank the relative quality of various choices, and they can evaluate the arduousness of various choice strategies, but trading-off the quality of a choice against the effort or cost of a decision is problematic both for consumers and for theorists.

One approach to the trade-off of search effort and choice quality involves invoking a linear cost structure for search components and then examining whether consumers act as if they are minimizing these costs. Johnson and Payne (1985), and Payne, Bettman and Johnson (1988) show that such a choice strategy selection is broadly consistent with an effort-accuracy model of search. The approach taken here is to construct a managerially-oriented task having a motivational structure similar to the naturalistic setting of a choice-based cutoff, with the provision that both costs and benefits are rendered explicit and comparable. This behavioral simulation allows assessment of the degree to which subjects are able to approach the known optimal policy, as well as the degree to which their adaptations to changes in the choice environment are in themselves consistent with optimal comparative statics. While this exploration cannot fully resolve the issue of whether naturalistic cutoffs are set optimally, it does address the question of whether humans are capable of recognizing and adapting toward optimal behavior, and helps suggest directions in which such dynamic decision-making can be improved.

The present study examines dynamic cutoff adjustment and thereby provides an empirical test of the normative model of cutoff selection developed by Feinberg and Huber (1996). Their model mirrors the cutoff process described above: a restrictive screen lowers evaluation costs, but may result in a lower-quality selection; by contrast, a less restrictive screen is more likely to uncover a better final alternative, but involves higher search costs. To render the underlying mathematical theory more analytically tractable, Feinberg and Huber invoke two structural assumptions. First, they posit the existence of an essentially costless indicator to screen alternatives and, second, they assume a constant cost must be paid for each alternative passing the screen. It is useful to consider the contexts in which these assumptions reflect actual decision-making. Examples of the first assumption, that of nearly costless indicators, are quite common: recommendations of friends, consumer magazine ratings, brand names, locations, prices and even preconceived notions can provide inexpensive, if fallible, ways to limit search.[1] A fixed cost assessed for each alternative passing the screen (with no explicit learning or scale effects) is reasonable, for example, in situations where one must pay for each evaluation (as in a credit check on potential applicants) or where each alternative passing the screen must undergo a more-or-less identical evaluation process (as in a preplanned college tour). This second assumption, that of a constant search cost, is not intrinsic to the mathematical theory, which can accommodate a general cost function. A convex function could reflect, for example, increasing marginal cost of time or cognitive resources, while a concave function might reflect learning curve

[1] More precisely, it is not the overall costlessness of the indicator that is important, but rather a negligible *marginal* cost for each alternative *rated*; a fixed cost, for example, to purchase a consumer magazine in order to profit from its suggestions is consistent with the theory proposed.

effects. In the experiments to follow, a linear cost structure was employed not so much for its conformity with the mathematical development in Feinberg and Huber (1996), but for its ability to be effectively communicated to experimental respondents.

Feinberg and Huber distinguish between two distinct cutoff strategies. A **Level** strategy specifies a minimally-acceptable *level* of the screening attribute, while a **Quota** strategy specifies the minimum *number* of alternatives (as ranked by the attribute) that will be evaluated. For example, in selecting automobile dealers to visit, a Level strategy might constrain visits to those within ten miles of home, while a Quota strategy might limit search to the three closest. The Quota and Level strategies are similar in that they each seek to balance the degree of effort involved in additional search against the eventually-diminishing marginal utility of the alternatives evaluated. However, they differ in the domain over which each allows decision-makers to have control. The Level strategy assures one of a minimal expected quality, but results in a variable evaluation cost, which depends on the number of alternatives surpassing the specified minimal rating. By contrast, the Quota strategy controls search cost (e.g., only three dealers will be visited), while the minimal quality rating of those evaluated is variable. Feinberg and Huber's analytic framework demonstrates that the Quota strategy provides higher expected profit, over a broad range of environmental conditions.

The present paper explores whether this same dominance of Quota over Level extends to the intuitive responses of human decision-makers. The answer is by no means clear, as a variety of factors that are mathematically transparent can appear rather opaque to human decision-makers. For example, a Quota strategy often offers the advantage of integer responses, which may be easily recalled and understood. By contrast, the less granular nature of the Level strategy may aid decision-making by encouraging hypothesis testing and subsequent incremental adaptation.[2]

To summarize, this paper will address the following issues, in decreasing order of emphasis:

1. *Relative Performance*: Does the theoretically superior Quota strategy result in greater profitability than the Level strategy among fallible decision-makers? Are there differences in the two strategies in terms of time required for decision?

2. *Adaptability*: How do respondents respond to different 'environmental' conditions – correlational (i.e., accuracy of the screening indicator), financial (i.e., the unit cost of evaluation), and

[2] While Level cutoffs do technically require 'integer' responses in many cases, it is the near-continuity of the possible screening rules that constitutes its distinguishing feature.

Default value (i.e., if no choice results from the screening process) – in terms of conformity with optimal comparative statics?

3. *Understanding* : To what extent are people able to surmise the rules governing effective use of cutoffs? Once respondents discern, by 'muddling through' the choice process, something approaching an optimal solution, can they distill such 'bottom-up' experience with the cutoff process into codified, 'top-down' knowledge?

The next section provides a formal description of the Quota and Level choice processes and summarizes several testable propositions suggested by the analysis of Feinberg and Huber. Then follows a description of a computer-based behavioral simulation for which a normatively optimal solution exists and can be specified. Finally, the results of the simulation, in regard to the three questions above, are discussed.

2 Formal Development of The Quota and Level Models

Consider a decision-maker who must choose from among n alternatives and has some ratings or beliefs that imperfectly assess their relative quality. On the basis of these ratings, X, the decision-maker may specifically exclude some alternatives from further evaluation, reducing the temporal, cognitive and financial cost involved with such evaluation. All those alternatives passing this initial X-based screen pass into the second stage, where the real utility or quality of the alternatives, Y, is revealed (without error), and a cost for each evaluated alternative is assessed. There is a known, stable correlation, ρ, between the prior ratings (X) and the revealed utilities (Y). If the decision-maker chooses to invoke a Quota cutoff, the n_c best-rated alternatives pass to the next stage to be evaluated, while with a Level cutoff, all alternatives rating at least as well as x_c are evaluated; in either case, a cost, c, is incurred for each alternative evaluated. Should no alternative pass, the process terminates, and the decision-maker receives the value of the average alternative in the set.[3]

Feinberg and Huber's (1996) solution to this cutoff problem begins by arranging the ratings values $\{X_i\}$ in ascending order, $X_{1:n}, X_{2:n}, ..., X_{n:n}$; corresponding utility values are given by $Y_{[1:n]}, Y_{[2:n]}, ..., Y_{[n:n]}$. Such Y-values are ordinarily called concomitants or induced order statistics (David 1981, Nagaraja

[3]This assumption of receiving the average is made initially for clarity of mathematical development, and is readily generalized to other types of 'default' payoffs, discussed below.

and David 1993), and are related to ratings by

$$Y_{[i:n]} = \rho X_{i:n} + \sqrt{1-\rho^2}\,\epsilon_{[i]}$$

The $\left\{\epsilon_{[i]}\right\}$ are independent, identically distributed and (for tractability) are assumed normal, so that without loss of generality (X, Y) can be taken to be $N(0,0,1,1,\rho)$.

For the Quota cutoff, n_c, one must maximize the following expected net gain function, the expected value of the greatest of the utilities less the total (known) evaluation cost:

$$E\left[max\left\{Y_{[n:n]}, Y_{[n-1:n]}, ..., Y_{[n-n_c+1:n]}\right\}\right] - cn_c$$

By contrast, for the Level cutoff, the exact number of alternatives evaluated is unknown, though its expected value is given by $n\left[1-\Phi\left(x_c\right)\right]$ (with Φ the standard normal cdf). Thus, one seeks to choose x_c to maximize a corresponding net gain function, representing the maximal utility of those alternatives whose ratings surpass x_c, less total expected evaluation cost:

$$E\left[max\left\{Y_i \mid i \ni X_i > X_c\right\}\right] - n\left[1-\Phi\left(x_c\right)\right]$$

Choosing the appropriate values of n_c and x_c for a given set of parametric values (ρ, n, c) is not a straightforward exercise. For the Quota, it can be shown that it is optimal to evaluate k alternatives, where k maximizes:

$$\max_k \binom{n}{k} \int_{y=-\infty}^{y=\infty} \int_{x=-\infty}^{x=\infty} y \frac{\partial}{\partial y} \left[\int_{u=x}^{\infty} \Phi\left(\frac{y-\rho u}{\sqrt{1-\rho^2}}\right) \phi\left(u\right) du\right]^k d\Phi^{n-k}\left(x\right) dy - kc \qquad (1)$$

Thus, the problem is reduced to comparing the values of at most n numerical integrations. By contrast, there is a closed-form solution of sorts for the optimal Level cutoff; x_c must be chosen to satisfy:

$$c = -\int_{u=-\infty}^{u=\infty} u \frac{\partial}{\partial u} \left[\Phi\left(\frac{\rho x_c - u}{\sqrt{1-\rho^2}}\right)\left(\Phi\left(x_c\right) + \int_{v=-\infty}^{v=u} \phi\left(v\right) \Phi\left(\frac{\rho v - x_c}{\sqrt{1-\rho^2}}\right) dv\right)^{n-1}\right] du \qquad (2)$$

Equations [1] and [2], while complex, allow one to draw a number of conclusions about optimal cutoff strategies. The most important conclusion from the perspective of the present paper is the greater profitability of the optimal Quota over the optimal Level strategy. This dominance does not occur universally; where cost is sufficiently high, one is better off with a very restrictive Level cutoff than any Quota. However, Feinberg and Huber provide evidence that when the decision-maker expects to examine a non-trivial number (typically two or more) of alternatives, the Quota method is superior; all but one of the experimental scenarios presented below exhibit optima for which this condition is satisfied. The theoretical model, however, makes no predictions regarding which method should allow superior performance – economically, time-wise or in terms of environmental adaptability – in actual practice.

Given a typical situation where multiple evaluations are to be expected, it is also possible to examine the comparative statics of the cutoff process, in effect determining how the optimal Level or Quota depends on the parameters of the choice environment. We define a more restrictive cutoff as one which reduces the number passing the cutoff in the Quota case or the expected number evaluated in the Level case. It can then be shown that, generally speaking, more restrictive cutoffs are justified when (1) the cost, c, of each evaluation increases, (2) the number of alternatives, n, available from which to screen decreases, and (3) the accuracy, measured by ρ, of the indicator increases.

In addition, Feinberg and Huber present generalizations of equations [1] and [2], modifying them to include a Default (or fallback) option, which will be essential in constructing the forthcoming behavioral simulation. Such an option represents what one would fall back upon if, for example, one visited a set of prospective apartments but, in the end, decided not to move. They demonstrate as well that the better the quality of the Default, the more restrictive the optimal cutoff should be. The Default offers the further advantage of serving as a proxy for the value of continued search. While the model with the Default set to the average utility of the choice set is strictly limited to the analysis of irreversible cutoffs, a specific Default value can be taken to represent the decision-maker's expectation with regard to discarding the original cutoff and engaging in additional search. To take but one example, should a candidate not be accepted into any of the 'first-tier' graduate programs applied to, the cutoff model need not require that he or she settle for the best of the remaining lot; rather, the Default represents the expected benefit (taking into account waiting time, another round of applications and the like) of waiting at least one additional year in order to repeat the process in the hope of a superior outcome. In this sense, the Default operates as a proxy for the cost of additional 'rounds' in the decision process.

Equations [1] and [2], modified as in Feinberg and Huber (1996) to include arbitrary Default values, enable the determination of optimal cutoff policies in the following behavioral simulation, permitting an assessment of the degree to which fallible subjects are able to initially intuit and ultimately approach an optimal solution, as well as the degree to which they are able to adapt their cutoff responses to changes in the choice environment. Specifically, we experimentally test the appropriateness of subjects' responses to changes in the accuracy (ρ) of the indicator, the unit evaluation cost (c) and the quality of the Default (D), as well as the more general question of whether the superiority of Quota over Level strategies implied by the formal model carries over to human decision-making.

3 Comparison of the Quota and Level Cutoff Strategies

Addressing the set of questions posed above requires a behavioral simulation with the following four properties. First, comparing the two cutoff strategies requires a realistic and involving problem that is equally amenable to a Quota or Level approach. Second, to assess not only the relative profitability of the two strategies, but the degree of absolute performance on the part of respondents, an optimal solution to the problem must exist and be estimable. Third, the task needs to employ multiple successive trials to assess any trends toward or away from the optimal policy, as well as to allow subjects to acclimate to between-trial variation. Finally, the behavioral exercise has to be flexible enough to allow credible manipulation of parameters such as processing cost or indicator accuracy.

The task developed for the simulation achieves these diverse goals by placing respondents in the role of purchasing agent who searches for the lowest bidder on a contract. A handling cost reflects the 'unit evaluation' or processing cost attached to each invited bid. To convey the role of the imperfect indicator, subjects can make use of a regression model to predict these bids for each contractor. Each subject operates in one of two separate experimental conditions. In the Level condition, respondents designate a maximum acceptable value of the predicted bid, while in the Quota condition, they specify the number of contractors (as ordered by the regression model's predictions) invited to provide bids.

Respondents consisted of 148 MBA and undergraduate Management concentrators. To involve them in the task and provide motivation to perform well, they were placed in groups of 10 - 20, and the three top scorers in each group were recognized with monetary rewards of \$50, \$35 and

\$25, respectively. Business students are appropriate for the task as they are conversant with the rudiments of forecasting (for the correlational manipulation described below), and are motivated by the managerial setting of the task.

The problem, implemented entirely by personal computer, was designed to be one to which the students might readily relate. Several factors guided the selection of a managerial / institutional decision-making setting. For example, it was necessary to choose a task that did not allow any respondents to capitalize on prior knowledge or differential expertise. Further, because 'cost' is not only one of the primitives of the cutoff model but one of the parameters manipulated, it was crucial to avoid any type of endogeneity, where some respondents would perceive search costs as significantly greater than others. Finally, a specific and unambiguous payoff had to be given after each trial, a type of ready feedback more common in managerial contexts than in consumer-related ones.

The task begins by placing the respondent in the role of Pat Swift, an MBA of indeterminate gender, who is trying to improve operations at a state purchasing department. Before Swift arrived, the Department of Transportation's own crews completed the final finishing of standardized road segments for \$20,000 each. Swift initiates a cost-cutting program, whereby 25 independent contractors bid on available jobs, resulting in an average low (winning) bid of \$18,000, thus saving \$2000 on each job. After Swift receives an award for cost-cutting, the State Employees Union counters by highlighting the \$100 cost for each bid processed. They argue that Swift's "cost saving measures" in fact lose the State \$500 per contract ($\$100 \cdot 25 - \$2000$). In response, Swift develops a regression model which predicts bids on the basis of each contractor's distance from the site and an index of construction activity. This regression ($r = .5$) allows Swift to establish a policy that limits the number (and thus processing cost) of contractors submitting bids. Under the Level cutoff strategy, the State sends out requests for bids to contractors whose predicted bids fall below a particular dollar amount (e.g., $x_c = \$19,500$). Under the Quota strategy, the decision rule specifies the number of contractors (e.g., $n_c = 3$), with bids solicited from contractors with the best predicted bids. Under either strategy, the lowest price among contractors submitting bids determines the cost of the job, and the difference between this and the State's \$20,000 cost, less the \$100 per-bid processing cost, represents the profit resulting from the bidding process.

Once familiar with these particulars, respondents assume the role of Pat Swift and are asked to determine cutoff points using either the Level or the Quota methods, as assigned. They are provided enough information so that in principle they could determine the optimal policy, although they are hardly expected to do so, given the complexities involved. Instead, respondents make 'backcasting' trials, examining how well they would have done had they used a particular cutoff in five past contracts. The backcasting trials serve several purposes, chief amongst them to allow the respondents to obtain a 'feel' for the way in which the task operates and the effect of various choices on profitability. Moreover, because they encounter different instantiations of the joint distribution between ratings and utility, subjects are exposed to the intrinsic variability from choice to choice, something which they might only have guessed at otherwise. Finally, because it will be important to compare performance across a number of conditions (described below), the backcasting trials enabled respondents to get up to speed quickly, so that any apparent effects not be attributable to artifacts of greater progressive experience with the mechanical aspects of the choice task.

After setting either a Level or Quota cutoff, the respondent receives feedback similar to that shown in Table 1. In this example, the bids are provided for the 6 contractors passing the chosen cutoff (i.e., $Quota = 6$ or $Level = \$19,110$). The fifth contractor is awarded the job, based on an actual bid of $18,425, producing a revenue of $1,575, relative to the $20,000 benchmark. Subtracting a cost of $600 for processing six bids results in a net profit of $975. Following the practice trials, the exercise begins in earnest with 10 actual contracts. For each, the respondent specifies a Quota or Level cutoff, and is given the predicted and actual bids for those passing the screen and the profit for the trial.

Table 1

Example of Feedback Provided Given
A Quota of 6 or a Level of $19,000

Number	Predicted	Actual Bid
1	$ 18,150	$ 18,500
2	$ 18,250	$ 19,757
3	$ 18,910	$ 19,145
4	$ 18,950	$ 19,685
5	$ 19,010	$ 18,425 ← lowest
6	$ 19,110	$ 20,665
Minimum Bid =		$ 18,425
Cost for Processing 6 Bids =		$ 600
Total Project Cost =		$ 19,025
Gain over $20,000 =		$ 975

After completing 10 rounds of the 'base' condition, respondents are told that certain aspects of the decision environment have changed, and that they will make additional cutoff decisions under these new conditions. Respondents were split into two groups at this stage, one which experiences a change related to the accuracy of the indicator (that is, involving ρ) and one which experiences two financial changes, related to evaluation cost, c, and Default value, D. Splitting the sample reduces the number of parametric manipulations to which respondents are subjected, simplifies the task considerably, and lessens the time typically required to complete the entire exercise.

The first group of respondents ($n = 95$), are subject to the *Improve Information* condition. Members of this group are informed that Pat Swift is now able to enhance the regression model by including an index of each company's past bids, increasing its predictive accuracy (fit) from $r = .5$ to $r = .8$ and reducing the (conditional) standard error of the estimates from \$886 to \$600 (that is, to $\$1000 \cdot \left[1 - (.8)^2\right]^{1/2}$).

The second group ($n = 53$) are subject to two more rounds of 10 cutoff decisions, each involving a change in financial parameters. In the *Decrease Cost* condition, they are also told that the evaluation process has been improved, because Pat Swift had been able to lower the evaluation cost from its previous value of \$100 to \$30 for each contract allowed past the screen. After 10 rounds with this new, reduced evaluation cost, they then experience the *Improve Default* condition. In this condition they are informed that the Department of Transportation's own crews have become more competitive, reducing the State's internal road finishing cost from its previous value of \$20,000 to \$18,400, a reduction in 'Default' of \$1600. Table 2 provides a summary of the four experimental conditions, optimal policies, expected number of evaluations performed, and payoffs at optimum.

Table 2: Summary of Experimental Conditions, Optima and Profits

	Corr. ρ	Cost c	Default D	Optimal Level	Exp'd Number	Exp'd Profit	Optimal Quota	Exp'd Profit
Base Condition:	0.5	\$100	\$20,000	\$19,000	4.0	\$1,194	3	\$1,252
Improve Information:	0.8	\$100	\$20,000	\$18,800	2.9	\$1,484	2	\$1,583
Decrease Cost:	0.5	\$30	\$20,000	\$19,500	7.7	\$1,584	7	\$1,597
Improve Default:	0.5	\$30	\$18,400	\$19,000	3.7	\$1,763	4	\$1,761

It must be stressed that respondents are never told whether cutoffs should become more or less restrictive with these 'improved' environments, but are merely informed that 10 more rounds will be run subject to them. It is, in fact, far from obvious whether the internal choice rules employed by respondents should change at all relative to the information they are given, and it is to be expected that some degree of inertia would be observed as they acclimate to altered environmental conditions.

For each contract, the actual bids offered were generated by randomly selecting a number of standard normal deviates from a master list of 100 possibilities. This randomization is important in that it produces a broad sampling of possible feedback streams, thus limiting the possibility that a particular feedback pattern, such as a declining profit trend, could account for any specific set of results. Figures 1 and 2 depict the expected profit associated with various cutoff rules in the four environmental conditions for the Level and Quota conditions, respectively. These curves are revealing, in that they illustrate the motivation for respondents to adjust their cutoffs in the various experimental conditions. We begin with an examination of the profit function for the Quota condition, followed by that of the Level and a comparison between the two.

The Quota profit curve for the Base Condition is shown by the bottommost of the curves in Figure 1, where a Quota of either 3 or 4 contractors results in a profit of about $1250. Deviations in either direction serve to reduce the payoff, quite severely for more restrictive cutoffs, so that a Quota of 1 results in a loss in expected payoff loss of nearly $300. An error in the opposite direction is less costly on the basis of per-unit deviation – a Quota of 10 is needed to effect a similar $300 loss. Those in the Improve Information condition then experience a gain in payoff for any cutoff invoked, presumably leading to the more restrictive cutoff of 2 bids, at optimum. By contrast, those in the Decrease Cost condition also experience a gain in payoff, but one leading to a less restrictive optimal Quota cutoff of 8. The contrast between the payoff functions in these two conditions provides a stringent test of adaptability. If subjects adapt appropriately to both these changes in condition, it cannot mean that they were following a simplistic, across-the-board trend towards more or less restrictive cutoffs, as they must increase restrictiveness in one condition and reduce it in the other. Finally, those in the Improve Default condition encounter a relatively flat payoff function, for which response to nearly-optimal cutoffs does not vary to a great degree throughout the relevant range. This condition thus provides a test of whether respondents can adapt in the appropriate direction despite a very poor signal-to-noise ratio.

Figure 2 depicts comparable payoff curves for the Level strategy. While they may appear somewhat different in terms of their shape and skewness, there are nevertheless important similarities between the Quota and Level payoffs functions. First, note that the direction of appropriate adjustment in the optimal strategies is the same as for the Quota case. Compared with the Base Condition, the Improve Information condition motivates a less restrictive strategy (optimal Level changing from $19,010 to $18,800), while the Decrease Cost condition prescribes a more restrictive strategy ($19,500), as does the Improve Default condition ($19,000). The expected payoffs for these optimal levels are listed in Table 2. Compared with the Level, the Quota strategy generates 5 - 7% greater payoffs, with the exception of the Improve Default condition, where the payoffs are essentially identical. While it may have seemed prudent to test the Level against the Quota only in conditions where the payoff functions have similar optima, the relative lack of profitability of the Level strategy renders this infeasible. To account for these differences in potential payoffs, analyses will be performed on 'raw' payoffs as well as on payoffs taken as a proportion of achievable optima, thereby adjusting for the superior theoretical performance of the Quota strategy.

One difference apparent from Figures 1 and 2 is that the Quota payoff curves exhibit a marked rightward skew, while that of the Level response is seemingly more symmetrical; however, this asymmetry is largely due to the different scalings of the horizontal axes. It is possible to render the Level and Quota payoff functions directly comparable by a rescaling in terms of the expected number of bids evaluated for any (dollar-based) Level cutoff. These revised Level response curves appear in Figure 3, superimposed with those of the Quota for the four experimental conditions. Note that for the Base and the Improve Information conditions, the optimal achievable payoff is not only noticeably greater for the Quota than for the Level, but prescribes that fewer alternatives be evaluated; by contrast, the payoff curves for the Decrease Cost condition are quite close to one another, while for the Improve Default condition they are nearly identical. Figure 3 is also telling in that it suggests that it may be possible to employ a Quota rule as a reasonable approximation to a Level one; that is, subjects may obtain approximately the same benefit from a Level cutoff for which the expected number of evaluations is near that of the optimal Quota. This insight is an important one, and will be discussed later in more detail.

In summary, the four payoff curves for each of the cutoff strategies not only admit of distinct optima (both in optimal cutoff and resulting profit), but display rather different sensitivity to deviations from near-optimal behavior. Additionally, while two of the conditions would be expected to afford significantly better performance for those respondents using a Quota cutoff, the remaining conditions should allow nearly equal performance. Finally, despite these differences, the directional prescriptions between experimental differences, in terms of optimal cutoff adjustment, are the same for both cutoff methods.

It is important to understand just how difficult it is in any of the conditions for respondents to settle upon the optimal strategy. Given the parameters of the simulation, determining optimal cutoffs is very difficult, even through extensive trial and error in a stable environment. In particular, given the large standard deviation ($\$500 - \$800$ around each contractor's bid), falling within \$100 of the optimum over the course of the ten trials would be quite a feat. The focal research issues, then, are which strategy, Quota or Level, affords greater profitability, speed and adaptability; and, secondarily, the extent to which decision-makers can approximate optimal cutoff usage in a very noisy (if reasonable) structured choice environment.

4 Results

The presentation of results follows the three research questions presented at the outset. First, we examine the issue of relative effectiveness of the Quota compared with the Level strategy, measured as a function of profitability and time taken. We next examine within and across strategies for evidence of respondent adaptability relative to the four environmental conditions. Finally, we assess the degree to which respondents are able to articulate the correct strategies as a result of their experiences in the simulation.

4.1 Relative Performance of Level vs. Quota Strategies

Subjects' performance using the Quota and Level strategies are first compared with respect to effectiveness, measured by profitability, and effort, measured by time taken. The results, shown in Table 3, are strong and unequivocal. In all but the Improve Default condition, the Quota cutoff is significantly more profitable than the corresponding Level ($p < .0001$). These results remain significant ($p < .0001$) even adjusting for the theoretically superior expected profit at optimum of the Quota strategy. For the Improve Default condition, adjusting in this manner for the greater

Quota optimum allows a marginally stronger test of its superiority, although the difference between the strategies is still not significant ($p \approx .14$ to $p \approx .12$)

Table 3
Profit and Time Taken for Level and Quota Strategies

Evaluation Cost	$c = \$100$				$c = \$30$			
Default Value	$D = \$20,000$						$D = \$18,400$	
Ratings Correlation	$\rho = .5$		$\rho = .8$		$\rho = .5$			
Strategy	Quota	Level	Quota	Level	Quota	Level	Quota	Level
Actual Mean Profit	\$1054*	\$927	\$1453*	\$1247	\$1447*	\$1272	\$1651^{ns}	\$1624
(stddev)	(38)	(52)	(35)	(47)	(73)	(73)	(90)	(85)
Optimal Profit	\$1252	\$1194	\$1583	\$1484	\$1597	\$1584	\$1761	\$1763
% Optimal Achieved	84%*	78%	92%*	84%	91%*	80%	95%^{ns}	92%
Time taken (minutes)	6.2*	7.5	2.8*	4.2	2.5*	5.5	1.4*	3.2
(stddev)	(1.0)	(1.5)	(0.7)	(0.7)	(1.2)	(2.0)	(0.7)	(1.1)
n	76	72	47	48	29	24	29	24

*Difference between Level and Quota significant at $p < .0001$

The Quota strategy also required less time in all four conditions. Decision time for the Level was, respectively, 20%, 50%, 115% and 130% greater than that for the corresponding Quota (all $p < .0001$). In the Base condition, where a non-trivial proportion of subjects' time was likely spent acclimating, despite the backcasting trials, to the decision environment, the relatively small difference is in its way understandable, even though it is still highly significant. Further, it is illuminating that, as the experiment wore on, the degree of disparity in time taken grows dramatically.

In short, based on profitability and speed of use in all parametric conditions tested, compared with a corresponding Level strategy, Quota dominates. This dominance in profitability exceeds considerably that predicted by theory alone, while there was no theoretical basis on which to anticipate the considerable differential in time necessary for decision.

4.2 Adaptability and Sampling Behavior

Insight into the process employed by respondents in the Quota and Level strategy groups can be gained from an examination of their average cutoffs, listed in Table 4. The focus here is on respondents' average responses compared with that theoretically-achievable, as well as how appropriately each adapts to changes in the environment. As the data will show, respondents exhibit a tendency to examine more alternatives than is theoretically optimal, although they also adapt reasonably well, considered both in terms of average response and within-respondent changes.

Table 4 **Adaptivity for Level and Quota Strategies**								
Evalution Cost	$c = \$100$				$c = \$30$			
Default Value	$D = \$20,000$						$D = \$18,400$	
Ratings Correlation	$\rho = .5$		$\rho = .8$		$\rho = .5$			
Condition	Base		Improve Accuracy		Decrease Cost		Improve Default	
Strategy	Quota	Level	Quota	Level	Quota	Level	Quota	Level
Average Cutoff Used	6.7	\$19220	4.8*	\$19030*	7.1*	\$19720*	6.8	\$18820*
Optimal Cutoff	3	\$19000	2	\$18800	7	\$19500	4	\$19000
n	76	72	47	48	29	24	29	24
*Difference from previous condition in Average Cutoff Used in correct direction and significant at $p < .05$								

Oversampling. Examination of the average Quota cutoffs invoked by respondents reveals a general tendency to be less restrictive than is optimal. Thus, in the Base Condition ($\rho = .5, c = \$100, D = \$20,000$) an average of 6.7 contracts were evaluated, while the optimal Quota was 3. Similarly, the optimal Level cutoff is $\$19,000$ while respondents, on average, solicited contractors with expected bids up to $\$19,220$. In seven out of eight cases the average cutoffs invoked by respondents are less restrictive than theory dictates. While this deviation reflects a consistent bias, it is hardly an unreasonable one. Allowing more contractors to bid affords a hedge of sorts in what is clearly a rather uncertain environment. Undersampling allows the possibility of only one alternative, or even none at all, passing into the evaluation stage, and respondents may feel justified in guarding against such an outcome, which leaves a good deal unknown about whether they might have fared better had they invoked a less restrictive cutoff. Thus, despite the transparency of the linear cost mechanism and the feedback received after each decision, it is possible that a certain degree of risk-aversion, with its implied asymmetric loss function, may have come into play. While the model of Feinberg and Huber (1996) can account analytically for such asymmetries at the individual level, the implementation of the model in the simulation does not accord any value to information gained by engaging in excessive (i.e., suboptimal) evaluation, and in fact care was to insure that respondents were explicitly aware of the linear-cost nature of the payoffs received.

The one case for which respondents' average cutoff behavior was more restrictive than the optimum was for the Improve Default condition ($\rho = .8, c = \$30, D = \$18,400$) of the Level cutoff. In that case, the State's automatic (Default) bid of $\$18,400$ provides a strong incentive to resist evaluating expected bids above that value. This intuitive strategy, viewed in light of the rather flat nature of the response curve in that condition (Figure 2), may have led respondents to be more restrictive than is optimal under the explicit expected profit framework employed.

Adaptiveness in Averaged Responses. Table 4 lists cutoffs, averaged over rounds and respondents, and reveals patterns of change, for both Level and Quota strategies, that trend consistently in the direction of the optimal adjustment. For example, those in the Improve Information condition became more restrictive in their Quota (6.7 to 4.8; $p < .01$) and their Level cutoffs (\$19,220 to \$19,030; $p < .01$). By contrast, those in the Decrease Cost condition became less restrictive in their Quota (6.7 to 7.1; $p < .05$) and their Level (\$19,220 to \$19,720; $p < .01$) cutoffs; when placed in the Improve Default condition, these respondents then restricted their Quota (7.1 to 6.8; *n.s.*) and their Level (\$19,030 to \$18,820; $p < .01$) cutoffs appropriately.

The pattern of results evident in Table 4 indicates that, on average, respondents adapted to changes in the choice environment in a manner consistent with directional changes in the theoretical optimum. It is therefore natural to ask, given this confluence between prescribed and actual choice behavior, whether respondents would be able to codify the apparent lessons regarding such environmental changes.

Ability to Articulate the Appropriate Strategy. Given that respondents trend, between conditions, in the direction of the theoretically-optimal strategy, one might ask to what degree adaptivity derives from what might be termed an understanding of the problem at hand. Respondents answered three questions assessing their understanding of the correct comparative statics for the Pat Swift problem. The first question concerned how a cutoff should change if processing cost were decreased (i.e., a smaller value of c), the second dealt with a worsening of the fit of the predictive model (i.e., a smaller value of ρ), and the third with a decrease in the State's internal road finishing cost (i.e., a lower/better value of D). Four answers were available: to raise, to lower, to keep the cutoff constant, or "do not know". The correct responses (Table 2) are, for both Level and Quota, in the first and second cases (c, ρ) to raise the cutoff used, and in the third (D) to lower.[4] Table 5 tallies the proportion of respondents answering correctly of those offering an answer; questions were asked regarding changes in all three relevant parameters, regardless of whether the respondents in a particular conditional assignment

[4]Note that, due to the bidding context of the problem, to 'raise' a Level cutoff is to invite *more* evaluation. Further, a decrease in the State's cost is a *increase* in the quality of the Default. To avoid confusion related to these scalings, the questions asked of respondents were phrased in terms of the bidding problem with which they were familiar, not a generalized cutoff setting.

(correlational or financial) were in fact exposed to such changes. Focusing momentarily on responses to parametric changes that the respondents actually experienced, the evidence is compelling that those in the Quota condition understood the correlation-based comparative statics better than their Level counterparts, those in the Level condition understood the Default-based comparative statics better, while neither understood the comparative statics regarding cost; these observations are based not only on significant differences between Level and Quota, but on whether the proportion of respondents in any cell was significantly greater than $\frac{1}{2}$.

Table 5

Evidence of Top-Down Understanding of Level and Quota Strategies

Percentage Specifying Correct Response to:	Correct Response Q or L	'Correlational' Condition Quota	'Correlational' Condition Level	'Financial' Condition Quota	'Financial' Condition Level
Lower processing cost (c)	Raise	54%	36%	48%	59%
Poorer regression indicator (ρ)	Raise	85%**†	29%	64%	39%
Decrease in State's cost (D)	Lower	31%	62%	31%**	73%†

* Difference between Level and Quota sig. at $p < .01$; †Proportion $> \frac{1}{2}$ at $p < .01$

The lack of understanding of the effect of lowering evaluation cost is rather striking. Elementary economic reasoning would indicate that as something beneficial (i.e., evaluation) becomes less costly, one should 'consume' more of it; yet in none of the four relevant cells of Table 5 do significantly more than half of respondents seem to be able to reason this through. This is all the more remarkable considering that the 'cost' manipulation, for both Level and Quota, was the one involving the largest optimal deviation, that is, the strongest signal to alter cutoffs in the direction of additional evaluation. Differences between Quota and Level for the remaining manipulations were more pronounced. Respondents in the Quota condition seem to have possessed a far better sense of appropriate comparative statics relative to changes in predictive accuracy of the initial ratings (ρ), and a poorer one regarding shifts in the quality of the Default (D). The relative ability of the Level respondents to adapt to the Default manipulation is not difficult to understand. As discussed earlier, the excellent Default of \$18,400 leads to a natural reluctance to evaluate a bid higher than that amount. Indeed, this Default appears to cause a mild overreaction, in which the average cutoff of \$18,800 is more restrictive than the optimal value of \$19,000; this is the only such example of screening over-restrictiveness in any parametric condition with either cutoff method.

Perhaps most telling is a comparison of the responses of those exposed to the relevant manipulations to the responses of those who were not; for example, is there any real difference between the $31\% - 73\%$ proportional split for those who experienced the Default manipulation and the analogous $31\% - 62\%$ split for those that did not, in terms of their prescriptions regarding changes in Default? Most surprisingly, for all three parameters, a standard χ^2 test demonstrates that being exposed to a parametric manipulation confers no statistical advantage in correctly surmising the appropriate reaction to changing that parameter (in all cases, $\chi^2_1 < 1$; $p > .5$). In short, although subjects displayed a considerable degree of between-parametric-condition adaptivity in the actual task, they were no better at codifying the lessons of that adaptivity than those who merely engaged in educated guessing.

Rather than focus on tentative differences between the methods, an altogether different conclusion to be drawn from Table 5 is that respondents in both conditions show a remarkable lack of insight into the inner workings of the problem, despite having spent the better part of an hour working on it and being offered a reward for good performance. If an external heuristic is not available to frame the problem easily, they appear to be unable to intuit the appropriate strategy. This lack of top-down knowledge stands in contrast to the comparative success in 'muddling through' the problem in a bottom-up fashion. An implication of this intriguing observation is raised in the following section.

5 Discussion

Behavioral simulation of a precisely-structured cutoff process has allowed not only a study of the differences between how the Level and Quota strategies are used, but also a working understanding of what drives cutoff behavior in a general setting. We first consider the contrasts between the Quota and Level strategies, and then discuss the boundaries of the demonstrated effects. Finally, we address the features the two cutoff strategies share, and resulting implications for improved decision-making.

5.1 The Superiority of the Quota over the Level Strategy

Within the context of the simulation, the performance of subjects employing the Quota strategy was superior to those using a Level one. This superiority appears along two dimensions: respondents using a Quota cutoff achieved greater profitability (even as a proportion of the optimum) and took significantly less time than those using a Level cutoff. Moreover, it should be stressed that there are other advantages associated with Quota cutoffs that were not examined in this study. In particular, Quotas are less susceptible to fluctuations brought on by shifts in the number of alternatives available or to how the ratings and utility values are scaled. For example, an optimal Quota of, say, 5 will remain optimal or nearly so in the face of even fairly large changes (e.g., 50%) in the number of alternatives or the standard deviations of ratings or utility, whereas using a Level cutoff will necessitate considerable alteration in response to changes of the same magnitude.

If Quotas are possessed of such commendable qualities, why are they not more in evidence? Several answers suggest themselves. The simplest of these may be that Quota cutoffs are not typically offered up by nature or even by common market mechanisms; goods are not arranged so one can easily locate the four most comfortable lounge chairs or the ten best colleges, even if the relevant attributes used to assess these qualities are decided upon in advance. Instead, to use a Quota cutoff requires that one fashion a system (a committee, a sorting scheme, a computer program) that ranks alternatives and presents them according to preordained criteria (c.f., Russo et. al. 1986).

A second reason why Quotas may not appear to be used more frequently owes to the fact that cutoffs in many, if not most, real life situations are reversible; they can be altered simply and quickly if the number of alternatives revealed is not appropriate. By contrast, the two-stage cutoff model posits a type of irreversibility, whereby one must get by with the evaluation set that initially emerges. This irreversibility naturally occurs in situations where temporal or organizational constraints prohibit the alteration of screening criteria following inspection of some of the alternatives. Such constraints appeared in the simulation because Pat Swift attempted to construct a decision system that could be easily understood and managed by workers, as well as one that followed the protocols of bid submission. Similar constraints also occur in individual choice contexts such as appointment-based visits to homes one might buy, employment interviews, or college admissions, where one may not be easily able to alter the evaluation set once search has begun. The important point here is that the contrast between Quotas and Levels is primarily applicable only to the subset of those decision problems where the initial evaluation set, once determined, is largely unchangeable, although, as discussed previously, reversibility and 'staged' decision processes can be accommodated through proper specification of the Default.

The final, and most speculative, answer to the question of why Quotas are not more in evidence is that decision-makers may be implicitly adjusting Level cutoffs in an effort to approximate a Quota strategy. This type of process was suggested by a respondent in the Level condition who indicated that his strategy was to "find a Level that resulted in 4-5 bids". The parallels in the Quota and Level payoff functions, as depicted on a common axis in Figure 3, indicates that such a strategy may be, normatively speaking, quite acceptable. A common example of using a Level cutoff to simulate an appropriate Quota occurs in key word data searches, which typically display the number of matches before offering them in full citation. A keyword that is too broad or general may turn up several thousand matches but, when combined with another keyword, may yield only a few dozen; yet another addition may reduce the number to a mere handful, or none at all. In this case, one gains information regarding the appropriateness of a Level cutoff by how many matches one obtains. The important point here is that the focus is on the *number* passing the screen, rather than the specific attribute value (Level) used in screening. Widing and Talarzyk (1993) have developed a decision aid in which respondents can set cutoffs and subsequently alter them, finding that respondents continually adapted their cutoffs until they obtained a choice set of manageable size, from which they were then comfortable making a final choice. Thus, a fruitful agenda for future research concerns the advantages and disadvantages of a Quota over a Level *orientation*. Our finding that Quota dominate Level strategies in this admittedly highly-structured context suggests that Quotas may also prove superior in a broader one, although additional work need be done to validate any such speculation.

5.2 Evidence of Adaptability

If, instead of contrasting the differences between the Level and the Quota strategies, one focuses on their similarities, then an important lesson derived from this study is the adaptability in both conditions to changes in the decision environment . It is important to note that these adaptations were achieved in an environment with a rather poor signal-to-noise ratio. This bottom-up adaptability is similar to the search adaptation found by Bettman et al. (1993) and by Huber and Klein (1993). By contrast, our finding of consistent oversampling in a dynamic decision-making context parallels the somewhat less optimistic view reported by Hutchinson and Meyer (1994), and evident in a number of other dynamic decision-making tasks (e.g., Cripps and Meyer 1984, Cox and Oaxaca 1992).

It seems quite clear, however, that this bottom-up adaptivity failed to generate an appreciable degree of top-down comprehension of the process. Even at the conclusion of 25 or 35 trials, well-educated respondents had very little consistent understanding of how their cutoffs ought to change given shifts in the accuracy of the indicator, processing cost, or the value of the Default option. Rather, respondents appeared to adjust their cutoffs in a manner similar to what consumers may do with market-based process and outcome feedback.

Such an adaptive process could take the following form: an overly-restrictive screen leaves a consumer uncertain as to whether the final alternative was optimal, leading to a less restrictive screen, and a greater number evaluated on successive trials; conversely, evaluating too many alternatives leads to frustration with the needless evaluation of lower-ranked alternatives, driving the consumer to restrict the screen in future searches. So it is that such a process can generate appropriate adaptive behavior without a corresponding understanding of the task-related structures driving that behavior (e.g., as documented extensively by Payne, Bettman and Johnson 1993).

This lack of top-down knowledge suggests that similar results might derive from less statistically-sophisticated respondents. It was felt that knowledgeable students were necessary for an understanding of the correlation and cost conditions. However, given that these respondents appeared not to know how to use the given information in an *a priori* sense, a future study might examine the extent to which statistically-naive but motivated respondents are able to perform as well as more knowledgeable ones.

While respondents were generally adaptive between conditions, it is not clear that with enough time they would have settled on anything like the ostensibly optimal solution. Instead, they appear to have stabilized at a point where they are consistently engaging in excess evaluation. This oversampling may be due to a rational need to continue to monitor the environment or an overemphasis on the quality of the final choice, ignoring the costs of unnecessary evaluation, regardless of how plain these were made. Another possibility is that, despite the explicit linearity of the loss function encoded in the payoff structure and feedback, respondents were somehow more sensitive to insufficient evaluation than to its opposite.

In conclusion, this paper began with a discussion of cutoffs broadly defined as part of common search behavior. It then demonstrated evidence of optimality and adaptivity for a carefully delineated choice problem. While the generalizations stemming from these results are intriguing, ultimately any generalizations can gain strength only from replication across a range of similar models.

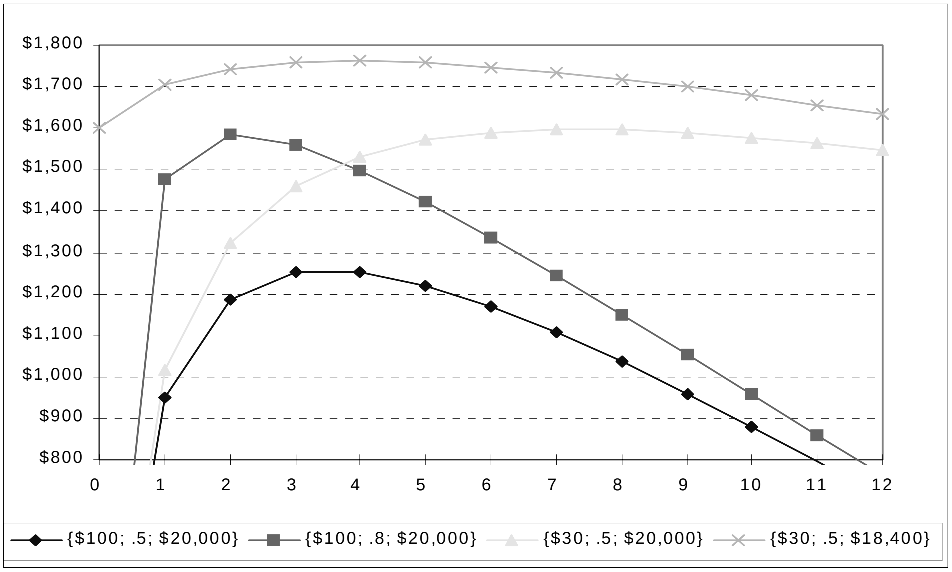

Figure 1: Quota Payoff Curves in Experimental $\{c, \rho, D\}$ Conditions

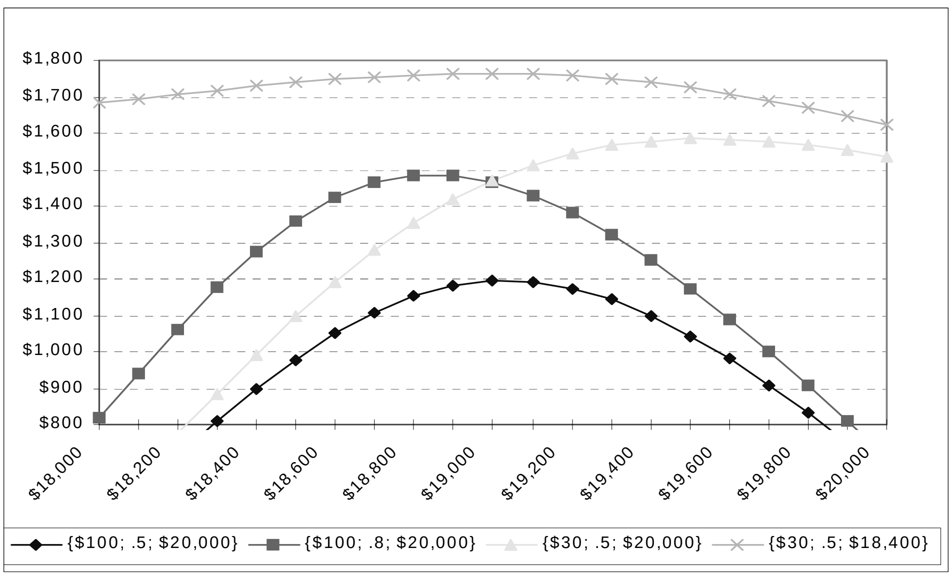

Figure 2: Level Payoff Curves in Experimental $\{c, \rho, D\}$ Conditions

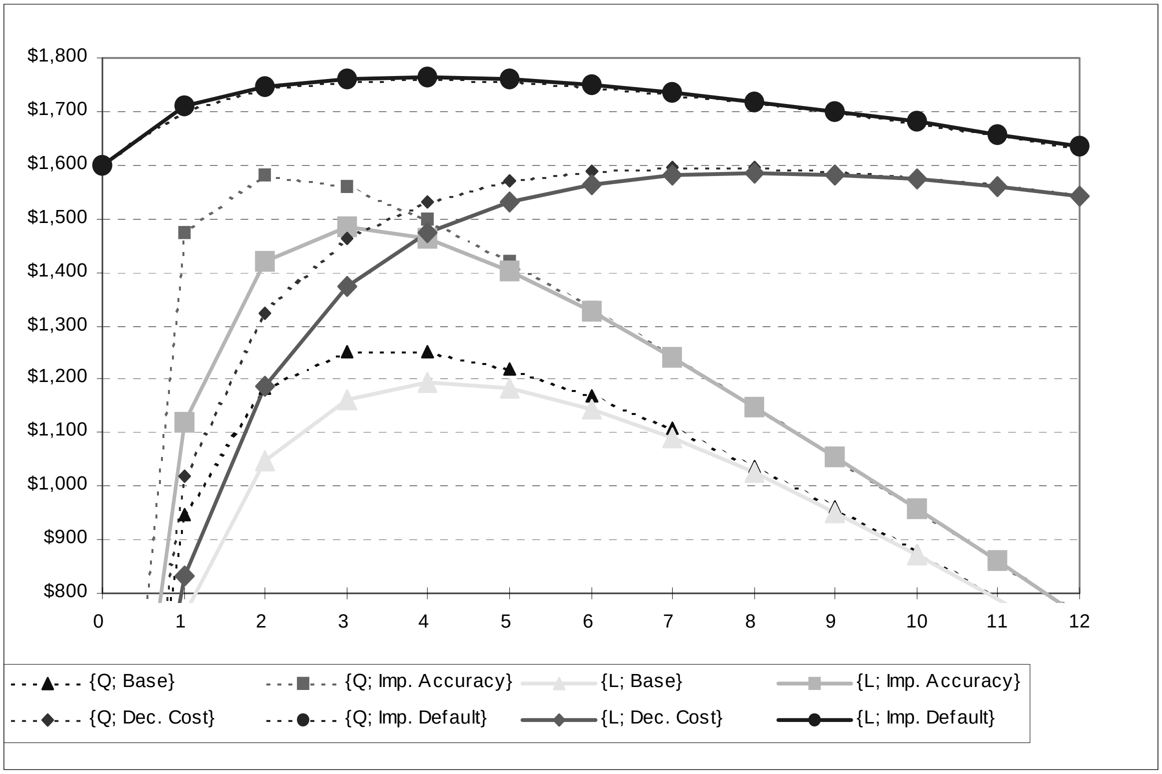

Figure 3: Payoff Curves for Quota and Level as a Function of Expected Number of Evaluations

References

[1] Bettman, James R., Eric J. Johnson, Mary Frances Luce and John Payne (1993) "Correlation, Conflict, and Choice," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(4), 1-21.

[2] Cox, James C. and Ronald L. Oaxaca (1992) "Direct Tests of the Reservation Wage Property," *Economic Journal*, 102 (November) 1423-1432.

[3] Cripps, John D. and Robert J. Meyer (1994) "Heuristics and Biases in the Timing and Replacement of Durable Products" *Journal of Consumer Research,* 21 (September) 304-318.

[4] David, H. A. (1981), *Order Statistics*, Second Edition, John Wiley, New York.

[5] Feinberg, Fred M. and Joel Huber (1996) "A Theory of Cutoff Formation Under Imperfect Information," *Management Science*, 2(1), 65-84.

[6] Gensch, David H. (1987) "A Two-Stage Disaggregate Attribute Choice Model," *Marketing Science*, 6(3), 223-239.

[7] Grether, David and Louis Wilde (1984) "An Analysis of Conjunctive Choice," *Journal of Consumer Research*, 10 (March) 373-385.

[8] Hauser, John, R. and B. Wernerfelt (1990) "An Evaluation Cost Model of Consideration Sets," *Journal of Consumer Research*, 16 (March), 393-408.

[9] Huber, Joel and Noreen M. Klein (1993) "Adapting Cutoffs to the Choice Environment: The Effects of Attribute Correlation and Reliability," *Journal of Consumer Research*, 18 (December) 346-357.

[10] Hutchinson, J. Wesley and Robert J. Meyer (1994) "Dynamic Decision-making: Optimal Policies and Actual Behavior in Sequential Choice Problems," *Marketing Letters*, 5(4), 369-382.

[11] Johnson, Eric J., and John W. Payne (1985), "Effort and Accuracy in Choice," *Management Science*, 31 (April), 395-414.

[12] Klein, Noreen M. and Steve W. Bither (1987), "An Investigation of Utility-Directed Cutoff Selection," *Journal of Consumer Research*, 14 (September) 240-256.

[13] Lussier, D. A. and R. W. Olshavsky (1979), "Task Complexity and Contingent Processing in Brand Choice," *Journal of Consumer Research*, 6, 154-165.

[14] Nagaraja, H. N. and H. A. David (1993), "Distribution of the Maximum of Concomitants of Selected Order Statistics," *The Annals of Statistics*, 22(1), 478-494.

[15] Nedungadi, Prakash (1990) "Recall and Consumer Consideration Sets Influencing Choice without Altering Brand Evaluations," *Journal of Consumer Research*, 17 (December) 263-276.

[16] Payne, John W. (1982), "Contingent Decision Behavior," *Psychological Bulletin*, 92(2), 382-402.

[17] Payne, John W., James R. Bettman and Eric J. Johnson (1988) "Adaptive Strategy Selecting in Decision-making," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14 (July), 534-552.

[18] Payne, John W., James R. Bettman and Eric J. Johnson (1993), *The Adaptive Decision Maker*, Cambridge University Press.

[19] Roberts, John H. and James M. Lattin (1991), "Development and Testing of a Model of Consideration Set Composition," *Journal of Marketing Research*, 28 (November), 429-440.

[20] Russo, J. Edward, Richard Staelin, Catherine A. Nolan, Gary J. Russell and Barbara L. Metcalf (1986), "Nutrition Information in the Supermarket," *Journal of Consumer Research*, 13(1), 48-70.

[21] Russo, J. Edward and France Leclerc (1994), "An Eye-Fixation Analysis of Choice Processes for Consumer Nondurables," *Journal of Consumer Research*, 21 (September) 274-290.

[22] Shugan, Steve M. (1980), "The Cost of Thinking," *Journal of Consumer Research*, 7 (September) 99-111.

[23] Stigler, George J. (1961), "The Economics of Information," *Journal of Political Economy*, 69 (June) 213-225.

[24] Tversky, Amos (1972), "Elimination by Aspects: A Theory of Choice," *Psychological Review*, 79 (4), 281-200.

[25] Widing, Robert E., and W. Wayne Talarzyk (1993) "Electronic Information Systems of Consumers: An Evaluation of Computer Assisted Formats in Multiple Decision Environments," *Journal of Marketing Research*, 30 (May) 125-141.

[26] Weitzman, M. (1979) "Optimal Search for the Best Alternative," *Econometrica*, 47 (3), 641-654.

[27] Wright, Peter . L. (1975) "Consumer Choice Strategies: Simplifying vs. Optimizing," *Journal of Marketing Research*, 12 (February), 60-67.