How Do Behavioral Assumptions Affect Structural Inference? Evidence From A Laboratory Experiment

Daniel Houser and Joachim Winter

05-2002

January 2002

mea – Mannheimer Forschungsinstitut Ökonomie und Demographischer Wandel
Gebäude L 13, 17_D-68131 Mannheim_Sekretariat +49 621 181-2773/1862_Telefax +49 621 181-1863_www.mea.uni-mannheim.de

HOW DO BEHAVIORAL ASSUMPTIONS AFFECT STRUCTURAL INFERENCE? EVIDENCE FROM A LABORATORY EXPERIMENT

DANIEL HOUSER
Department of Economics & Economic Science Laboratory
University of Arizona

dhouser@econlab.arizona.edu
http://w3.arizona.edu/~econ/houser.html

JOACHIM WINTER
Mannheim Research Institute for the Economics of Aging (MEA) &
Sonderforschungsbereich 504
University of Mannheim

winter@uni-mannheim.de
http://mea.uni-mannheim.de/winter

First draft: September, 2000
This revision: December, 2001

Abstract: We use a laboratory experiment to investigate the effect that assuming rational expectations has on structural inference in a dynamic discrete choice decision problem. Our experimental design induces preferences up to each subject's subjective rates of time preference, leaving unrestricted only this parameter and the decision rule that the subject uses in solving the problem. We analyze the data under the assumption that all subjects use the rational expectations decision rule, and also under weaker behavioral assumptions that allow for heterogeneity in the way people form decisions. We find no evidence that assuming rational expectations distorts inferences about the cross-sectional distribution of discount rates.

Keywords: search; discounting; heuristics; heterogeneity; statistical classification

JEL classification: D83; C44; D90

We would like to thank seminar participants at the University of Arizona, UC Berkeley, Mannheim, Purdue, Tilburg and Yale, as well as audiences at the 2001 European Meeting of the Econometric Society for useful comments. A substantial part of this research was completed while the first author was visiting SFB 504, University of Mannheim. Financial support from the Deutsche Forschungsgemeinschaft via SFB 504, the University of Arizona Foundation, the Russell Sage Foundation grant no. 98-00-01 and the International Foundation for Research in Experimental Economics are gratefully acknowledged.

1. Introduction

The latter half of the 20th century saw the emergence of a large literature documenting apparently "irrational" behavior by subjects in economics and psychology experiments. This has led theorists, particularly over the last 10 years, to begin to assess the equilibrium implications of certain types of rule-of-thumb behavior (see, e.g., Ellison and Fudenberg 1993, Krusell and Smith, 1995, Lettau and Uhlig, 1999, or Anderlini and Canning, 2000). Interestingly, however, there has been relatively little effort made by empirical economists to understand whether, as a practical matter, assuming rational expectations leads to misspecified models that substantially distort inferences about structural parameters. A reason is that both the analysis of field data and simulation procedures are limited in the extent to which they can address this issue. This paper uses a laboratory experiment to investigate the sensitivity of inferences about one structural parameter, the subjective rate of time preference, to the rational expectations assumption.

The experimental laboratory has proved a useful tool for investigating statistical questions that are not easily answered through the analysis of field data or Monte Carlo simulations. For example, an interesting recent paper by Cox and Oaxaca (1999) employs lab experiments to assess whether supply and demand parameters can be recovered from data generated by various common market institutions. They point out that standard procedures to recover these parameters require assumptions about the data generating process (DGP), and, to the extent that these assumptions generate specification errors, inferences can be erroneous. While Monte Carlo simulations can be used to study the effects of any particular misspecification, the DGPs considered in such exercises might not correspond to any naturally occurring DGP. Cox and Oaxaca (1999) point out that one advantage of the experimental laboratory is that it can be used to study the effects of the specific misspecification that arises when the true DGP stems from human behavior. Our study, which examines the effects of the rational expectations specification, takes advantage of this feature of laboratory experiments as well.

Rational expectations is the standard assumption in virtually all applied structural econometric work, and this has generated some criticism. As Blundell (1994, p. 4) points out, "structural models are often accused (in many cases quite rightly) of imposing untenable and unbelievable restrictions on behavior." Nevertheless, there seems to be very little evidence suggesting that the rational expectations assumption has had detrimental effects in any positive sense. Research using field data to estimate life-cycle labor supply models with expectations-robust procedures (see, e.g., Hotz and Miller, 1988, or Houser, 2001) typically obtains results that line up well with findings based on rational expectations specifications.

Unfortunately, there is a fundamental limitation to the amount one can learn about the effect of expectation specification errors from the analysis of field data. In particular, Rust (1994) shows that, under very general conditions, individuals' preferences and their expectations processes are jointly nonparametrically unidentified. Because peoples' "true" period return functions are not generally known, and because the identification problem implies that multiple period return function and expectation specification combinations will fit the data equally well in-sample, examining the consequences of expectation misspecification with field data presents a substantial challenge. In contrast, the experimental laboratory provides an environment in which period return functions can be controlled, and the identification problem consequently circumvented. The result is that inferences about the effect of expectation misspecification on structural inference becomes possible.

This paper reports data from a price search experiment that induces preferences (but not the way people form expectations) up to a single structural parameter: the rate of subjective time preference (or the "discount rate."). Our experimental design, which builds on Rapoport and Tversky (1970) and Seale and Rapoport (1997), provides information on discount rates because it incorporates the possibility that a subject's experimental earnings might be paid to them two weeks after they complete their laboratory session. Note that inferences about time preference have been drawn in many other laboratory experimental studies; see Chesson and Viscusi (2000) for a recent contribution and further references. Also, an interesting study by Warner and Pleeter (2001) takes advantage of a natural experiment, providing data on real choices with large stakes, to estimate time discount rates. Finally, there are also recent econometric studies which estimate discount rates from field data using indirect inference methods (e.g., Samwick, 1998; Gourinchas and Parker, 2000; and Cagetti, 2000).

We compare three approaches to inference about subjects' discount rates. One assumption common to each approach is that subjects are risk neutral. This assumption is standard in the price search literature, and using it here allows our results to be more easily compared to that literature. Moreover, while preferences regarding risk and intertemporal substitution can be separately identified conceptually, say using the framework of Kreps and Porteus (1978), they are typically difficult to distinguish in practice. Still, the assumption of risk neutrality might fail. Because risk averse subjects should stop search sooner on average, and since it turns out that more impatient subjects should also stop search sooner on average, the failure of risk-neutrality would tend to bias downwards our discount rate estimates.

Our first results derive from a simple revealed preference analysis that is based only on subjects' accepted payoffs. This procedure enables us to establish discount rate bounds under the weak assumption that subjects prefer higher to lower payoffs (and risk neutrality). Although this procedure is robust, it ignores all continuation decisions which, in general, also contain information about discount rates. In particular, a subject's stopping rule within the sequential search task will vary with her rate of time preference, both under rational expectations and under many other decision rules (which we call heuristics in the sequel). Incorporating stronger behavioral assumptions, therefore, typically allows more precise inference about rates of time preference, but at the possible cost of bias due to model misspecification.

Our second results stem from the assumption that subjects have rational expectations, while the final set of results are based on a relatively more robust approach. This latter approach requires that we pre-specify a set of candidate decision rules, which includes the rational expectations decision rule as well as sophisticated and naive heuristics, and then use a classification procedure, related to El-Gamal and Grether (1995), to determine which combination of heuristic and discount rate is most likely for each subject. This procedure relaxes both the rational expectations assumption and the assumption that all subjects use the same decision rule.

We find that inferences about discount rates are substantially sharpened, relative to the revealed preference analysis, by imposing behavioral restrictions. We also show that the data is described less well by rational expectations behavior than by simpler reservation price heuristics. The reason is that subjects tend to stop searching sooner, on average, than is consistent with rational expectations behavior and low rates of time discounting (i.e., high willingness to accept delayed payoffs). At the same time, rational expectations behavior under high time discounting is inconsistent with the frequency of delayed payoffs that are accepted. It turns out that some of some of the heuristics we consider accommodate both relaxed stopping criteria along with little time discounting, and it is these that explain our data best. Nevertheless, our inferences about the cross-sectional distribution of discount rates are not statistically significantly different under rational expectations and heuristic specifications. The mean and median from the rational expectations analysis are 0.65 and 0.68, respectively, while under weaker behavioral assumptions the statistics are both 0.70.

Although population-level inference seems robust to behavioral specifications, inferences about any given individual's discount rate can vary substantially between the behavioral specifications. While we cannot know any subject's "true" rate of time preference, we can nevertheless investigate the relationship between each subject's estimated discount rate and an external measure of the extent to which he/she is forward looking. The external measure we use is based on responses to the "consideration for future consequences" (CFC) personality survey (Strathman, *et al.*, 1994). We administered this instrument as part of the our price-search experiment. We find that neither revealed preference nor rational expectations estimates are statistically significantly correlated with this measure, while our heuristic-based discount rate estimates and the CFC measure are statistically significantly positively correlated.

The remainder of this paper is structured as follows. In section 2, we the introduce our experimental design. We then present evidence of discounting and derive bounds on subjective discount rates (Section 3). In sections 4 and 5, we discuss structural inference about discount rates under rational expectations and under weaker behavioral assumptions, respectively. In section 6, we discuss differences between the estimates at both the aggregate and individual levels. Section 7 concludes.

2. Experimental design

Subjects make decisions in a standard price search problem (with recall) modified to incorporate delayed payments. Each subject's goal is to purchase a "widget" that they value at 100 tokens. The widget is sold at many different locations, and visiting a new location costs one token. Each location posts both a price and availability for the widget that are revealed if the subject visits that location. Subjects are told that the price at each location is drawn independently from a normal distribution with mean and standard deviation of 100 and 10, respectively. Subjects are also told that the availability at each location is drawn independently (both across locations and with respect to the price draw) and is immediate with probability 0.5 or 14 days with probability 0.5. Subjects may stop their search at any time and purchase the widget from any of the locations they visited. An exact transcript of the instructions for this experiment is provided in the appendix.

One reason we chose a price search design is that allows inference about time preference within an environment where past research has shown decisions to be reasonably consistent with rational expectations (see, e.g., Hey (1982, 1987)). This might be because price search is a task that most people experience from time-to-time, and this familiarity might build an intuition that improves their laboratory performance. Like Cox and Oaxaca (1999), we believe it is important to conduct our study in a laboratory environment that has an obvious analogue in the naturally occurring world. The reasons we chose to include the possibility of recall and to give subjects full information about the price distribution are (1) the possibility of recall together with the time-differentiated availability of the product (and possibility of a delayed payment of the search surplus) allows us to observe within-game trade-offs between different points in time, and (2) giving subjects full knowledge about the price distribution, along with the opportunity to become experienced with the task, avoids problems arising from "learning" that might distort inference about time preference.

The perception that delayed payment involves risks or transaction costs can bias inference about a subject's rate of time preference. We made every effort to mitigate this confound. Subjects were told that at the end of the experiment their earnings would be put into an envelope. If they chose an immediate payment then the envelope would be given to them immediately. If they chose a delayed payment their envelope would be stored in a safe until it was delivered or picked up in 14 days. The instructions read (see the appendix): "We absolutely guarantee that 14 days from today you will be able to obtain your envelope with your earnings. Moreover, we will pay any costs (such as mailing costs) required to get your earnings to you after 14 days."

The experiment was conducted at the University of Arizona's Economic Science Laboratory in the Spring and Summer of 2000. Subjects were recruited from the general student population. All experiments were run entirely on computers using Visual Basic software written by the authors. This software is available from either author on request.

Upon arrival, each subject was given an instruction sheet, a transcript of which is provided in the Appendix, and seated privately at a computer terminal. To ensure subjects were experienced with the task and comfortable with the computer interface, they were allowed to play an unlimited number of practice games before playing ten games that were relevant to their final payment. Subjects earned $5 for participating in the experiment (paid to everybody at the end of the session regardless of their choices) and an additional amount, possibly time-delayed, determined by selecting randomly the outcome of one of the ten payment-relevant games. Subjects were paid privately at the end of the experiment. On average, subjects were in the laboratory for about an hour and earned about $20.

Our results are based on the play of 68 subjects. This yields a total of 680 payoff-relevant games and a total of 4527 rounds (continuation decisions). On average, subjects played 10 practice games before starting their 10 payoff-relevant games. During practice subjects on average played more rounds per game, accepted higher prices, obtained lower payoffs, and also obtained zero payoffs more often than in the payment-relevant games.

3. Evidence of discounting and bounds on subjective discount rates

Table 1 contains an overview of the intertemporal choices made by subjects. The top panel contains data for all 680 games. In 371 games, subjects chose a product with instant delivery (corresponding to an instant payoff), and in 309 games subjects chose a product with delayed delivery (corresponding to a payoff in 14 days). On average, accepted prices are higher, and the resulting payoffs lower, for chosen products with instant delivery. This is consistent with the notion that at least some of our subjects discount future payments. Discounting is also suggested by Figure 1 which shows the distributions of accepted prices for both delivery dates.

Our first approach to drawing inferences about subjects' discount rates, which can be viewed as *revealed intertemporal preference* analysis, requires only the assumption that subjects prefer higher to lower payoffs (and risk neutrality). With this assumption, we are able to derive bounds on each subject's discount rate by observing their location choices. For example, suppose a subject stops searching, and her best payoffs for immediate and delayed delivery widgets are 88 and 90, respectively. If the subject chooses the location that offers immediate delivery, it follows that her discount rate can be no larger than 88/90, while if she chooses the location offering delayed delivery, her discount rate can be no smaller than 88/90. This fact allows us to construct bounds on a subject's discount rate, given that the subject has a meaningful choice between delivery dates. In particular, only those games in which the subject has prices which would result in positive payoffs at both delivery dates contain information about her discount rate. This is the case for 604 of the total 680 games. The bottom panel of Table 1 provides descriptive statistics for the sample that excludes the 76 zero-payoff games from the revealed preference analysis.

Table 2 lists the six situations a subject can face after she decides to stop searching. These situations are defined by comparing payoffs for the chosen delivery date with the payoffs corresponding to the best price available at the alternative delivery date (of which many will be zero as well). Observations for cases 1, 2, and 3 (delivery now) are shown in the top panel of Figure 2, observations for cases 4, 5, and 6 (delivery later) are shown the bottom panel. Observations that correspond to cases 1 and 4 are above the 45 degree line, observations that correspond to cases 2 and 5 are on the 45 degree line and observations below the 45 degree line correspond to cases 3 and 6, respectively. Table 2 also contains the counts of these observations.

From these observed decisions on intertemporal tradeoffs, we can compute bounds on subjects discount rates. We exclude those cases in which subjects chose combinations of price and date that were strictly dominated in terms of the associated payoffs (for example, all case 4 observations) and all observations with zero payoff for the best available price at the alternative delivery date.

We end up with observed choices that allow us to obtain lower and/or upper bounds that lie in the [0, 1] interval for most of our subjects. Note that, because all subjects play 10 games, we have more than one lower or upper bound for some subjects. In these cases, we compute the intra-subject maximum of the lower bound, or the intra-subject minimum of the upper bound, respectively. There were a several cases where a subject's decisions did not reveal an upper or lower bound. In these cases we set the lower bound to 0 and/or the upper bound to 1. Once the bounds were determined we computed the "revealed preference" estimate of the subjective discount rate for all subjects as the mean of the lower and upper bounds.

This revealed preference estimation procedure can be interpreted as follows. We impose a prior for the subjective discount rate which is a uniform over the [0, 1] interval. Observed lower and upper bounds away from zero or one, respectively, allow us to narrow our prior. For those subjects for whom we cannot construct bounds away from zero and one, we end up with an estimate of 0.5 (this is the case for 38 subjects, more than half of our 68 subjects). Although using the mean of the set of possible discount rates is arbitrary, it is reasonable because we will use this estimate within each inferential approach, and our primary interest is in whether differences appear between approaches.

The distribution of the bounds on subjective discount rates and the estimate of the subjective discount rate that we obtain are summarized in Table 3. In general, the bounds on discount rates are relatively wide, and there is a fair amount of heterogeneity in these bounds; this corresponds to the scattered price observations in Figure 2. Figure 9 shows the distribution of estimated discount rates (together with estimates obtained using other behavioral assumptions and procedures, to be discussed in the following sections). It is evident that the revealed preference approach without stronger behavioral assumptions does not substantially narrow the flat prior.

4. Inference under rational expectations

Figure 3 suggests that the revealed intertemporal preference approach provides quite imprecise inference about our subjects' discount rates. Although it is attractive in that it requires very weak behavioral assumptions, this also means that only the final decision in each game, the location decision, informs the discount rate. Under stronger behavioral assumptions it is usually the case that continuation decisions inform the discount rate as well. We adopt this approach in this and the next sections.

A natural starting point is to assume that subjects follow the rational expectations stopping rule. Under this assumption, the subjective discount rate is identified and can be estimated as the only free parameter of the corresponding decision process. To see this, it is useful to derive the optimal solution to the search task. Since the value of the "widget" in the experimental task is 100 tokens, and since each search costs one token, it is reasonable to derive the solution to the problem under the restriction that the number of searches cannot exceed 100. For the remainder of this paper, we label a subject whose behavior is consistent with this rule as a "rational expectations" agent.

Let $t \in \{1, 2, ..., 100\}$ denote the number of searches that the subject has made. (For notational convenience, we suppress a subject index as long as no confusion can arise.) After making the tth search the agents' state vector is $S_t = \{t, p_0^b, p_1^b\}$ where p_0^b is the lowest price the subject has encountered for immediate delivery, and p_1^b is the lowest price encountered for delayed delivery. After the tth search the agent may either stop searching and choose a location from which to buy the item or continue to search. If the agent stops searching then she chooses to buy the item from the location with the highest payoff. The available payoffs are:

$$\begin{aligned}\Pi_0(S_t) &= \max\{0, 100 - p_0^b - t\} \\ \Pi_1(S_t) &= \max\{0, \delta(100 - p_1^b - t)\}\end{aligned} \qquad (1)$$

where δ is the agent's rate of time preference. It is convenient to define the chosen payoff, conditional on stopping, as $\Pi(S_t) = \max(\Pi_0(S_t), \Pi_1(S_t))$. The agent will stop searching only if the stopping value is higher than the continuation value. The recursive formulation of this decision problem is therefore:

$$V_t(S_t) = \max\{\Pi(S_t), E[V_{t+1}(S_{t+1}) \mid S_t]\} \qquad (2)$$

where E represents the mathematical expectations operator, and the expectation is taken with respect to the distribution of $S_{t+1} \mid S_t$.

It is immediate from (1) and (2) that this problem has, at every t, the reservation price property. That is, there is some best available price for immediate delivery items, which depends both on the number of completed searches and the best available price for delayed-delivery items, such that at this and any lower price it is optimal to stop searching while at every higher best available price it is optimal to continue. The same is true for delayed delivery widgets. To see this, note that both arguments that enter the maximization operator in (2) are weakly decreasing (in both prices), and that when prices are sufficiently low the continuation value is equal to the stopping payoff less one, while at sufficiently high prices the stopping payoff is zero while the continuation payoff is necessarily bounded above zero (since all price draws are independent and the measure of prices that generate positive payoffs is strictly positive). The intermediate value theorem then implies that the stopping and continuation value cross at least once. That this crossing is unique follows from the fact that the slope of the stopping payoff at the crossing point must be –1 (since the crossing occurs at a positive value for the stopping and continuation payoff), and the slope of the continuation payoff function is evidently bounded from below by –1 (because a one unit improvement generated by the current price draw can never do better than improve the continuation payoff by one unit).

From (1) and (2) it is also easily seen that both reservation prices depend on the rate of time preference, δ. One way to see this is simply to observe that the continuation value incorporates both immediate and delayed delivery events, and the value of delayed delivery clearly depends on the rate of time preference. Figure 4 plots the path of reservation prices, calculated by solving the dynamic discrete choice problem implied by (1) and (2) numerically, for the cases $\delta = 1.0$ and $\delta = 0.1$. Note that by "reservation price" we mean that the subject should stop searching if the price draw is less than *or equal to* the reservation value, and that, consistent with our experiment, we report only integer valued prices.

When the discount rate is unity the problem reduces to standard search with recall with a single reservation price. This reservation price begins at 90, decays slowly to about 87 by round 10, reaches 80 by the 19^{th} search and decays at a rate of about one per round from that point forward. When the discount rate is 0.1, the reservation price for the current payoff follows roughly the same pattern as the earlier case, but is slightly higher. The reason is that the continuation value clearly falls with the discount rate, so the person is willing to accept smaller stopping payoffs. Reservation prices for the delayed delivery follow an inverted "U" shaped pattern. The very low initial values arise because the continuation value incorporates the chance of finding an item for immediate delivery, and to offset this expected value under heavy time discounting requires very low prices.

With reservation prices in hand, it is a simple matter to estimate the discount rate for each individual. Their decision task requires them to decide when to stop searching and, once stopped, the location from which to purchase the "widget". Given any discount rate, each period each state-contingent continuation decision is either consistent or inconsistent with the reservation price associated with that discount rate. When the subject stops searching and chooses a location, that choice is either consistent or inconsistent with the given discount rate. Hence, if T_j represents the number of locations searched before the subject stopped in game j, then there are $T_1 + ... + T_{10} + 10$ payoff-relevant decisions made by the subject during the course of the experiment. We associate each subject with the discount rate that is consistent with the greatest number of those decisions (this procedure is formalized in the next section).

In this research we consider only 11 discount rate values: {0, 0.1, …, 1.0}. Such a coarse grid reduces the computational burden, and prior checks confirmed that choices for most subjects in our sample are not sensitive to smaller changes in the discount rate. Figure 5 describes our discount rate inferences for each subject. The solid circle in the figure represents the point estimate of the discount rate. For 36 of our subjects the best discount rate is not unique: several discount rates explain the subject's choices equally well. In these cases the point estimate is assumed to be the mean of the candidate values, which is described in the figure by a solid circle in the middle of a line-segment. The mean of the estimates is 0.65 and the median is 0.68.

5. Inference under heuristics

An alternative to assuming that all subjects use the rational expectations decision rule is to impose the weaker assumption that each subject's decision rule is a member of a pre-specified, finite “candidate” decision-rule set. This approach relaxes the rational expectations assumption, and allows different subjects to use different decision rules. As long as the candidate rules in the set generate different state-contingent decisions for at least one state that occurs with positive probability, then the decision rule used by each subject is asymptotically (in number of decision problems played by the subject) identified. Moreover, if the choices predicted by the decision rules also depend on an intertemporal trade-off, then it is possible to estimate jointly the subject’s subjective discount rate and the decision rule she uses. The primary advantage of this procedure is that inferences about discount rates are less likely to be biased by misspecification of the decision rule.

5.1 Specification of the set of admissible heuristics

Following the approach taken in earlier experimental work on behavior in search tasks by Hey (1982) and Moon and Martin (1990), and in a slightly different context by Müller (2001), we base our analysis on a pre-specified set of candidate decision rules. We assume in our statistical analysis that subjects use one of these rules to determine their stopping decision in every round of the search problem, conditional on observable state variables and unobservable preference parameters (in our case, these are the discount rate and, possibly, parameters of the decision rules themselves).

The decision rules we consider fall in three broad classes. The first rule is the solution to the search problem derived under rational expectations (rule 1). The second class comprises six sophisticated heuristics which reflect some properties of the optimal solution but are still simple to apply (rules 2 through 7), and the third class consists of 19 naive heuristics which are not directly related to the optimal solution but extremely simple to apply (rules 8 through 26). Table 4 contains an overview of all 26 rules.

Decision rule 1 is the *optimal search rule*. Subjects who follow this rule stop searching if the best available price for delivery now is smaller than the reservation price for delivery now, or if the best available price for delivery later is smaller than the reservation price for delivery later. The reservation prices are given by the solution to the search problem under rational expectations, as discussed in Section 4 above.

The next class of decision rules are sophisticated heuristics; they have the reservation price property. We have designed these rules to take account of two important determinants of the optimal reservation price, namely, the known standard deviation of the price distribution and the accumulated search cost. Subjects who use one of these heuristics stop searching when the best available present-value payment exceeds a certain threshold level. In the first version (rules 2, 3 and 4), this threshold is a multiple (0.5, 1.0, or 1.5) of the standard deviation of the price distribution; note that this threshold is constant over time. In contrast, in the second version we assume that the threshold falls linearly in the accumulated search cost (rules 5, 6 and 7). In that version, subjects always stop after a maximum number of searches (since the standard deviation of the price distribution is 10, this happens in rounds 5, 10, or 15, depending on the threshold parameter).

An intuitive interpretation of these heuristics is that subjects understand that a "good" realization of the random price process is a price that is well below the mean of the known distribution, and natural focal points for the distance between a price draw and the mean of the distribution are multiples of the standard deviation. Following these sophisticated heuristics does not require solving the underlying dynamic decision problem. In particular, subjects do not form expectations over future price draws.

Figure 6 compares the reservation prices for various decision rules, discount rates, and delivery dates. Figure 6a shows reservation prices for items with immediate availability implied by rules 2 and 5 (i.e., parameter 0.5) and a discount rate of 1.0, to the corresponding optimal reservation price path. The path of rule 2 begins above the optimal path, is about the same at round 5 and then falls and stays below. Hence, following this rule would see people either stopping too easily in early rounds or continuing their search when they should stop in later rounds. Rule 5, on the other hand, raises the reservation price at all search durations. Under rule 5, searchers stop with any positive payoff beginning in round 5, while this is not true for the optimal rule until about round 20. Figure 6b provides a similar comparison for delayed delivery reservation prices under a discount rate of 0.5. The reservation price rules are qualitatively similar.

Decision rules 8 through 12 have been suggested by Moon and Martin (1990), following earlier work by Hey (1982). Subjects who follow rules 8, 9, or 10, the *one bounce rule*, search at least 2, 3, or 4 rounds, respectively, and then stop as soon as a price draw results in a present-value payment that is lower than the best payment available from earlier rounds. Rules 11 and 12 are even simpler – subjects just do a *fixed number of searches* (2 or 3, respectively) and then take the best available price. These rules were designed for an environment with unknown price distributions.

The remaining decision rules are based on *winning streaks* (rules 13 through 20) and *loosing streaks* (rules 21 through 26). Subjects who follow these rules stop searching if they receive two or three consecutive price draws that are above or below some fixed threshold level. Both winning and loosing streak heuristics are formulated in two versions that differ in how these threshold levels are constructed. The first version is based on the price draws themselves while the second version is based on the present-value payments implied by the current price draw. The threshold levels are, once again, formulated as multiples (0.5 or 1.0) of the standard distribution which we interpret as focal points; details can be found in Table 4. All rules based on winning or loosing streaks take into account knowledge of the price distribution – the idea is that subjects use the standard deviation to determine what a "good" or a "bad" draw is. However, subjects who follow these rules falsely assume, at least implicitly, that there is some dependence in the price process.

5.2 Statistical classification procedure

Our statistical classification procedure involves nesting the maximum likelihood classification algorithm proposed by El-Gamal and Grether (1995) within a discount rate grid search. To implement this procedure we assume that each of our subjects follows exactly one of the heuristics in our universe of heuristics, and that they use the same heuristic within each game that they play (for an alternative approach that does not require one to pre-specify the universe of alternatives see Houser, Keane and McCabe (2001)). This seems reasonable in view of the fact that all subjects are experienced when they begin the payoff relevant games. Of course, a subject's behavior need not necessarily be perfectly consistent with any of the heuristics in our universe of alternatives. Following El-Gamal and Grether (1995), we circumvent this problem by assuming that subjects play their heuristic with error, and that the error rate is the same across all subjects.

More formally, to implement the classification procedure we assume that each subject has a fixed subjective discount rate $\delta_i \in [0,1]$. We also assume that each subject follows a single decision rule $c_i \in C$, where C is a pre-specified universe of candidate heuristics. Finally, we assume that a subject plays her decision rule with error rate ε and that this error rate is constant over time and across subjects.

Each heuristic/discount rate combination we consider provides a unique map from a subject i's state (information set) S_{it} to their continuation decision $d_{it} \in \{0,1\}$. Let d_{it}^*

denote the observed decision of subject i in period t, and let $d_{it}^{c_i,\delta_i}(S_{it})$ be the unique decision implied by heuristic c_i and discount rate δ_i under information S_{it}. Then, define the indicator function $x_{it}^{c_i,\delta_i}(S_{it}) = 1(d_{it}^* = d_{it}^{c_i,\delta_i}(S_{it}))$. This variable takes value one if the subject i's tth continuation decision agrees with the prediction, and is zero otherwise. Note that no difficulty arises if an error leads a subject to continue their search when they should have stopped, since the subject's state and corresponding continuation decision remain defined. It is, therefore, not inconsistent to assume that a subject who violates her decision rule in one period can still act according to this rule during the remaining search periods.

Finally, the last decision that a subject makes in every game is the location decision. We assume that the location decision is made independently of the continuation decisions and with the same error rate. In particular, we assume that a subject who stops with J location alternatives chooses the "right" location with probability $(1-\varepsilon)$ and a member of the set of "wrong" locations with probability ε, independent of J.

We observe T_i decisions for subject i. Following El-Gamal and Grether (1995), define the sufficient statistic $X_i^{c_i,\delta_i} = \sum_{t=1}^{T_i} x_{it}^{c_i,\delta_i}(S_{it})$. Then the likelihood function for subject i is

$$f^{c_i,\delta_i,\varepsilon}(x_{i1}^{c_i,\delta_i},\ldots,x_{iT_i}^{c_i,\delta_i}) = (1-\varepsilon/2)^{X_i^{c_i,\delta_i}}(\varepsilon/2)^{T_i - X_i^{c_i,\delta_i}}. \qquad (3)$$

If there are I subjects, each indexed by i, then a natural way to estimate jointly the discount rate for each subject and their heuristic c_i is to solve:

$$\left(\hat{c},\hat{\delta},\hat{\varepsilon}\right) = \arg\max_{\left\{\begin{smallmatrix} c_i \in C \\ \delta_i \in [0,1] \\ \varepsilon \end{smallmatrix}\right\}_{i=1,I}} \sum_{i=1,I} \log(f^{c_i,\delta_i,\varepsilon}(x_{i1}^{c_i,\delta_i},\ldots,x_{iT_i}^{c_i,\delta_i})), \qquad (4)$$

where $\hat{c} = (\hat{c}_1, ..., \hat{c}_{T_i})$ and $\hat{\delta} = (\hat{\delta}_1, ..., \hat{\delta}_{T_i})$.

To avoid over-fitting with respect to the number of heuristics used to explain the subject population's behavior, El-Gamal and Grether (1995) suggest that, if there are k heuristics used to explain all of the subjects' decisions, then the log-likelihood should be penalized by an amount $R(k)$. Adopting their suggestion for our environment implies that the penalty factor is

$$R(k) = k \log(2) - 68 \log(k) - k \log(26). \tag{5}$$

Log-likelihood (4), minus penalty factor (5), is maximized over the set of all possible k-tuples (k=1,…,26) that can be formed from our universe of 26 decision rules.

El-Gamal and Grether (1995) suggest the penalty factor given by (5) for the following reason. Let C^k denote a set of k heuristics. The penalized log likelihood implied by (4) and (5) is the Bayesian posterior that arises under the priors that (i) the probability that the population includes exactly k heuristics is $1/2^k$, (ii) all possible k-tuples of heuristics in any C^k are equally likely (in our case each has probability $1/26^k$), (iii) all allocations of heuristics to subjects are equally likely (in our case each has probability $1/k^{68}$) and (iv) all error rates (between zero and one) are equally likely and do not depend on the number of rules used in the population or on the way those rules are assigned.

It is worthwhile to point out that an alternative approach for classifying subjects according to their decision rules is based on the EM algorithm (for an application in this context, see Little and Rubin, 1983). As discussed by El-Gamal and Grether (1995, section 6), using an EM classification algorithm has two advantages over our approach. First, since we classify each subject to the rule that maximizes their contribution to the likelihood function, we are in essence minimizing the number of errors attributed to each subject. This results in a downward bias in our estimate of the error rate. Second, small-sample misclassification errors are not taken into account when we estimate the discount rate. However, as the number of observations per subjects tends to infinity, both estimation algorithms will coincide with probability one and share the same consistency properties (El-Gamal and Grether, 1995, p. 1144).

5.3 Inference about discount rates

Table 5 reports the aggregate fraction of choices that are consistent with each heuristic for each discount rate in {0, 0.1, …, 1.0}. The numbers of the heuristics are defined in Table 4. Heuristics 13 through 26 are streak-based and provide a relatively poor and relatively discount-rate insensitive fit. The worst fits are provided by heuristics 22 through 26, the bounce and fixed search heuristics. Recall that the latter heuristics have been designed for search tasks in which the distribution of prices is unknown, so it is not surprising that in our experiment, they are dominated by heuristics which account for knowledge of the price distribution. The best aggregate fits are provided by the six reservation price heuristics (rules 2 through 6) and the optimal, rational expectations behavior (rule 1). That the optimal rule is consistent with up to 88% of all choices is consistent with others' findings that behavior is often near optimal in search experiments.

Although behavior is near optimal, it is clear that the reservation price heuristics provide a superior aggregate fit. Heuristic 5, whose reservation price profiles are plotted in figure 6, matches 96% of all choices made in the experiment when the discount rate is above 0.6. Note also that while aggregate consistency tends to be higher at higher discount rates for the reservation price heuristics, it is highest for the optimal rule when the discount rate is 0.6. This suggests that subjects are, on average, stopping to search sooner than is consistent with rational behavior at high discount rates. While reducing the discount rate helps to fit stopping decisions when a current payoff is chosen, it becomes inconsistent with the decision to accept a delayed payoff. A relaxed stopping criterion along with a high discount rate, which reconciles the decision to accept delayed payoffs, is consistent with our reservation price heuristics.

Table 5 also reports the number of subjects whose behavior is best explained by the heuristic. These counts appear in parentheses next to the heuristic number, unless the heuristic was not assigned to any subject. Five heuristics appear in subjects' "best" sets: 1, 2, 5, 6 and 7. The reason the number of heuristic assignments exceeds the number of subjects is that, for 24 subjects, multiple heuristics are able to explain their behavior equally and maximally well (in other words, the likelihood function is flat with respect to some decision rules). Heuristics 2, 5 and 6 occur most frequently: at least one of these heuristics appears in all but two subjects' "best" sets. Two subjects' choices are uniquely best explained by heuristic 7. The rational expectations rule is assigned to five subjects, but always in conjunction with a sophisticated heuristic. This last finding has an interesting interpretation. Even though observed behavior is consistent with optimal, rational expectations behavior in five cases, it is also consistent with behavior that follows some heuristic that is much easier to use. This finding is consistent with earlier studies which found that in search experiments, observed behavior is often close to optimal behavior (as discussed in the introduction).

Figure 7 summarizes the outcome of the El-Gamal and Grether procedure when applied to our data. Each dot corresponds to a single subject. The label above the dot reports the heuristic to which the subject was assigned. In 19 cases there are multiple assignments. For example, "256" indicates that the subject's behavior is equally well explained, at potentially different discount rates, by heuristics 2, 5 and 6. Note that only these three heuristics survive the selection procedure. The vertical position of each dot measures the fraction of choices consistent with that heuristic at an optimal discount rate, and that this number is not less than 92% for any subject, and for several subjects all of their choices are predicted by their assigned decision rule(s).

Inferences about the subjective discount rates are provided in Figure 8. In the cases where multiple heuristics are assigned, the line extends from the lowest to highest discount rate among all of those in the optimal set. For example, if a subject's behavior is maximally consistent with rule 5 at discount rates .8 and .9, and also maximally consistent with rule 6 at discount rates .6 and .7, then the line would extend from 0.6 to 0.9. In relation to the inferences drawn from the rational case, discount rates are slightly higher on average and clearly less precise. One reason for the reduced precision is that, since there is no continuation value incorporated into the reservation prices, variation in the discount rate only affects the reservation price path for delayed payment items. Hence, changes in the discount rate have relatively fewer behavioral implications, and therefore carry less information, than in the rational case.

The mean estimated discount rate, where for each subject the estimate is the midpoint of the interval of possible discount rates, is 0.70 (se=0.03) in the present case as compared to the value of 0.65 (se=0.04) found in the rational case. The median is 0.70 which compares well with the median of 0.68 obtained by assuming rational expectations. The densities of discount rates estimated under rational expectations restrictions and under the weaker heuristics and revealed-preference assumptions are compared in Figure 9. The density in the heuristics case is tighter, with less mass at both extremes of the distribution.

At the individual level, the point estimate of the discount rate changes for 56 of 68 subjects when we give up rational expectations restrictions in favor of heuristics. It is larger under heuristics for 29 subjects, with an average increase of 0.30, and smaller under heuristic for 27 subjects, with an average decrease of 0.19. The largest increase in the estimated discount rate is 0.75 (from 0 under optimal behavior to 0.75), and the largest decrease is 0.45 (from 1 to 0.55).

6. Effect of behavioral assumptions at the aggregate and individual level

6.1 Effects on inference about aggregates

Figure 9 compares the cross-sectional distribution of discount rate point estimates that arises under each of the three models considered above. This distribution is of interest for policy experiments. For example, if one wanted to stimulate savings behavior with a tax incentive, the efficient policy could depend on the mean of the population's actual discount rate distribution. Consequently, a typical goal of structural inference is to learn about the central tendency of the population distribution of some parameter of interest.

The means of the distributions of point estimates that arise under rational expectations and heuristic behavior are 0.65 (se=0.04) and 0.70 (se=0.03), respectively, and it turns out that these are not statistically significantly different. The medians are also not statistically significantly different, and it turns out that the estimates' distributions are not statistically significantly different. (A two-sample, two-tailed t-test for equality of means cannot reject the null that the means of the distributions are the same at standard significance levels (p=0.26). A two-sample Wilcoxon rank-sum test cannot reject the null that the medians of the two distributions are the same (p=0.54). A Kolomogorov-Smirnov test for equality of the distribution functions cannot reject the null that they are the same at the 5% significance level (p=0.07).

The mean of the revealed preference estimates, 0.59 (se=0.2), is not statistically different from the mean of the rational expectations discount rate distribution. However, the median (0.50) and overall distribution both differ significantly from the others. (The p-values associated with the t-test, rank-sum test and Kolomogorov-Smirnov test are, respectively, 0.17, 0.01 and 0.0 against the rational expectations distribution, and 0.0, 0.0 and 0.0 against the heuristic distribution.)

If one is interested in the central tendency of the cross-sectional distribution of discount rates, then our findings suggest that there is little cost to imposing rational expectations instead of using a more robust, but in our case cumbersome, heuristics approach. Often, however, an important additional goal of structural estimation is "policy analysis" that depends on functions of the structural parameters. To investigate whether policy analysis might be affected by different behavioral assumptions, we examine how inferences about the price elasticity of demand for immediately available widgets compares under rational expectations and heuristic decision rules.

We begin by simulating the "baseline" quantity demanded under the restriction of rational expectations. Given a discount rate, simulated search proceeds sequentially with price and availability draws made according to the distributions used in our experiment. Rational expectations is imposed by (a) using the rational expectations stopping rules and (b) weighting discount rates according to their rational expectations density as described in Figure 9. The "high-price" simulation proceeds identically to the baseline except that the location of the price distribution *for immediately available widgets only* is shifted to the right by one unit, becoming N(101,100), and the rational expectations decision rule is accordingly recalculated for every discount rate. Then, the rational expectations price elasticity is determined by comparing the quantity of immediately available widgets purchased in the two simulations in the usual way. The price elasticity under weaker behavioral assumptions is computed analogously, with discount rates and heuristics weighted according to their in-sample estimated distribution. In cases where multiple heuristics were assigned to the same subject, our simulations chose one of the candidates arbitrarily. A total of 16 subjects were assigned to heuristic two, and 31 and 21 to heuristics five and six, respectively. Hence, to compute the desired elasticities we need to obtain four measures of quantity demanded, and to obtain each we conducted 68,000 separate simulations. Specifically, we have 68 discount rate-heuristic joint assignments. We simulated under each of these joint assignments 1000 times.

Under the rational expectations baseline, the immediately available widget accounts for 60% of all purchases in simulations with non-zero payoffs. This is higher than the 54% found in the experimental data (see Table 1). Under the high-price simulation the immediately available widgets are purchased only 55% of the time, suggesting a price elasticity of about eight. Under the heuristics baseline it turns out that 53% of the widgets purchased in games with non-zero payoffs are of immediate availability, and that this number falls to about 49% under the high-price simulation, implying a price elasticity of just over seven.

These calculations ignore games with zero payoffs. It is interesting to note that in the data 11% of the games resulted in zero payoffs, while in the heuristic and rational expectations simulations the numbers are 7% and 18%, respectively. The overall average payoff from rational expectations simulations, under the heuristic discount rate distribution and including zero payoff games, is 1.3% higher than the overall average payoff from the heuristic simulations. If one includes half of the zero payoff games in both the immediate and delayed delivery categories, then the rational expectations elasticity estimate becomes

6.4 while the heuristic elasticity changes to 7.0. Taken together, these results seem to suggest that, at least in our case, elasticity estimates are not very sensitive to the rational expectations assumption.

6.2 Validity at the individual level

Although the cross-sectional distributions of the discount rate estimates are quite similar, we pointed above that inference about any given individual's discount rate often varies substantially with behavioral assumptions. In many cases relevant for policy analysis, particularly those where interest centers on a proper subset of the sample, it is important to have confidence in individual estimates. In this section we show that, in a narrow sense, the estimates under heuristics seem to have greater validity at the individual level than those under rational expectations or revealed preference. Because we cannot know any subject's "true" discount rate, our approach to assessing validity is to compare the estimated discount rates with an independent, external measure of the extent to which one considers future consequences when making decisions.

Immediately following the experiment all subjects were asked to answer a questionnaire which contained a set of 12 questions. This repertoire was developed as part of a study of personality traits by Strathman *et al.* (1994). Based on these 12 items, Strathman *et al.* constructed and validated a scale labeled *consideration for future consequences* (CFC). The same questions were used by Daniel and Webley (1997) in their mail survey on savings attitudes. They argue that an individual's consideration of future consequences should be related to her intertemporal choice behavior. Daniel and Webley report positive correlations of the CFC measure with age, with saving (controlling for age), and also with other psychological constructs related to intertemporal preferences.

Based the answers to the 12 questionnaire items, we constructed the CFC measure for the 67 subjects who completed the survey. Using standard factor analysis methods, we obtained a scale reliability coefficient (Cronbach's alpha) of 0.765. Table 7 reports the correlations of the CFC measure with our three measures of the subjective discount rate (i.e., revealed preference, optimal behavior, and heuristics estimates) and the correlations among the discount rates estimates. We report Kendall's tau, a rank correlation measure that is more robust than ordinary correlations, given that our discount rate estimates take on only few values.

The strongest correlation is between the two discount rate estimates obtained under relatively strong behavioral assumptions. The correlations of these two estimates with the revealed-preference estimates are much weaker. The estimated correlations between the CFC measure and the discount rates under revealed-preference, rational expectations and heuristics are –0.07, 0.03 and 0.18, respectively. The CFC measure is not statistically significantly correlated with either the rational expectations or revealed preference discount rate estimates. However, the null hypothesis that there is no correlation between the CFC measure and the discount rates under heuristics is rejected ($p = 0.042$). While other psychological surveys might produce different results, our results suggest one externally validated dimension in which the most reasonable discount rate estimates are those implied by the heuristics model, and these coincide with the decision rules that provide the best fit according to our statistical classification procedure.

7. Conclusions

This research aimed to assess the impact of assuming rational expectations on inference about preferences within the context of a price search experiment. Our novel experimental design induced preferences up to a single free structural parameter: the rate of time preference. We drew inferences about this parameter under the assumption that subjects behaved according to rational expectations and under weaker behavioral assumptions. We found that the cross-sectional distributions of discount rate point estimates implied by rational expectations and heuristic behavior are not statistically significantly different. However, greater differences exist between the distribution implied by a revealed preference analysis and the other two cases. The reason is that our revealed preference analysis was not informative with respect to discount rates.

Although the rational expectations and heuristics discount rate distributions agree, we showed that inference at the individual level differs substantially between the two cases. One approach to assessing which set of estimates is "better" is to compare each with an external measure of forward looking behavior. We did this, using as our external measure the outcome of a personality survey that has been previously validated in the psychology literature. We found that the estimates implied by heuristics behavior are statistically significantly positively correlated with the external measure, while the estimates stemming from the other procedures are not.

The assumption of rational expectations has been criticized primarily because it seems unlikely that people actually perform rational expectations calculations when making decisions. While decades of experimental studies support this contention, it does not follow that the rational expectations assumption is not useful. Relatively more robust behavioral specifications can be cumbersome to estimate and difficult to interpret, and such costs must be carefully balanced against their expected benefits. This research suggests that the benefits of more robust specifications are subtle, but might lie primarily in improved individual-level inference.

Appendix: Transcript of instructions for the experiment

WELCOME

This is a study of individual decision making for which you will earn cash. The amount of money you earn depends on your decisions, so it is important to read and understand these instructions. All the money that you earn will be awarded to you in cash and paid to you privately. The funding for this experiment has been provided by a research foundation.

GENERAL OVERVIEW

Your goal is to purchase a "widget". This widget is sold at many different locations. Different locations sell the widget for different prices and might have different delivery schedules. At some locations the widget might be available today, while at other locations you may have to wait some time before the widget is delivered to you. In order to determine the price and availability at a location you must pay a "search" cost. You may search as many locations as you like, and when you are finished searching you may purchase the widget from any location that you have visited. Your total earnings are equal to your value for the widget minus the price you pay minus the total of the search costs you incurred.

VALUES, PRICES AND AVAILABILITY

Your goal is to purchase an widget that has a value to you of 100 tokens. This widget is sold at many different locations. Different locations sell the widget for different prices and have different availabilities for the widget. It costs you one (1) token to visit a new location and determine the price it charges for widgets and whether widgets are available.

The price and availability at each location are determined randomly by the computer. The price at any location is drawn from a Normal distribution with mean of 100 and a standard deviation of 10. *This means that half of the locations charge a price that is greater than 100 (your value for the widget) and that about 95% of all prices are between 80 and 120.*

A location's availability for the widget has nothing to do with the price they charge for the widget. Each location's availability is either now or in 14 days with equal chance. *This means that half of the locations offer immediate availability, and the other half offer delivery in 14 days, and their availability has nothing to do with their price.*

YOUR SPECIFIC TASK AND HOW EARNINGS ARE DETERMINED

Your goal is to purchase a widget that has a value to you of 100 tokens. This widget is sold at many different locations. Different locations sell the widget for different prices and have different availabilities for the widget. It costs you one (1) token to visit a new location and determine its price and availability for the widget.

You may visit as many locations as you like, paying one (1) token for each location. When you are done searching you will be given the opportunity to purchase the widget from any of the locations that you have visited. Your total token earnings are equal to your value for the widget (100) minus the price that you choose to pay minus your total search costs. These tokens are available either today or in 14 days, depending on the availability at the location where you purchased the widget.

> EXAMPLE 1: Suppose you search three (3) locations and decide to buy the widget for 80 tokens at a location with availability 14 days. Then your total earnings are: 100 – 80 – 3 = 17 tokens in 14 days.

> EXAMPLE 2: Suppose you search 12 locations and decide to buy the widget for 90 tokens at a location with immediate availability. Then your total earnings are: 100 – 90 – 12 < 0 tokens immediately. However, you never earn less than zero tokens! So in this sort of case your earnings will be set to exactly zero tokens.

You will complete ten "money" rounds of this task. The number of tokens that you earn and the availability of those tokens for each round will be recorded on your computer screen. When you have finished the 10 money rounds exactly one of them will be selected at random to be the "payoff" round. For this round only, the number of tokens that you earned will be converted to dollars at a rate of one token = one U.S. dollar. *These dollars will be paid to you either today or in 14 days, depending on the availability of the widget in the "payoff" round.* All of the tokens in the remaining rounds become valueless.

> EXAMPLE 3. Suppose that the situation described in EXAMPLE 1 is chosen as the "payoff" round. This means that you have earned $17 and that these $17 will be paid to you in 14 days.

PAYOFFS 14 DAYS FROM NOW

You have earned $5 by participating, and this will be paid to you in cash before you leave today. The remainder of your earnings will be paid to you today or in 14 days, depending on your choices. Before you begin the experiment you will be asked to address an envelope to yourself. Before you leave today all of your earnings will be placed in that envelope. If you are to be paid today, then the envelope will be given to you today. If you are to be paid in 14 days then we will mail the sealed envelope to you 14 days from now, or you can pick it up at your convenience 14 days from now. Between the time you leave today and 14 days from now your envelope will be stored in a safe at the University of Arizona. It will not be touched. *We absolutely guarantee that 14 days from today you will be able to obtain your envelope with your earnings. Moreover, we will pay any costs (such as mailing costs) required to get your earnings to you after 14 days.*

TODAY'S PROCEDURE

You may practice as much as you like. You should practice until you develop a good "feel" for the way the task works and become comfortable with the computer interface. When you are finished practicing press the "Play for Money" button which appears on the bottom left corner of your screen. You will then play 10 times in a row for money. You will draw a playing card (Ace through 10) to determine which money round becomes the payoff round. The amount of your dollar payoff, as well as its timing, will then be determined as described above.

We are interested in the decisions you make while working alone. *Please, no talking.* If you have any questions please raise your hand. When you are finished please raise your hand and someone will assist you.

References

Anderlini, L. and D. Canning (2000), Structural stability implies robustness to bounded rationality. *Journal of Economic Theory,* forthcoming.

Blundell, R. (1994), Evaluating structural microeconometric models of labor supply. In C. A. Sims (ed), *Advances in Econometrics, Sixth World Congress*. Volume II. Cambridge, UK and New York: Cambridge University Press, 3–48.

Cagetti, M. (2000), Wealth accumulation over the life cycle and precautionary savings. Manuscript, University of Virginia.

Carbone, E. and J. D. Hey (1997), How do people tackle dynamic decision problems? Manuscript, University of York.

Chesson, H. and W. K. Viscusi (2000), The heterogeneity of time-risk tradeoffs. *Journal of Behavioral Decision Making*, 13, 251–258.

Cox, J. C. and R. L. Oaxaca (1999), Can supply and demand parameters be recovered from data generated by market institutions? *Journal of Business and Economics Statistics*, 17, 285–297.

Daniel, T. and P. Webley (1997), The economic psychology of saving: The role of individual differences associated with intertemporal choice. Unpublished manuscript, University of Exeter.

El-Gamal, M. A. and D. M. Grether (1995), Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Association*, 90, 1137–1145.

Ellison, G. and D. Fudenberg (1993), Rules of thumb for social learning. *Journal of Political Economy*, 101, 612-643.

Gourinchas, P.-O. and J. A. Parker (1999), Consumption over the life cycle. Manuscript, Princeton University and University of Wisconsin, Madison.

Hey, J. D. (1981), Are optimal search rules reasonable? And vice versa? (And does it matter anyway?) *Journal of Economic Behavior and Organization*, 2, 47–70.

Hey, J. D. (1982), Search for rules for search. *Journal of Economic Behavior and Organization*, 3, 65–81.

Hey, J. D. (1987), Still searching. *Journal of Economic Behavior and Organization*, 8, 137–144.

Hotz, V. Joseph and Robert A. Miller (1988), An Empirical Analysis of Life Cycle Fertility and Female Labor Supply. *Econometrica*. 56:1, 91-118.

Houser, D. (2000), Bayesian analysis of a dynamic model of labor supply and saving. Manuscript, University of Arizona.

Houser, D., M. Keane and K. McCabe. How do people actually solve dynamic decision problems? Manuscript, University of Arizona.

Kreps, D. M. and E. L. Porteus (1978), Temporal resolution of uncertainty and dynamic choice theory, *Econometrica*, 46, 185-200.

Krusell, P. and A. A. Smith, Jr. (1996), Rules of thumb in macroeconomic equilibrium: A quantitative analysis. *Journal of Economic Dynamics and Control*, 20, 527–558.

Lettau, M. and H. Uhlig (1999), Rules of thumb versus dynamic programming. *American Economic Review*, 81, 148–174.

Little, R. and D. Rubin (1983), On jointly estimating parameters and missing data. *The American Statistician*, 37, 218–220.

Moon, P. and A. Martin (1990), Better heuristics for economic search: Experimental and simulation evidence. *Journal of Behavioral Decision Making*, 3, 175–193.

Müller, W. (2001): Strategies, heuristics, and the relevance of risk aversion in a dynamic decision problem. *Journal of Economic Psychology*, 22, 493-522.

Rapoport, A. and A. Tversky (1970), Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5, 105–120.

Rust, J. (1994), Structural estimation of Markov decision processes. In R. F. Engle and D. L. McFadden: *Handbook of Econometrics*, Volume 4, Amsterdam: North-Holland, 3081-3143.

Samwick, A. (1998), Discount rate heterogeneity and social security reform. *Journal of Development Economics*, 57, 117–146.

Seale, D. A. and A. Rapoport (1997), Sequential decision making with relative ranks: An experimental investigation of the “secretary problem”. *Organizational Behavior and Human Decisions*, 69, 221–236.

Strathman, A., F. Gleicher, D. S. Boninger, and C. S. Edwards (1994), The consideration of future consequences: Weighing immediate and distant outcomes of behavior. *Journal of Personality and Social Psychology*, 66(4), 742–752.

Warner, J. T. and S. Pleeter (2001), The personal discount rate: Evidence from military downsizing programs. *American Economic Review*, 91, 33-53.

Table 1: Accepted prices and resulting payoffs

	Delivery date of the chosen product			
	Now		Later	
Including games with zero payoff				
Number of games (680 total)	371		309	
Mean accepted price	86.329	(6.514)	85.350	(5.639)
Mean resulting payoff	9.199	(6.517)	9.731	(6.252)
Excluding games with zero payoff				
Number of games (604 total)	329		275	
Mean accepted price	85.228	(5.457)	84.705	(5.175)
Mean resulting payoff	10.374	(5.974)	10.935	(5.554)

Notes: Standard deviations in parentheses.

Table 2: Combinations of payoffs for accepted and best alternative prices

The payoff associated with the best price available for a product with the alternative delivery date is ...	Delivery date of the chosen product	
	Now	Later
... higher than the payoff for the chosen product	Case 1 109 games	Case 4 92 games
... equal to the payoff for the chosen product	Case 2 3 games	Case 5 4 games
... less than the payoff for the chosen product	Case 3 217 games	Case 6 179 games

Note: Figures exclude games with zero payoff.

Table 3: Alternative estimates of subjective discount rates

Behavioral assumption	Mean	Min	Max	Std. dev.
No restriction (revealed preference)				
Lower bound	0.222	0	1	0.307
Upper bound	0.967	0.250	1	0.148
Mean of bounds	0.595	0.125	1	0.181
Optimal behavior	0.635	0	1	0.334
Assigned heuristic	0.704	0	1	0.222

Notes: The total number of subjects is 68.

Table 4: Candidate decision rules for the sequential search problem

Number	Description	Decision rule is a function of delta	Parameter values
1	*Optimal search* Stop searching when the best available price for delivery now is lower than the reservation price for delivery now, or when the best available price for delivery later is lower than the reservation price for delivery later.	yes	-
2, 3, 4	*Reservation price heuristic (version 1)* Stop searching when the best available payoff exceeds γ times the standard deviation of the price distribution.	yes	γ = 0.5, 1.0, 1.5
5, 6, 7	*Reservation price heuristic (version 2)* Stop searching when the best available payoff exceeds γ times the standard deviation of the price distribution less the accumulated search cost.	yes	γ = 0.5, 1.0, 1.5
8, 9, 10	*One bounce heuristic* Search at least γ rounds and then stop as soon as a price draw would result in a present-value payment that is lower than the best payment available from earlier rounds.	yes	γ = 2, 3, 4
11, 12	*Fixed number of searches heuristic* Stop searching after γ rounds.	no	γ = 2, 3
13, 14, 15, 16	*Winning streak heuristic (price version)* Stop if you receive more than γ_1 consecutive price draws below E(p) - $\gamma_2\sigma$(p).	no	γ_1 = 2, 3 γ_2 = 0.5, 1.0
17, 18, 19, 20	*Winning streak heuristic (payoff version)* Stop if you receive more than γ_1 consecutive price draws that would result in a payment over $\gamma_2\sigma$(p).	yes	γ_1 = 2, 3 γ_2 = 0.5, 1.0
21, 22, 23, 24	*Loosing streak heuristic (price version)* Stop if you receive more than γ_1 consecutive price draws above E(p) + $\gamma_2\sigma$(p).	no	γ_1 = 2, 3 γ_2 = 0.5, 1.0
25, 26	*Loosing streak heuristic (payoff version)* Stop if you receive more than γ consecutive price draws that would result in a zero payment.	yes	γ = 2, 3

Notes: We assume that after stopping according to one of these decision rules, subjects accept the price that results in the best present-value payment.

Table 5: Fraction of correctly explained choices, by discount rate (means across subjects)

	assumed value of the discount rate										
Heuristic*	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0
1 (5)	0.86	0.87	0.87	0.87	0.88	0.87	0.86	0.85	0.85	0.85	0.85
2 (21)	0.95	0.95	0.95	0.94	0.94	0.93	0.92	0.90	0.90	0.89	0.89
3	0.91	0.90	0.89	0.89	0.88	0.87	0.87	0.87	0.86	0.86	0.86
4	0.87	0.86	0.86	0.85	0.85	0.85	0.85	0.85	0.84	0.84	0.84
5 (46)	0.96	0.96	0.96	0.96	0.95	0.95	0.95	0.94	0.93	0.92	0.92
6 (29)	0.95	0.95	0.94	0.94	0.93	0.92	0.91	0.91	0.90	0.90	0.89
7 (4)	0.90	0.89	0.88	0.88	0.87	0.87	0.87	0.86	0.86	0.86	0.86
8	0.46	0.46	0.46	0.46	0.45	0.46	0.46	0.46	0.47	0.50	0.58
9	0.51	0.51	0.51	0.51	0.51	0.51	0.51	0.51	0.52	0.55	0.62
10	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.56	0.56	0.58	0.65
11	0.41	0.41	0.41	0.40	0.40	0.40	0.40	0.40	0.40	0.39	0.39
12	0.48	0.48	0.48	0.48	0.48	0.48	0.48	0.47	0.47	0.47	0.47
13	0.83	0.83	0.83	0.82	0.82	0.82	0.82	0.82	0.82	0.81	0.81
14	0.83	0.83	0.83	0.82	0.82	0.82	0.82	0.82	0.82	0.81	0.81
15	0.84	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.82	0.82
16	0.84	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.82	0.82
17	0.84	0.84	0.84	0.84	0.83	0.83	0.83	0.83	0.83	0.82	0.82
18	0.84	0.84	0.84	0.84	0.83	0.83	0.83	0.83	0.83	0.82	0.82
19	0.84	0.84	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.82	0.82
20	0.84	0.84	0.83	0.83	0.83	0.83	0.83	0.83	0.83	0.82	0.82
21	0.79	0.79	0.79	0.79	0.79	0.78	0.78	0.78	0.78	0.78	0.78
22	0.79	0.79	0.79	0.79	0.79	0.78	0.78	0.78	0.78	0.78	0.78
23	0.82	0.82	0.82	0.82	0.82	0.82	0.82	0.82	0.82	0.81	0.81
24	0.82	0.82	0.82	0.82	0.82	0.82	0.82	0.82	0.82	0.81	0.81
25	0.47	0.47	0.47	0.47	0.47	0.47	0.46	0.46	0.46	0.46	0.46
26	0.47	0.47	0.47	0.47	0.47	0.47	0.46	0.46	0.46	0.46	0.46

*** Numbers in parentheses indicates the number of subjects for whom the heuristic is a "best" heuristic. No number indicates the heuristic is not assigned to any subject. The sum of the assignments is greater than the number of subjects because some subjects are assigned multiple heuristics.**

Table 6: Correlations of the CFC scale and three discount rate estimates

	CFC scale	Revealed preference	Rational expectations
Revealed preference	-0.0699		
Rational expectations search	0.0294	0.1624*	
Heuristics-based search	0.1787**	0.2696***	0.3250***

Notes: The table reports are Kendall's robust correlation coefficients. *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Figure 1: Distribution of accepted prices

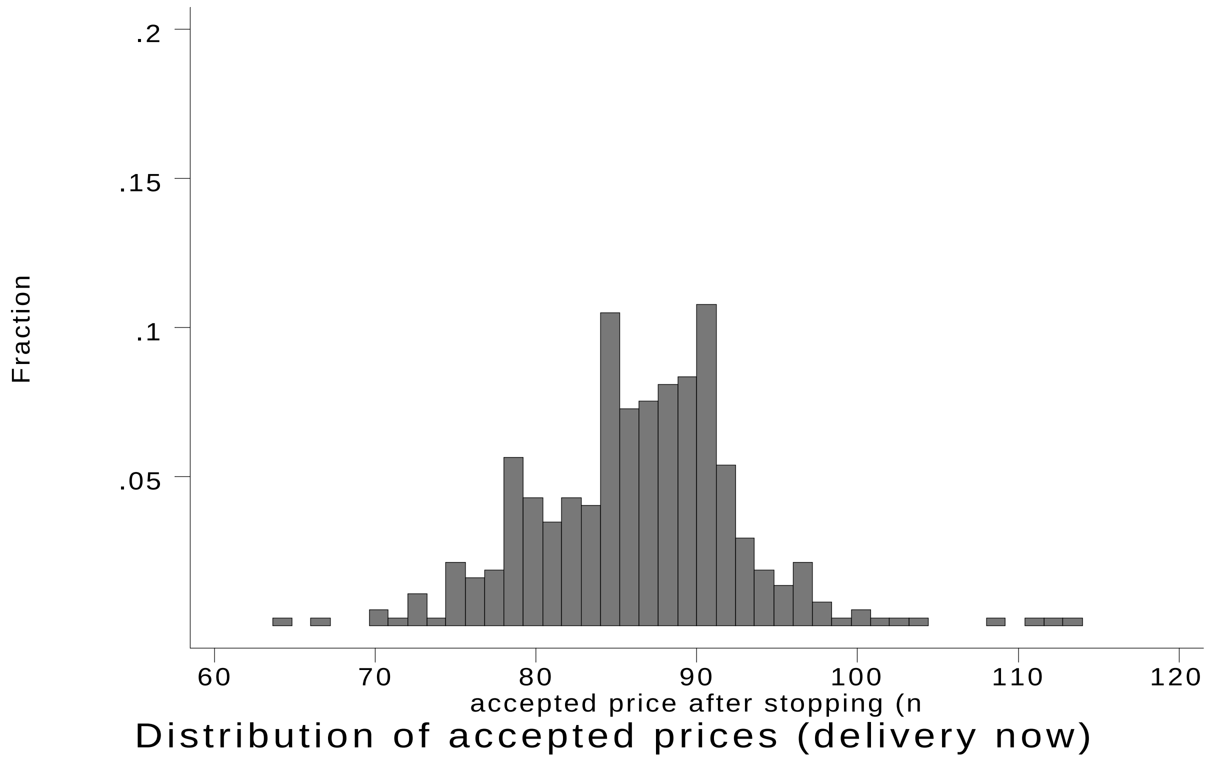

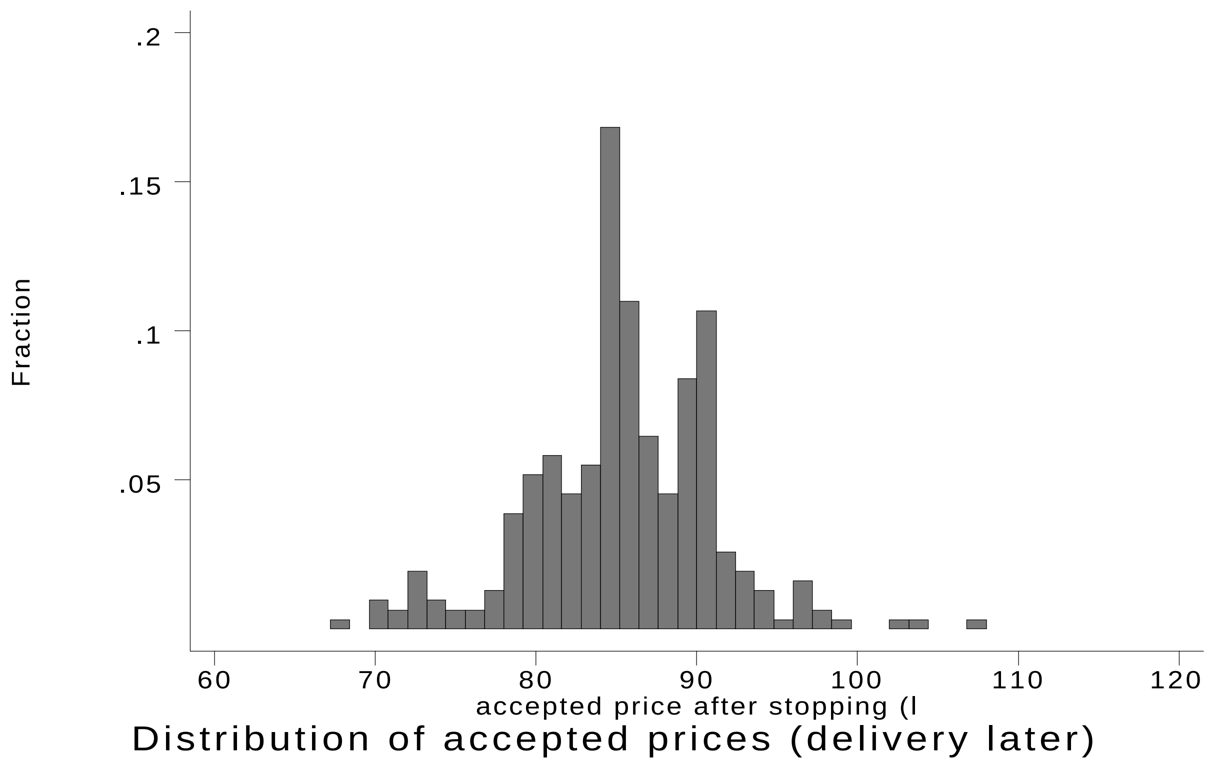

Figure 2: Payoffs for accepted price and best alternative price

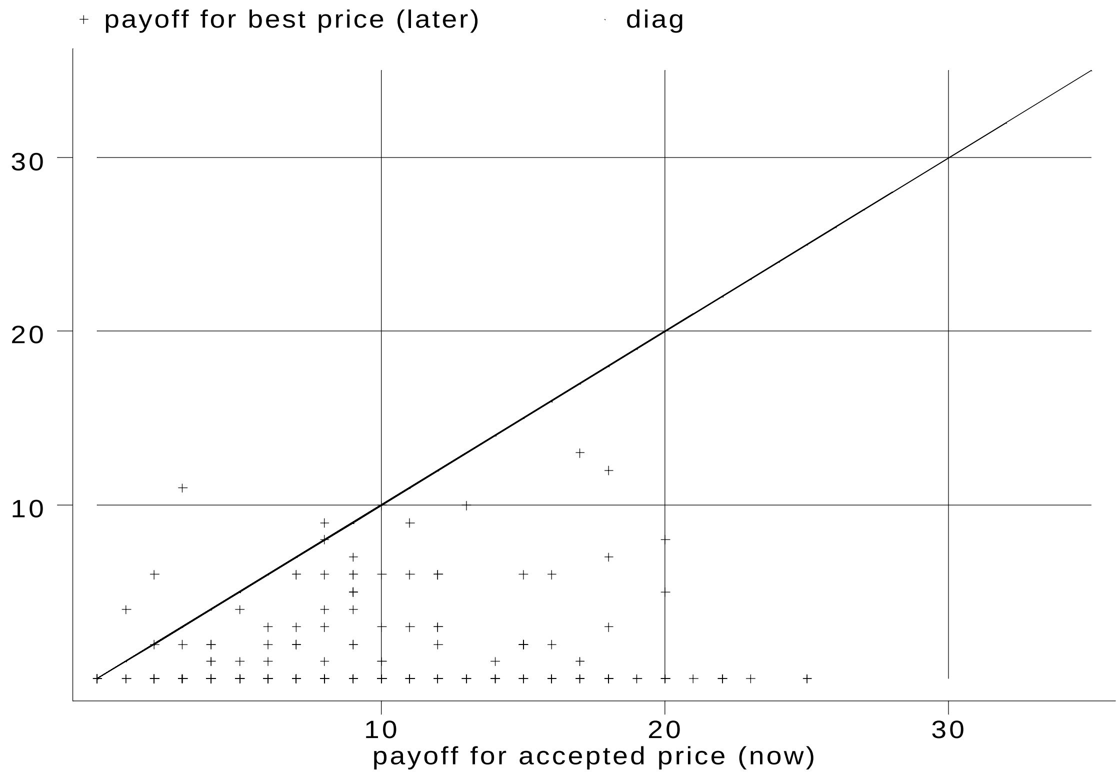

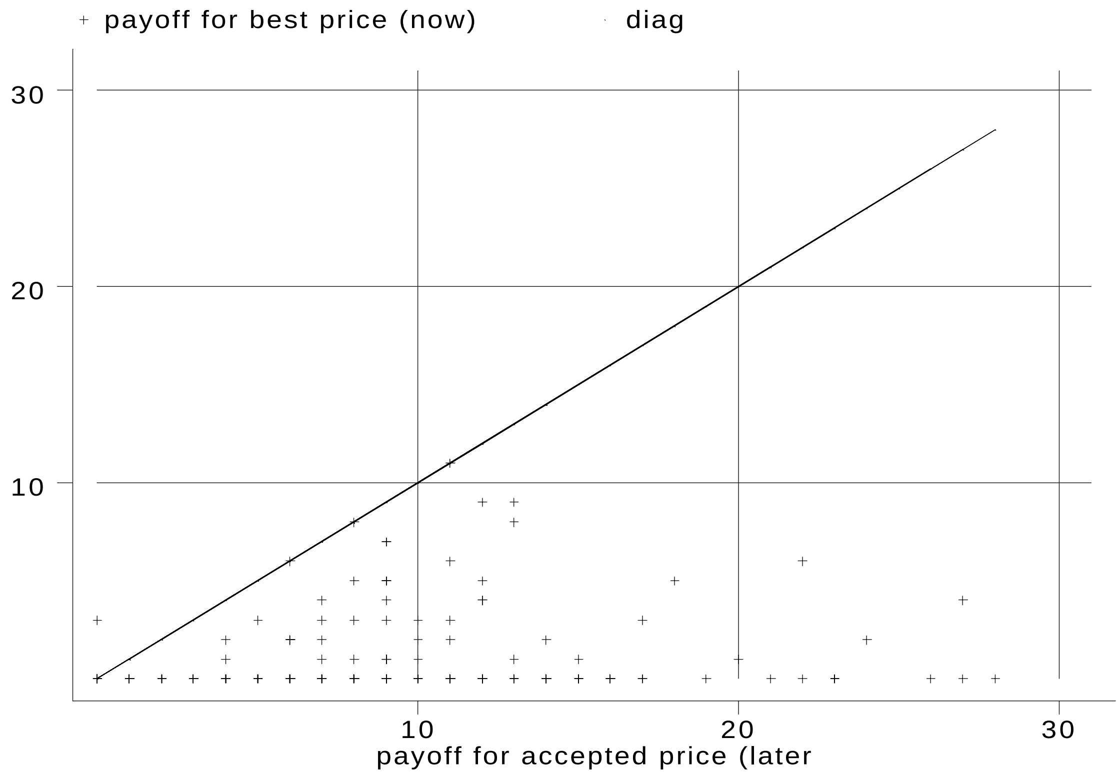

Figure 3: Estimates of discount rates without behavioral assumptions

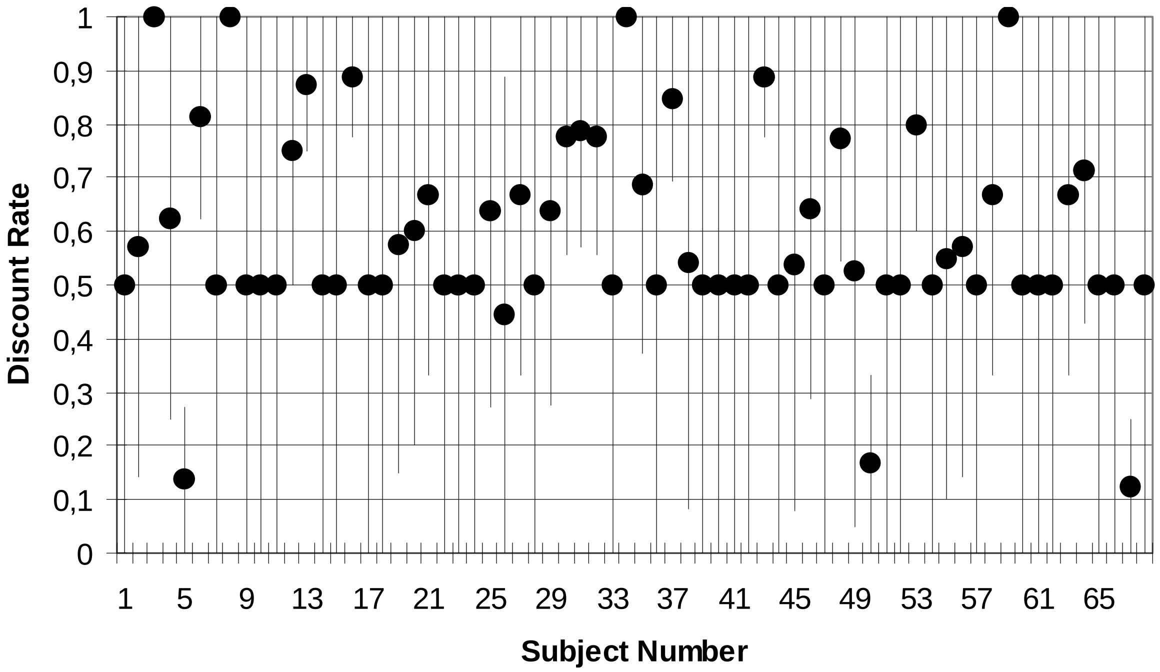

Figure 4: Rational expectations reservation prices for alternative discount rates and delivery dates.

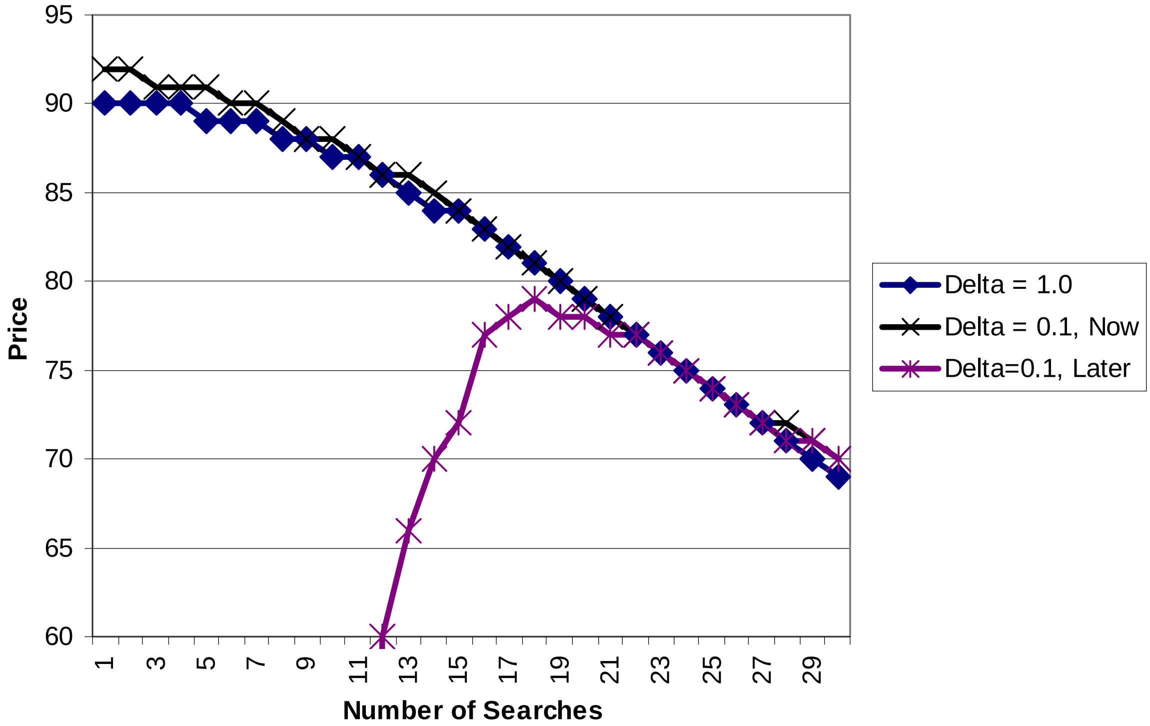

Note: The reservation prices for delivery now assume that the best available price for delayed delivery is 100. The reservation prices for delayed delivery assumes that the best available price for immediate delivery is 100.

Figure 5: Estimates of discount rates under rational expectations

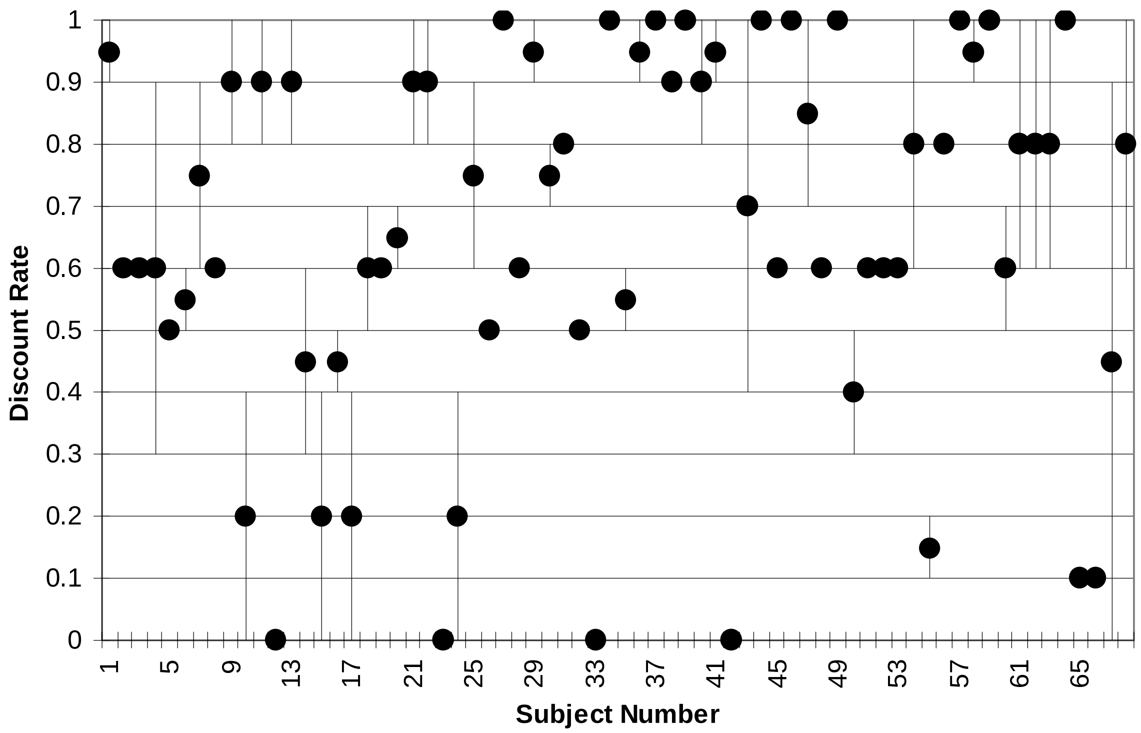

Figure 6a: Reservation prices for rational expectations and heuristic decision rules, immediate delivery.

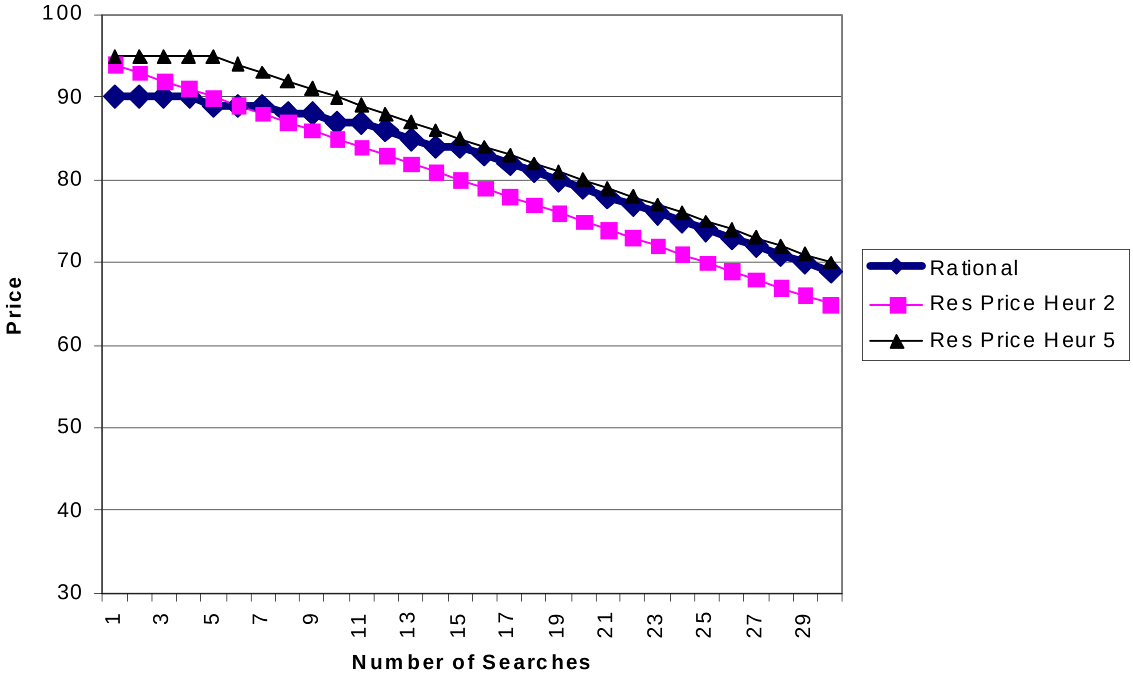

Figure 6b: Reservation prices for rational expectations and heuristic decision rules, delayed delivery.

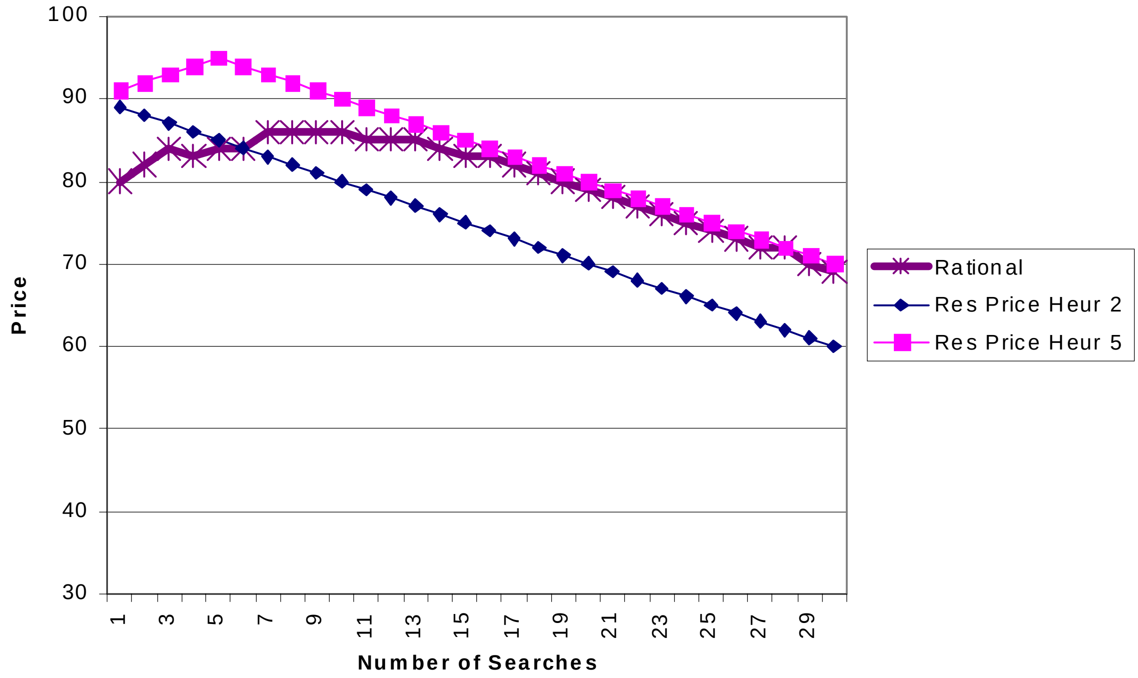

Note: The discount rate is set to 0.5 in Figure 6b. The rational expectations reservation prices were calculated as described in the note to Figure 4.

igure 7: Fit of the classification model for the assigned heuristics

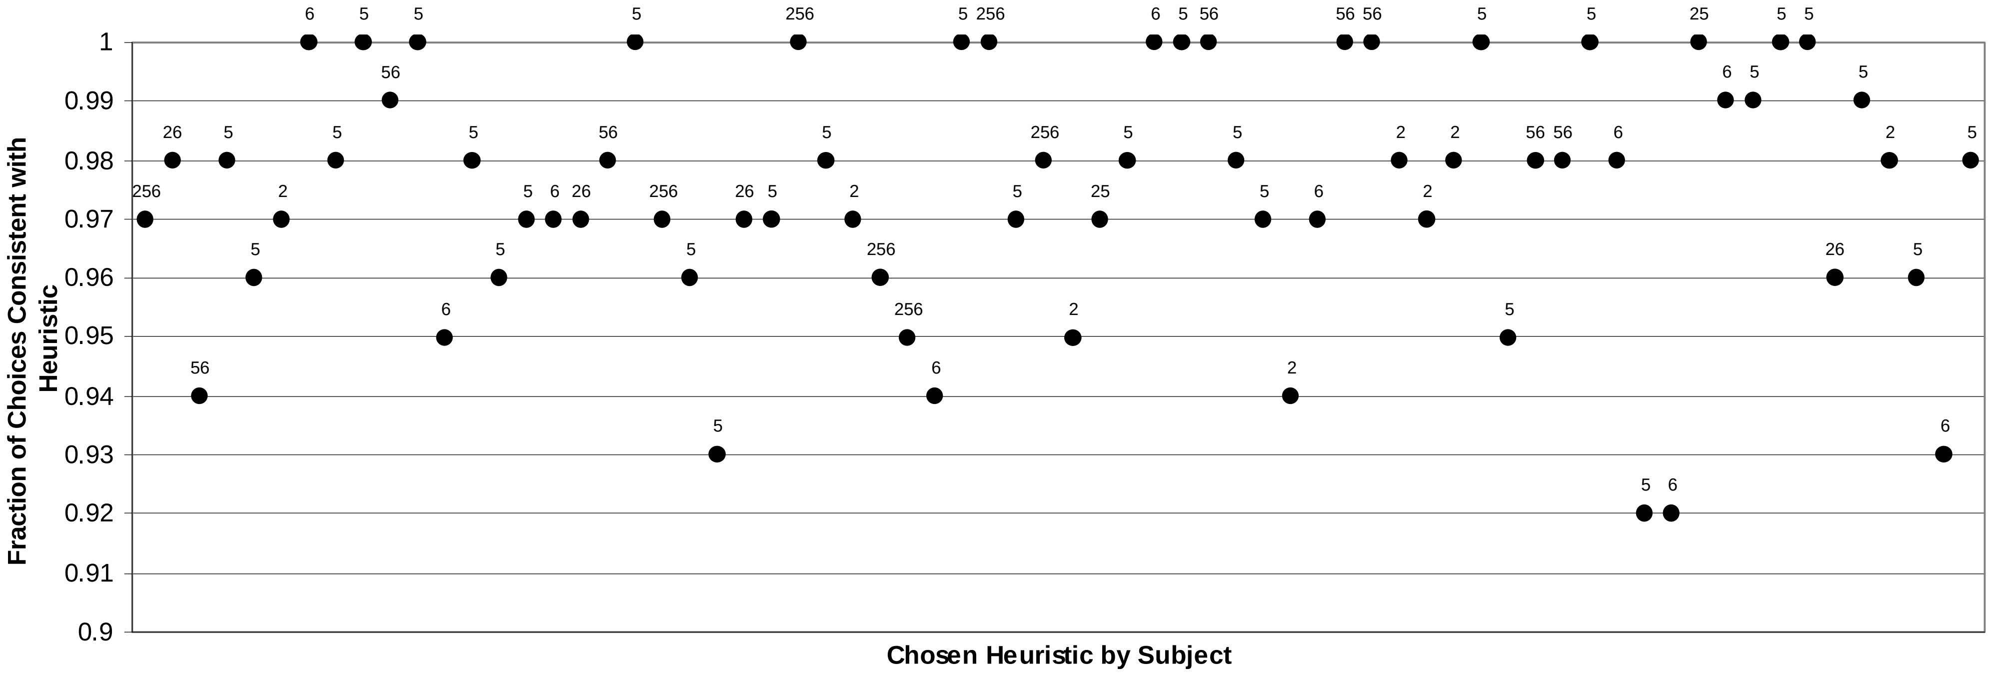

Figure 8: Estimates of discount rates under the assigned heuristic

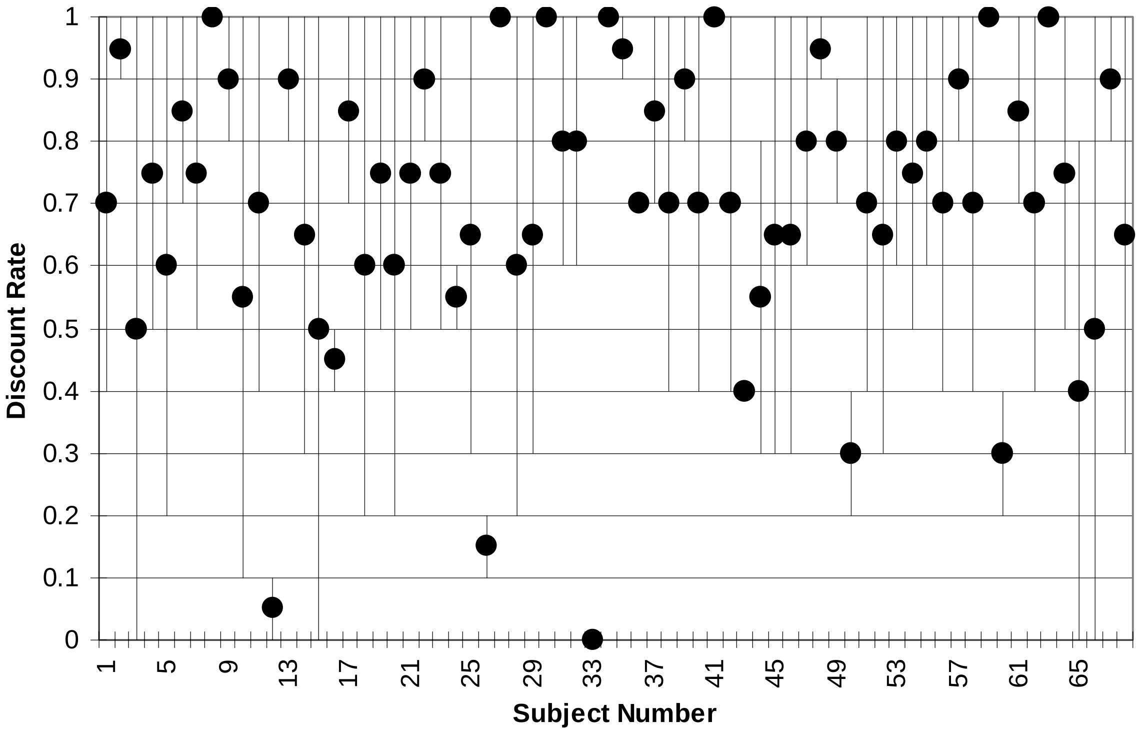

Figure 9: Density of discount rates under alternative behavioral assumptions

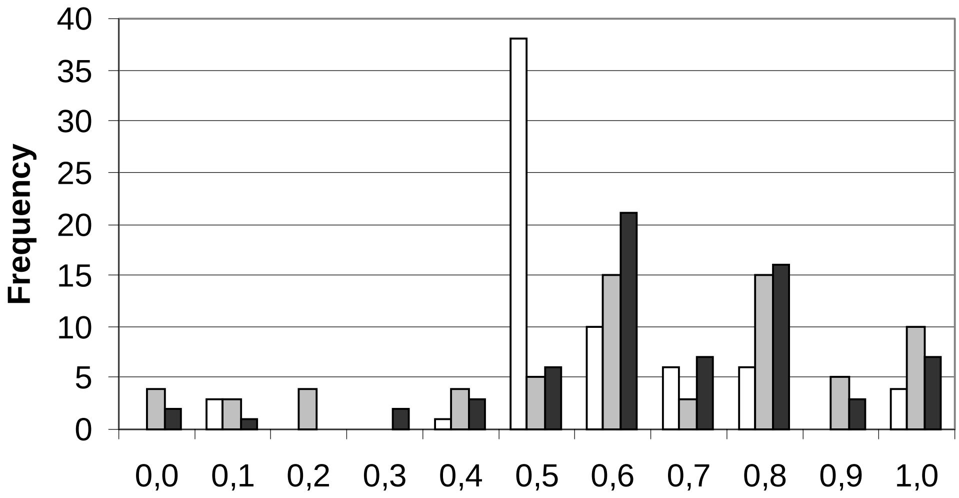

Note: This figure shows the distribution of discount rates estimated under three behavioral assumptions. White bars: revealed preference without restrictions on behavior; tan bars: heuristics assigned to each subject by our classification procedure; black bars: optimal behavior. The total number of subjects is 68.