TESTING JOB SEARCH MODELS: THE LABORATORY APPROACH

James C. Cox and Ronald L. Oaxaca

In the theoretical literature on optimal job search, models have been developed for infinite and finite search horizons, costly and subsidized search, zero and positive discount rates, search with and without recall of past wage offers, and known and unknown wage offer distributions. Much of the literature involves models of search by individual agents, although some papers embed job search in a market context. This body of theory can be tested with laboratory and field experiments and with econometric analysis of data from the historical record. Each approach has its advantages and disadvantages. This paper explains how the theory can be tested with controlled experiments, reviews the literature on laboratory job search experiments, and reports some new experiments on search with recall opportunities.

We consider experiments that are designed to test finite horizon models of search by an individual agent that have the formal structure of stochastic dynamic programming models. The decision task faced by our subjects is much more complicated than those in most individual choice experiments because of the dynamic nature of the search environment. Consider, for example, an experiment designed to test a sequential job search model. Suppose that the subject receives an offer in some period that is not the last period in the search horizon. The offer is a certain amount of money. However, if the subject accepts the offer then he must forgo the

Research in Labor Economics, Volume 15, pages 171–207.
Copyright © 1996 by JAI Press Inc.
All rights of reproduction in any form reserved.
ISBN: 0-7623-0111-2

opportunity to continue the search and, possibly, receive a more remunerative offer. This opportunity cost of accepting an offer consists of the uncertain payoffs that are forgone by relinquishing the opportunity to draw from the wage offers distribution in all of the periods that remain in the search horizon. The distribution of these forgone payoffs is complicated by the fact that the dollar value of a specific wage offer changes from one period to the next during an experimental trial because of the finite length of the (search and earnings) horizon.

The rest of the paper contains the following. In section I we discuss the role of controlled experiments in testing job search models. We explain the important role that experiments have in testing these models because several of the postulated determinants of search behavior are not observable in nonexperimental data sources. Section II contains a discussion of the literature that reports laboratory experimental research on job search models. Section III develops a finite horizon search model with recall opportunities. This model differs from many in the literature because it specifies that economic variables are discrete (rather than continuous). This distinction is essential to the numerical solutions of the model that provide the basis of the most stringent tests of it. The numerical solutions are for the linear (risk neutral) special case of the concave (risk averse or risk neutral) search model. They are used in two-sided tests of the risk neutral special case of the model. The same numerical solutions are used in one-sided tests of the general concave (risk neutral or risk averse) model because the reservation wages for a risk averse agent never exceed those for a risk neutral agent. We explain our experimental design in section IV. In section V we present our analysis of the results of the experiments. Section VI contains some concluding remarks.

I. THE ROLE OF CONTROLLED EXPERIMENTS

A definitional characteristic of job search models is that they are dynamic models of optimal search under conditions of uncertainty. Most such models incorporate the assumption that a searcher's only incentive in the search process is provided by the income stream that results from search. Search behavior must be consistent with two behavioral hypotheses in order for search models to be reasonably accurate predictors of job search. The two hypotheses are the following.

H-SOLVE: Normally-intelligent human beings are capable of making choices in a dynamic, uncertain decision environment as if they were finding the optimal solutions to stochastic dynamic decision problems.

H-PREFERENCE: The income stream that results from search dominates other possible determinants of job search behavior.

That hypothesis H-SOLVE is necessary for the empirical validity of job search theory can be explained as follows. If, using whatever heuristics they may, people do not make job-search choices that are (approximately) optimal then the models will not (reasonably) accurately predict their behavior. Next consider hypothesis H-PREFERENCE. Our socialization may endow us with much richer preferences for the outcomes of job search than simple wealth maximization. If that is the case, then job search models might not accurately predict job search behavior even if individuals behave as if they are good stochastic dynamic decision makers. The distinction between H-SOLVE and H-PREFERENCE, and the importance of both to empirical evaluation of job search models, sheds light on the role of laboratory experiments in such research.

Consider some of the difficulties in attempting to use nonexperimental data to test job search theory. The models imply that the feasibility of recalling (inventorying) past wage offers, the length of the search horizon, and agent information on the distribution of wage offers are central determinants of an optimal search strategy. But possibilities of wage offer recall, the length of search horizons, and agent information on wage offer distributions are not usually reported in nonexperimental data sources. Thus it is difficult, if not impossible, to use such data even to test hypothesis H-SOLVE, let alone to test the conjunction of H-SOLVE and H-PREFERENCE that would be the question in nonexperimental environments. In contrast, controlled experiments have some unique advantages for empirical evaluation of the models. Laboratory experiments and field experiments both have advantages and disadvantages and thus we will discuss both types.

It is possible to design laboratory experiments that are intended to test the conjunction of hypotheses H-SOLVE and H-PREFERENCE. One can do this by incorporating emotive terms such as "job," "unemployment," and "income from working," in the experimental instructions and asking the subjects to engage in job market "role-playing." Interpretation of the experimental results then requires the additional behavioral hypothesis that laboratory role-playing is equivalent to actual labor market participation. If the theory fails the experimental test, it is then not clear whether this is due to the failure of hypothesis H-SOLVE or to the failure of hypothesis H-PREFERENCE. If the theory survives the experimental test, it is then not clear whether: (a) the results support the conjunction of hypothesis H-SOLVE, hypothesis H-PREFERENCE, and the behavioral equivalence hypothesis; or (b) the subjects were not motivated by the role-playing instructions, hence hypothesis H-PREFERENCE was not really tested, and the results only support hypothesis H-SOLVE.

In order to obtain results with an unambiguous interpretation, one can design experiments to test only hypothesis H-SOLVE. This can be done by carefully avoiding any use of possibly emotive terms such as "job" or "employment" in the experiments. The experimental subjects can be given the opportunity to make sequential choices in a stochastic environment about accepting or not accepting "points" that have a known conversion rate into dollars. In addition, the experimenter can control, and thereby observe, the possibility of recalling past wage offers, the length of the search horizon, and agent information about the wage offer distribution. Theoretically-hypothesized determinants of search behavior such as the discounting rate of interest and the cost or subsidy to search can also be controlled in a laboratory experiment. Thus experiments of this type are well-suited for yielding empirical information about hypothesis H-SOLVE. If the results of such experiments are inconsistent with H-SOLVE then job search models are called into question, if not rejected, as acceptable predictors of job search behavior. If the results are consistent with H-SOLVE then the research can proceed to examine H-PREFERENCE, or the conjunction of the two hypotheses.

The conjunction of hypotheses H-SOLVE and H-PREFERENCE can be tested in field experiments. Such experiments could be done with offers in field labor markets in which the subject role-playing would be actual labor market participation. One disadvantage of such experiments, relative to laboratory experiments, is that they are much more expensive. Other disadvantages include the researcher's inability to control the subjects' search horizons and difficulty in controlling their information about the wage offer distribution and their discount rates. The reason for this is that the experimental field labor market could not be isolated from other uncontrolled markets. However, if laboratory experiments had previously found that subject behavior was consistent with hypothesis H-SOLVE then the control problems in field experiments intended to incrementally test hypothesis H-PREFERENCE might not be prohibitive.

II PREVIOUS EXPERIMENTAL TESTS OF JOB SEARCH MODELS

Experimental tests of job search theory are reported in Braunstein and Schotter (1981, 1982), Cox and Oaxaca (1989, 1992a, 1992b), Harrison and Morgan (1990), and Hey (1987). Cox and Oaxaca (1989) presents a discrete wage, finite horizon version of the theory of sequential search that focuses upon the opportunity costs associated with the foregone earnings from rejecting an offer when it is received. The risk neutral model is a special case of the general concave search model in which income preferences can be represented by some strictly increasing, concave utility function. The reservation wages of a risk averse agent never exceed those of a risk neutral agent. Hence, numerical solutions for the risk neutral special case of the model can be used to test both that special case and the general concave (risk averse or risk neutral) search model. Braunstein and Schotter (1981, 1982) present experimental tests of many hypotheses derived from the survey by Lippman and McCall (1976a, 1976b). The Braunstein-Schotter papers report results that are generally consistent with the theory except for one major exception: the inconsistency between the observed decreasing reservation wages of the subjects and the predicted constant reservation wage path of the infinite horizon search model that they use to interpret the data (Braunstein and Schotter, 1981, p. 20). In most of their experiments, Braunstein and Schotter did not limit the number of searches that their subjects could undertake; however the experiments clearly had to have a finite (and rather short) length. Braunstein and Schotter's data show decreasing reservation wages for their experiments where points had a one-for-one transformation into dollars (their "risk neutral" experiments). This result is inconsistent with the constant reservation wage implication of the infinite horizon search model but it is consistent with the decreasing reservation wage path of finite horizon search models. What conclusion should one draw about consistency with search theory? Our approach makes possible an unambiguous interpretation of the experimental results because we develop a finite horizon search model and then run experiments that are clearly finite horizon.

A. Tests with Search Duration Data

Cox and Oaxaca (1989) reports sequential search experiments in which there is no opportunity for recall of past offers. These experiments consist of a baseline and several treatments that were administered to two sets of subjects. The design of the baseline (BL1 & BL2) trials was as follows. The length of the search and earnings horizon was set at 20 periods, meaning that subjects could receive no more than 20 offers and that an offer that was accepted in period t would yield earnings for $20 - t + 1$ periods. The probability of obtaining an offer in any period was 0.5, which was operationalized by asking subjects to draw from an urn containing two black balls and two white balls. Conditional upon receiving an offer, the probability was 0.1 of receiving any particular wage offer between 1 and 10; this was operationalized by drawing a ball from a bingo cage containing ten balls numbered from 1 to 10. In the baseline trials, the interest rate and the net search subsidy/costs were set equal to zero. The baseline experimental design served as a reference from which the treatment experimental designs differed in only one aspect at a time.

The interest rate treatment (TR1) set the interest rate at 10%. The subsidy treatment (TR2) provided a net search subsidy of 5 points per period. The risk treatment (TR3) was a mean-preserving contraction experiment: conditional upon receiving an offer, the probability was 0.25 of receiving any particular wage offer between 4 and 7 by drawing from a bingo cage containing four balls numbered from 4 to 7. In the cost treatment (TR4), subjects incurred a net search subsidy of −10 points per period. The probability treatment (TR5) reduced the probability of receiving an offer in any period to 0.25 by asking subjects to draw from an urn holding three black balls and one white ball. Finally, the horizon treatment (TR6) shortened the search horizon to 10 periods.

There were 30 distinct individuals in the first subject group. They were given the following sequence of baseline and treatment experimental trials: BL1, BL1, TR1, TR1, BL1, BL1, TR2, TR2. This experimental design reflects three objectives. First, we wanted the experiments to be relatively short. It is important in any experiment to not exceed the attention (or patience) span of the subjects. However, we thought that this was especially important in sequential search experiments in order not to bias the search duration (and other) results: if the subjects had become tired, impatient, or bored, they might have been inclined to accept low offers in order, simply, to end the experiment. Second, we wanted to repeat the baseline trials after an intervening set of treatment trials in order to ascertain whether the subjects "returned to baseline" behavior. If the subjects' behavior in the second set of baseline trials differed significantly from their behavior in the first group of baseline trials (i.e., they failed to return-to-baseline), that would indicate that learning and/or sequencing effects were being confounded with treatment effects in the experiment. The third design objective was to run the baseline and treatment trials in pairs, both as a further check for sequence and/or learning effects and to increase the number of observations.

The second subject group consisted of 30 individuals that had not been in the first group of experiments. They encountered the following sequence of baseline and experimental trials: TR3, TR3, BL2, BL2, TR4, TR4, BL2, BL2, TR5, TR5, TR6, TR6. This experimental sequence was shaped by the same design objectives as the first sequence; however our decisions about the second sequence were informed by our experience in running the preceding group of experiments. Subjects in the first group typically only took 30 to 35 minutes to complete that experimental sequence. Thus we concluded that we could increase the number of experimental treatments in the second sequence and still have relatively short experiments. As it turned out, we doubled (from 2 to 4) the number of treatments in the second sequence but held constant the number of baseline trials; the result was a 50% increase (from 8 to 12) in the total number of experimental trials. In addition, for the second subject group, the money conversion factor ($/point) was increased by 50% in order to ensure dominance of the payoffs in the presence of four additional experimental trials.

Our reported results were based on the responses of 60 distinct subjects who participated in these search experiments. Each experiment included four baseline trials and two trials for each treatment; consequently we had a total of 120 observations for each baseline design and 60 observations for each treatment. The results were as follows. Subjects terminated search in the period predicted by the risk neutral model, conditional on the draws from the wage offers distribution, in 77% of the 600 trials. Disaggregating by treatment, we found that the success percentage for the risk neutral model varied from a low of 68% of each of the 60 trials in the subsidy treatment (TR2) and the risk treatment (TR3) to a high of nearly 87% of the 60 trials in baseline 1A (the first sequence of the baseline design for subject group one). The success rates for the concave (risk averse or risk neutral) model were, of course, even higher. Thus, 94% of the search terminations in 600 trials were consistent with the concave search model. The concave model's success percentage varied from a low of 87% of each of 60 trials in the subsidy treatment (TR2) and the horizon treatment (TR6) to a high of 98% of each of 60 trials in baseline 1B (the second sequence of the baseline design for subject group one), baseline 2B (the second sequence of the baseline design for subject group two), and the probability treatment (TR5).

We conducted both parametric and nonparametric tests on the consistency of the observed laboratory behavior with the predictions of the risk neutral and concave models, conditional on the actual draws from the wage offers distribution. The nonparametric test that we employed was the Kolmogorov-Smirnov full distribution test. This test allowed us to take into account all of the moments of the distribution of search terminations. We tested whether or not the observed distributions of search terminations by the subjects were significantly different from the theoretically-predicted distributions of search terminations conditional on the draws. The results of one-sided and two-sided tests were reported. The interest rate treatment exhibited the closest agreement between the observed distribution of search terminations by subjects and the theoretically-predicted distribution for the risk neutral model. The probability treatment was the only one for which we could reject the hypothesis that the distribution of search durations by the subjects was the same as the theoretically-predicted distribution for the risk neutral model. The one-sided tests indicated that the concave model could not be rejected for any of the treatments.

A parametric test compared the actual mean durations of search with the theoretically-predicted means for the risk neutral model. For subjects for whom search terminated on or after the period predicted by the risk neutral model conditional on the draws, it was possible to infer the exact theoretically-predicted outcomes by comparing the subjects' responses with the optimal reservation wages for the given search problem. The theoretically-predicted means were obtained as the average of the theoretically-predicted outcomes for the values of the random variables drawn in each treatment by the subjects who did not terminate search prematurely. We concluded that a subject terminated search prematurely if the theoretically-predicted reservation wage exceeded the offer accepted by that subject and no offers were received in prior periods that equalled or exceeded the optimal reservation wages for those periods. If a subject prematurely terminated search in a particular trial, the theoretically-predicted values of search duration were specified as the theoretically-predicted values conditional on continuing the search beyond the termination period selected by the subject. Tests for significant differences between observed and theoretical means were based on the asymptotic normal distribution of the ratio of mean differences to estimated asymptotic standard errors. In five of the ten treatments, the mean difference in search duration was statistically significant at the 5% level on a two-tailed test. In all but one treatment, the observed mean search duration was less than the predicted mean according to the risk neutral model. The concave (risk averse or risk neutral) model was accommodated by conducting the appropriate one-tailed test on mean differences. Only the horizon treatment yielded a statistically significant difference for the one-tailed test. This inconsistency with the concave model could conceivably have resulted from the fact that the (shortened) horizon treatment was the last one in the sequence for subject group two. Thus the habituation effects of the 20 period horizons from the preceding 10 trials may not have had time to dissipate.

Another set of statistical tests was conducted to determine whether or not treatment effects were statistically significant. Again we appealed to asymptotic normality of the ratio of mean differences to estimated asymptotic standard errors. First, we tested whether the mean responses to the treatments were significantly different from mean responses in the baseline trials. Four of the twelve treatment/baseline comparisons were statistically significant.

We were also interested in the question of whether the treatment effects observed in the laboratory were significantly different from the treatment effects predicted by the risk neutral model. We found that the observed treatment effects were significantly different from the theoretically-predicted treatment effects in only three of the twelve search duration comparisons. Our test for treatment effects was also applied to sequencing effects in the baseline trials. Recall that interest rate treatment (TR1) trials intervened between the two sets of baseline trials with the first subject group, and the costly search treatments (TR4) intervened between the baseline trials with the second subject group. Did the order of the baseline trials make a difference? We found that differences in subject mean search duration between baseline 1A and baseline 1B and between baseline 2A and baseline 2B were not statistically significant. Furthermore, mean discrepancies between subject responses and the theoretical predictions were not significantly different between the first and second baseline trials for both subject groups. It is in this sense that we concluded that subjects returned to their baseline behavior after an intervening treatment.

B. Direct Tests with Observed Reservation Wages

Cox and Oaxaca (1992a) advances our understanding of the performance of finite horizon job search theory by subjecting the theory to direct tests with observed reservation wages and by comparing the performance of the theory with that of a naive rule. We found that a naive rule performed as well as the risk neutral model in (indirect) tests that used search duration data. In contrast, (direct) tests using observed reservation wages implied rejection of both the naive rule and the risk neutral model but not the general concave model. Our interpretation of this finding was that indirect tests based on search duration are weak tests of the theory.

Procedures for conducting these experiments were as follows. Each trial consisted of a maximum of 20 periods; that is, the length of the finite horizon was 20 periods. During each period in which a trial was in progress, a subject first drew a realization of a random variable that determined whether he/she received an offer. If an offer was received, the amount of the wage was determined by the realization of a second random variable. The wage was then multiplied by the (prespecified) annuity factor that converted it into the end-period capital value of the income stream implied by receiving the wage each period until the end of the horizon. The capital value was then added to the (zero or positive) prespecified earnings (or search subsidy) for the period to determine the total earnings from stopping the search at that time.

Direct tests of job search models are based on observed reservation wages. We were able to perform such tests by conducting experimental trials in which the subject responses consisted of stated minimum acceptable offers for which they were willing to make binding precommitments of acceptance. Economic theory does not make a distinction between job search decisions that are based on precommitments to reservation wages and decisions based on acceptance or rejection of known offers. Nevertheless, we were concerned that behavior might be different in these two environments. In order to test for the presence of a "precommitment effect" on search decisions, we paired experimental trials involving precommitment with trials in which subject responses were acceptance or rejection of known wage offers.

These experiments consisted of two parts. In part I, a subject was asked to record either an S (for "stop") or a C (for "continue") during each period. The first S recorded during an experimental trial ended that trial; that is, the S response indicated job acceptance. The C response, during any period except the last period in a trial, indicated that the subject had decided to continue the search, which meant that the offer, if any, in that period was rejected. Part II of an experiment produced observations of reservation wages by asking subjects to record the minimum acceptable offer for each period of a trial. Draws were then made from the wage offers distribution. If, in any period, a subject received an offer that equaled or exceeded the pre-recorded reservation wage, the trial was terminated and the income earned in that trial was added to the subject's accumulated earnings in the experiment. If the subject did not receive an offer that equaled or exceeded the pre-recorded reservation wage, he/she was permitted to continue the search (but was not allowed to accept the offer, if any).

There were eight trials in each of the two parts of an experiment. The first two of eight trials were "baseline" trials with the following parameterization: (a) the interest rate was zero; (b) the net search subsidy was zero; (c) the probability of receiving an offer in any period was 0.5; and (d) the conditional probability that the amount of an offer was any integer from 1 through 10 was 0.1. The third and fourth trials in each part were "subsidy treatment" trials which differed from the baseline trials only by including a net subsidy to search of 5 cents per period. The fifth and sixth trials in each part were more baseline trials. The seventh and eighth trials in each part were "probability treatment" trials which differed from the baseline trials only by using an offer probability of 0.25 in place of 0.5.

This two-part design, with eight trials in each part, was used in order to implement the following objectives. The second pair of baseline trials was included in each part of the experiment so that we could test for whether the subjects "returned to baseline" (behavior) after the subsidy treatment. If they did not, then the subsidy (and probability) treatment effects on behavior would be confounded in the data with the effects of learning or other sequencing phenomena. The entire sequence of baseline and treatment trials in part I (without precommitment) was repeated in part II (with precommitment); thus all of part I was a baseline control for part II. This enabled us to test for the presence of a precommitment treatment effect on search behavior.

The 30 subjects in these experiments were University of Arizona undergraduates with no known previous experience with manually-run, individual choice experiments. Some of the subjects did have previous experience with computerized laboratory market experiments. We used the search model in Cox and Oaxaca (1989) to select parameters for the experimental design that implied that the unconditional expected payoff to a risk neutral, perfectly-optimizing subject from participating in one of these experiments was $24.06. The expected payoff conditional on the actual draws of the random variables was $22.44. Subjects actual earnings in the experiments varied from $17.56 to $24.54, and averaged $21.44. Experiments lasted about one hour.

The precommitment-induced *observable* reservation wages are the central feature of the experiments reported in Cox and Oaxaca (1992a). Since precommitment is *not* generally a feature of job search decision making in non-experimental environments, an important question was whether these observed reservation wages are the same as the implicit reservation wages subjects use when they are not required to precommit to a minimum acceptance wage. Of course one cannot merely assume that the two types of reservation wages are the same. Therefore, we employed a variety of tests to address this question. The experimental design question is whether the precommitment requirement itself is a treatment that influences subject behavior.

Both nonparametric median (Fisher sign) tests and parametric means tests were used in Cox and Oaxaca (1992b) to test for the presence of a precommitment treatment effect. Since these experiments used each subject as his own control, the draws will vary for that subject across trials. We were able to control for variation in the draws by comparing deviations from theoretically-predicted search durations conditional on the actual draws. Thus both tests compare the difference between search duration in the precommitment and no-precommitment trials with the difference between the theoretically-predicted risk neutral search duration (given the actual draws) in the precommitment and no-precommitment environments.

The conclusions reported in Cox and Oaxaca (1992b) follow from the possibility for making both within-part and across-part tests for return-to-baseline in our experimental design. Using data from the *second set* of two baseline periods and the probability treatment periods from both parts of the experiments leads to the conclusion that there was *no* significant precommitment treatment effect on search duration. In contrast, using data from the *first set* of two baseline periods and the subsidy treatment periods from both parts leads to the conclusion that there *was* a significant precommitment treatment effect. This supports the conclusion that the precommitment requirement affected subject search duration for the first two experimental treatments and then was "shaken off" for the remaining treatments.

The results from *within-part* tests reinforce this interpretation. Recall that the third (two-trial) treatment in the experimental sequence in each part was a repeat of the first (baseline) treatment. We performed tests of the null hypothesis of a return-to-baseline behavior after the intervening subsidy treatment. The results showed relatively high p values for the no-precommitment baseline experiments and relatively small p values for the precommitment baseline experiments. Consequently, one could not reject return-to-baseline for the no-precommitment experiments but could reject return-to-baseline for the precommitment experiments. This was consistent with subjects "shaking off" the influence of the precommitment requirement by the time they reached the second baseline treatment. The fact that the same baseline treatment was used when a precommitment effect was detected and when it disappeared made it difficult to argue that it was merely the sequential arrangement of the four basic treatments that somehow had caused a precommitment effect to appear and then disappear.

Since the experiments reported in Cox and Oaxaca (1992a) were designed to elicit (binding) reservation wage messages from the subjects, it was possible to conduct direct tests of the reservation wage path. Our data analysis began with an informal inspection for consistency of observed search terminations with the predictions of a naive rule and the risk neutral and concave (risk averse or risk neutral) search models. About 94% of the search terminations were consistent with the concave model. Furthermore, 76% of the search terminations were consistent with the risk neutral model and 75% of them were consistent with the naive rule. Thus the naive rule had about the same success rate in predicting search duration as did the risk neutral model. One interpretation of this finding is that indirect tests, based on reservation wage implications such as predicted search duration, are weak tests of search models. The results of the direct tests of reservation wages supported this interpretation.

We also reported likelihood ratio tests of mean reservation wage paths. The risk neutral model and the naive rule were rejected for every experimental treatment in both tests using all the data and tests using data for periods in which there were at least two searches. In contrast, the concave model was not rejected in any test using the at-least-two-searches subsample of the data. It was rejected only for the first baseline in the tests using all of the data.

Recall of past offers was not possible in any of the experiments reported in our previous papers. The risk neutral (or linear) model survives tests based on search duration and the general concave (risk averse or risk neutral) model survives those tests remarkably well. Tests with reservation wages lead to different results. The risk neutral search model is rejected by tests with reservation wage data; however, both the general concave model and a (specific) logistic concave search model are not rejected by the reservation wage tests. We next examine whether risk neutral and concave search models with perfect and stochastic recall can survive experimental tests.

III. A MODEL OF FINITE HORIZON SEARCH WITH RECALL

In this section we generalize the finite horizon sequential search model in Cox and Oaxaca (1989) to include both perfect and stochastic recall opportunities. Besides its finite horizon, the other feature of that model that is essential to experimental testing is its inclusion of a discrete wage offers distribution. Discrete wages are essential to derivation of the numerical solutions that make possible quantitative tests of the theory.

In the absence of recall opportunities, a wage offer that is not accepted when it is received is no longer available. In contrast, in a model with perfect recall opportunities, once an offer is received it is always available. Stochastic recall is an intermediate case in which a wage offer that is not accepted when it is received may be available later. Recall of past offers will never occur even if it is possible in an optimal infinite-horizon search from a known distribution by a risk neutral agent (Lippman and McCall, 1976, pp. 168–172). This is an immediate implication of the constant reservation wage in that basic search model. But past offers may be recalled in an optimal finite-horizon search because of decreasing reservation wages. Since actual (field or laboratory) search will be finite in length, search with recall opportunities is an interesting topic of study.

Assume that the finite search horizon is T periods. Assume that the probability that a job offer is received in any period in the search horizon is p. If an offer is received, the wage is a realization of the random variable W with discrete density function,

$$g(w) = \begin{cases} Prob(W = w) & \text{for } w^l \le w \le w^h \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

Define the cumulative conditional wage offers distribution function:

$$G(w) = \sum_{w^l}^{w} g(z) \tag{2}$$

The searcher will have made t draws from the unconditional wage offers distribution if he is still searching at time t. Minimal rationality implies that only the highest available wage offer will be considered for acceptance. In time period 1, the probability that job search yields a wage offer that is less than or equal to w (denoted $F(w)$) is the sum of the probability of not getting an offer, $1 - p$, and the probability of receiving an offer that does not exceed w, $pG(w)$. Thus

$$\begin{aligned} F_1(w) &= 1 - p + pG(w) \\ &= 1 - p(1 - G(w)) \end{aligned} \tag{3}$$

If there are no recall opportunities then the cumulative distribution function for wage offers is the same in all periods in the search horizon: $F_t(w) = F(w) = 1 - p(1 - G(w))$, $t = 1, 2, ..., T$. If there are recall opportunities, the model is a little more complicated. Let $\theta^{t-\tau}$ be the probability that an offer which was received in period τ can be recalled in period t when $\tau < t$. Accordingly, θ is the probability that an offer made in the preceding period, $\tau = t - 1$, can be recalled in the current period. The probability that job search in period t yields a wage offer in period t that is less than or equal to w is $1 - p(1 - G(w))$, where the explanation is similar to that for $F_1(w)$. The probability that job search in period τ (where τ < t) does *not* yield a wage offer greater than w, that is available in period t, is

$$1 - p\theta^{t-\tau}(1 - G(w)) = 1 - p + p(1 - \theta^{t-\tau}) + p\theta^{t-\tau} G(w) \tag{4}$$

The right-hand-side of equation (4) is the probability that no offer is received in period τ (which is $1 - p$), plus the probability than an offer is received in period τ but that it is not available in period t (which is $p(1 - \theta^{t-\tau})$, plus the probability that an offer is received in period τ and it is available in period t and it does not exceed w (which is $p\theta^{t-\tau}G(w)$).

The probability that the highest wage offer available in period t from job search in periods $\tau = 1, 2, ..., t$, is less than or equal to w is

$$F_t(w) = \prod_{\tau=2} [1 - p\theta^{\tau-1}(1 - G(w))][1 - p(1 - G(w))], \text{ for } t = 2, 3, \quad T \tag{5}$$

Equations (3) and (5) are the probability distribution functions for the highest available wage offers in period 1 and periods 2 through T. Perfect recall is the special case where $\theta = 1$. Stochastic recall is where $0 < \theta < 1$.

The discrete density function associated with $F_t(\cdot)$ is

$$f_t(w) = F_t(w) - F_t(w - 1), \text{ for } t = 1, 2, ..., T. \tag{6}$$

The discrete density function $f_t(\cdot)$ will be used to derive optimal reservation wages below.

Assume that if an offer is accepted in period t then wage payments begin in period $t + 1$ and continue through period $T + 1$. Let w_t be a possible reservation wage for period t. Define $I_t(w_t)$ as the time period-one discounted value of the expected utility of income from job search in period t. We consider the case in which a net subsidy to search in period t, denoted s_t, is paid in period t whether or not an offer is accepted. Thus, if the agent searches in period t, he receives the net subsidy, s_t, with certainty. With reservation wage w_t, the probability of not having an acceptable offer at time t is $F_t(w_t - 1)$. If an acceptable offer is not available in period t, then the agent continues the search in period $t + 1$, with discounted expected utility $I_{t+1}(w_{t+1})$. Therefore, the expected utility from job search satisfies the equation of motion,

$$I_t(w_t) = (1+r)^{-(t-1)}\, u(s_t) + R_t \sum_{w_t}^{w^h} u(w) f_t(w) + F_t(w_t - 1) I_{t+1}(w_{t+1}), \tag{7}$$

where $u(\cdot)$ is the agent's von Neumann-Morgenstern utility function normalized such that $u(0) = 0$, r is the rate of interest, and

$$\begin{aligned} R_t &= \sum_{\tau=t+1}^{T+1} (1+r)^{-(\tau-1)} \\ &= \sum_{\tau=t}^{T} (1+r)^{-\tau}. \end{aligned} \tag{8}$$

We define the *concave* model as the one in which $u(\cdot)$ is assumed to be strictly increasing and concave. The *linear* model is the special case where $u(\cdot)$ is linear. Thus, in the concave model the agent is assumed to be risk averse or risk neutral while, in the linear model, he is assumed to be risk neutral.

Optimal reservation wages can be derived by backward recursion in the following way. The returns to search are zero after the end of the search horizon; hence

$$I_{T+1}(w_{T+1}) = 0, \text{ for all } w_{T+1}. \tag{9}$$

Equations (7) and (9) imply

$$I_T(w_T) = (1+r)^{-(T-1)} u(s_T) + R_T \sum_{w_T}^{w^h} u(w) f_T(w).$$

$I_T(\cdot)$ is decreasing in w_T; therefore the optimal reservation wage in period T equals zero:

$$w_T^* = 0.$$

Optimal reservation wages for other periods are derived as follows.[1] Collecting terms in equation (7), noting that

$$F_t(w_t - 1) = 1 - \sum_{w_t}^{w^h} f_t(w),$$

and setting $w_{t+1} = w_{t+1}^*$ yields

$$I_t(w_t) = I_{t+1}(w_{t+1}^*) + (1+r)^{-(t-1)}\, u(s_t) + \sum_{w_t}^{w^h} [R_t\, u(w) - I_{t+1}(w_{t+1}^*)] f_t(w). \tag{12}$$

The optimal reservation wage for period t, w_t^*, is the discrete value of w_t that maximizes $I_t(w_t)$. The bracketed term in equation (12) will be positive, zero, or negative as $w \gtreqless u^{-1}(I_{t+1}(w_{t+1}^*)/R_t)$. Therefore, the optimal reservation wage, w_t^*, $t = 1, 2, ..., T-1$, is the discrete value that is not less than, and closest to, the greater of 0 and $u^{-1}(I_{t+1}(w_{t+1}^*)/R_t)$.

IV. EXPERIMENTAL DESIGN

Each experiment for a single subject is divided into four parts. Part I consists of two baseline trials without the possibility of recall. Part II consists of two trials with perfect recall opportunities. Part III consists of two trials without recall. Part IV consists of two trials with stochastic recall; in these trials an offer received in period τ has a probability of being available in period t, where $\tau < t$, equal to $(3/4)^{t-\tau}$, that is, $\theta = 3/4$.

Each trial contains 20 periods (the finite horizon). During each period in which a trial is in progress, a subject is asked to record either an S (for "stop") or a C (for "continue"). If a subject chooses C during any period (except the last one) in a trial then the trial continues on to the next period. That is, the C response indicates that the subject has decided to continue the search, which means that no offer is accepted in that period. The S response indicates that the subject has decided to stop the search and accept an available offer. Of course, the first S recorded in a trial ends it, and the subject then begins the next trial in the experiment.

In all parts of the experiments the probability of receiving an offer in any period, p, is $1/2$ and the conditional probability that the amount of an offer is any one integer from 1 through 10 is $1/10$. There is no net subsidy to search and the rate of interest is zero.

Procedures for conducting these experiments are as follows. During each period in which a trial is in progress, a subject first draws a ball from urn 1 which holds two white balls and two black balls. If a subject draws a black ball from urn 1, he or she does not receive an offer during that period. If a subject draws a white ball from urn 1, he or she does receive an offer, the amount of which will be determined by a subsequent draw from urn 2. Urn 2 is a bingo cage that contains ten balls that are distinctly numbered with the integers from 1 to 10.

A third urn is used in the Part IV trials with stochastic recall (but not in the other parts of the experiment). Urn 3 holds 3 white balls and 1 black ball. If a subject indicates that he would like to recall an offer received one period earlier, he is permitted to make one draw from urn 3. If he draws a white ball then the offer is available but if he draws a black ball then the offer is not available. If the subject desires to recall an offer received two periods earlier, he must draw two white balls (with replacement between draws) and no black balls from urn 3 in order to be allowed to recall the offer. An offer received τ periods earlier is only available if a subject draws white balls (with replacement) and no black balls in τ draws from urn 3.

At the beginning of an experiment, the ten bingo balls for urn 2 are placed on a table in ascending order of the numbers on them. The subject is asked to inspect the numbers on the balls and then to put them in the bingo cage (urn 2). The subject is also asked to inspect the contents of urn 1 before the start of each trial. In Part IV, the subject is also asked to inspect the contents of urn 3. In addition, the subjects are informed that they are welcome to reinspect the contents of any of the containers at any time during the experiment. Furthermore, after each draw from any of the containers the experimenter brings to the subject's attention the fact that he or she is returning to the urn the ball that was drawn. Finally, a subject is occasionally shown the contents of urn 1 between draws even if he does not request it, most usually after he has drawn a run of black balls.

Complete subject instructions, including some sample record sheets, are contained in Appendix 1. In writing these instructions, we were careful to avoid the use of terms, such as "job" or "employment" or "unemployment," that might have motivational effects on subjects that are distinct from the motivational effects of the monetary earnings from participation in an experiment. The reason for this is to ensure that the experiment is *controlled*. The methodological views that lead to our decision to avoid the use of emotive terms, and to other features of the experimental design, are explained in section I.

There were 30 distinct subjects. Each subject participated in an experiment at a distinct time and had no contact with the other subjects that is known to the experimenters. The search model developed above can be used to calculate the expected payoff to a risk neutral, perfectly-optimizing subject from participating in one of these experiments. The unconditional expected payoff is $17.82. The expected payoff conditional on the actual draws of the random variables is $16.72. The actual average subject payoff was $15.76, with a range from $12.25 to $18.75. Experiments lasted 41 minutes on average. The fastest subject completed the experiment in 30 minutes and the slowest subject took 65 minutes. The subjects were University of Arizona undergraduates with no known previous experience with manually-run, individual choice experiments. Some of the subjects did have experience with computerized laboratory market experiments.

V. EXPERIMENTAL RESULTS AND TESTS OF THE MODELS

We report results for the 30 distinct subjects who participated in the experiments. Each subject participated in the two initial baseline trials, two perfect recall trials, two subsequent baseline trials, and two stochastic recall trials; hence we have 60 observations for each baseline and treatment.

The subjects exercised recall in 9 out of the 60 search terminations in the perfect recall trials. They successfully exercised recall in 2 out of the 60 search terminations in the stochastic recall trials. Furthermore, 10 out of the (total of) 11 recalled offers were: (a) less than the risk neutral theoretical reservation wages when they were initially rejected; and (b) still less that the risk neutral reservation wages when they were accepted. Thus, all 10 of these offer recalls are consistent with the decreasing reservation wages of *risk averse* optimal agents. The other (single) recalled offer was equal to the (same) risk neutral reservation wage both when it was initially rejected and when it was subsequently accepted. This offer recall is not consistent with the reservation wages of an optimal agent with any given attitude toward risk.

Table 1 presents the overall success rate of the linear (risk neutral) and concave (risk averse or risk neutral) models in predicting the subjects' search durations. Overall, the subjects terminated search in the periods predicted by the linear model in 61% of the 240 trials. This is quite a bit lower than the risk neutral model's 76% success rate in the experiments that we conducted without any opportunities for recall of past offers (Cox and Oaxaca, 1992a, Table 1).[2] The first baseline and second baseline risk neutral success rates of 73% and 75% are lower than most of the baseline success rates for the linear model reported in Cox and Oaxaca (1989). However, the really dramatically lower success rates that occur are in the perfect recall trials. The 35% success rate of the linear model for the perfect recall trials is about one-half the *minimum* success rate observed in all earlier experiments. The 60% success rate of the linear model in the stochastic recall trials is, perhaps not surprisingly, between the perfect recall and no recall (or baseline) success rates.

The success rates for the concave model vary between 93% (in the first baseline) and 98% (in the perfect recall trials). These very high success rates are comparable to our earlier 1,080 observed search terminations in experimental trials without recall opportunities (Cox and Oaxaca, 1992a, Table 1).

More can be learned about the models' performance from the mean search duration data in Table 2. In all four treatments the observed mean search duration is less than the risk neutral model's predicted search duration conditional on the draws from the wage offers distribution. This is consistent with risk aversion. Comparing the figures in the first three rows of Table 2 reveals why the linear model's success rate (reported in Table 1) for the perfect recall trials is so much lower than it is for other treatments. Although mean observed search duration is higher in the perfect recall trials than in either baseline, as predicted by theory, the recall-induced increase in duration is not nearly as large as predicted by the risk neutral model. For example, the risk neutral model predicts that the subjects will search on average 1.64 (= 6.18 – 4.54) periods longer in the perfect recall trials than in the first baseline trials but the observed difference of 0.81 (= 4.98 – 4.17) periods is less than one-half the predicted difference. Subjects behave as if they are more risk averse in the perfect recall trials than in the baseline trials. But the theory assumes that risk attitude is a characteristic of the agent that is independent of experimental treatments. Furthermore, the wage offers distribution is less risky in the presence of (certain or stochastic) recall opportunities than in their absence.[3] Hence, we conclude that the theory is failing to predict how subjects will respond to the introduction of perfect recall opportunities.

Let $\overline{D}^i$ be the observed average search duration for treatment i, and let $\overline{D}^{in}$ be the risk neutral (linear) model's theoretical average search duration for treatment i

Table 1. Theoretically consistent search terminations

	Linear Model		*Concave Model*	
Treatment	*Ratio*	*%*	*Ratio*	*%*
	44/60	73.3		93.3
	21/60	35.0		98.3
	45/60	75.0		96.7
	36/60	60.0		95.0
	146/240	60.8		95.8

conditional on the actual draws from the wage offers distribution, where *i* denotes first baseline, perfect recall, second baseline, or stochastic recall. A hypothesis test for the linear model is given by

$$H_0 : E\,(\overline{D}^i - \overline{D}^{in}) = 0,\, H_1 : \sim H_0.$$

A hypothesis test for the concave model is given by

$$H_0 : E\,(\overline{D}^i - \overline{D}^{in}) \leq 0,\, H_1 : \sim H_0.$$

Table 3 presents the p values for mean/median tests of the linear and concave models. The means test is simply the matched-pairs test with the normal distribution. The nonparametric test is the Fisher sign test for the median. At conventional significance levels, the linear model would be rejected for all of the experimental treatments except the first baseline trials. Here the means test would not lead to rejection of the linear model at conventional levels of significance, although the Fisher sign test would lead to rejection at the 7.6% level of significance. In contrast, the concave model would not be rejected for any of the treatments.

Table 4 presents the p values for mean/median test comparisons of the baseline, perfect recall, and stochastic recall trials. The first column of Table 4 lists the treatments that are being compared in each row of the table. The second and third columns report the results of tests for differences between treatments, that is,

Table 2. Mean search duration

Treatment	*Linear Model*	*Observed*
First Baseline		
Perfect Recall		
Second Baseline		
Stochastic Recall		

Table 3. Mean/Median search duration tests of the linear and concave models (p values)

	H_0: $E(\overline{D}^i - \overline{D}^{in}) = 0$		H_0: $E(\overline{D}^i - \overline{D}^{in}) \leq 0$	
Treatment	*Means Test*	*Sign Test*	*Means Test*	*Sign Test*

$$H_0 : E(\overline{D}^i - \overline{D}^j) = 0, H_1 : \sim H_0.$$

The p values in columns 2 and 3 are helpful in answering the question: Is there a significant difference between the observed average search durations in two treatments? Any such differences would be attributed to the joint effects of the experimental treatments and the variation in the draws from the wage offers distribution. The fourth and fifth columns report the results of tests for differences between risk neutral theoretical predictions conditional on draws from the wage offers distribution, that is,

$$H_0 : E(\overline{D}^{in} - \overline{D}^{jn}) = 0, H_1 : \sim H_0.$$

The p values in columns 4 and 5 address the question: Does risk neutral search theory predict a significant difference between average search durations in two treatments? This information provides a way of determining how well the experimental design succeeded in distinguishing between the theoretical treatment effects given the offers drawn in the experimental trials. The sixth and seventh columns report the results of tests for differences between deviations of observed average search durations and deviations of risk neutral theoretical average search durations conditional on the draws, that is,

$$H_0 : E\,[(\overline{D}^i - \overline{D}^j) - (\overline{D}^{in} - \overline{D}^{jn})] = 0, H_1 : \sim H_0.$$

The p values in columns 6 and 7 are relevant to answering the question: Is the observed difference in average search durations in two treatments significantly different from the difference predicted by the risk neutral model? Consistency with the risk neutral model requires that these differences not be statistically significant.

The p values in the first row of Table 4 do not lead us to reject any of the null hypotheses; thus we conclude that the subjects did "return-to-baseline" after the intervening perfect recall treatment. From the second row of Table 4 we conclude that there was no significant difference in average search duration between the perfect recall and first baseline trials (p = 0.244 or 0.285); however, theory predicts

Table 4. Mean/Median Search Duration Tests for Treatment Effects (p values)

	$H_0 : E(\overline{D}^i - \overline{D}^j) = 0$		$H_0 : E(\overline{D}^{in} - \overline{D}^{jn}) = 0$		$H_0: E[(\overline{D}^i - \overline{D}^j) - (\overline{D}^{in} - \overline{D}^{jn})] = 0$	
Comparison	*Means Test*	*Sign Test*	*Means Test*	*Sign Test*	*Means Test*	*Sign Test*
First Baseline/Second Baseline						
Perfect Recall/First Baseline						
Perfect Recall/Second Baseline						
Stochastic Recall/First Baseline						
Stochastic Recall/Second Baseline						
Perfect Recall/Stochastic Recall						

a difference (p = 0.012 or 0.003) and the observed difference is significantly different from the predicted difference (p = 0.043 or 0.005). Taken together, the p values in row 2 support the conclusion that the subjects did not respond to the availability of perfect recall as predicted by theory.

The third row of Table 4 presents the p values for comparison of the perfect recall trials with the second baseline trials. At the 1% level of significance we would not reject any of the null hypotheses; that is, we would conclude that there was no significant difference between the perfect recall and second baseline trials and that this is predicted by theory given the draws from the wage offers distribution. Our conclusions would be somewhat different if we based them on a 10% significance level. Based on the means test we would conclude that the risk neutral model predicts a significant difference between the perfect recall and second baseline trials (p = 0.064), although none is observed (p = 0.267, p = 0.138). Based on the Fisher sign test we would conclude that there is a significant difference between perfect recall and second baseline (p = 0.031) even though none is observed (p = 0.691) and none is predicted by theory given the draws from the wage offers distribution (p = 0.431).

The p values in the fourth and fifth rows of Table 4 would not lead one to reject any of the null hypotheses; thus, we conclude that the risk neutral model predicts no significant difference between the baseline and stochastic recall trials and that is what we observed.

The last row of Table 4 presents a second case in which the inferences are sensitive to the distinction between tests with 1% and 10% significance levels. Based on the means test, at the 1% level of significance we would not reject any of the null hypotheses. On the other hand the Fisher sign test would lead us to conclude that the risk neutral model predicts a significant difference between the perfect recall and stochastic recall trials (p = 0.003), although none is observed (p = 0.022, p = 0.189). However, from the means test at the 10% level of significance we would conclude that the risk neutral model predicts a significant difference between perfect recall and stochastic recall trials (p = 0.058), although that is not observed (p = 0.182, p = 0.119). Furthermore, the Fisher sign test would lead us to conclude that there is no significant difference between perfect recall and stochastic recall (p = 0.189) even though there is an observed difference (p = 0.022) and the risk neutral model predicts a significant difference (p = 0.003) conditional on the draws from the wage offers distribution.

Let $F_i(x)$ and $F_{in}(x)$ represent the population distribution function and the risk neutral distribution function (conditional on the wage offers) of search terminations by period of search for treatment *i*. We use the Kolmogorov-Smirnov (KS) goodness-of-fit tests to test the linear and concave models on the basis of the full distribution of search terminations. The corresponding linear and concave model hypotheses are given by

$$H_0 : F_i(x) = F_{in}(x),\ H_1 \quad F_i(x) \neq F_{in}(x), \text{ and}$$

Table 5. Kolmogorov-Smirnov Goodness-of-Fit tests of the linear and concave models (p values)

Treatment	$H_0 : F_i(x) = F_{in}(x)$	$H_0 : F_i(x) > F_{in}(x)$
First Baseline	$p > 0.20$	$p > 0.10$
Perfect Recall	$p < 0.01$	$p > 0.10$
Second Baseline	$0.05 < p < 0.10$	$p > 0.10$
Stochastic Recall	$p < 0.01$	$p > 0.10$

$$H_0 : F_i(x) \geq F_{in}(x), H_1 : F_i(x) < F_{in}(x) .$$

Table 5 reports the p value ranges for the KS tests. At the 10% level of significance we would reject the risk neutral model for every treatment except the first baseline trials. In contrast, we would not reject the concave model for any treatment.

VI. CONCLUSIONS

Cox and Oaxaca (1989) reported experimental tests of a finite horizon sequential search model in which there were no opportunities for recalling previous wage offers. That paper reported experiments which included variations in the interest rate, search costs and subsidies, riskiness of the wage offers distribution, the probability of receiving an offer, and the length of the search horizon. On the basis of several parametric and nonparametric tests applied to search duration data, we concluded that the risk neutral model survived the experimental treatments and that the concave model survived them remarkably well. These results caused us to look further for the boundaries of predictive failure of finite horizon search theory.

We are subjecting search theory to further tests in two different ways: (a) complicating the search environment; and (b) changing the data used in the tests in the original search environment. We have pursued research strategy (a) with two different experimental designs. The experiments first reported in the present paper complicate the search environment by introducing opportunities for certain and stochastic recall of previous offers. Another paper, in progress, complicates the environment by introducing search from an unknown distribution with treatments in which the theoretical reservation wage property alternatively exists or does not exist. We have pursued research strategy (b) by subjecting the theory to direct tests using reservation wage data (Cox and Oaxaca, 1992a). That approach began with a test (for an experimental artifact) to check whether the experimental technique needed to elicit reservation wages (precommitment) was itself a significant treatment (Cox and Oaxaca, 1992b). We concluded that the reservation wage data was usable and that the direct tests with that data were stronger than the (relatively weak) tests based on search duration. Direct tests with reservation wage data lead to rejection of the risk neutral model but they did not imply rejection of either the general concave search model or a specific (logistic) concave model.

A *concave* search model represents search by an agent that is risk neutral or risk averse. Thus the direct tests with reservation wages required us to postulate that experimental subjects may be risk averse in order to maintain consistency between reservation wage data and theoretical reservation wage paths. Furthermore, in all of the many experimental treatments in our research program, when there have been deviations from the predictions of the *risk neutral* search model they have been overwhelmingly on the side of early search terminations. The data have been subject to a variety of parametric and nonparametric tests for consistency with the concave model. As discussed above, there was only one instance in which the concave model was rejected. This was the case of the horizon treatment reported in Cox and Oaxaca (1989). In these experiments there were six non-baseline treatments and four baseline treatments. The median percentage of deviations from the predictions of the risk neutral model that were early terminations was 74%. With the exception of the horizon treatment in which only 20% of the deviations were early, the early search termination rate ranged from 58% to 94%. In Cox and Oaxaca (1992a) there were 2 baseline and 2 non-baseline treatments run both with and without precommitment. The median percentage of deviations that were early terminations in the no-precommitment experiments was 64%. The early deviation percentages ranged from 59% for the first baseline treatment to 71% for the second baseline experiment. In the precommitment experiments, the median early deviation rate was 90%. The early deviation percentages ranged from 63% in the second baseline treatment to 92% for the subsidy and probability treatments. The median early deviation rate for the recall experiments is 88%. The early deviation rate ranges from 75% for the first baseline treatment to 97% for the perfect recall treatment.

From the preponderance of evidence, there can be no doubt as to the one-sided nature of deviations from the predictions of the risk neutral search model. In seeking an explanation for this phenomenon, it is natural for an economist to appeal to risk aversion. Indeed, we are unaware of any alternative model that predicts early search terminations and that has testable implications that could distinguish it from the concave search model. However, could the incidence of early search terminations be an artifact of the experimental design? An examination of the risk neutral solution to the perfect recall reservation wage path in Table A1 of Appendix 2 reveals that the optimal reservation wage is 9 for the first 13 periods. This does not leave much room for random errors to lie above the risk neutral reservation wage path. As Table 1 reveals, the risk neutral model fares rather poorly for the perfect recall treatment. Only 35% of the search terminations occurred when predicted by the risk neutral model. This is by far the worst outcome for the risk neutral model that we have found thus far in our research program. At the same time, the perfect recall treatment reveals the highest early deviation rate that we have observed. The risk neutral reservation wage path under stochastic recall is seen from Table A1 of Appendix 2 to be also relatively high in the early periods (the optimal risk neutral reservation wage is 8 for the first 12 periods). Only 60% of the search terminations (Table 1) for this treatment coincided with the predictions of the risk neutral model and the early deviation rate is 88% for this treatment. One can see that the higher the reservation wage path, the greater would be the tendency for *random* deviations from the risk neutral model's predictions to lie on the side of lower reservation wages and early search terminations.

At the same time, the evidence strongly suggests that experimental design artifacts cannot fully account for the well-documented tendency for search to terminate earlier than predicted by the risk neutral model. Table A2 of the Appendix 2 shows that the optimal risk neutral reservation wage paths for the interest, risk, cost, and probability treatments begin at 6. In the cases of the cost and probability treatments, the optimal reservation wage drops to 5 in period 3. Therefore, these wage paths begin near the median of the conditional wage offers distribution. The lowest early deviation rate detected for these treatments was 68%, for the risk treatment. The highest early deviation rate without precommitment was 94%, for the probability treatment reported in Cox and Oaxaca (1989). From these results it is clear that deviations still tend to be one-sided from what one would expect from purely random errors.

This suggests the question as to whether there is any independent data that supports the conclusion that many of our subjects are risk averse for the level of payoffs usually used in laboratory experiments. Cox, Smith, and Walker (1988) reports auction market experiments for more than 200 subjects from the same subject pool of University of Arizona undergraduate students that we use in our search experiments. They report that almost all of their subjects' bidding behavior is consistent with a risk neutral or risk averse (i.e., "concave") bidding model and, furthermore, that the estimated bid functions of 68% of them are significantly risk averse. This risk aversion explanation of bidding behavior is corroborated by the experiments reported by Harlow and Brown (1990a,b) using subjects from the same subject pool. They report results for more than 100 subjects who bid in auctions, completed two psychological survey tests, and submitted to biochemical blood tests. Harlow and Brown conclude that the psychological survey tests support the risk aversion explanation of bidding and, furthermore, that there is a measurable biochemical basis for their interpretation of subjects' experimental economic behavior.

We interpret the test results in the present paper as follows. The risk neutral model's search duration predictive accuracy is lower in the stochastic recall trials than in all of our earlier experimental trials without recall. Furthermore, the risk neutral model fails decisively to predict search duration in the perfect recall trials. In contrast, the general concave model's success rate for both stochastic and certain recall trials remains at the very high levels observed in earlier experiments without recall. But the testable implications of the general concave model are weak. For example, this concave model "succeeds" whenever subjects search *no longer* than is predicted by the risk neutral model. A closer look at the data leads to a different conclusion about the concave model. Subjects search as if they are *more* risk averse in the presence of perfect recall opportunities than in their absence. But the introduction of perfect recall *decreases* the riskiness of the wage offers distribution. Hence we conclude that the general concave model also fails to predict how subjects will respond to the introduction of perfect recall opportunities.

APPENDIX A

Instructions

General

This is an experiment in the economics of decision making. Various research support agencies have provided funds for the conduct of this research. The instructions are simple and, if you follow them carefully, you might make a considerable amount of money which will be paid to you in cash at the end of the experiment. The experiment is divided into four parts. Each part contains two trials and each trial consists of twenty periods. The instructions for Part I of the experiment are contained in these pages. You will be given the instructions for the other parts when you are ready to begin them.

Part I

In Part I of this experiment, you will participate in two trials, each of which has 20 periods. During each period of a trial, you will be asked to decide whether to Stop or to Continue that trial. If you decide to Stop a trial in some period *t*, then the experimenter will add to your cash earnings the amount written in the Earnings from Stopping column for period *t*. If you decide to Continue a trial in period *t*, you will then have the opportunity to Stop that trial in later periods when the Earnings from Stopping may be higher or lower than they were in period *t*.

Each period, Earnings from Stopping will be determined by a random process. This random process works as follows. You will first draw a ball from Container 1. If the ball you draw is black, you get 0 points for that period. If you draw a white ball from Container 1, you may then draw a ball from Container 2. The number on the Container 2 ball is then your Points Drawn for that period. Your Earnings from Stopping for that period are calculated by multiplying your Points Drawn by the Money Conversion Factor for that period. You will be informed of the amounts of these Money Conversion Factors for all 20 periods at the beginning of each trial.

The following examples should help to clarify these instructions.

Specific Examples

Consider the sample Trial A record sheet. The first line of the sample Trial A record sheet contains the sentence: "Container 1 contains 2 white balls and 2 black balls." This sentence introduces the following feature of the experiment. You will be presented with Container 1, a coffee can, which contains 2 white balls and 2 black balls. You will be asked to inspect the contents of Container 1 to verify that it does, in fact, contain 2 white balls and 2 black balls. You will then be asked to draw one ball from Container 1 without looking inside. (Note that your chance of drawing a white ball is 2 out of 4 and your chance of drawing a black ball is 2 out of 4.) If you draw a white ball from Container 1, you may then proceed to draw from Container 2. If you draw a black ball from Container 1, you may *not* draw from Container 2 during that period of the trial; instead, you may either Stop the

Trial *A*

Container 1 contains 2 white balls and 2 black balls.
Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10.

A Period Number	*B* Points Drawn	*C* Money Conversion Factor (cents/point)	*D* Earnings from Stopping (col. *B* × col. *C*)	Your Decisions
1		30		
2		28.5		
3		27		
4		25.5		
5		24		
6		22.5		
7		21		
8		19.5		
9		18		
10		16.5		
11		15		
12		13.5		
13		12		
14		10.5		
15		9		
16		7.5		
17		6		
18		4.5		
19		3		
20		1.5		

In this trial, you earned _____ cents.

trial at this point or Continue on to the next period of the trial and draw again from Container 1, after having replaced the black ball.

Suppose that, in some period, you draw a white ball from Container 1 and thus may draw from Container 2. What does Container 2 contain? The second and third lines of the sample Trial A record sheet contain the sentence: "Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10." This sentence now conveys the information that Container 2 contains a total of ten balls. One and only one ball has the number 1 on it, one and only one ball has the number 2 on it, and so on through one and only one ball having the number 10 on it. (Note that this means that your chance of drawing any whole number from 1 through 10 from Container 2 is 1 out of 10.) You will be asked to inspect the contents of Container 2, which is a bingo cage, to verify that it does, in fact, contain one each of balls with whole numbers from 1 through 10.

Let's continue the explanation by considering an example of what might actually happen during one of the two trials in Part I of the experiment. We will continue to make use of the sample Trial A record sheet. This record sheet now contains the information that Container 1 contains 2 white balls and 2 black balls and that Container 2 contains one each of balls with whole numbers from 1 through 10. The record sheet also contains a Money Conversion Factor for each of the 20 periods of the trial. The Money Conversion Factor begins at 30 cents per point and decreases by 1.5 cents per point each period.

Let's now consider what might happen in sample Trial A, beginning with period 1. Suppose that in period 1 you draw a white ball from Container 1. You may then draw a ball from Container 2 *in period 1*. Next, suppose that the ball that you draw from Container 2 has the number 6 on it. PLEASE WRITE the number 6 in the period 1 row of the Points Drawn column of the record sheet. Next, multiply the 6 points by the Money Conversion Factor of 30 cents per point to get 180 cents. Please write the number 180 in the period 1 row of the Earnings from Stopping column. You now must decide whether to Stop sample Trial A in period 1 or to Continue this trial. If you were to decide to Stop in period 1, your earnings for this trial would be 180 cents and you would then begin the next trial. If your period 1 decision is to Continue sample Trial A, you then proceed to period 2. Note that, in continuing from period 1 to period 2, the Money Conversion Factor decreases from 30 to 28.5 cents per point.

Suppose that your period 1 decision is to Continue. Please write the letter C for Continue in the period 1 row of the Your Decisions column. You may now proceed to period 2. First, you return the white ball drawn in period 1 to Container 1 and the 6 ball to Container 2. Then you may make your draws for period 2. Suppose that you draw the black ball from Container 1 in period 2. Then you may not draw from Container 2 in that period. Please write the number 0 in the period 2 row of the Points Drawn and Earnings from Stopping columns. You now must decide whether to Stop in period 2 or to Continue. If you were to decide to Stop in period 2, your earnings for this trial would be 0 cents and you would then begin the next trial. If your period 2 decision is to Continue, you then proceed to period 3. Note that, in continuing from period 2 to period 3, the Money Conversion Factor decreases from 28.5 to 27 cents per point.

Suppose that your period 2 decision is to Continue. Please write the letter C in the period 2 row of the Your Decisions column. You now proceed to period 3. First, you return the black ball drawn in period 2 to Container 1. Now you may make your draws for period 3. Suppose that you draw a white ball in period 3. You may then draw from Container 2 in period 3. Suppose that your period 3 draw from Container 2 yields the ball with the number 8 on it. Please write the number 8 in the period 3 row of the Points Drawn column. Then multiply the 8 points by the Money Conversion Factor of 27 cents per point to get 216 cents. Please write the

TRIAL *B*

Container 1 contains 2 white balls and 2 black balls.
Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10.

A	*B*	*C*	*D*	*E*
Peroid Number	*Points Drawn*	*Money Conversion Factor (cents/point)*	*Earnings from Stopping (col. B × col. C)*	*Your Decisions*
1	3	30	90	C
2	6	28.5	171	C
3	0	27	0	C
4	4	25.5	102	C
5	5	24	120	C
6	0	22.5	0	C
	7	21	147	S
8		19.5		
		18		
10		16.5		
11		15		
12		13.5		
13		12		
14		10.5		
15		9		
16		7.5		
17		6		
18		4.5		
19		3		
20		1.5		

In this trial, you earned 147 cents.

number 216 in the period 3 row of the Earnings from Stopping column. You must now decide whether to Stop this trial in period 3 or to Continue this trial.

Suppose that you decide to Stop this trial in period 3. Please write the letter S for Stop in the period 3 row of the Your Decisions column. Since you have decided to Stop, this ends Trial A. Your earnings in Trial A are the Earnings from Stopping in period 3, which are 216 cents. Please write the number 216 in the blank space in the sentence at the bottom of the page. That sentence then becomes: "In this trial, you earned *216* cents."

You have now completed sample Trial A and are ready to begin sample Trial B. Consider the sample Trial B record sheet. The numbers in the Points Drawn column imply that you drew white balls from Container 1 in periods 1, 2, 4, 5 and 7, and that you drew black balls in periods 3 and 6. Thus, you were not eligible to draw from Container 2 in periods 3 and 6 but did draw numbered balls from Container 2 in periods 1, 2, 4, 5 and 7.

The numbers drawn from Container 2 are recorded in the Points Drawn column. These Points Drawn numbers have been multiplied by the corresponding Money Conversion Factors to get the numbers written in the Earnings from Stopping column. The information recorded in the Your Decisions column indicates that you decided to Continue the trial in periods 1, 2, 3, 4, 5 and 6, and that you decided to Stop the trial in period 7. The money earnings that would result from these decisions are recorded at the bottom of the record sheet. Note that, in this example, you would have earned more if you had stopped the trial in period 2 rather than period 7 even through you drew more points in period 7 than in period 2. Why? Because the decrease in the Money Conversion Factor (from 28.5 to 21) more than offset the increase in Points Drawn (from 6 to 7) in its effect on the Earnings from Stopping.

To calculate the amount that you earn in this experiment, we will first add the amounts written at the bottom of the record sheets *for Trials 1–8*. We will next round that amount *up* to the closest multiple of 25 cents. This amount will be paid to you in cash at the end of the experiment.

DO YOU HAVE ANY QUESTIONS?

You are now ready to begin the experiment. We will first conduct Trial C, which is a practice trial with *no* money payoff. Then we will begin the trials with money payoffs.

Part II

In Part II of this experiment, you will participate in two trials, each of which has 20 periods. Part II differs from Part I in only one way: whether or not the number of Points Drawn in an earlier period can be substituted for the number of Points Drawn in a later period in calculating the Earnings from Stopping for the later period. In Part I, if you decided to Continue in some period then the points that you drew in that period were not available to calculate Earnings from Stopping in any later period. In contrast, if in Part II you decide to Continue in some period then

TRIAL *C*

Container 1 contains 2 white balls and 2 black balls.
Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10.

A	*B*	*C*	*D*	*E*
Period Number	*Points Drawn*	*Money Conversion Factor (cents/point)*	*Earnings from Stopping (col. B × col. C)*	*Your Decisions*
1		30		
		28.5		
3		27		
4		25.5		
5		24		
6		22.5		
7		21		
8		19.5		
9				
10		16.5		
11		15		
12		13.5		
13		12		
14		10.5		
15		9		
16		7.5		
17		6		
18		4.5		
19		3		
20		1.5		

In this trial, you earned cents.

the points that you drew in that period can be used instead of the Points Drawn in *any* later period to calculate Earnings from Stopping in the later period. However, any calculation of Earnings from Stopping for a period must use the Money Conversion Factor *for that period*, regardless of when the points were drawn.

The following example should help to clarify these instructions.

Specific Example

Consider the sample Trial D record sheet. The numbers in the Points Drawn column imply that you drew white balls from Container 1 in periods 2, 3 and 6, and that you drew black balls in periods 1, 4 and 5. Thus, you were not eligible to draw from Container 2 in periods 1, 4 and 5, but did draw numbered balls from Container 2 in periods 2, 3 and 6.

TRIAL *D*

Container 1 contains 2 white balls and 2 black balls.
Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10.

A *Period Number*	*B* *Points Drawn*	*C* *Points Chosen*	*D* *Money Conversion Factor (cents/point)*	*E* *Earnings from Stopping (col. C × col. D)*	*F* *Your Decisions*
1	0	0	30	0	C
2	3	3	28.5	85.5	C
3	7	7	27	189	C
4	0	7	25.5	178.5	C
5	0	3	24	72	C
6	5	7	22.5	157.5	S
7			21		
8			19.5		
9			18		
10			16.5		
11			15		
12			13.5		
13			12		
14			10.5		
15			9		
16			7.5		
17			6		
18			4.5		
19			3		
20			1.5		

In this trial, you earned 157.5 cents.

The numbers drawn from Container 2 are recorded in the Points Drawn column. Note that there is a new column labeled Points Chosen. The addition of this column reflects the additional choices that are now available to you. In calculating the Earnings from Stopping for any period, you may now choose to use the Points Drawn in that period or the Points Drawn in *any earlier period.*

The numbers in the Points Chosen column indicate that you chose 3, rather than 0, in period 2 and that you chose 7, rather than 0 or 3, in periods 3 and 4. Also, you chose 3, rather than 0 or 7, in period 5. The C's in Your Decisions column indicate that you decided to Continue in periods 1–5.

Let's now consider what happened in period 6 of sample Trial D. The number 5 in the Points Drawn column indicates that in period 6 you drew a white ball from Container 1 and the 5 ball from Container 2. Since you had drawn 0, 3, and 7 in earlier periods, these numbers, in addition to 5, could have been your Points Chosen for period 6. In this case, you selected 7 as your Points Chosen. You then calculated the Earnings from Stopping using the Money Conversion Factor *for period 6* (157.5 = 7 × 22.5) and decided to Stop in that period.

Note the following features of the numbers in the sample Trial D record sheet. By selecting 7 as your Points Chosen for period 6, rather than the 5 that you drew in that period, you increased your earnings from 112.5 (= 5 × 22.5) to 157.5 (= 7 × 22.5). However, if you had decided to Stop in period 3 rather than period 6 you would have earned more money (189 cents) because the Money Conversion Factor was higher in period 3. In contrast, if you had been lucky enough to draw a 9 in period 6, and selected it as your Points Chosen, then you would have earned 202.5 (= 9 × 22.5) from deciding to Stop in period 6.

DO YOU HAVE ANY QUESTIONS?

You are now ready to begin Part II of the experiment.

Part III

Part III is exactly the same as Part I. You will participate in two trials, each of which has 20 periods. As in Part I, the only points that you can use to calculate your Earnings from Stopping in a period are the points (if any) that you draw *in that period.*

DO YOU HAVE ANY QUESTIONS?

You are now ready to begin Part III of the experiment.

Part IV

In Part IV of this experiment, you will participate in two trials, each of which has 20 periods. Part IV is similar to Part II but somewhat different. In Part II, if you decided to Continue in a period then the points that you drew in that period could always be used instead of the Points Drawn in any later period to calculate Earnings from Stopping in the later period. In comparison, if in Part IV you decide to Continue in a period then the points that you draw in that period *might* be available for use in calculating Earnings from Stopping in some later period. Whether or not the Points Drawn in one period will be available in some specific later period will be determined by your draws from Container 3 (which is a container that we did *not* use in Parts I–III of the experiment).

Containers I and 2 will be used in Part IV to determine Points Drawn in a period in the same way they were used in Parts I–III. Container 3 will be used in the following way. Suppose that, in some period, you decide that you do *not* want to use the Points Drawn in that period to calculate Earnings from Stopping; you decide, instead, to find out whether Points Drawn in some earlier period are available to make this calculation. If you choose the Points Drawn *one* period earlier, you would have to draw a white ball from Container 3 *one* time in order for them to be available. If you choose the Points Drawn *two* periods earlier, you would have to draw a white ball from Container 3 on each of your first *two* draws (with replacement of the white ball between draws) in order for the points to be available. In general, Points Drawn x periods earlier are only available if you are able to draw white balls from Container 3 on each of your first x draws (with replacement between draws).

The following example should help to clarify these instructions.

A Specific Example

Consider the sample Trial E record sheet. The three lines above the table list the contents of Containers 1, 2, and 3. As before, Container 1 contains 2 white balls and 2 black balls and Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10. The last line above the table contains the sentence: "Container 3 contains 3 white balls and 1 black ball." This means that your chance

TRIAL *E*

Container 1 contains 2 white balls and 2 black balls.
Container 2 contains one each of balls with integer (or whole) numbers from 1 through 10.
Container 3 contains 3 white balls and 1 black ball.

A Period Number	*B* Points Drawn	*C* Points Wanted	*D* Points Available	*E* Money Conversion Factor (cents/point)	*F* Earnings from Stopping (col. D × col. E)	*G* Your Decisions
	3	3	3	30	90	C
	0	3	3	28.5	85.5	C
		7	7	27	189	C
4	5	7	0	25.5	0	C
	0	7	7	24	168	S
6				22.5		
7				21		
8				19.5		
9				18		
10				16.5		
11				15		
12				13.5		
14				10.5		
15						
16				7.5		
17				6		
18				4.5		
19				3		
20				1.5		

In this t al, you earn d 168 cents.

TRIAL *F*

				E	*F*	G
Period Number	*Points Drawn*	*Points Wanted*	*Points Available*	*Money Conversion Factor (cents/point)*	*Earnings from Stopping (col. D × col. E)*	*Your Decisions*
1				30		
2				28.5		
3				27		
4				25.5		
5				24		
6				22.5		
7				21		
8				19.5		
9				18		
10				16.5		
11				15		
12				13.5		
13				12		
14				10.5		
15				9		
16				7.5		
17				6		
18				4.5		
19				3		
20				1.5		

In this trial, you earned cents.

of drawing a white ball from Container 3 is 3 out of 4. You will be asked to inspect the contents of Container 3, a plastic tub, to verify that is does contain 3 white balls and 1 black ball.

Consider, again, the sample Trial E record sheet. The numbers in the Points Drawn column imply that you drew black balls from Container 1 in periods 2 and 5 and that you drew white balls in periods 1, 3 and 4. The numbers drawn from Container 2 in periods 1, 3 and 4 are recorded in the Points Drawn column. Note that there are two new columns labeled Points Wanted and Points Available. We will next explain the meaning of the numbers in these columns.

In period 1 your number of Points Drawn was 3. Since this was the first period in the trial, no other points had been drawn. Hence, your only possible choice for the Points Wanted column was 3. Since you drew the 3 points *in period 1*, they are necessarily available *in that period.* Multiplying the 3 Points Available times the

Money Conversion Factor of 30 yields the Earnings from Stopping of 90. Your decision was to Continue.

Things get more interesting in period 2 of Trial E. The 0 in the Points Drawn column indicates that you drew a black ball from Container 1. Either the 3 Points Drawn in period 1 or the 0 Points Drawn in period 2 can now be listed in the Points Wanted column. You decided to list 3 as your Points Wanted in period 2. Since you drew the 3 in period 1, which is *one* period earlier than period 2, you must draw *one* white ball from Container 3 if it is to be available. The number 3 written in the Points Available column for period 2 indicates that you did draw a white ball from Container 3. However, after calculating your Earnings from Stopping you decided to Continue.

Now consider the period 3 row of the sample Trial E record sheet. Your number of Points Drawn was 7 in period 3. In choosing from among all Points Drawn in this trial (3, 0, and 7) you chose 7 as your Points Wanted. Since 7 was drawn *in this period* it was necessarily available. Multiplying the Points Available of 7 times the Money Conversion Factor of 27 yields the Earnings from Stopping of 189. The C in the last column indicates that you decided to Continue.

In period 4, your number of Points Drawn was 5. You decided that your number of Points Wanted was 7. Since the 7 was drawn *one* period earlier, you would have to draw a white ball *one* time from Container 3 in order to have the 7 be available. The 0 in the Points Available column indicates that you drew a black ball from Container 3.

Finally, in period 5 your Points Drawn were 0. Once again, you chose 7 as your Points Wanted. Since the 7 was drawn *two* periods earlier, you would have to make *two* successive draws of white balls from Container 3 (with replacement) to have the 7 be available. The 7 in the Points Available column indicates that you drew two white balls as follows. Your first draw was a white ball. This ball was returned to Container 3. Then you drew a second time. Your second draw was also a white ball. You then decided to Stop this trial and claim your earnings of 168 cents.

DO YOU HAVE ANY QUESTIONS?

You are now ready to begin Part IV of the experiment. We will first conduct Trial F, which is a practice trial with *no* money payoff. Then we will conduct two trials with money payoffs.

APPENDIX B

Table A1. Integer dynamic programming solutions

	Baseline		*Perfect Recall*		*Stochastic Recall*	
t	w_t^*	D_t	w_t^*	D_t	w_t^*	D_t
1	7	4.6	9	3.8	8	3.6
2	7	5.5	9	4.2	8	4.1
3	7	6.4	9	4.8	8	4.8

	Baseline		Perfect Recall		Stochastic Recall	
	w_t^*	D_t	w_t^*	D_t	w_t^*	D_t
4	7	7.3	9	5.4	8	5.6
5	7	8.1	9	6.2	8	6.5
6	6	8.8	9	7.0	8	7.4
7	6	9.8	9	7.8	8	8.3
8	6	10.7	9	8.7	8	9.2
9	6	11.6	9	9.6	8	10.2
10	6	12.5	9	10.5	8	11.1
11	6	13.3	9	11.4	8	12.1
12	5	14.1	9	12.4	8	12.9
13	5	15.0	9	13.3	7	13.7
14	5	15.8	8	14.1	7	14.7
15	4	16.6	8	15.1	7	15.6
16	4	17.4	8	16.1	6	16.5
17	3	18.2	8	17.1	6	17.5
18	3	18.9	7	18.0	5	18.3
19	2	19.6	5	19.0	4	19.2
20	0	[illegible]	0	20.0	0	20.0

Table A2. Integer dynamic programming solutions: Previous treatments

	Interest		Subsidy		Risk		Cost		Probability		Horizon	
t	w_t^*	D_t	w_t^*	D_t	w_t^*	D_t	w_t^*	D_t	w_t^*	D_t	w_t^*	D_t
3	6											
4	6											
5	6											
6	6											
7	5											
8	5											
9	5											
10	5											
11	5											
	5											
13	5											
14	4											
15	4											
16	4											
17	3											
18	3											
19	2										—	
20	0										—	

ACKNOWLEDGMENT

We are grateful to the National Science Foundation for financial support (grant numbers SES-8404915, SBR-9311351). Capable research assistance was provided by Lucy Atkinson Eldridge and Yvonne Durham. Valuable computer programming was supplied by Shawn LaMaster.

NOTES

1. If negative net search subsidies (positive net costs of search) should cause the expected present value of income from continuing search to become negative in some period $\tau + 1$, then $w_\tau^* = 0$.
2. The linear model's success rates for the four baseline treatments reported in Cox and Oaxaca (1989, Table 3) are 87%, 83%, 80%, and 73%.
3. There is a first-order stochastic dominance ordering of the wage offers distributions in the presence and absence of recall opportunities. For example, in the absence of recall opportunities the c.d.f. for wage offers is $F(w) = 1 - p(1 - G(w))$ in all periods in the search horizon. With perfect recall, the c.d.f. is $F_t(w) = [1 - p(1 - G(w))]^t$ for $t = 1, 2, ..., T$. Clearly, $F_t(w) < F(w)$ for $w^l \le w^h$ and $t = 2, 3, ..., T$.

REFERENCES

Braunstein, Yale M., and Andrew Schotter (1981), "Economic Search: An Experimental Study." *Economic Inquiry* 19: 1–25.

Braunstein, Yale M., and Andrew Schotter (1982), "Labor Market Search: An Experimental Study." *Economic Inquiry* 20: 133–144.

Cox, James C., and Ronald L. Oaxaca (1989), "Laboratory Experiments with a Finite Horizon Job Search Model." *Journal of Risk and Uncertainty* 2: 301–329.

Cox, James C., and Ronald L. Oaxaca (1992a), "Direct Tests of the Reservation Wage Property." *Economic Journal* 102: 1423–1432.

Cox, James C., and Ronald L. Oaxaca (1992b), "Tests for a Reservation Wage Effect." In *Decision Making under Risk and Uncertainty: New Models and Empirical Findings*, edited by John Geweke. Dordrecht: Kluwer Academic Publishers.

Cox, James C., Vernon L. Smith, and James M. Walker (1988), "Theory and Individual Behavior of First-Price Auctions." *Journal of Risk and Uncertainty* 1: 61–99.

Harlow, W. V., and Keith Brown (1990a), "Understanding and Assessing Financial Risk Tolerance: A Biological Perspective." *Financial Analysts Journal* 46: 50–62, 80.

Harlow, W. V., and Keith Brown (1990b), *The Role of Risk Tolerance in the Asset Allocation Process: A New Perspective.* Charlottesville, VA: Research Foundation of the Institute of Chartered Financial Analysts.

Harrison, Glenn W., and Peter Morgan (1990), "Search Intensity in Experiments." *Economic Journal* 100: 478–486.

Hey, John D. (1987), "Still Searching." *Journal of Economic Behavior and Organization* 8: 137–144.

Lippman, Steven A., and John J. McCall (1976), "The Economics of Job Search: A Survey." *Economic Inquiry* 14: 155–189.