

Brown, Meta; Flinn, Christopher J.; Schotter, Andrew

Working Paper

Real-time search in the laboratory and the market

Staff Report, No. 410

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Brown, Meta; Flinn, Christopher J.; Schotter, Andrew (2009) : Real-time search in the laboratory and the market, Staff Report, No. 410, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60947>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Real-Time Search in the Laboratory and the Market

Meta Brown
Christopher J. Flinn
Andrew Schotter

Staff Report no. 410
December 2009

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Real-Time Search in the Laboratory and the Market

Meta Brown, Christopher J. Flinn, and Andrew Schotter

Federal Reserve Bank of New York Staff Reports, no. 410

December 2009

JEL classification: C91, J64

Abstract

While widely accepted models of labor market search imply a constant reservation wage policy, the empirical evidence strongly suggests that reservation wages decline in the duration of search. This paper reports the results of the first real-time-search laboratory experiment. The controlled environment that subjects face is stationary, and the payoff-maximizing reservation wage is constant. Nevertheless, subjects' reservation wages decline sharply over time. We investigate two hypotheses to explain this decline: 1) searchers respond to the stock of accruing search costs, and 2) searchers experience nonstationary subjective costs of time spent searching. Our data support the latter hypothesis, and we substantiate this conclusion both experimentally and econometrically.

Key words: job search, consumer search, reservation wage, experiment

Brown: Federal Reserve Bank of New York (e-mail: meta.brown@ny.frb.org). Flinn: New York University (e-mail: christopher.flinn@nyu.edu). Schotter: New York University (e-mail: andrew.schotter@nyu.edu). This research is supported in part by a grant from the C.V. Starr Center for Applied Economics at New York University and the Center for Experimental Social Science. Judy Goldberg provided valuable research assistance. Raj Advani and Jay Ezrielev were responsible for the experiment programming. An editor and three anonymous referees provided many valuable comments, as did seminar participants at Queen's University, Université Laval, and the Midwest Economics Association. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

The standard model of labor market search posits repeated (infinite horizon) search from a fixed distribution of wages in a stationary environment. Under these circumstances the optimal search policy has a reservation wage property, i.e., the decision to accept an offer involves a comparison of the offer with a time-invariant constant, w^* , which itself is a function of the parameters characterizing the search environment. A wage offer greater than or equal to w^* is accepted, while any other is rejected and search is continued.

Some early empirical job search studies, however, accumulated evidence of a declining trend in reservation wages over the course of an unemployment spell.¹ Since then several theoretical explanations for the declining reservation wage profiles have emerged, all of which relax the stationarity assumption of the standard model. For example, Gronau (1971) assumes searchers have finite lives; Burdett (1977), van den Berg (1990), Albrecht and Vroman (2005) and others model short-term unemployment benefits; Albrecht, Gautier and Vroman (2006) permit directed search to heterogeneous prospects;² Danforth (1979) and Rendon (2006) consider general liquidity constrained search; and Burdett and Vishwanath (1988) model learning about the job offer distribution. Behavioral explanations for declining reservation wages have been suggested by Braunstein and Schotter (1981, 1982) and others.

In this paper we present the results of the first ever (to our knowledge) real-time-search laboratory experiment performed in an effort to investigate this falling reservation wage phenomenon. Where field studies of unemployed search often rely on strong identifying assumptions to recover searchers' unobservable reservation wages, in the laboratory we are able to elicit subjects' complete reservation wage profiles in an incentive-compatible manner.³ Further, in the laboratory we are able to abstract from market sources of non-stationarity. We present subjects with a stationary environment in which they encounter offers at a stochastic rate and in which they may accumulate search costs as time progresses. We find that, despite the stationarity of their environment, there is a strong tendency of subjects in our experiment to lower their reservation wages over time. The importance of this finding is that it demonstrates a behavioral cause of declining reservation wages, which may operate in the field alone or in conjunction with the market factors discussed above.

While there are a wide variety of possible behavioral explanations for declining searcher standards in a stationary environment, we find that most can be categorized in one of

¹See, for example, Kasper (1967), Keifer and Neumann (1979), and Lancaster and Chesher (1983). Braunstein and Schotter (1981, 1982) first generated the result in the context of a search experiment.

²This is related to the model of Salop (1973).

³In the partial-partial equilibrium search context, Flinn and Heckman (1982) and Wolpin (1987) present the first rigorous investigations of identification and estimation of these models using unemployment duration and wage information.

two ways. First, searchers may respond to accumulating monetary costs by lowering their acceptance standards. For example, the sunk cost fallacy demonstrated by Thaler (1980) could lead searchers who fail to recognize previously accumulated search costs as sunk to accept lower wage offers in response to larger accumulated costs. Second, searchers may respond to passing search time itself. A lengthy search spell may be inherently discouraging, or workers may accumulate *subjective* search costs not captured in standard models.⁴

To distinguish between these two explanations we employ both econometric and experimental methods. We first attempt to identify the impact of both search time and accumulated search cost on the declining reservation wage in a standard framework. However, since these two variables are highly colinear in the experimental data, it is hard to separate their influence using econometric methods alone. As a result, we run two additional experimental treatments, one in which searchers wait uncertain amounts of time for offers to arrive but accumulate no search costs, and one in which offers arrive immediately but with a stochastic cost. Estimates of the time dependence of stated reservation wages in these two new treatments both show significant declines. Nonetheless, we find that only the waiting treatment is sufficient to generate a downward trend in reservation wages as steep as that observed in the standard framework. Based on these results, it appears that the primary factor determining the time path of reservation wages in our experiments is the uncertain wait.

The three treatments each impose discounting through the structure of the experimental payoff calculation. One might be concerned that subjects otherwise aware of the incremental nature of the search problem respond to the "shrinking pie" of discounted payoffs by lowering their reservation wage choices. We field a final treatment in which discounting is imposed not structurally but through the threat of termination of game play. This treatment allows us to compare the two different methods of discounting. We find strikingly similar time patterns in reservation wages, and in this sense the experiment cross-validates the use of structural discounting and termination risk-based discounting in laboratory search experiments. We know of no other experiment to have done this.

The experimental hazard into offer acceptance displays both increasing and decreasing regions. Our non-monotonic experimental hazard demonstrates the ability of an extensive downward trend in reservation wages in a majority of environments to overwhelm the decreasing hazard generated by environment heterogeneity and produce regions of increasing employment hazards. As an examination of the external validity of our experimental results, we present evidence of increasing and decreasing regions of the raw employment hazard in

⁴One referee has pointed out the possibility that subjects seek a positive net payout in each spell. This may lead to non-stationary cost responses, but may be less relevant to unemployed search in the field than other possible behavioral responses. We also find declining reservation wage patterns in an experimental treatment that guarantees positive net payouts, described below.

recent data on young US job seekers, and we review evidence from Ham and Rea (1987) of a non-monotonic hazard despite extensive controls for market sources of non-stationarity. The experimental and field data are consistent in this sense, and our experimental results suggest a possible behavioral component to the observed reservation wage declines and non-monotonic hazards in U.S. employment data.

The paper proceeds as follows. In Section 2 we recount the standard search model. Section 3 presents our experimental design and Section 4 our results. In Section 5 we consider a non-stationary modification of the standard search framework, and we discuss external validation of our laboratory results. Section 6 concludes.

2 The Standard Search Model

The theoretical foundation for the experimental analysis is a standard partial-partial equilibrium continuous time search model. Wage offers arrive according to a Poisson process with time homogeneous parameter λ , which is the (constant) rate at which the searcher encounters new wage offers. When a job offer arrives the searcher receives a wage offer w , which is an independently and identically distributed draw from the distribution F . The searcher accepts or rejects any arriving offer, and the accepted wage is received in perpetuity. We do not permit the searcher to recall previous offers,⁵ and we assume that, once hired, a worker cannot be fired or search on-the-job.

In Appendix A.1 we derive the expected wealth-maximizing searcher’s decision rule in this stationary environment. As is well-known, the agent will use a decision rule that possesses the reservation wage policy, that is,

$$\begin{aligned} \text{Accept } w &\Leftrightarrow w \geq w^* \\ \text{Reject } w &\Leftrightarrow w < w^*, \end{aligned}$$

where

$$w^* = b + \frac{\lambda}{\rho} \int_{w^*} (w - w^*) dF(w), \quad (1)$$

b represents the instantaneous value of search, and ρ is the instantaneous discount rate.

Due to the time homogeneity of all primitive parameters, the value w^* is constant. If the primitive parameters are non-stationary, if agents suffer from sunk-cost fallacies, or if their objective is not expected wealth maximization, a reservation wage decision rule may not

⁵In a stationary environment with full information on the offer distribution, an expected wealth maximizing agent would never exercise a recall option. We explicitly preclude it here since we cannot assume that our experimental subjects are expected wealth-maximizers nor “correctly” solving the optimization problem.

exist, or, if one does exist, the reservation wage may not be time-invariant. Non-stationary reservation wage policies are considered below and in van den Berg (1990).

3 Experimental Design

Three separate search experiments (or treatments) were performed in the Experimental Economics Laboratory of the Center for Experimental Social Science at New York University. Each treatment included 38 inexperienced undergraduate students recruited from the undergraduate population of New York University and brought into a laboratory where they sat in front of computer terminals and read computerized instructions. Each experiment lasted approximately one hour (by the time the last searcher completed the last search) and the average payoff was approximately \$16. Subjects were paid a \$5 show-up fee, and were paid all contingent compensation in experimental dollars. These were converted into U.S. dollars at the rate of 1 experimental dollar to \$.07 U.S. for the baseline treatment and 1 experimental dollar to \$.05 U.S. for the additional treatments. Based on their laboratory behavior and questions, the subjects' motivation and understanding of instructions appeared to be high.

3.1 Laboratory Implementation of the Standard Search Model

The decision task of the subjects in our baseline experiment was identical to the standard search problem described in Section 2. Subjects searched in continuous real-time search environments in which time was measured in seconds. When the subject clicked his or her mouse indicating that he or she was ready to start, the clock started. Wage offers, drawn from a two parameter Weibull distribution, $F(w \mid \alpha, \beta)$, arrived at a rate determined by a Poisson distribution with parameter λ .⁶ The cost of search was constant at b experimental dollars per second, where $b \leq 0$ in all treatments.⁷ Each subject searched until he or she accepted a wage. He or she completed five consecutive search spells in each of three assigned search environments.⁸ Search environments were randomly determined. Elements of the parameter vector $X = (\alpha, \beta, \rho, \lambda, b)$ had bounded supports $\alpha \in [0.015, 0.6]$, $\beta \in [0.5, 2]$,

⁶Appendix B discusses the implementation of continuous time offer arrival in the laboratory and the offer distribution parameterization.

⁷It is important to note that the existence of a reservation wage in the canonical partial-partial equilibrium model of search does not impose any sign restriction on the utility flow while searching, b . For purposes of the experiment, we restricted the flow value of search to be nonpositive. Consequently, we refer to b as a search cost in what follows.

⁸Subjects were told that they would search 5 times in a row in each of three environments. They were informed when environments changed. Each new environment began with the tutorial described below, relevantly parameterized.

$\rho = 0.03$, $\lambda \in [0.05, 0.5]$, and $b \in [-0.05, -0.03]$. Each environment was selected by drawing parameter values independently from uniform distributions defined over these supports.⁹

Before search began, subjects had an opportunity to experience the exact environment they would be searching in. This was done through a tutorial program.¹⁰ The tutorial first described the wage distribution by presenting a histogram representing its properties and explaining its salient features. Rather than presenting the actual wage offers, we transformed each offer into the current time discounted present value of receiving the wage to the horizon, so that subjects saw instead the distribution of $\frac{w}{\rho}$. This was done so that subjects would not have to convert each wage offer into a discounted present value on their own. It leaves the problem unaltered. In addition, the notion of a discount rate was explained both in words and in terms of graphs depicting how wages and costs were discounted to the beginning of the search period.

The tutorial contained a dynamic demonstration of the temporal features of the problem. Subjects sat in front of their terminals for two minutes and watched discounted wage offers arrive at the rate defined by the environment. As each offer arrived a beep was heard and the words << **PAYMENT OFFER** >> were listed on the screen.

Table 1 replicates a screen which was filled as the two minute demonstration proceeded. The table columns contain the elapsed time, the present value of the offer ($\frac{w}{\rho}$), the value of the offer discounted to the beginning of the spell, cumulative discounted search costs when the offer arrived, and the payoff net of search costs if the offer were to be accepted. From the table we see that the first wage arrived after 20 seconds and had an associated lifetime income of 4.8 experimental dollars which, when discounted to the present, was worth only 2.63. The discounted search cost incurred during those 20 seconds (0.42 experimental dollars) is then included in the net discounted payoff value minus search cost of 2.21. After the first offer arrived this particular subject had to wait 5 seconds for the next and 31 seconds for the third offer. Such a table was designed to meet several objectives. First, it gave the subject an idea of the waiting time process, and, most importantly, the randomness of it. It then showed him or her the trade-offs involved in stopping or waiting for a better wage in

⁹Appendix Table 1 shows the means and standard deviations of the environment parameter values used in the experiment. The parameter range for offer arrival rate λ is chosen to be analogous to the range of values estimated in this standard search model in Flinn and Heckman (1982). Given the demonstration in Flinn and Heckman (1982) that the discount rate ρ and unemployment flow utility b are not separately identified in field data, we fix ρ at what we believe to be a reasonable level. Instantaneous discount rates in field research imply annual discounts of, for example, 0.05 in Dey and Flinn (2005) and 0.12 in van den Berg (1990). Elapsing seconds in our experiment are parameterized to behave similarly to elapsing weeks of search in comparable models estimated in the field. Thus a discount of 0.03 is not far from standard discounts in continuous time search models estimated using field data.

¹⁰The computer tutorial is available for viewing at http://cess.nyu.edu/exp_data_programs.html. The data are available on request from the authors.

the future, net of search costs. For example, notice that while the best payment offer was received at the 96 seconds mark, its discounted value was far less than the much smaller payment offer received at the 20 second mark.¹¹ At the end of the tutorial search began in the relevant environment. Before a subject began to search in a new environment, he or she was presented with a tutorial for the specific environment drawn.

Search took place as follows. Subjects waited at their terminals until an offer arrived. During this time the screen displayed only the elapsing search time in seconds. When an offer arrived the words << **PAYMENT OFFER** >> appeared on the screen and time was stopped (i.e., we allowed subjects to deliberate over their reservation wage without accumulating further search costs). The value of the arriving offer was not shown to the subject. Rather, without seeing the offer, the subject was asked to state a value such that if the arriving offer were greater, the offer would be accepted automatically. If the arriving offer were less, the offer would be rejected and search would continue. This procedure is incentive compatible in that revealing the truth about one's reservation value is a dominant strategy. It is equivalent to the Becker, DeGroot, and Marschak (1964) procedure.¹² Thus we are able to observe subjects' true reservation values, an object rarely available in field data.¹³ The subject entered a reservation value and then clicked an "OK" box to move on. If the offer was below the chosen reservation value, the screen announced, "Your reservation price was greater than the payment. Search will continue." The offer was displayed, along with its discounted value and the accumulated cost. If the offer was above the chosen reservation value, the screen announced, "Your reservation price was less than the payment. The payment was accepted." The screen then displayed the offer and the subject's payoff for the spell, and the search spell ended. A subject's final payoff equalled the sum of her or his payoffs across all fifteen search rounds.

3.2 Additional Treatments

The colinearity of elapsed time and accumulated costs in the standard search framework limits our ability to distinguish between these two influences on searcher behavior in the baseline experiment. However, we were able to isolate time and cost effects in the laboratory. In a second experimental treatment, the No Wait-Cost treatment, we removed all waiting from game play. Subjects received offers without waiting at a monetary cost of bt experimental

¹¹Note that the ex post payoff maximizing offer varied across tutorials.

¹²It is similar to the manner of eliciting credibly reported reservation wages used by Cox and Oaxaca (1992).

¹³In field data, the analyst only knows whether an offer $w \in A$, where A is the acceptance set. We force individuals to use a critical value strategy, i.e., one in which the set A is connected. Assuming A is in fact connected, we learn everything about the subject's decision rule.

dollars, where t was drawn from the same λ -parameterized distribution as in the baseline experiment. All other aspects of the environment were parameterized as before.¹⁴

In the third treatment, the Wait-No Cost treatment, we removed the monetary search cost altogether but continued to impose stochastic laboratory waiting periods. This treatment is precisely equivalent to performing the baseline experiment with $b = 0$.

Cox and Oaxaca (1989, 1996) note that subjects are aware of the finiteness of laboratory sessions. Subjects therefore may not believe an experimenter’s promise of unlimited search time, and this may influence their search behavior in purportedly infinite horizon search experiments. With an eye to this concern, our intent was that a binding search horizon never influence a searcher’s decision during the course of the experiment. We believe that the parameterization of environments in the experiment achieves this objective. The predicted average length of a completed search spell using an expected payoff maximizing strategy and facing the slowest possible job offer arrival rate ($\frac{1}{\lambda} = 20$ seconds) is just under one minute. Observed search spell durations averaged just 22 seconds, with a maximum of almost 4 minutes.¹⁵ Thus a searcher who always took the maximum observed time would complete 15 trials in just under an hour of active search, easily within any reasonable time allotment for a university laboratory experiment.

It is important to point out that, in each of our new treatments, expected payoff maximizing reservation wage $w^*(X)$ is still constant over time. In the No Wait-Cost experiment the reservation wage is identical in each environment to that in the baseline experiment, since the standard theory of search does not model non-monetary waiting costs. In the Wait-No Cost experiment, $w^*(X)$ is higher than in the Wait-Cost and No Wait-Cost experiments but still constant.

4 Results and Estimation

Since we run an experiment in real time where observed reservation wages may follow any time path, the question of whether subjects adhere to the expected payoff maximizing reservation wage profile is not well defined. However, we can pose the following narrower questions.

¹⁴This new experimental treatment also facilitates comparison to several previous search experiments, such as Braunstein and Schotter (1981, 1982), Hey (1987), Cox and Oaxaca (1989, 1992), and Harrison and Morgan (1990), in which the subject indicates the decision to search and receives an immediate offer.

¹⁵The duration of the experiment was not announced to subjects during either the recruitment phase or the experiment.

4.1 Do subjects select the expected payoff maximizing reservation wage on average?

For each randomly chosen environment in which subjects search, we calculate w^* using expression (1). We were surprised to find that baseline treatment subjects' first stated reservation wages, in the first trial of the first environment, averaged 101.92 percent of their environments' optimal reservation wages. In other words, a standard search model applying a payoff-maximizing assumption does a surprisingly good job of explaining subjects' initial reservation wages.

After this, however, the fit of the standard search model to subjects' behavior falls off. Pooling observations across environments, trials and offers, the average $w^*(X)$ is 13.42. The mean stated reservation wage is 9.90.¹⁶ A 95 percent confidence interval on the mean difference between the stated and the payoff maximizing reservation wage runs from -5.33 to -2.44, easily excluding zero, and thus we reject the hypothesis that baseline treatment subjects selected $w^*(X)$ on average.¹⁷

The results for the Wait-No Cost treatment are similar. Subjects initially select a reservation wage whose average is reasonably near $w^*(X)$. In the pooled data the stated reservation wage average is 13.48, the $w^*(X)$ average is 15.19 and the confidence interval on the difference is -2.39 to -1.05. Again subjects under bid the expected payoff maximizing reservation wage. However, in the No Wait-Cost treatment the average stated reservation wage is 12.61, the $w^*(X)$ average is 11.18 and the confidence interval on the difference is 0.73 to 2.15. Subjects set higher reservations, on average, than the payoff maximizing wage. This is the first of several indications that subjects in the No Wait-Cost treatment behave differently from those in the waiting treatments.

The above confidence intervals are constructed by resampling individual-trial observations. However, if there is individual heterogeneity in the distance between the payoff maximizing and stated wages that persists across trials, then our resampling method should account for this. Constructing confidence intervals by resampling individual subjects' full experiment histories generates much broader confidence intervals, which exclude zero only for the baseline treatment. This sensitivity of the approximated confidence intervals to individual resampling suggests substantial individual heterogeneity in reservation wage levels given environment characteristics.

¹⁶All empirical "reservation wages" reported here and throughout the rest of the paper are expressed as the total $\frac{w}{p}$ value derived from the wage, as shown to the experimental subjects.

¹⁷Confidence intervals reported in this and the following two subsections are each based on relevant statistics from 5000 bootstrap samples.

4.2 Do subjects follow a constant reservation wage strategy?

To this point we have avoided confronting a fundamental problem with descriptive analysis of our experimental data. Searchers willing to accept lower wages finish searching sooner, on average, and therefore are under-represented in data on second and later offers. This renders our discussion of average reservation values less informative, and is a particular challenge as we turn to wage dynamics. An additional problem is that the rich environment heterogeneity that is a part of our experimental design generates extensive variation in payoff maximizing reservation wages, potentially obscuring other relevant sources of variation.

One measure we can take to address both searcher and environment heterogeneity is to difference reservations from offer to offer within a given searcher and environment.¹⁸ Table 2 shows reservation wage changes from the first offer to the accepted offer within a spell. Spells in which the first offer is accepted are assigned zero differences, and differences are averaged across all search spells and environments for each treatment. We find significant declines in stated reservation wages from the first to the accepted offer for each of the three treatments. Average differences for the Wait-Cost, No Wait-Cost and Wait-No Cost treatments, respectively, are 2.24, 1.83 and 2.71. The bootstrapped 95 percent confidence intervals around these means are (1.48, 2.82), (1.04, 2.59) and (1.35, 4.50), ruling out the null hypothesis of no change.¹⁹ Average declines conditional on two or more offers are substantially larger, at 4.16, 5.44 and 6.25, and, again, their confidence intervals rule out no change. Finally, in analysis of reservation wage transitions within a search spell, we find that 64.05 percent of baseline treatment offer-to-offer transitions involve reservation wage decreases.²⁰ On average and at the level of the individual search spell, the predicted flat reservation wage profile is a poor description of what we see subjects choosing in the experiment.

Where we do not compare repeated observations on the same searcher in the same environment, we see evidence of both dynamic selection and the clouding effect of environment heterogeneity. Table 3 reports descriptive statistics for groups of searchers who received the same final number of offers. Among searchers who accepted their first offers, the average

¹⁸Note that this assumes that the relevant searcher and environment heterogeneity is time-constant. This is an assumption that we maintain throughout the empirical section, and formalize in the model in subsection 4.4. Searcher heterogeneity in trends would not be accounted for by our approach.

¹⁹All bootstrapped confidence intervals from this point forward are constructed by resampling individuals' full search histories. We find minimal sensitivity of the confidence intervals on reservation wage measures that are differenced within individuals to individual or individual-trial resampling, suggesting comparatively limited heterogeneity in reservation wage trends.

²⁰21.14 percent show no change, and 14.81 percent are increases. The pattern is similar for the Wait-No Cost treatment; however, in 58.87 percent of No Wait-Cost treatment transitions the reservation wage is unchanged.

stated reservation is 5.71.²¹ Average first stated reservations for those who received exactly two, three, four and five or more offers are 8.79, 11.05, 15.89 and 19.03. The increasing pattern in the mean first stated reservations across offer groups is similar in the other two treatments, and reflects dynamic selection.²² We may be less concerned about such selection within groups in Table 3. In the baseline treatment, the two-offer group's stated reservations fall from 8.79 to 5.54 on average. The three-offer group's reservations fall from 11.05 through 8.93 to 6.04, and the four-offer group's from 15.89 through 14.33 and 12.64 to 9.79.²³ Though the declines are large, 95 percent confidence intervals around group-offer averages overlap heavily from the first to the second offer, the second to the third and so on. It appears that uncontrolled environment heterogeneity is able to mask significant declines in reservation wages over time.

4.3 Spell-to-spell transitions and learning

One possible explanation for an observed decline in subjects' reservation profiles is that conventionally rational subjects learn about the wage offer over time. Those who draw lower offers remain in the search pool and revise their wage offer distribution priors downward. We attempted to remove learning dynamics from the experiment via pre-search tutorials and practice. Here we investigate the behavior of stated reservation wages between search spells in order to gauge the role of learning in the experiment.

Table 3 "recoveries" are calculated as the difference between the reservation wage that led to acceptance in the previous trial and the first stated reservation wage in the current trial. The recovery average is calculated using only trial-to-trial transitions in which the previous search and the current search occur in the same environment. Our question is whether searchers beginning a new round in the same environment behave as they did at the end of the previous round in that environment, as they did at its beginning, or neither. A recovery average of zero would indicate that, on average, searchers continue where they left off in a new search spell in the same environment. A recovery average equal to the decline average would indicate that, on average, searchers begin the process with the same reservation wage every time. In the former case, within-spell reservation wage declines are carried over to the next round; in the latter, they are not.²⁴

²¹Note that this is the modal group. 310 of the 570 completed search spells in the baseline experiment, 289 No Wait-Cost spells and 265 Wait-No Cost spells ended at the first offer.

²²As in Ham and LaLonde (1996), a randomization of participants precedes the dynamic decision problem, and nonetheless the comparison of outcomes is hindered by dynamic selection. Thus, like Ham and LaLonde, we turn to panel techniques in an experimental context.

²³The small five or more offer group has a flatter profile. By and large the reservation wage averages by offer number behave similarly in the other treatments.

²⁴Nothing in the exercise guarantees that the result will fit either of these descriptions.

Recovery averages in Table 2 roughly balance the previous trials’ average reservation wage declines for each treatment. Where the mean Wait-Cost decline is 2.24, the mean recovery is 1.92; comparable figures for the other treatments are (1.97, 1.83) and (2.61, 2.72). Bootstrapped 95 percent confidence intervals on the mean decline-recovery differences fail to reject the null of zero. This finding would seem to argue against a learning explanation. If subjects learn about the wage offer distribution during active search, we would expect what they have learned to carry over into the next search round. However, we see an approximately complete recovery of searchers’ reservation wages between spells.

In summary, we have found decidedly non-stationary searcher behavior in the laboratory which cannot be explained by the standard market factors cited above. Some concern regarding dynamic selection in our descriptive analysis remains. The suggestive evidence of non-stationarity raises some new questions. For example, as we argue in section 1, many intuitive explanations for non-stationary search behavior can be ascribed to one of two features of the search problem: elapsed search time and accumulated costs. What influence does each factor have on searchers’ reservation wage profiles?

In the next subsections we take two approaches to addressing the above questions. First, we develop the more formal analysis of reservation wages, elapsed search time, and accumulated search costs that follows. Second, we return to our experimental variation of the roles of time and costs in the search problem.

4.4 “Out of Equilibrium” Estimation of Reservation Wage Functions in the Baseline

As above, denote an environment by X , which includes parameters describing the offer distribution function F , the instantaneous net utility flow from searching, b , the rate of arrival of offers, λ , and the discount rate ρ . As we know from Section 2, under the assumptions that (1) the environment is constant and (2) the objective of the agent is to maximize expected wealth, an individual i searching in trial j within environment X should use the offer acceptance rule

$$r_{ij}(t; X) = w^*(X).$$

If the environment is not constant, or if the agent is not an expected wealth-maximizer, the decision rule may be a function of t , or even more importantly, may not be uniquely characterizable in terms of a reservation wage. The experiment as conducted forces individuals to use reservation wage policies, which may not be optimal given their objective function. This caveat aside, to characterize the policies individuals use when we allow for

time dependence, we assume that

$$r_{ij}(t; X) = m_{ij}(X) + \phi(t, bt) + \varepsilon_{ij}(t; X), \quad (2)$$

where $m_{ij}(X)$ is an individual- and trial-specific function of the constant environmental parameters, $\phi(t, bt)$ is a function of elapsed time and accumulated cost and $\varepsilon_{ij}(t; X)$ is the state, at time t , of a driftless²⁵ Brownian motion process with constant instantaneous variance parameter $\sigma_{ij}^2(X)$, potentially specific to an individual, environment, and trial. The driftless Brownian motion process is described by an initial condition, $\varepsilon_{ij}(0; X)$, which is unrestricted, and the stochastic law of motion

$$\varepsilon_{ij}(t; X) = \varepsilon_{ij}(0; X) + \sigma_{ij}(X)Z(t),$$

where $Z(t)$ is the standard Brownian motion (or Wiener) process. This implies that $u_{ij}(t; X) \equiv \varepsilon_{ij}(t + \tau; X) - \varepsilon_{ij}(\tau; X) \sim N(0, \sigma_{ij}^2(X)t)$, for $t > 0$ and $\tau \geq 0$.

There is no presumption that

$$m_{ij}(X) = m(X) = w^*(X),$$

that is, we do not assume that the m function is common across individuals and trials, nor do we assume that it equals the reservation wage function under expected wealth maximization. On the other hand, it is necessary for us to assume that the function $\phi(t, bt)$ is common to all population members.

Let subject i 's sequence of stated reservation wages in trial j within environment X_k be given by $\{r_{ijkl}\}_{l=1}^{n_{ijk}}$, and let the time of each offer be given by $\{t_{ijkl}\}_{l=1}^{n_{ijk}}$, where $0 < t_{ijk1} < \dots < t_{ijkn_{ijk}}$ and where n_{ijk} is the number of offers observed in (i, j, k) . Denote the set of every (i, j, k) with an associated $n_{ijk} \geq 2$ as S_2 . Within this set, consider the difference sequence $\Delta r_{ijkl} = r_{ijkl+1} - r_{ijkl}$, $l = 1, \dots, n_{ijk} - 1$. The specification (2) implies that

$$\Delta r_{ijkl} = \phi(t_{ijkl+1}, bt_{ijkl+1}) - \phi(t_{ijkl}, bt_{ijkl}) + u_{ij}(t_{ijkl+1} - t_{ijkl}; X_k). \quad (3)$$

Expression (3) characterizes contiguous reservation wage differences Δr_{ijkl} , within set S_2 of trials that contain two or more offers, in terms of their dependence on elapsed time, accumulated cost and an error term of known distribution.

²⁵The requirement that the process be driftless stems from it serving as a disturbance term in what is essentially a time series/cross sectional regression model. With no restriction on the initial condition of the process, $\varepsilon_i(0)$, differencing is required in order to consistently estimate the function ϕ . The differenced ε process will be mean 0 only if all individual specific trends are equal to 0.

With an assumption regarding the functional form of ϕ , estimation of the parameters characterizing (3) is reasonably standard given the strong, but standard, assumptions made on the distribution of the disturbance terms $u_{ij}(t; X)$. We assume that the function ϕ belongs to the set of polynomial functions in $(t, -bt)$, where $\mathbb{Q}(t, bt; \theta_d)$ denotes the polynomial in t and bt (including interactions) of order d , which is characterized by the parameter vector θ_d . We can only hope to precisely estimate low dimensional polynomials. As a result, we restrict attention to the cases of $d = 1, 2$.

Given our assumptions, investigation of time and accumulated cost effects on reservation wages involves estimation of a linear (in the parameters) regression function with a mean zero, potentially heteroskedastic disturbance term, with experimental data from the set S_2 . By dividing equation (3) by $\sqrt{t_{ijkl+1} - t_{ijkl}}$, we produce a covariance matrix of the disturbances which is diagonal, with elements containing $\sigma_{ij}^2(X_k)$. The OLS estimator of θ_d in this transformed equation is unbiased and consistent. If $\sigma_{ij}(X_k) = \sigma \forall (i, j, k)$, then the OLS standard errors are consistent as well. However, if the instantaneous variance parameters do vary across subjects, trials, and/or environments, consistent estimates of the standard errors can be obtained by use of the White (1980) estimator. We test for the presence of heteroskedasticity in the differenced and transformed data, and choose the OLS or White estimator of the standard errors accordingly. The nature of the relationship of the reservation wage process and the two arguments t and bt can be investigated using standard tests for linear restrictions on the parameters of the polynomial functions.

In an experimental auction study, Ham, Kagel, and Lehrer (2005) discuss the difficulty of “formulat(ing) and test(ing) non-optimizing behavioral models...” Our solution, like theirs, involves specifying a fairly agnostic econometric model of subject behavior and comparing its results to the relatively restrictive predictions of the existing theory.

4.5 Estimation Results

Point estimates for the model with \mathbb{Q} functions of orders $d = 1$ and 2 are reported in Table 4. In response to a high degree of correlation between t and bt in the baseline treatment data, we report estimates for time-only specifications for this group. While the collinearity problem limits the precision of coefficient estimates in t and bt specifications, we are able to estimate time-only coefficients with reasonable precision. Since there is no variation in cost in the Wait-No Cost treatment, or in wait time in the No Wait-Cost treatment, we avoid the collinearity problem and are able to isolate the effect of waiting in the former and cost in the latter treatment. In addition, we find that some features of the estimation are sensitive to the inclusion of a handful of extreme outliers among the observed reservation wage changes.

The estimates represent samples in which the largest and smallest 2.5 percent of reservation wage changes have been trimmed.²⁶

Consider first the estimated time dependence of reservation wages in baseline specification (1). Elapsed time is measured in seconds, so that the point estimates indicate the effect of one second of elapsed time on the change in the reservation wage. In (1), the coefficient on t is -0.2090. This implies that an additional 4.8 seconds of elapsed time is associated with a one experimental dollar decrease in the reservation wage. A White test rejects the null hypothesis of homoskedasticity in the transformed data, and so we focus on the White standard errors in brackets. The point estimate is significantly different from zero at the 1 percent level. Thus we find that the subjects' reservation wages decline significantly over time. Specification (2) again shows a very precisely estimated negative coefficient on elapsed time, and a mild though insignificant convexity in the time path. The net marginal effect of time on the reservation wage change remains negative for all search durations observed in the sample. Again the White test rejects the null of homoskedasticity and we rely on the White standard errors. An F-test of the null hypothesis that the coefficients on t and t^2 in specification (2) are jointly zero does reject the null at the 5 percent level of confidence. Specification (4) includes time and cost regressors, with squared terms and an interaction. Though the collinearity problem does not allow us to identify time and cost coefficients separately, an F-test easily rejects the null that coefficients on all five regressors are jointly zero. Once cost data are included in the estimation, the White test fails to reject the null that $\sigma_i = \sigma \forall i$ in the differenced and transformed data.

In the analysis of the No Wait-Cost treatment data, we replace $\phi(t, bt)$ with $\phi(bt)$; similarly, we replace $\phi(t, bt)$ with $\phi(t)$ in the Wait-No Cost estimation. Estimates using first and second order \mathbb{Q} approximations to ϕ for these two treatments are also reported in Table 4. The experimental separation of the wait and search cost effects of passing time on the search problem allows us to estimate the effects of the wait and the cost on reservation wage dynamics separately. This time we find decisive evidence of negative effects of both elapsed time and accumulated cost on the change in the reservation wage. Wait-No Cost treatment coefficient estimates are similar to those in the Wait-Cost time-only specification: the linear coefficient implies that a six second wait is associated with a one experimental dollar decrease in the reservation wage, and the quadratic specification demonstrates a decreasing and mildly convex time path for all sample members. The significant No Wait-Cost treatment estimates of the cost dependence of reservation wages imply, in the linear case, that a one experimental dollar increase in accumulated cost is associated with a 1.78 experimental

²⁶The results are qualitatively very similar where the dependent variable is specified as the change in log reservation wages and the sample is untrimmed.

dollar decrease in the reservation wage. The cost dependence of the reservation wage is estimated to be decreasing and convex.²⁷

It is difficult to compare the contributions of wait time and accumulated costs to the time dependence of the reservation wage based on the magnitudes of these coefficients alone, or even to interpret curvature parameters, so we move to graphical comparisons. Figure 1 plots the predicted change in the reservation wage based on the quadratic specification estimates for all three treatments.²⁸ The reservation wage declines over time, and a 95 percent confidence band on each prediction demonstrates that we can rule out zero or positive changes in the reservation wage throughout the observation period for all three treatments.²⁹

The No Wait-Cost treatment displays the weakest time dependence of reservation wages. This may be unsurprising, given that 59 percent of observed offer-to-offer reservation wage transitions here involve no change. We find that the No Wait-Cost trajectory lies everywhere above the Wait-Cost and the Wait-No Cost trajectories, and no overlap of its confidence band with those of the other treatments. The Wait-Cost and Wait-No Cost trajectories are considerably steeper, and their confidence bands show substantial overlap.³⁰ These experimental results provide one answer to our question of which feature of search dynamics is responsible for the non-stationarity in searchers' choices. While the laboratory wait is sufficient to generate a decreasing reservation wage trend statistically indistinguishable from the trend in the more standard search problem, the search cost is not.

4.6 Robustness Check: Discounting via Termination Risk

Some may fear that the implementation of discounting in our experiment, as represented in Table 1, might induce a fear of losses in searcher behavior as future payments, which are discounted back to present value, appear to shrink and payments appear to approach negative values.³¹ This may distract subjects from the stationarity of the incremental payoffs in the search problem that may lead to a lowering of stated reservation wages. Such problems are inescapable in an experimental setting, however, since the subject must leave the lab with a single payoff and that payoff must be discounted to some fixed point while in the stationary search setting payoffs are in effect flow values. While searching a searcher receives

²⁷White tests fail to reject homoskedasticity of the differenced and transformed data in all Table 4 Wait-No Cost and No Wait-Cost specifications, and so we report only OLS standard errors.

²⁸The Wait-Cost prediction curve is based on the $d = 2$ time and cost specification.

²⁹The confidence bands are generated using the asymptotic distribution theory associated with the maximum likelihood estimator.

³⁰In fact, the slope of the predicted time path under laboratory waiting is more than twice the slope for costs alone throughout Figure 1.

³¹See Shin and Ariely (2004) for a direct demonstration of the shrinking pie effect. Their subjects took costly actions in order to maintain options of questionable value.

flow b ; following offer acceptance a searcher receives flow w . But the ex post value of the total search spell must depend on the time of offer acceptance, τ . Since the experimental subject is rewarded based on the ex post value of the spell, there seems to be no way to represent the full stationary search problem to an experimental subject without bringing in the time dependence of payoffs in a manner that may induce her to perceive the payoff "pie" as shrinking.

However, the experimental literature does present one viable and convenient means of inducing necessary discounting in our search experiment without introducing the need for payoff discounting per se and that is random termination.³² Roth and Murnighan (1978), among other experimental studies, induce discounting in laboratory subjects through the random termination of experimental game play. Terminated rounds or sessions under this approach yield zero payoffs to subjects. In order to investigate the sensitivity of our results on the time dependence of reservation values to the structurally imposed discounting used in the previous treatments, we field a final laboratory treatment that induces discounting through termination of the search spell. Search spell termination times are distributed according to a Poisson distribution with parameter $\phi = 0.03$. The expected time to termination is thus $1/0.03 = 33.3$ seconds. If the spell is terminated before an offer is accepted the spell yields a zero payoff. No other discount is imposed (i.e., $\rho = 0$). In order to maintain comparable payoffs to previous treatments, an accepted offer pays $w/0.03$, with w distributed according to a Weibull distribution. We introduced this "Terminated Risk Treatment" under the Wait-No Cost conditions where the instantaneous search cost is removed. This allows us to isolate the pure discounting effect since there is no possibility of losses resulting from termination and hence we can compare the results of this treatment directly to the Wait-No Cost treatment done with payoff discounting. This new treatment is theoretically identical to our Wait-No Cost treatment run before with monetary discounting and was the treatment that we felt explained the decreasing reservation wage phenomenon the best. In all other aspects we parameterize the environments subjects face identically to those of the Wait-No Cost treatment.

The presentation of the experiment to subjects in the Termination Risk treatment requires no explanation of mounting search costs and no representation of dwindling discounted net payoffs. The tutorial screen presented to subjects before the start of game play was in this case analogous to the first two columns of Table 1, reporting simply the time of offer arrival and the payoff. As before, this tutorial screen evolved in real time. One new element was added to the tutorial: a termination announcement once the stochastic termination time arrived. Subjects were shown three search spells from origination to termination, to give

³²Thanks to the referees for suggesting this application of the random termination method.

them an idea of the effect of termination on the search process. They were also shown the values, in seconds, of 30 draws from the distribution of termination times. Game play proceeded precisely as it did in the three prior treatments, with two exceptions. First, no information on discounted payoffs or accumulated costs was presented during or after game play, as these were now irrelevant. Second, if the stochastic termination time arrived before an offer was accepted, the termination was announced, the search spell ended and the subject received a payoff of zero.

Thirty-six subjects each performed 5 completed search spells in each of 3 environments in the Termination Risk treatment. Estimates of the empirical model described at the start of this section using the Termination Risk treatment data are reported at the bottom of Table 4. Where the time coefficient estimated for the Wait-No Cost treatment was -0.1592 with a standard error of 0.0171, we find a time coefficient for the Termination Risk treatment of -0.1478 with a standard error of 0.0215. The similarity in the magnitude and precision of the two estimates is striking. The predicted time paths are also quite similar for the quadratic specification, with the exception that the Termination Risk estimates show only a linear effect of time on the reservation value where the Wait-No Cost estimates indicated a slight decrease in the effect of time on the reservation value as the spell continues.³³ Repeating the Figure 1 confidence interval predictions based on the linear specifications for each treatment again produces a No Wait-Cost confidence interval that lies everywhere below the intervals for the (now three) waiting treatments, supporting the findings in section 4.5.³⁴

Hence our reservation wage time path results are not sensitive to whether discounting is imposed structurally in the calculation of final payoffs or through the inclusion of a termination hazard in the experiment. Given the identical parameterization of the search environments in the Wait-No Cost and Termination Risk treatments, perhaps this is unsurprising. One interpretation of our results is that the experiment cross-validates the structural and termination risk methods of inducing discounting in a laboratory search experiment.

4.7 The Experimental Hazard

One additional dimension of search behavior of central interest in the empirical job search literature is the hazard rate out of unemployment, into offer acceptance, over time. Define

³³As in the Wait-No Cost treatment, White tests fail to reject the null hypothesis of homoskedasticity for both the linear and the quadratic Q specifications and therefore we report only OLS standard errors.

³⁴However, exogenous terminations lead to smaller sample sizes and larger standard errors in the Termination Risk treatment. Hence the new treatment yields quite a broad confidence interval under the quadratic specification, and this interval does in fact overlap the No Wait-Cost quadratic specification confidence interval for some values of t .

$h(t; X)$ as the employment hazard at time t conditional on searching in environment X , and $h(t)$ as the unconditional hazard at time t . If reservation wage $w^*(t; X)$ is decreasing in t , as implied by the experimental results, then we should expect to see $h(t; X)$ rise over time in environment X . However, when evaluating the experimental hazard we must pool across the wide range of environments implemented in the experiment, constructing unconditional hazard $h(t)$. It is well known that if $w^*(t; X)$ is constant over time for all environments, and if there is any heterogeneity in $h(t; X)$, then unconditional hazard $h(t)$ will be nonincreasing over time.³⁵ As searchers in high $h(t; X)$ environments exit and searchers in low $h(t; X)$ environments remain, mixing leads to decreasing hazards.

The unconditional hazard rate for the experiment, $h(t)$, is constructed from the marginal densities of duration times in the population.³⁶ Figure 2 graphs a nonparametric estimate of the unconditional hazard using data from the baseline treatment. Figure 2.a contains the unsmoothed estimate of the hazard, and Figures 2.b and 2.c contain smoothed versions. The general pattern observed is not that of a steady increase, but instead includes substantial decreasing regions. This may seem puzzling given the strong evidence of a declining tendency in reservation wages in the experiment, and the increasing conditional hazards that the estimated reservation wage profiles would seem to imply. As we can see from any of the plots, the mixing phenomenon has obliterated much of the positive duration dependence in the conditional hazards. Because we have access to the reservation wages, we know that the lack of an overall increasing trend in the marginal hazard rate is purely an artifact of the selection process. Our experiment thus demonstrates the ability of downward pressure from mixing across search environments to overwhelm a positive trend in an employment hazard.

5 Beyond the Laboratory

Given the finding that searchers systematically lower their acceptance standards over time, even in a controlled stationary environment, we examine the role of time-varying subjective costs of search in the context of the theory and the field. First we modify the stationary search model of section 2 to accommodate time-varying search costs of reasonably general origin, and we ask whether this modification is able to capture features of the reservation profile and exit from search observed in the laboratory. Second, we turn to data on young job seekers in the U.S. and ask whether the searcher behavior we observe in the laboratory is also evident in the field.

³⁵See, for example, Barlow and Proschan 1981, Section 4.4.

³⁶The precise hazard estimator and smoothing method are described in Appendix C.

5.1 The Case of Time-Varying b

It is not difficult to extend the standard stationary search model to include time dependent utility flows in the state of unemployed search. Keeping all other parameters in expression (1) fixed, we merely allow the instantaneous value of search, $b(t)$, to be a differentiable, nonincreasing function of t . In this case, it is reasonably straightforward to demonstrate that the job acceptance decision possesses a reservation wage property, with the critical value given by $w^*(t)$, which is formally derived in Appendix A.2. Moreover, we have the following result.

Proposition 1 *If $b'(t) \leq 0$ for all $t \geq 0$, then $w^*(t) \leq 0$.*

Proof. See Theorem 2 of van den Berg (1990). ■

We illustrate the path of the reservation wage given a falling flow utility of being in the search state with the following example. As was true in the case of the experiments, we assume that the distribution of wage offers is Weibull, with parameters $\alpha = 0.25$ and $\beta = 1.3$. The offer arrival rate, λ , is set at 0.2, and $\rho = 0.01$. We assume that the flow utility in the unemployment state at search duration t is given by $b(t) = \underline{b} + a\delta^t$, with $\underline{b} = 3$, $a = 5$, and $\delta = 0.95$. Given that $a > 0$, this function attains a maximum value at $t = 0$, where $b(0) = \underline{b} + a$, and the infimum of the function is \underline{b} . The first derivative is $a \ln(\delta)\delta^t < 0$ and the second derivative is $a(\ln(\delta))^2\delta^t > 0$.

Figure 3.a contains a plot of the path of reservation wages over the course of an unemployment spell. The path of reservation wages inherits the general properties of the b function. Since the environment in which search is conducted is assumed to be stationary, the declining reservation wages result in a nondecreasing hazard rate out of the search spell. The hazard function for the example is plotted in Figure 3.b. Thus we find that this simple modification of the standard search model captures the salient features of behavior we observe in the laboratory: reservation wages decline with increasing search duration, and the model is able to produce an increasing hazard (and thus, we infer, increasing hazard regions under mixing). The fit of the model to the observed reservation wage profile may be modified through the manipulation of the $b(t)$ function.

5.2 The Examination of Experimental Results using Field Data

The analysis of the experimental data strongly suggests a tendency for searchers in stochastic but stationary environments to reduce their reservation wage values as the search episode increases in length. The objective of this section is to determine whether this result has validity in data gathered outside the lab.

The main empirical implications drawn from the experimental evidence are the following:

1. Reservation wages are decreasing in the duration of search.
2. Individual hazard rates out of the search state are increasing in duration.

The field data we use in this exercise are taken from the 1997 National Longitudinal Survey of Youth (NLSY-97). In the year 2004, sample members were between 20 and 24 years of age, which is the approximate age range of most of our experimental subjects. From the weekly labor market history file, we constructed unemployment spells that began after the interview conducted in 2004 and before the interview conducted in 2005. We used only the first unemployment spell in this period for each respondent when there was more than one unemployment spell during this period. Unemployment spells still in progress at the time of the interview date in 2005 were considered to be right-censored. By constructing spells in this manner, we hoped to minimize the types of measurement problems that commonly plague duration-based analyses, problems which make it difficult to accurately determine the shape of the hazard function.³⁷

Our estimated hazard functions are based on a sample of 1921 unemployment spells; 1402 were complete (i.e., not right-censored). Figure 4 contains plots of the “raw” nonparametric hazard function estimator, and two smoothed versions. The hazard estimator is identical to that employed in the experimental data and described in Appendix C. In Figure 4.a, we see that there is a general tendency for the hazard to decline with duration of the unemployment spell, but there is considerable “noise” in the estimated function, as we observed in the experimental data. There is no clear indication of an increasing hazard.

When we smooth the hazard, this is no longer the case. With a moderate degree of smoothing, as in Figure 4.b, there are three increasing hazard regions ($t < 3$, $27 < t < 31$ and $37 < t < 41$). At a high degree of smoothing (Figure 4.c), however, the hazard contains almost exclusively flat and decreasing regions.

We believe that these results provide some evidence that, in recent U.S. employment data, the marginal hazard rate is increasing over some nonnegligible subset of the sample space.³⁸ According to the proposition, this is consistent with the increasing conditional

³⁷In a previous draft of this paper, we utilized unemployment spell information from the Survey of Income and Program Participation to conduct this exercise. The SIPP data suffer from a well-known “seam” problem, which is the tendency for events, such as an unemployment spell, to be reported as ending at sampling frequencies, which are approximately four months in the SIPP. (See, for example, Ham, Li and Shore-Sheppard 2007.) The new data don’t suffer from this type of reporting problem in any obvious way.

³⁸Other contexts in which nonmonotone employment hazards have been demonstrated include Swedish employment record data (Bergstrom and Edin 1992 and Korpi 1995) and U.S. CPS job seekers (Flinn and Heckman 1983) and unemployment beneficiaries (Addison and Portugal 1987).

hazard rates observed in the experimental evidence. In general our field data results suggest that observed hazards out of unemployment are not consistent with the (weakly) decreasing profile predicted by standard search models with constant reservation wages, whether or not there is mixing over search environments.

A drawback to our analysis of employment hazards in these recent field data is that, unlike the experimental subjects, NLSY-97 job seekers may face non-stationary search environments resulting from time-limited unemployment benefits, changing demand conditions, and, eventually, binding credit constraints. These market factors alone might explain increasing hazard regions. Ham and Rea (1987) analyze employment hazards while accounting for market sources of non-stationarity. They study prime-aged, unemployed Canadian men from 1975-1980, who are less likely to face meaningful credit constraints than our young NLSY-97 cohort, they include arguably complete detail on unemployment benefits, and they control for changing demand conditions via both regional and industry unemployment measures. With and without modeling heterogeneity among searchers, Ham and Rea find that the direct effect of duration on the employment hazard is negative for shorter search durations, but positive for longer durations.³⁹

Both our experimental results and the field studies described above indicate the presence of increasing segments of the offer acceptance hazard. The common practice of assuming stationary search in multiple distinct environments is unable to fit this pattern. Our results suggest that accounting for the behavioral effects of elapsed search time may be an appropriate way of reconciling the conventional search framework with the evidence of increasing hazards.

6 Conclusion

This paper describes the first real-time laboratory search experiment of which we are aware. The experiment also differs from most previous search experiments, excepting that of Cox and Oaxaca (1992), in that it elicits searchers' underlying reservation wages in an incentive-compatible manner, and differs from all prior search experiments that we are aware of in its use of random variation in the search environment. We designed the experiment to exclude non-stationary phenomena that could mimic market features such as budget constrained or directed search. The search environments are stationary in terms of their monetary features, though they may be non-stationary in other, experiential features. Thus the experiment seeks to isolate the influence of behavioral features of the search problem on searchers' decision

³⁹A potentially important caveat is that search durations longer than 40 weeks, where the estimated hazard turns up, are less common.

making over time.

We find decisive evidence of declining reservation wages in the duration of search in our baseline experiment, in which subjects confront the standard, stationary search environment. This leads to the question of whether it is the elapsed time in the laboratory or the accumulating search costs that drives the time dependence of reservation wages. To answer this question we repeat our experiment, varying the roles of time and costs in the search process. We find a significantly steeper decrease in reservation wages over time for all treatments that include waiting; combining wait time and accumulating costs leads to the steepest reservation wage decline over time. Hence it appears to be the uncertain laboratory wait, and not the accumulating search cost, that contributes most to the observed reservation wage decline, though the effects of time and costs are cumulative.

Given these new insights into behavioral aspects of searchers' decisions, we modify the standard continuous-time stationary search framework to accommodate non-stationary, potentially subjective costs of time spent searching. We find that a fairly flexible specification of the time dependence of the value of search easily permits the determination of a unique, time-varying acceptance rule that retains the reservation property. We demonstrate a non-monotonic experimental acceptance hazard, and report similar properties of the employment hazard in NLSY-97 data and in leading research on unemployment duration from previous decades. We conclude that behavioral responses to features of job search such as the length of the search spell and accumulating search costs as studied here may be of use in explaining evidence of decreasing reservation wages and increasing employment hazards observed in the field.

References

- [1] Addison, John T., and Pedro Portugal. 1987. "On the Distributional Shape of Unemployment Duration." *Review of Economics and Statistics*, 69: 520-526.
- [2] Albrecht, James, Pieter Gautier, and Susan Vroman. 2006. "Equilibrium Directed Search with Multiple Applications." *Review of Economic Studies*, 73(4): 869-891.
- [3] Albrecht, James, and Susan Vroman. 2005. "Equilibrium Search with Time-Varying Unemployment Benefits." *Economic Journal*, 115(505): 631-648.
- [4] Banks, Jeffrey, Mark Olson, and David Porter. 1997. "An experimental analysis of the bandit problem." *Economic Theory*, 10: 55-77.
- [5] Barlow, Richard E., and Frank Proschan. 1981. *Statistical Theory of Reliability and Life Testing*. Silver Spring, MD: To Begin With.
- [6] Becker, Gordon, Morris DeGroot, and Jacob Marschak. 1964. "Measuring utility by a single-response method." *Behavioral Science*, 9 (July): 226-232.
- [7] Bergstrom, R., and Pers-Anders Edin. 1992. "Time aggregation and the distributional shape of unemployment duration." *Journal of Applied Econometrics*, 7: 5-30.
- [8] Braunstein, Yale M., and Andrew Schotter. 1981. "Economic Search: An Experimental Study." *Economic Inquiry*, 19: 1-25.
- [9] Braunstein, Yale M., and Andrew Schotter. 1982. "Labor Market Search: An Experimental Study." *Economic Inquiry*, 20: 133-144.
- [10] Burdett, Kenneth. 1977. "Unemployment Insurance Payments as a Search Subsidy: A Theoretical Analysis." Unpublished.
- [11] Burdett, Kenneth, and Tara Vishwanath. 1988. "Declining Reservation Wages and Learning." *Review of Economic Studies*, 55: 655-665.
- [12] 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.
- [13] Cox, James C., and Ronald L. Oaxaca. 1992. "Direct Tests of the Reservation Wage Property." *Economic Journal*, 102: 1423-1432.
- [14] Cox, James C., and Ronald L. Oaxaca. 1996. "Testing Job Search Models: The Laboratory Approach." *Research in Labor Economics*, 15: 171-207.

- [15] Danforth, John. 1979. "On the Role of Consumption and Decreasing Absolute Risk Aversion in the Theory of Job Search." In *Studies in The Economics of Search*, ed. Steven A. Lippman and John J. McCall, 109-131. New York: North-Holland.
- [16] Dey, Matthew, and Christopher Flinn. 2005. "An Equilibrium Model of Health Insurance Provision and Wage Determination." *Econometrica*, 73: 571-627.
- [17] Flinn, Christopher J., and James J. Heckman. 1982. "New Methods for Analyzing Structural Models of Labor Force Dynamics." *Journal of Econometrics*, 18: 115-168.
- [18] Flinn, Christopher J., and James J Heckman. 1983. "Are Unemployed and Out of the Labor Force Behaviorally Distinct Labor Force States?" *Journal of Labor Economics*, 1: 28-42.
- [19] Gronau, Reuben. 1971. "Information and Frictional Unemployment." *American Economic Review*, 61: 290-301.
- [20] Ham, John, John Kagel, and Steven Lehrer. 2005. "Randomization, Endogeneity and Laboratory Experiments: the Role of Cash Balances in Private Value Auctions." *Journal of Econometrics*, 125: 175-205.
- [21] Ham, John, and Robert LaLonde. 1996. "The Effect of Sample Selection and Initial Conditions in Duration Models: Evidence from Experimental Data on Training." *Econometrica*, 64(1): 175-205.
- [22] Ham, John C., Xianghong Li, and Lara Shore-Sheppard. 2007. "Correcting for Seam Bias when Estimating Discrete Variable Models, with an Application to Analyzing the Employment Dynamics of Disadvantaged Women in the SIPP." http://www-rcf.usc.edu/~johnham/documents/Ham_Seam_Bias_aug_31_final.pdf.
- [23] Ham, John C. and Samuel A. Rea. 1987. "Unemployment Insurance and Male Unemployment Duration in Canada." *Journal of Labor Economics*, 5(3): 325-353.
- [24] Harrison, Glenn W., and Peter Morgan. 1990. "Search Intensity in Economic Experiments." *Economic Journal*, 100: 478-486.
- [25] Hey, John D. 1987. "Still Searching." *Journal of Economic Behavior and Organization*, 8: 137-144.
- [26] Kasper, Hirschel. 1967. "Asking Price of Labor and the Duration of Unemployment." *Review of Economics and Statistics*, 49: 165-172.

- [27] Kiefer, Nicholas M., and George R. Neumann. 1979. "An Empirical Job Search Model with a Test of the Constant Reservation Wage Hypothesis." *Journal of Political Economy*, 87: 89-108.
- [28] Korpi, Tomas. 1995. "Effects of Manpower Policies on Duration Dependence in Re-employment Rates: The Example of Sweden." *Economica*, 62: 353-371.
- [29] Lancaster, Tony, and Andrew Chesher. 1983. "An Econometric Analysis of Reservation Wages." *Econometrica*, 51: 1161-1176.
- [30] Rendon, Silvio. 2006. "Job Search and Asset Accumulation under Borrowing Constraints." *International Economic Review*, 47: 233-263.
- [31] Roth, A. E. and Murnighan, J. K. 1978. "Equilibrium Behavior and Repeated Play of the Prisoners' Dilemma." *Journal of Mathematical Psychology*, 17:189-197.
- [32] Salop, Steven. 1973. "Systematic Job Search and Unemployment." *Review of Economic Studies*, 40: 191-201.
- [33] Shin, Jiwoong, and Dan Ariely. 2004. "Keeping Doors Open: The Effect of Unavailability on Incentives to Keep Options Viable." *Management Science*, 50: 575-586.
- [34] Thaler, Richard. 1980. "Toward a Positive Theory of Consumer Choice." *Journal of Economic Behavior and Organization*, 1: 39-60.
- [35] van den Berg, Gerard. 1990. "Nonstationarity in Job Search Theory." *Review of Economic Studies*, 57: 255-277.
- [36] Wang, Jane-Ling. 2005. "Smoothing Hazard Rates." *Encyclopedia of Biostatistics*, 2nd Edition, 7: 4986-4997.
- [37] White, Halbert. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*, 48: 817-838.
- [38] Wolpin, Kenneth. 1987. "Estimating a Structural Search Model: The Transition from School to Work." *Econometrica*, 55: 801-817.

A Derivation of Reservation Wage Policies in the Stationary and Nonstationary Case

A.1 Stationary Environment

The payoff-maximizing searcher's decision rule in a continuous time stationary environment is derived as follows. The value of search over an arbitrarily small time period ε is given by

$$V_n = (1 + \rho\varepsilon)^{-1} \{b\varepsilon + \lambda\varepsilon \int \max(\frac{w}{\rho}, V_n) dF(w) + (1 - \lambda\varepsilon)V_n + o(\varepsilon)\},$$

where $(1 + \rho\varepsilon)^{-1}$ serves as an “instantaneous” discount factor, $b\varepsilon$ is the flow value of search accumulated over the period ε , $\lambda\varepsilon$ is the “approximate” probability of receiving one offer over the period, and $o(\varepsilon)$ represents all events that can occur over the period ε that involve more than one event, with $\lim_{\varepsilon \rightarrow 0} o(\varepsilon)/\varepsilon = 0$. The first term in the max operator is the value of accepting an offer of w , w/ρ , and the second argument is the value of rejecting it, which is the value of continuing search. The idea behind this representation is that all changes in the choice set and decisions are made at the end of the period ε .

Rewrite the expression as

$$\begin{aligned} (1 + \rho\varepsilon)V_n - V_n &= b\varepsilon + \lambda\varepsilon \int \max(\frac{w}{\rho} - V_n, 0) dF(w) + o(\varepsilon) \\ \Rightarrow \rho\varepsilon V_n &= b\varepsilon + \frac{\lambda\varepsilon}{\rho} \int_{\rho V_n} (w - \rho V_n) dF(w) + o(\varepsilon). \end{aligned}$$

Dividing both sides by ε and taking limits as $\varepsilon \rightarrow 0$, we have

$$\rho V_n = b + \frac{\lambda}{\rho} \int_{\rho V_n} (w - \rho V_n) dF(w). \quad (4)$$

Given that V_n is independent of current offer w , the choice to accept wage w in perpetuity, valued at $\frac{w}{\rho}$, or to continue search, valued at V_n , displays reservation property

$$\begin{aligned} \text{Accept } w &\Leftrightarrow w \geq w^* \\ \text{Reject } w &\Leftrightarrow w < w^*. \end{aligned} \quad (5)$$

The reservation wage, $w^* \equiv \rho V_n$, is defined by implicit function 4.

A.2 Nonstationary Environment

We consider the case of optimal search with a monotonic, time-varying function $b(t)$, where t is the duration of the search spell. Suppose that $b'(t) \leq 0$, with $\lim_{t \rightarrow \infty} b(t) = B$. The value of search is now time-varying, and is denoted $V_n(t)$. Then we have

$$\begin{aligned} V_n(t) = & (1 + \rho\varepsilon)^{-1} \left\{ \int_t^{t+\varepsilon} b(u) du + \lambda\varepsilon \int \max\left(\frac{w}{\rho}, V_n(t + \varepsilon)\right) dF(w) \right. \\ & \left. + (1 - \lambda\varepsilon)V_n(t + \varepsilon) + o(\varepsilon) \right\}, \end{aligned}$$

which can be rewritten as

$$\begin{aligned} (1 + \rho\varepsilon)V_n(t) - V_n(t + \varepsilon) &= \int_t^{t+\varepsilon} b(u) du \\ &+ \lambda\varepsilon \int \max\left(\frac{w}{\rho} - V_n(t + \varepsilon), 0\right) dF(w) + o(\varepsilon) \\ \Rightarrow V_n(t + \varepsilon) - V_n(t) &= \rho\varepsilon V_n(t) - \int_t^{t+\varepsilon} b(u) du \\ &- \frac{\lambda\varepsilon}{\rho} \int_{\rho V_n(t+\varepsilon)} (w - \rho V_n(t + \varepsilon)) dF(w) + o(\varepsilon). \end{aligned}$$

Dividing by ε and taking limits, we have

$$V'_n(t) = \rho V_n(t) - b(t) - \frac{\lambda}{\rho} \int_{\rho V_n(t)} (w - \rho V_n(t)) dF(w),$$

or

$$w^{*'}(t) = \rho w^*(t) - \rho b(t) - \lambda \int_{w^*(t)} (w - w^*(t)) dF(w).^{40} \quad (6)$$

B Laboratory Implementation of Continuous Time Offer Arrival

The realized duration between two successive offers was generated by pseudo-random number draws from a Uniform distribution defined on $[0, 1]$ using the inverse-CDF method. That is, in an environment in which the rate of arrival of offers was λ , draw a (pseudo) random number, x . Treat x as a probability, and find the value of t (the duration time) that corresponds to

⁴⁰In order to establish the existence and uniqueness of solution $w^*(t)$ we require further conditions. Define $b_0 = b(0)$ and $\lim_{t \rightarrow \infty} b(t) = b_1$. Assuming that $b(t)$ is monotonic and continuously differentiable, $b_0 > b_1$, $b_0 < +\infty$ and $b_1 > -\infty$, we are able to demonstrate the existence and uniqueness of $w^*(t)$ that satisfies (6). A proof of existence and uniqueness for a more general case can be found in van den Berg (1990).

that probability. Since

$$P(T \leq t | \lambda) = 1 - \exp(-\lambda t)$$

when T has a negative exponential distribution, then define the (pseudo) random waiting time associated with the draw x as

$$t = -\frac{\ln(1-x)}{\lambda}.$$

The specific functional form for the probability density of arriving offer w is

$$f(w | \alpha, \beta) = \alpha \beta w^{\beta-1} e^{-\alpha w^\beta}.$$

This density implies a probability mass of wages exceeding threshold w of

$$1 - F(w) = \exp(-\alpha w^\beta).$$

C Estimation of the Hazard Function

The hazard function estimation is performed as follows. The nonparametric estimator of the hazard function is given by

$$\hat{h}(t) = d(t)/R(t), \quad t = 1, 2, \dots, T,$$

where T is the last time in which a sample member is observed to exit a search spell, $d(t)$ is the number of individuals who exited from search at duration t , and $R(t)$ is the number who had not completed their search spell by duration t . $R(t)$, which is the number of individuals at risk of leaving unemployment at duration t , contains both individuals with right-censored spells of search (whose exact exit dates are not observable) that lasted at least t weeks, as well as those individuals with completed spells whose exit dates are t or greater. In the experimental data, there are no right-censored search spells, while in the NLSY97 data there are a large number.

The function \hat{h} can exhibit a substantial amount of instability, particularly over areas in the sample space where the risk set is small (i.e., larger values of t). This makes regions of increasing and decreasing duration dependence difficult to identify. Following the methods described in Wang (2005), we smooth the function $\hat{h}(t)$ using a local linear smoother. The

estimate of the hazard at duration t is given by

$$\hat{a}_0(t) = \arg \min_{a_0, a_1} \sum_{i=1}^T R(i) K((t-i)/q) \{ \hat{h}(i) - (a_0(t) + a_1(t)(i-t)) \}^2,$$

where K is the standard normal kernel, q is a bandwidth parameter that we have chosen as reported in Figure 2, and $a_0(t)$ and $a_1(t)$ are parameters chosen by the smoothing algorithm to fit a local linear polynomial to the raw hazard. (See Wang 2005 for a motivation for this functional form). The smoothed hazard value at time t is simply approximation $\hat{a}_0(t) + \hat{a}_1(t)(i-t)$ evaluated at $i = t$, hence $\hat{a}_0(t)$. In the NLSY97 data, $T = 52$ weeks, capturing up to one year of unemployed search. In the experimental data, $T = 52$ seconds. This time limit captures the vast majority of exits from the search state in the experimental data. For both the experimental and nonexperimental data, we present plots of the estimated hazard for two values of the smoothing parameter b , which were 4 and 8.

Table 1: Tutorial Screen [$\lambda = 1$ every 15 seconds, $b = 0.028/\text{second}$, $\rho = 3\%/\text{second}$]

Time Received	Current Value of Offer	Discounted Value	Total Search Cost	Net Payoff
20	4.8	2.63	0.42	2.21
25	10.5	4.96	0.49	4.47
56	12.97	2.42	0.76	1.66
58	14.60	2.56	0.77	1.79
82	12.41	1.06	0.85	0.21
86	14.47	1.10	0.86	0.23
96	15.74	0.88	0.88	0.00

Table 2: Reservation Wage Declines Within Trials and Recoveries Across Trials

Wait-Cost	Mean	SD	Median	Min	Max
Decline	2.24	6.23	0.00	-10.40	63.00
Recovery	1.92	5.97	0.50	-19.00	60.00
Decline - Recovery 95% CI	(-0.86, 0.84)				
No Wait & Cost	Mean	SD	Median	Min	Max
Decline	1.83	7.45	0.00	-16.00	89.00
Recovery	1.97	7.66	0.00	-27.00	82.00
Decline - Recovery 95% CI	(-1.42, 1.28)				
Wait & No Cost	Mean	SD	Median	Min	Max
Decline	2.71	10.45	0.00	-32.00	135.00
Recovery	2.62	8.73	0.20	-21.00	81.90
Decline - Recovery 95% CI	(-2.07, 1.74)				

Table 3: Reservation wage descriptives by number of offers and experimental treatment, pooling over search environments

Offer #	Wait, Cost					No Wait, Cost					Wait, No Cost				
	Total offers in trial					Total offers in trial					Total offers in trial				
	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
1: mean	5.17	8.79	11.05	15.89	19.03	6.48	10.80	11.05	18.34	16.81	7.25	11.87	11.19	11.58	24.72
<i>median</i>	3.79	7.10	10.00	14.50	14.00	5.00	7.75	9.25	18.50	15.00	5.60	10.00	10.00	9.45	12.00
<i>SD</i>	5.27	8.17	7.95	12.17	17.62	6.60	10.82	8.80	15.92	15.47	8.72	12.89	8.04	7.95	31.56
<i>n=</i>	310	146	57	29	28	289	122	54	30	75	265	142	62	42	59
2: mean	-	5.54	8.93	14.33	17.86	-	8.32	9.87	17.66	16.36	-	7.80	9.20	11.69	24.96
<i>median</i>	-	4.50	10.00	12.90	13.50	-	7.00	8.50	15.00	15.00	-	6.90	7.00	9.45	10.50
<i>SD</i>	-	5.06	5.92	9.52	15.71	-	6.79	7.23	16.15	14.11	-	6.16	7.29	9.58	31.93
<i>n=</i>		146	57	29	28		122	54	30	75		142	62	42	59
3: mean	-	-	6.04	12.64	18.14	-	-	7.52	17.45	16.94	-	-	7.48	9.75	23.34
<i>median</i>	-	-	4.00	12.00	13.00	-	-	5.00	15.50	14.90	-	-	4.75	8.10	12.00
<i>SD</i>	-	-	5.73	9.72	18.94	-	-	6.72	16.43	15.97	-	-	6.73	7.77	31.19
<i>n=</i>			57	29	28			54	30	75			62	42	59
4: mean	-	-	-	9.79	18.16	-	-	-	13.85	16.07	-	-	-	8.47	20.60
<i>median</i>	-	-	-	8.00	10.00	-	-	-	11.00	15.00	-	-	-	6.90	12.00
<i>SD</i>	-	-	-	8.43	21.65	-	-	-	9.37	15.34	-	-	-	6.30	24.96
<i>n=</i>				29	28				30	75				42	59
5: mean	-	-	-	-	17.53	-	-	-	-	14.76	-	-	-	-	16.91
<i>median</i>	-	-	-	-	9.00	-	-	-	-	10.00	-	-	-	-	10.00
<i>SD</i>	-	-	-	-	24.85	-	-	-	-	15.14	-	-	-	-	21.32
<i>n=</i>					28					75					59

Table 4: Estimates of reservation wage changes, (Wait, Cost),
(No Wait, Cost), & (Wait, No Cost) treatments

Treatment	Independent Variable						N ,
		t	bt	t^2	b^2t^2	bt^2	R^2
Wait, Cost	(1)	-0.2090 (0.0316) [0.0228]	-	-	-	-	492, 0.0819
	(2)	-0.2515 (0.0476) [0.0424]	-	0.0005 (0.0004) [0.0005]	-	-	492, 0.0845
	(3)	-0.3724 (0.2458) [0.1973]	4.1316 (6.1655) [5.0061]	-	-	-	492, 0.0827
	(4)	-0.1947 (0.4023)	-1.3656 (9.9139)	-0.0374 (0.0203)	-22.4585 (11.3506)	1.8592 (0.9627)	492, 0.0926
No Wait, Cost	(1)	-	-1.7760 (0.1997)	-	-	-	897, 0.0811
	(2)	-	-2.2563 (0.2732)	-	0.0836 (0.0326)	-	897, 0.0878
Wait, No Cost	(1)	-0.1592 (0.0171)	-	-	-	-	723, 0.1066
	(2)	-0.2215 (0.0256)	-	0.0008 (0.0002)	-	-	723, 0.1196
Termination Risk	(1)	-0.1478 (0.0215)	-	-	-	-	276, 0.1467
	(2)	-0.1409 (0.0404)	-	-0.0002 (0.0008)	-	-	276, 0.1469

OLS standard errors are reported in parentheses and White standard errors in brackets.

Appendix Table 1: Parameters drawn in the baseline experiment

Parameter	Mean	Standard Deviation	Min	Max
α	0.342	0.175	0.023	0.598
β	1.203	0.410	0.524	1.989
λ	0.139	0.113	0.050	0.500
b	-0.039	0.005	-0.050	-0.030

Figure 1: Predicted change in reservation wages over time




Figure 2.a
Nonparametric Hazard Estimate
Experimental Data




Figure 2.b
Smoothed Hazard Estimate
 $q = 2$




Figure 2.c
Smoothed Hazard Estimate
 $q = 4$




Figure 3.a
Reservation Wage Function




Figure 3.b
Hazard Rate Out of Search




Figure 4.a
Nonparametric Hazard Estimate
NLSY97 Data




Figure 4.b
Smoothed Hazard Estimate
 $q = 2$




Figure 4.c
Smoothed Hazard Estimate
 $q = 4$

