

Estimating Individual Discount Rates With Field Experiments

by Glenn W. Harrison

Inferences about individual discount rates are drawn from two sets of field experiments. One was a natural field experiment in the United States, undertaken by the U.S. Department of Defense to encourage military personnel to retire early. The other was an “un-natural” field experiment in Denmark, where laboratory methods were applied in the field by the Danish Ministry of Business and Industry in a deliberate effort to elicit discount rates for policy purposes.

In the Danish study, Harrison, Lau and Williams [2000] estimate individual discount rates with respect to time streams of money using controlled laboratory experiments. These discount rates are elicited by means of field experiments involving real monetary rewards, and employ the laboratory design developed by Coller and Williams [1999]. The experiments were carried out across Denmark using a representative sample of 268 people between 19 and 75 years of age. Individual discount rates are estimated for various households differentiated by socio-demographic characteristics such as income and age. Our conclusions are that discount rates are constant over the 12-month to 3-year horizons used in these experiments, and that discount rates vary substantially with respect to several socio-demographic variables. Hence we conclude that it would be *reasonable to assume constant discount rates for specific household types, but not the same rates across all households*.

In the U.S. study, Harrison and Johnson [2002] reconsider the conclusions of Warner and Pleeter [2001], who find that the average individual discount rates implied by the observed choices were high relative to *a priori* expectations. Their linear model predicts average rates of 10.4% and 35.4% for officers and enlisted personnel, respectively. We extend their analysis in four respects. First, we use this field experiment to test the traditional hypothesis that discount rates are constant over different time horizons. Second, we consider the robustness of the inferences about average discount rates, taking into account the statistical uncertainty of the calculation. Third, we account for the possibility of censoring with respect to market interest rates. Censoring is a potential issue if individuals have access to alternative investment opportunities. Fourth, we augment their analysis to estimate the individual discount rates for the U.S. adult population as a whole, rather than just the military component of that population. We do so by constructing a statistical model of the decision to join the military, and using that statistical model to adjust the responses estimated for those in the military.

Conclusions are drawn about the complementarity of lab and field experiments, and remaining gaps in our empirical knowledge about individual discounting.

Coller, Maribeth, and Williams, Melonie B., “Eliciting Individual Discount Rates,” *Experimental Economics*, 2, 1999, 107-127.

Harrison, Glenn W., and Johnson, Eric, “Estimating Individual Discount Rates for the United States: Inferences from a Natural Field Experiment,” *Unpublished Manuscript*, Department of Economics, Moore School of Business, University of South Carolina, 2002.

Harrison, Glenn W.; Lau, Morten Igel, and Williams, Melonie B., “Estimating Individual Discount Rates for Denmark: A Field Experiment,” *Economics Working Paper 00-11*, Moore School of Business, University of South Carolina, 2000; forthcoming, *American Economic Review*.

Warner, John T., and Pleeter, Saul, “The Personal Discount Rate: Evidence from Military Downsizing Programs,” *American Economic Review*, 91(1), March 2001, 33-53.

Estimating Individual Discount Rates in Denmark: A Field Experiment

by

Glenn W. Harrison, Morten I. Lau and Melonie B. Williams[†]

November 2001

Forthcoming, *American Economic Review*, v.92, 2002.

Abstract. We estimate individual discount rates with respect to time streams of money using controlled laboratory experiments. These discount rates are elicited by means of field experiments involving real monetary rewards. The experiments were carried out across Denmark using a representative sample of 268 people between 19 and 75 years of age. Individual discount rates are estimated for various households differentiated by socio-demographic characteristics such as income and age. Our conclusions are that discount rates are constant over the 12-month to 3-year horizons used in these experiments, and that discount rates vary substantially with respect to several socio-demographic variables. Hence we conclude that it would be *reasonable to assume constant discount rates for specific household types, but not the same rates across all households.*

[†] Department of Economics, The Darla Moore School of Business, University of South Carolina, USA (Harrison); Centre for Economic and Business Research, Danish Ministry of Trade and Industry (Lau); and Office of Policy, Economics and Innovation, U.S. Environmental Protection Agency, USA (Williams). E-mail contacts: HARRISON@DARLA.BADM.SC.EDU, MOL@CEBR.DK, and WILLIAMS.MELONIE@EPAMAIL.EPA.GOV. We are grateful to the Danish Ministry of Business and Industry for funding this study, to Maribeth Collier, Ben Heijdra, Orley Ashenfelter and two referees for comments. None of our employers or overlapping generations are responsible for our conclusions.

Discount rates are often used in cost-benefit analysis. Whenever costs and benefits for a household or individual are spread over time, it is essential that one calculate present value equivalents in order to undertake meaningful comparisons. In most cases welfare analysts use market rates as the basis for these present value calculations. Sensitivity analysis often consists of varying the scalar discount rate up or down in relation to market interest rates.

Since discount rates are a reflection of subjective time preferences, one would expect *a priori* that they could differ across different individuals.¹ However, standard practice in inter-temporal welfare analyses is to assume that those rates are (i) the same across households, and (ii) the same for all time horizons. We elicit individual discount rates from subjects in order to test these two hypotheses. The first hypothesis is that discount rates for a *given time horizon* do not differ with respect to socio-demographic characteristics that characterize households in our sample. The second hypothesis is that discount rates for a *given individual* do not differ across time horizons.

We use survey questions with real monetary rewards to elicit individual discount rates and demonstrate the methodological complementarity between lab and field experiments. The survey questions are designed by Collier and Williams [1999], who elicit nominal individual discount rates for university students using controlled laboratory experiments.² We apply their experimental procedures, but employ subjects that are normally encountered in field surveys. Our experiments were carried out across Denmark for the Danish government, using a nationally representative sample of 268 people between 19 and 75 years of age.

Our results indicate that nominal³ discount rates are constant over the 1-year to 3-year horizons used in these experiments, and that discount rates vary significantly with respect to several socio-demographic variables. *On the basis of these results one can assume constant discount rates for specific household types, but not the same rates across all households.*

In section 1 we review the logic of our experimental design. Section 2 explains the field experiments conducted, and section 3 examines the results and relates them to those found in the existing literature.

¹ We elicit discount rates for individuals. To the extent that the characteristics of individuals are used to define “representative households,” we can refer to the individual and the household interchangeably. However, we remain agnostic concerning the way in which the individual discount rates of individual household members are aggregated into one household discount rate, akin to a social discount rate for the household as a small society.

² Collier and Williams [1999] explain how their design relates to findings in the extant experimental literature. We review this discussion in our working paper, available at [HTTP://DMSWEB.BADM.SC.EDU/GLENN/IDR/DKIDR.HTM](http://DMSWEB.BADM.SC.EDU/GLENN/IDR/DKIDR.HTM). This web page also contains links to all experimental instructions, data, and software to replicate our results. For the convenience of Danish-challenged readers we also provide on this web site an English translation of all instructions and questionnaires.

³ At the time of the experiments the inflation rate in Denmark was just under 2% per annum, and had been steady for several years. It rose to 3% per annum by the end of the longest horizon used in our experiments. The realized rates of inflation, taking the front-end delay into account, were 0.3%, 1.2%, 3.2% and 6.1% for the 6, 12, 24 and 36 month horizons, respectively..

1. Experimental Design

The basic question used to elicit individual discount rates is extremely simple: do you prefer \$100 today or \$100+ x tomorrow, where x is some positive amount? If the subject prefers the \$100 today then we can infer that the discount rate is higher than $x\%$ per day; otherwise, we can infer that it is $x\%$ per day or less. The format of our experiment modifies and extends this basic question in six ways.

First, we pose a number of such questions to each individual, each question varying x by some amount. When x is zero we would obviously expect the individual to reject the option of waiting for no rate of return. As we increase x we would expect more individuals to take the future income option. For any given individual, the point at which they switch from choosing the current income option to taking the future income option provides a bound on their discount rate. That is, if an individual takes the current income option for all x from 0 to 10, then takes the future income option for all x from 11 up to 100, we can infer that their discount rate lies between 10% and 11% for this time interval. The finer the increments in x , the finer will we be able to pinpoint the discount rate of the individual.

Second, we simultaneously pose several questions with varying values of x , selecting one question at random for actual payment after all responses have been completed by the individual. In this way the results from one question do not generate income effects which might influence the answers to other questions. Although one could allow for these effects in the later analysis, they could easily cause more statistical problems than the extra data is worth.

Third, we provide two future income options rather than one “instant income” option and one future income option. For example, we offer \$100 in one month and \$100+ x in 7 months, interpreting the revealed discount rate as applying to a time horizon of 6 months. This avoids the potential problem of the subject facing extra transactions costs⁴ with the future income option. If the delayed option were to involve greater transactions costs, then the revealed discount rate would include these subjective transactions costs. By having both options entail future income we hold these transactions costs constant.

Fourth, we consider four possible time horizons: 6 months, 12 months, 24 months and 36 months. In one series of experiments subjects were randomly assigned to a session in which they were asked to consider one of these time horizons. In these sessions we only elicited discount rates pertaining to that horizon. In another series, with different subjects, we ask the subject to state preferences over all four time horizons, knowing that we will select one time horizon at random for possible payment. A comparison of these two series will allow some evaluation of the effect of explicitly asking subjects to consider multiple time horizons. It is plausible that this could mitigate any tendency for subjects to reveal time-inconsistent

⁴ Including the possibility of default by the experimenter.

discount rates.

Fifth, we elicit information from subjects to help us identify what market rates of interest they face. This information will be used to allow for the possibility that their responses in our surveys are *censored* by market rates. To explain the censoring problem, assume that you value a cold beer at \$3, which is to say that if you had to pay \$3 for one beer you would. If I ask you whether or not you are willing to pay \$2.50 for a *lab* beer, your response to me will depend on whether or not there is a market price of *field* beer⁵ lower than \$2.50. If the market price of the field beer is \$2.00, and you know that you can buy a beer outside the lab at this price, then you would never rationally reveal to me that you would pay \$2.50 for my lab beer. In this case we say that your response is censored by the market price (Harrison [1992; p.1432]). Fortunately, there are simple statistical procedures for allowing for this possibility, and we employ those in our statistical analysis.

It is easy to see how this censoring problem applies here. Consider a subject with a true individual discount rate (IDR) of 30%. In the absence of field substitutes for lab incentives, we would expect this subject to choose to save in the lab when the lab instrument provides a rate of return of 30% or higher. Now assume that this subject can *borrow* in the field at a rate of 14%. Although she demands at least 30% interest to delay consumption and save in the lab, at rates between 14% and 30% she is better off borrowing in the field at 14% and not delaying consumption in the field, leaving the money in the lab earning 14% or more, and repaying the field debt at the time she collects from the experimenter. In this case, the subject should rationally choose to invest in the lab when the lab instrument provides a rate of return of 14% or more. Hence, censoring would imply that the true IDR *could* actually be greater than or equal to the observed borrowing rate when we observe lab investment responses that suggest that the IDR is equal to the borrowing rate.⁶ In other words, if we ignored the possibility of censoring of lab responses we would incorrectly infer that this subject had an IDR of 14%. Instead, we can only infer from these lab responses that the subject has a true IDR between 14% and ∞ . The problem is symmetric for censoring with respect to savings rates, although less significant empirically.⁷

The implication of allowing for censoring is that we cannot presume that the “raw” responses in the

⁵ Assume further that a beer in the lab is the same product as a beer in the field.

⁶ When the subject reports an IDR interval that exceeds the borrowing rate that we calculate for the subject, we assume that there are subjective and unobserved transactions costs such that the true (unobserved) market rate for the subject is equal to the lower bound of the reported interval. The subject’s responses are then treated statistically as being censored at that inferred borrowing rate.

⁷ Consider, for example, a subject with a true IDR of 3%. In the absence of field substitutes for lab incentives, we would again expect this subject to choose to invest in the lab instrument as long as it provides a return of 3% or higher. Now suppose that this subject can *save* in the field at a rate of 10%. Although she would be willing to save at 3%, at rates between 3% and 10% she is better off investing in the field and refusing to invest in the lab. Hence censoring would imply that the true IDR could actually be less than or equal to the observed savings rate when we observe lab investment responses that suggest that the IDR is close to the savings rate.

lab are unbiased indicators of the true IDR of the subject. Moreover, if we ignored field censoring then we could easily be led to think that we were measuring responses with more precision than would be warranted.

Sixth, we provide respondents with the interest rates associated with the delayed payment option. This is an important control feature if field investments are priced in terms of interest rates. If subjects are attempting to compare the lab investment to their field options, this feature may serve to reduce comparison errors since now both lab and field options are priced in the same metric.⁸

2. The Danish Experiments

2.1 Sample

In 1996 the Danish Ministry of Business and Industry contracted with the Danish Social Research Institute (SFI, after the Danish name *Socialforskningsinstituttet*) to undertake the field surveys.⁹ The final surveys were conducted between June 16 and July 8, 1997, throughout Denmark.

The sample population consisted of a random selection from individuals 19-75 years old who had participated all three times in the European Community Household Panel Survey (ECHP) previously conducted by SFI. These persons were chosen because they had some experience with respect to economic surveys, and because we could expect a high response rate. The sample was constructed in two steps.

The 275 municipalities in Denmark were proportionally stratified with respect to the number of persons between 19 and 75 years of age on January 1, 1997. Copenhagen and Aarhus, the two largest municipalities, had their own stratum due to their size. Most of the other municipalities were divided into 23 strata. Some remote municipalities, primarily tiny islands, were not represented in the sample because the population is relatively small and the subjects would spend too much time on traveling to the experimental session.

The 27 sessions were divided equally across geographic locations with 5, 10 and 15 participants in each experiment. In turn, the 27 sessions were located such that the number of participants at the experiments correspond to the relative size of the population in the given stratum. For example, approximately 11% of the population between 19 and 75 years of age live in Copenhagen, and 3 sessions with

⁸ Collier and Williams [1999] suggest that behavior in these studies may be affected by uncontrolled factors other than time preferences that may help explain observed anomalies. They suggest that subjects may attempt to *arbitrage* between lab and field investment opportunities, but may make mistakes in comparing these opportunities because the lab and field investments are “priced” in different terms. Lab investments are priced in *dollar* interest (the difference between the early and later payments), while field investments are priced in terms of annual and effective interest *rates*. A rational subject should never choose to postpone payment in the laboratory at interest rates lower than those she can receive in the external market, for example, but she may make mistakes in converting dollar interest to an interest rate (or vice versa) for the purposes of comparison. The use of hypothetical or small payments is likely to exacerbate this problem because of the cognitive costs associated with the subject’s arbitrage problem; at lower stakes subjects are likely to expend less cognitive effort on getting the comparison right.

⁹ At the time, Harrison was Director of the MobiDK Project, within the Ministry. Lau was a Senior Researcher with the MobiDK Project.

a total of 30 participants were held in Copenhagen which corresponds to 11.1% of the total sample size.¹⁰

Most strata consist of several municipalities, and the strata were constructed according to traffic connections. The sessions were held in the evening to facilitate attendance by working subjects. It was important that the participants not spend too much time on traveling in order to join the experiments. In some cases, it was necessary to divide a given stratum into two subgroups, since the distance between some potential participants and the location of the session would otherwise be too great. Accordingly, a random draw from the subgroups was made, weighting the two subgroups with respect to the relative size of the population between 19 and 75 years of age.

The interviewers initially contacted 6, 12 or 17 persons, the number depending on the specific session and assuming a show-up rate of approximately 80%. If a respondent declined to participate, the interviewers contacted a “stand-in” roughly the same age. Hence, either 6, 12 or 17 persons were confirmed before the experiment took place. However, some persons did not show up at the sessions and the actual number of participants varied accordingly.

A total of 268 subjects participated in the experiments. The sample was designed to be equally split between single-horizon and multiple-horizon treatments, and then equally split by time horizon within the single-horizon treatments. All subjects were randomly assigned to treatment condition.¹¹ The sample was representative of the adult population of Denmark, due to the stratified sampling methods employed.

2.2 Primary Experimental Instructions

Apart from logistical correspondence between SFI and the subject concerning attendance at the session, the only information that the subject received was from the survey instrument administered in the experiment. The initial contact letter to the subjects posed the general nature of the task, and informed subjects that they would be paid 500 DKK after participating in the survey and that one subject would receive at least 3000 DKK. No other details of the experiment were provided until the subjects arrived at the session.

Upon arrival at the experimental session, subjects were given the following information:

One person in this room will be randomly chosen to receive a large sum of money. If you are the individual chosen to

¹⁰ It is possible that some subjects were confused as to whether they lived in Copenhagen or Greater Copenhagen, so we have tended to lump these together in the statistical analysis. The area called Copenhagen in the survey covers 3 communes: Copenhagen, Frederiksberg and Gentofte. The total population in this area is 600,000 people, which is around 11% of the total population. Three sessions in Copenhagen with 27 subjects in total matches this share well. Some of the sessions referring to Zealand cover some of the suburbs in Copenhagen. The population in Copenhagen, including all suburbs, is 1.35 million, which is around 26% of the total population. We suspect that some subjects who live in the suburbs write that they live in Copenhagen instead of the Greater Copenhagen area.

¹¹ Due to the vagaries of no-shows, the actual sample differs slightly from this design. There were 118 subjects in the 15 single-horizon experiments, and 150 subjects in the 12 multiple-horizon experiments. Within the single-horizon experiments there were 26, 32, 31 and 29 subjects, respectively, in the 6 month, 12 month, 24 month and 36 month treatments.

receive this money (the "Assignee"), you will have a choice of *two* payment options; option A or option B. If you choose option B you will receive a sum of money 7 months from today. If you choose option A, you will receive a sum of money 1 month from today, but this option (A) will pay a smaller amount than option B.

Subjects were given payoff tables as illustrated in Table 1. They were told that they must choose between payment Options A and B for each of the 20 payoff alternatives. Option A was 3000 DKK in all sessions. Option B paid 3000 DKK + X DKK, where X ranged from annual rates of return of 2.5% to 50% on the principal of 3000 DKK, compounded quarterly to be consistent with general Danish banking practices on overdraft accounts. The payoff tables provided the annual and annual effective interest rates for each payment option and the experimental instructions defined these terms by way of example. Subjects were then told that a single payment option would be chosen at random for payment, and that a single subject would be chosen at random to be paid his preferred payment option for the chosen payoff alternative. The payment mechanism was explained as follows:

HOW WILL THE ASSIGNEE BE PAID?

The Assignee will receive a certificate which is redeemable under the conditions dictated by his or her chosen payment option under the selected payoff alternative. This certificate is guaranteed by the Social Research Institute. The Social Research Institute will automatically redeem the certificate for a Social Research Institute check, which the Assignee will receive given his or her chosen payment option under the selected payoff alternative. Please note that all payments are subject to income tax, and information on all payments to participants will be given to the tax authorities by the Social Research Institute.

Finally, prior to the choice task, the experimenter illustrated the randomization devices in a trial experiment which utilized different quantities of candies as payoffs. The trial Assignee was paid his candies at the end of the trial experiment, to illustrate the concrete nature of the payoffs.

The instructions for the 12-month, 24-month and 36-month horizon experiments were identical except for the obvious changes. The instructions for the multiple-horizons sessions were similar, with the single change that the subject was asked to provide responses for all four time horizons. All four time-horizons were presented simultaneously to the subject, who could respond to them in any order.¹² One time horizon was then selected for possible payment, and the remaining procedures were identical to the single-horizon sessions.

Across all time horizons, payoffs to any one subject could range from 3,000 DKK up to 12,333 DK. The exchange rate in mid-1997 was approximately 6.7 DKK per US dollar, so this range converts to \$450 and \$1,840.

2.3 Additional Experimental Questionnaires

In addition to the primary elicitation task, we collected information from subjects on a variety of socio-demographic characteristics. Specifically, we collected information on age, gender, size of town the

¹² The literal sequence of the time-horizon payoff tables in the survey instrument was the natural one, with the 6-month horizon coming first.

subject resided in, type of residence, primary occupation during the last 12 months, highest level of education, household type (viz., marital status and presence of younger or older children), number of people employed in the household, total household income before taxes, disposable household income, whether the subject is a smoker, and the number of cigarettes smoked per day.

We also elicited information on a number of financial variables to help us identify the market circumstances within which the discount rate responses should be viewed. Specifically, we collected information on whether the subject had various accounts (e.g., checking account, credit card, line of credit), the annual interest rate on those accounts, and the current balance. We also collected information on the subject's perception of his or her chances of obtaining a loan, line of credit or credit card.

3. Results

Our null hypotheses are that the discount rates for given time horizons do not differ across households, and that the discount rates for given households do not differ across time horizons.

3.1 Statistical Analysis

After removing subjects that gave incomplete or inconsistent responses, the final sample consists of 109 observations spread across the four single-horizon sessions, and 132 observations on the multiple-horizon sessions.¹³ The statistical analysis takes into account four features of these data.¹⁴ First, we account for the fact that we observe only interval-censored responses, rather than precise values of the IDR. Thus a subject that switched from A to B in option 8 would be viewed as choosing an annual effective rate in the interval (18.68%, 21.55%]. Second, we account for the stratification of our national sample, as described earlier. Third, we account for the “panel data” feature of our experiments in which some subjects provided four sets of responses rather than just one.¹⁵ Finally, we account for the possibility that market responses are censored by market savings and borrowing rates.¹⁶

The explanatory variables included in our statistical model are defined as follows:

¹³ An inconsistent response is one in which the subject switched between A and B more than once. This occurred in only 3% of the responses, reflecting 4% of the subjects. The remaining sample reductions are from subjects that neglected to answer some core demographic question.

¹⁴ Because of these statistical issues, we refer to the discount rates that are *predicted* by the regression model as the *elicited* discount rates. That is, some statistical analysis is needed to infer the discount rate that is implied by the raw response to the experimental instrument.

¹⁵ This feature amounts to a multi-stage sampling design in which there are up to four observations for each “primary sampling unit,” which in our case is the individual subject. The regression procedure we use allows for any amount of correlation within the observations for each primary sampling unit. See StataCorp [2001; User's Guide, p.324].

¹⁶ We estimate an interval regression model recognizing the features of the complex survey design used, employing version 7 of *Stata* documented in StatCorp [2001].

- ❑ T6, T12, T24 and T36: binary indicators¹⁷ of the 6-month, 12-month, 24-month, and 36-month time horizons, respectively;
- ❑ FEMALE: binary indicator if the subject was a female;
- ❑ YOUNG: binary indicator if the subject was less than 30 years old;
- ❑ MIDDLE: binary indicator if the subject was between 40 and 50 years old;
- ❑ OLD: binary indicator if the subject was greater than 50 years old;
- ❑ LOWER MIDDLE: disposable household income in 1996 between 100,000 and 199,999 Danish kroner;
- ❑ UPPER MIDDLE: disposable household income in 1996 between 200,000 and 299,999 Danish kroner;
- ❑ RICH: disposable household income in 1996 greater than or equal to 300,000 Danish kroner;
- ❑ SKILLED: binary indicator that the subject has completed more than the basic primary and secondary education in Denmark (i.e., completed more than “Basic school, General upper secondary education, and/or Vocational upper secondary education”);
- ❑ LONGEDU: binary indicator that the subject has completed some substantial higher education (referred to in Denmark as “medium-cycle or longer-cycle higher education”);
- ❑ COPEN: binary indicator that the subject lives in Copenhagen, including “Greater Copenhagen and its suburbs”;
- ❑ TOWN: binary indicator that the subject lives in a town with 10,000 or more inhabitants other than Copenhagen;
- ❑ OWNER: binary indicator that the subject lives in an apartment or house that they own;
- ❑ RETIRED: binary indicator that the subject is retired;
- ❑ UNEMP: binary indicator that the subject is unemployed;
- ❑ SINGLE: binary indicator that the subject lives alone, where the subjects were told that a “household is an economic unit, defined as a group of persons who live in the same residence where each person contributes to general expenditures”;
- ❑ KIDS: binary indicator that the subject lives with children;
- ❑ MULTIPLE: binary indicator that the subject gave responses in a multiple-horizon session;
- ❑ GSIZE: variable indicating the size of the group that attended the session that the subject participated in;
- ❑ BALANCE: binary indicator that the subject carries a positive balance in a line of credit¹⁸ or credit card; and

¹⁷ As a matter of convention we code all binary indicators with the Boolean interpretation in which a 1 denotes “true” and 0 denotes “false.” For example, T6=1 if the observation pertains to the 6-month horizon, and 0 otherwise.

¹⁸ It is common for Danes to carry a pre-arranged personal line of credit at a bank, so we view this as being similar to the credit card balances that Americans might carry in terms of convenience of access.

- CHANCES: binary indicator that the subject believes that the chances of getting a line of credit or credit card approved if they went to a bank are poor (less than 75% likely).

The characteristics employed in our statistical analysis are generally those also used by *Denmark's Statistics* in their household expenditure surveys.¹⁹

The regression results are presented in Table 2. The overall significance of the regression equation is provided by an adjusted Wald test statistic of the null hypothesis that all coefficients other than the constant are equal to zero. We reject this null hypothesis at any standard level.

The average discount rate elicited over all subjects is approximately 28%. Before examining how these rates vary with the experimental treatments, the absolute level of the elicited rate should be noted. Relative to the extensive experimental literature in which discount rates are elicited with a variety of hypothetical questions, this average is actually quite low. On the other hand, compared to discount rates popularly used in welfare analysis (roughly between 3% and 10%) these rates seem relatively high. Several factors might account for the absolute magnitude of the elicited rates.

First, despite our extensive attempts to encourage credibility, the subjects might have doubted that we would actually follow through on the payments.²⁰ These are, after all, artificial and constructed payment options. This uncertainty could plausibly have encouraged subjects to view these as “risky” prospects, in turn encouraging them to require a higher rate of return before investing for any longer time period. This particular credibility effect would likely be additive on the elicited discount rates over all time horizons, increasing all elicited discount rates by some fixed amount (e.g., 10 percentage points) to offset the “default risk.” The reason that this effect would be constant across time horizons is that the risk of default would not be likely to vary with the time horizon.

Second, since we elicited discount rates over real monetary amounts and operated with a finite budget, we were forced to constrain the amounts of money involved. Compared to many laboratory experiments with real payments, our field experiments use quite large amounts. Nonetheless, the subjects may have perceived these as small amounts of money. Whether or not that leads to a change in revealed discount rates is an open question, but *a priori* folklore amongst experimenters suggests that subjects might not take foregone income seriously if it falls below some subjective threshold. This could lead the subjects not to respond to the incentives offered by foregoing near-term consumption in our experiments.

¹⁹ These are standard classifications, but also have the advantage of allowing us to map the results into other databases and models that use these classifications for welfare analyses. Specifically, we plan to use these elicited rates to extend the calibration of “generational accounts” for Denmark and computable general equilibrium models for Denmark that represent households as inter-temporal utility maximizers.

²⁰ It is true that the Ministry of Business and Industry changed its name to the Ministry of Trade and Industry within the time horizon of the instruments being proffered, but this would not have been known at the time the experiments were conducted, and was largely a superficial change.

We attempt to control for the effect of varying incentives by including the variable *G*SIZE in our regression model. Expected payments to subjects varied with the size of the group they participated in, since this (inversely) scaled the probability that the subject would be selected as the one person to actually play out their choices for real payment. By controlling for this variable in the regression model, and generating predictions for the case in which group size was counter-factually assumed to be one, we can ascertain what the regression model predicts would be the elicited discount rate if the probability of being selected was one.

3.2 Elicited Discount Rates

The regression results are reported in Table 2. The coefficients in Table 2 may be directly interpreted as the marginal effect of each variable. In particular, the coefficients on the first four variables show the estimated discount rate for different horizons, since we do not have a constant term. An alternative way to view the effects of demographics is to generate predicted discount rates for everyone in the sample and then to stratify these predicted rates. These results are shown in Table 3. The demographic results in Table 3 show the effect of varying the indicated variable and all other characteristics that are associated with it. Thus, if women are better educated on average than men in Denmark, the effect of sex in Table 3 will include the effect of this difference in education whereas the marginal effect on that coefficient in Table 2 will not. We report both sets of demographic breakdowns since each is of policy interest.

Table 2 indicates that there was some difference in the estimated discount rates for the 6-month horizon compared to the others. The 6-month rate is roughly 6 percentage points higher, at 34.9%, whereas the estimated discount rates for other horizons are virtually identical at around 28%. An F-test confirms these claims. The only demographic characteristics that appear to matter in Table 2 are (i) the length of education, which is associated with a discount rate over 9 percentage points lower than otherwise; (ii) retirement, which is associated with a discount rate over 12 percentage points higher than otherwise; and (iii) unemployment, which is associated with a discount rate just over 7 percentage points lower than otherwise. In addition, if the individual perceives that they have a *poor* chance of getting a loan or credit card approved at a bank, their discount rate is over 7 percentage points higher. In each case we can plausibly entertain hypotheses that allow the causality to go both ways. In fact, one of the motivating policy forces behind our survey was a concern that Danes did not invest enough in education. Our results suggest that those that do invest in education may do so because they simply have a lower discount rate, and are more willing to trade off near-term costs for longer-term payoffs.

Although the individual coefficients do not indicate statistical significance at the conventional levels, we also observe a lowering of the estimated discount rate as individuals get richer. Again, we only observe an association; it is possible that individuals are richer because they are willing to invest more than others. In

fact, this could be correlated with investments in education. For this reason it is appropriate to examine the fully stratified results in Table 3, which show the joint effects of each demographic characteristic and those other characteristics correlated with it.

Table 3 generates several interesting results,²¹ complementing the marginal effects of Table 2:

- The overall individual discount rate in Denmark is estimated to be 28.1%. This reflects the stratification of our sample in order to obtain an efficient estimate of the national average. Figure 1 displays the distribution of estimated discount rates, which is roughly normal.
- The discount rates for men and women appear to be identical, confirming the marginal effects in Table 2. This result is particularly notable since many other characteristics vary with sex.
- Discount rates appear to decline with age, at least after middle age.
- There does appear to be a significant lowering of the discount rate for richer individuals, when we allow for richer households to “carry with them” the other characteristics they typically have, such as more education. The richest household have discount rates that are over 10 percentage points lower than the poorest households. Again, we must remain agnostic on the matter of causality, but this strong finding fuels the debate as to whether the poor are poor because they have a higher discount rate or have a higher discount rate because they are poor. This difference between Table 3 and Table 2 illustrates the potential importance of examining demographic effects both ways.
- Table 3 also shows a large difference between the discount rates of skilled and unskilled individuals, with those that have skills having a significantly lower discount rate.
- Perhaps surprisingly, students have a higher discount rate than non-students.
- The importance of the extent of education from Table 2 is confirmed in Table 3: those with longer investments in education are also those with substantially lower discount rates.
- Ownership of a house is associated with having a lower discount rate, which is consistent with the notion of owners as being more settled into their lifestyle than apartment dwellers.
- Retired individuals have higher discount rates, confirming the marginal effect from Table 2.
- The unemployed have lower discount rates than the employed.
- Finally, poor chances of being turned down for a loan or credit card by a bank are associated with the individual having much higher discount rates, as one would expect. This result also appeared in Table 2, and seems to cut across other demographics.

²¹ The total sample in Table 3 is listed as 696, even though some observations were deleted in the regression analysis in Table 2. The reason is that the complete sample is utilized when adjusting the standard errors for the sample stratification.

3.3 Comparison to the Literature

There have been several attempts to estimate discount rates for individuals in field settings using financial instruments.²² All of them find relatively high discount rates.

Ausubel [1991; Table 11, p.70] shows that nearly three-quarters of those holding credit cards in banks he surveyed do not pay off their balance on time and avoid finance charges, despite the fact that those finance charges amount to roughly 19% per annum. We find that subjects in our experiments that hold comparable balances in Denmark have essentially the same discount rates as those that do not hold such balances (see Tables 2 and 3).

Warner and Pleeter [2001] estimate individual discount rates for a large number of U.S. military personnel who were offered voluntary separation options. One option was an initial lump-sum payment, and the other was an annuity. They estimate (Table 6, p. 48) that officers had an average discount rate of between 10% and 19%, depending on the statistical specification assumed, and that enlisted personnel had discount rates between 35% and 54%.

Although these findings refer to selected segments of the population, albeit large segments with a diverse range of socio-demographic characteristics, they suggest that the level of discount rates that we find for the Danish population is consistent with other field evidence.

4. Conclusions

We demonstrate that it is possible to elicit discount rates from individuals in the field using real economic commitments, and that those discount rates are in an *a priori* plausible range. There are variations in discount rates across some socio-demographic characteristics of the Danish population, implying that inter-temporal welfare evaluations for those household groups should take these differences into account. On the other hand, elicited discount rates do not vary with respect to the time horizon used here beyond one year, consistent with the use of constant discount rates for *given* household types for those horizons.

²² There are also numerous studies estimating large discount rates implicit in the purchase of alternative consumer durables, and numerous laboratory studies using student subjects that utilize financial instruments also find large discount rates. Ruderman, Levine and McMahon [1986] review the former, and Coller and Williams [1999] review the latter. The only laboratory experiments with lower discount rates, that we are aware of, are those of Coller and Williams [1999], whose design we employed here. They find rates for American college students in the 15% to 18% range.

References

- Ausubel, Lawrence M., "The Failure of Competition in the Credit Card Market," *American Economic Review*, 81(1), March 1991, 50-81.
- Coller, Maribeth, and Williams, Melonie B., "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 1999, 107-127.
- Harrison, Glenn W., "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review*, 82, December 1992, 1426-1443.
- Ruderman, Henry; Levin, Mark, and McMahon, James, "Energy-Efficiency Choice in the Purchase of Residential Appliances," in W. Kempton and M. Neiman (eds.), *Energy Efficiency: Perspectives on Individual Behavior* (Washington, DC: American Council for an Energy Efficient Economy, 1986).
- StataCorp, *Stata Statistical Software: Release 7.0* (College Station, TX: Stata Corporation, 2001).
- Warner, John T., and Pleeter, Saul, "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review*, 91(1), March 2001, 33-53.

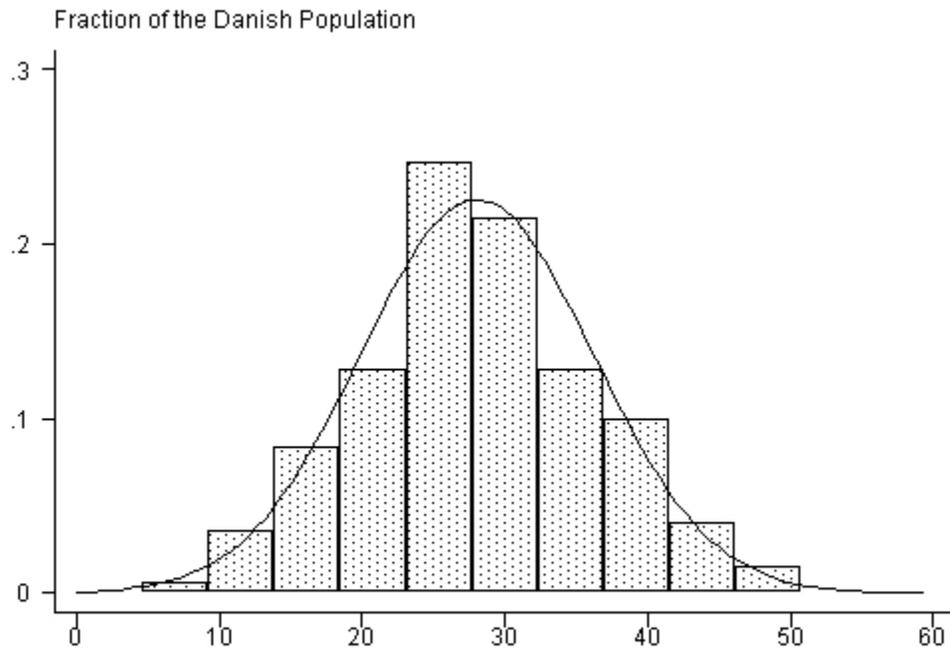


Figure 1: Estimated Discount Rates for the Danish Population

Table 1: Payoff Table for 6 Month Time Horizon

Payoff Alternative	Payment Option A (pays amount below in 1 month)	Payment Option B (pays amount below in 7 months)	Annual Interest Rate (AR)	Annual Effective Interest Rate (AER)	Preferred Payment Option (Circle A or B)
1	3,000 DKK	3,038 DKK	2.50%	2.52%	A B
2	3,000 DKK	3,075 DKK	5.00%	5.09%	A B
3	3,000 DKK	3,114 DKK	7.50%	7.71%	A B
4	3,000 DKK	3,152 DKK	10.00%	10.38%	A B
5	3,000 DKK	3,190 DKK	12.50%	13.10%	A B
6	3,000 DKK	3,229 DKK	15.00%	15.87%	A B
7	3,000 DKK	3,268 DKK	17.50%	18.68%	A B
8	3,000 DKK	3,308 DKK	20.00%	21.55%	A B
9	3,000 DKK	3,347 DKK	22.50%	24.47%	A B
10	3,000 DKK	3,387 DKK	25.00%	27.44%	A B
11	3,000 DKK	3,427 DKK	27.50%	30.47%	A B
12	3,000 DKK	3,467 DKK	30.00%	33.55%	A B
13	3,000 DKK	3,507 DKK	32.50%	36.68%	A B
14	3,000 DKK	3,548 DKK	35.00%	39.87%	A B
15	3,000 DKK	3,589 DKK	37.50%	43.11%	A B
16	3,000 DKK	3,630 DKK	40.00%	46.41%	A B
17	3,000 DKK	3,671 DKK	42.50%	49.77%	A B
18	3,000 DKK	3,713 DKK	45.00%	53.18%	A B
19	3,000 DKK	3,755 DKK	47.50%	56.65%	A B
20	3,000 DKK	3,797 DKK	50.00%	60.18%	A B

Table 2: Regression Analysis of Discount Rate Responses

Variable	Coef.	Std. Err.	t	P> t	[90% Conf. Interval]	
t6	34.86076	7.908359	4.41	0.000	21.8014	47.92012
t12	28.95233	7.976701	3.63	0.000	15.78012	42.12454
t24	27.44078	8.018661	3.42	0.001	14.19928	40.68228
t36	27.87162	8.046035	3.46	0.001	14.58491	41.15832
multiple	.8359218	2.228436	0.38	0.708	-2.843975	4.515818
female	1.014945	2.713695	0.37	0.709	-3.466278	5.496168
young	-1.094671	3.934629	-0.28	0.781	-7.592065	5.402722
middle	.1785973	3.446215	0.05	0.959	-5.512261	5.869455
old	-.4595653	3.754661	-0.12	0.903	-6.659771	5.740641
middle1	-1.305936	3.674648	-0.36	0.723	-7.374014	4.762143
middle2	-3.214197	4.309141	-0.75	0.456	-10.33004	3.901641
rich	-5.341135	4.102213	-1.30	0.194	-12.11527	1.432997
skilled	.7426614	3.275909	0.23	0.821	-4.666965	6.152288
student	4.204929	5.285858	0.80	0.427	-4.523798	12.93366
longedu	-9.202757	3.174322	-2.90	0.004	-14.44463	-3.960884
copen	-1.13076	3.209827	-0.35	0.725	-6.431263	4.169742
town	3.171888	2.845343	1.11	0.266	-1.52673	7.870505
owner	-3.764708	3.030948	-1.24	0.215	-8.769821	1.240406
retired	12.37832	5.048285	2.45	0.015	4.041905	20.71473
unemp	-7.769304	4.437314	-1.75	0.081	-15.0968	-.4418082
single	-2.401655	3.009327	-0.80	0.426	-7.371065	2.567755
kids	.2497801	3.11824	0.08	0.936	-4.899481	5.399041
gsize	.0238708	.3650134	0.07	0.948	-.5788889	.6266305
balance	1.829445	2.61292	0.70	0.485	-2.485364	6.144253
Chances	7.648062	3.996732	1.91	0.057	1.048115	14.24801

Table 3: Average Elicited Discount Rates Stratified by Major Demographics

Demographic Characteristic	Estimate	Std. Err.	[90% Conf. Interval]		Obs
ALL	28.1464	.53537	27.26233	29.03048	696
Male	28.06626	.76262	26.80692	29.3256	336
Female	28.22121	.7667374	26.95507	29.48735	360
Young	28.71521	.9551633	27.13791	30.2925	146
Middle	28.35924	.8708021	26.92125	29.79722	199
Old	25.05474	1.065985	23.29444	26.81503	158
Magnificent	30.02767	1.256172	27.95331	32.10203	193
Poor	32.92452	1.014352	31.24948	34.59955	171
Lower_Middle	30.08146	.676202	28.96482	31.19809	280
Upper_Middle	22.68201	.7520371	21.44014	23.92387	126
Rich	22.51315	1.251744	20.4461	24.5802	119
Unskilled	31.42633	.7387784	30.20636	32.6463	295
Skilled	25.73349	.6889163	24.59586	26.87113	401
Not a Student	27.48244	.5661343	26.54756	28.41732	621
Student	33.64402	1.291917	31.51063	35.7774	75
Less_Educated	30.9838	.547016	30.0805	31.88711	506
More_Educated	20.58996	.7659382	19.32514	21.85479	190
Not Copenhagen	28.50351	.5887187	27.53133	29.47568	531
Copenhagen	26.99719	1.236626	24.9551	29.03927	165
Not in a Town	26.79067	.7654368	25.52668	28.05466	388
Town	29.85428	.7091534	28.68323	31.02533	308
Not an Owner	31.66546	.7322497	30.45627	32.87465	291
Owner	25.6179	.6893576	24.47953	26.75626	405
Active	26.52264	.4946091	25.70587	27.3394	603
Retired	38.67471	1.029985	36.97386	40.37557	93
Working	28.38739	.5465463	27.48486	29.28992	655
Unemployed	24.29656	1.699674	21.48983	27.1033	41
Married	27.47882	.7236279	26.28387	28.67377	453
Single	29.39091	.8189498	28.03855	30.74328	243
No Children	28.89642	.7119442	27.72076	30.07208	431
Have Children	26.92657	.8296289	25.55658	28.29657	265
No Balance	28.19078	.7941139	26.87943	29.50213	387
Use a Balance	28.09083	.7535453	26.84647	29.33518	309
Good Chances	27.10611	.5349895	26.22266	27.98956	611
Poor Chances	35.62428	1.365615	33.36919	37.87937	85

Estimating Individual Discount Rates for the United States: Inferences from a Natural Field Experiment

by

Glenn W. Harrison and Eric Johnson[†]

INCOMPLETE DRAFT #7: January 2002

Abstract. Warner and Pleeter [2001] use the observed behavior from a large sample of early retirement decisions by U.S. military personnel to infer individual discount rates from what amounts to a natural field experiment. The stakes were huge, the sample large and diverse, the financial choices relatively simple, and the time horizon varied between 14 and 30 years. We extend the inferences to take into account the uncertainty from estimating a model of individual responses, and show that the estimated discount rates for enlisted personnel are relatively imprecisely estimated while those for officers are relatively precisely estimated. In neither case do we see evidence of discount rates that vary with the time horizons considered, although in the case of enlisted personnel this could just be due to the imprecision of the estimates. We also extend the inferences about military personnel to be able to make formal inferences about the discount rates of the general U.S. adult population.

[†] Department of Economics, Moore School of Business, University of South Carolina, USA. E-mail contacts: HARRISON@MOORE.SC.EDU and EJOHNSON@MOORE.SC.EDU. We are grateful John Warner for providing access to data.

In 1992 the United States Department of Defense started offering substantial early retirement options to nearly 300,000 individuals in the military. This voluntary separation policy was instituted as part of a general policy of reducing the size of the military as part of the “Cold War dividend.” Warner and Pleeter [2001] (WP) recognize how the options offered to military personnel could be viewed as a natural experiment with which one could elicit information on individual discount rates. In general terms, one option was a lump-sum amount and the other option was an annuity. The individual was told what the cut-off discount rate was for the two to be actuarially equal, and this concept was explained in various ways. If an individual is observed to take the lump-sum, one could infer that his discount rate was greater than the threshold rate. Similarly, for those individuals that elected to take the annuity, one could infer that his discount rate was less than the threshold.¹

Four features of this natural experiment make it particularly compelling for the purpose of estimating individual discount rates. First, the stakes were real. Second, the stakes were substantial, and dwarf anything that has been used in laboratory experiments with salient payoffs in the United States. The average lump-sum amounts were around \$50,000 and \$25,000 for officers and enlisted personnel, respectively. Second, the military went to some lengths to explain to everyone the financial implications of choosing one option over the other, making the comparison of personal and threshold discount rate relatively transparent. Third, the options were offered to a wide range of officers and enlisted personnel, such that there are substantial variations in key demographic variables such as income, age, race and education. Fourth, the time horizon for the annuity differed in direct proportion to the years of military service of the individual, such that there are annuities between 14 and 30 years in length. This facilitates evaluation of the classical hypothesis that discount rates are stationary over different time horizons.

WP recognize that one problem of interpretation might arise if the very existence of the scheme signaled to individuals that they would be forced to retire anyway. As it happens, the military also significantly tightened up the rules governing “progression through the ranks” so that the probability of being involuntarily separated from the military increased at the same time as the options for voluntary separation were offered. This background factor could be significant, since it could have led to many individuals thinking that they were going to be separated from the military anyway, and hence deciding to participate in the voluntary scheme even if they would not have done so otherwise. Of course, this background feature could work in any direction, to increase or decrease the propensity of a given individual

¹ This design is essentially the same as one used in a long series of laboratory experiments studying the behavior of college students. See Coller and Williams [1999] for a recent review of those experiments. Comparable designs have been taken into the field, such as the study of the Danish population by Harrison, Lau and Williams [2002]. The only difference is that the field experiment evaluated by WP only offered each individual one discount rate: Harrison, Lau and Williams [2002] offered each subject 20 different discount rates, ranging between 2.5% and 50%.

to take one or other option. In any event, WP allow for the possibility that the decision to join the voluntary separation process itself might lead to sample selection issues. They estimate a bivariate probit model, in which one decision is to join the separation process and the other decision is to take the annuity rather than the lump-sum.

WP conclude that the average individual discount rates implied by the observed separation choices were high relative to *a priori* expectations. Their linear model predicts average rates of 10.4% and 35.4% for officers and enlisted personnel, respectively.

We extend their analysis in four respects. First, we use this field experiment to test the traditional hypothesis that discount rates are constant over different time horizons.² Second, we consider the robustness of the inferences about average discount rates, taking into account the statistical uncertainty of the calculation. Third, we account for the possibility of censoring with respect to market interest rates. Censoring is a potential issue if individuals have access to alternative investment opportunities. Fourth, we augment their analysis to estimate the individual discount rates for the U.S. adult population as a whole, rather than just the military component of that population. We do so by constructing a statistical model of the decision to join the military, and using that statistical model to adjust the responses estimated for those in the military.

In section 1 we replicate the analysis of WP, reviewing the essential features of their calculations. In section 2 we correct for uncertainty about the statistical model used to infer the discount rate, for possible market censoring in this inference, and for the sample selection effects of just having data on the responses of military personnel. We study the implied level of discount rates, how precisely they are inferred, and how they vary with the time horizon of the annuity offered.

1. Replication and Recalculation

We obtained the raw data from John Warner, and were able to replicate the main results with a reasonable tolerance using alternative statistical software.³

We use the same method as WP [2001; Table 6, p.48] to calculate estimated discount rates.⁴ After

² This hypothesis have been rejected in many hypothetical settings, and even in some settings in which there were real economic consequences. Frederick, Loewenstein and O'Donoghue [2002] review the literature.

³ The single probit regression results reported by WP were implemented using *SAS*, and the bivariate probit results implemented using *LIMDEP*. It turns out that the specific bivariate probit model they implemented is a probit model with sample selection modeled as a probit equation as well (Greene [1995; p.466/7]), as their discussion suggests. We replicated all of their findings in version 7 of *Stata*, using the *PROBIT* and *HECKPROB* routines to implement their two types of models. All data and code for our replication is available from [HTTP://DMSWEB.BADM.SC.EDU/GLENN/IDR/USA/](http://DMSWEB.BADM.SC.EDU/GLENN/IDR/USA/). We are grateful to John Warner for answering several questions of detail and providing unpublished computer runs.

⁴ In their Table 3, WP calculate the mean predicted discount rate from a single-equation probit model, using only the discount rate as an explanatory variables, employing a shortcut formula which correctly evaluates the mean discount rate. Specifically, the predicted mean is equal to the estimated intercept divided by the coefficient on the discount rate offered.

each regression model is estimated it is used to predict the probability that each individual would accept the lump-sum alternative at discount rates varying between 0% up to 100% in increments of 1 percentage point. For example, consider a 5% discount rate offered to officers, and the results of the single-equation probit model. Of the 11,212 individuals in this case, the *average* probability of accepting the lump-sum at this rate is 0.79 according to the estimated model. The lowest predicted probability of acceptance for any individual at this rate is 0.207, and the highest is 0.992. There is a standard deviation in the predicted probabilities of 0.13. This standard deviation is taken over all 11,212 individual predictions of the probability of acceptance. It is important to note that it assumes that the estimated coefficients of the probit model are exactly correct; we evaluate this assumption in section 2.

Similar calculations are undertaken for each possible discount rate between 0% and 100%, and the results tabulated. The results are shown in Figure 1. The vertical axis shows the probability of acceptance for the sample, and the horizontal axes shows the (synthetically) offered discount rate. The average, minimum, maximum, and 95% confidence intervals are shown. Again, this is the distribution of predicted probabilities for the sample, assuming that the estimated coefficients of the probit regression model have no sampling error.

Once the predicted probabilities of acceptance are tabulated for each of the 11,212 officers, we loop over each officer and identify the *smallest* discount rate at which the lump-sum would be accepted by that officer. This provides a distribution of estimated *minimum* discount rates, one for each individual in the sample.

In Figure 2 we report the results of this calculation, showing the distribution of personal discount rates initially offered to the subjects and then the distributions that are implied by the single-equation probit model used by WP.⁵ The top panels show the after-tax discount rates that were offered, and the bottom panels show the discount rates inferred from the estimated model. The first column shows the results for officers, and the second column shows the results for enlisted personnel. The horizontal axes in all four charts are identical, to allow a visual reference across the. The main result is that the distribution of estimated discount rates is much wider than the distribution of offered rates. Indeed, for enlisted personnel the distribution of estimated rates is almost entirely out-of-sample in comparison to the offered rates above it. There is nothing “wrong” with these differences, although they will be critical when we start to consider standard errors on these estimated discount rates. Again, the estimated rates in the bottom charts of Figure 2 are based on the logic of Figure 1: no prediction error is assumed from the estimated statistical model when it is applied at the level of the individual to predict the threshold rate at which the lump-sum would be accepted.

⁵ Virtually identical results are obtained with the model that corrects for possible sample-selection effects.

Probability Using Single Equation Probit Model

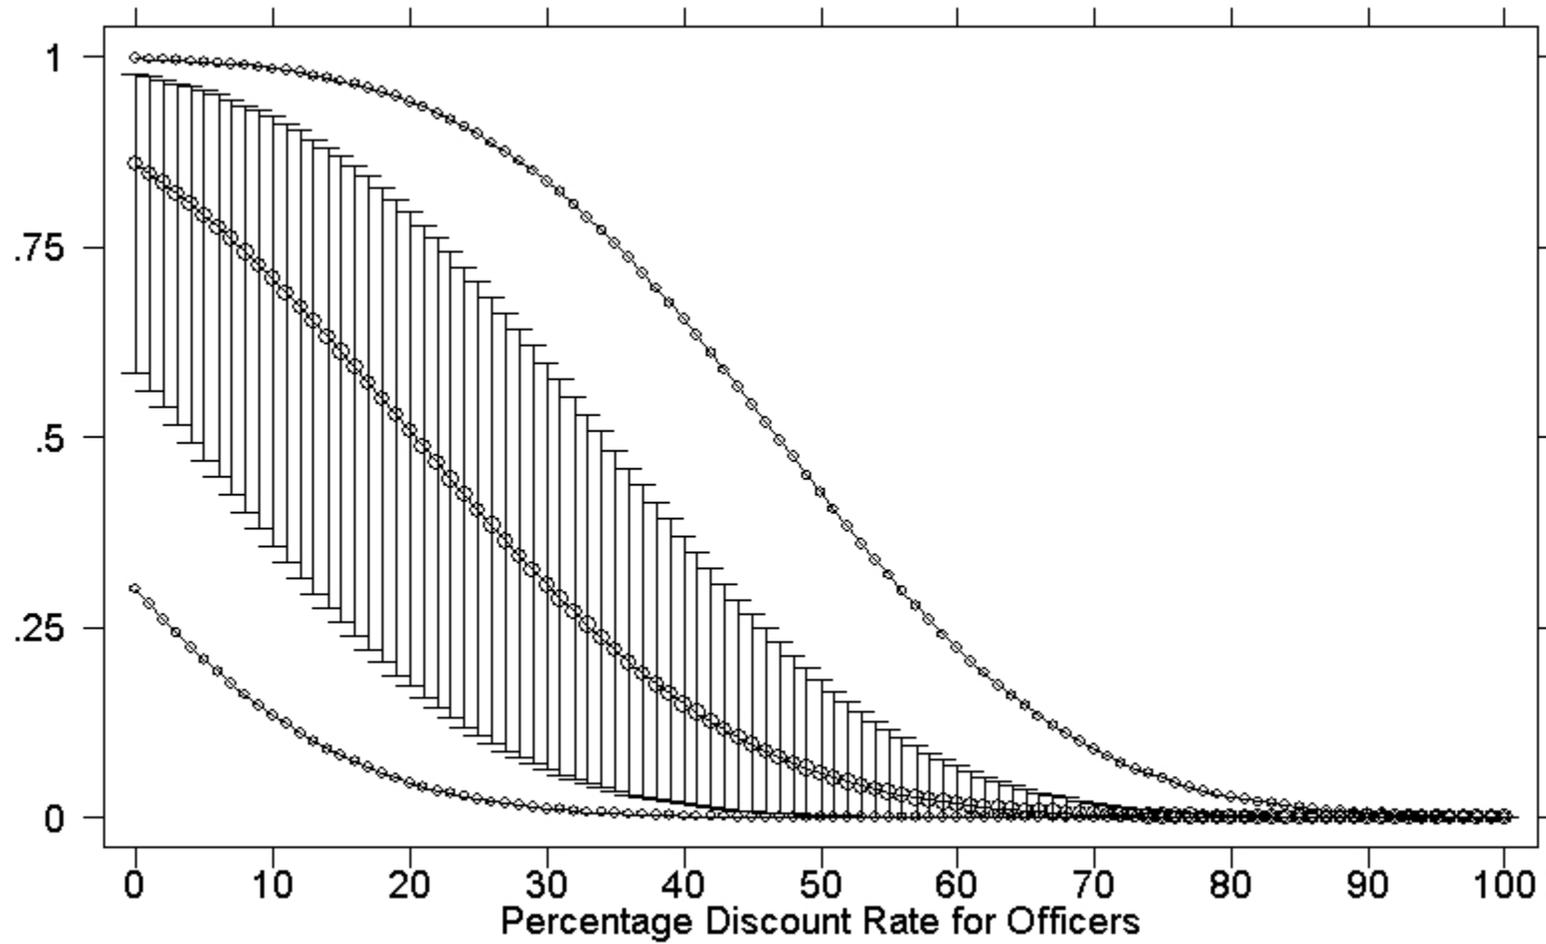
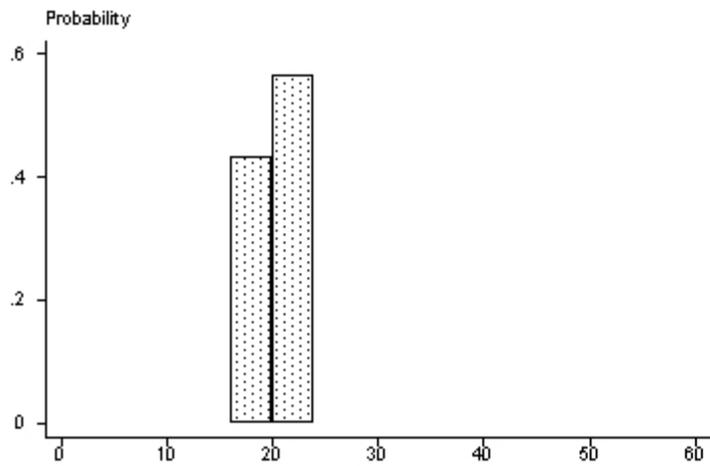
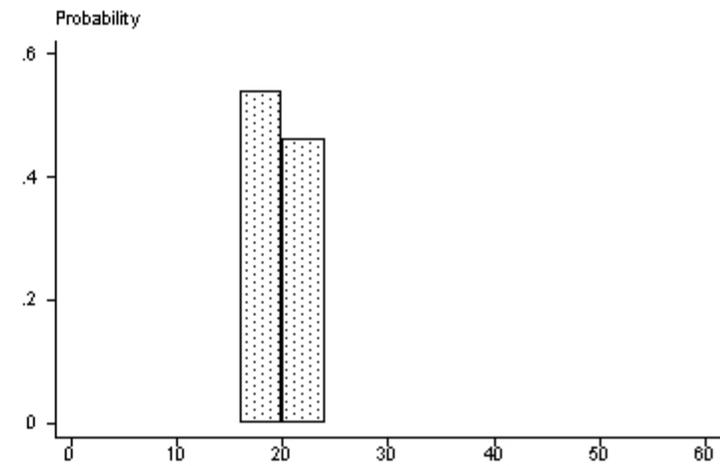


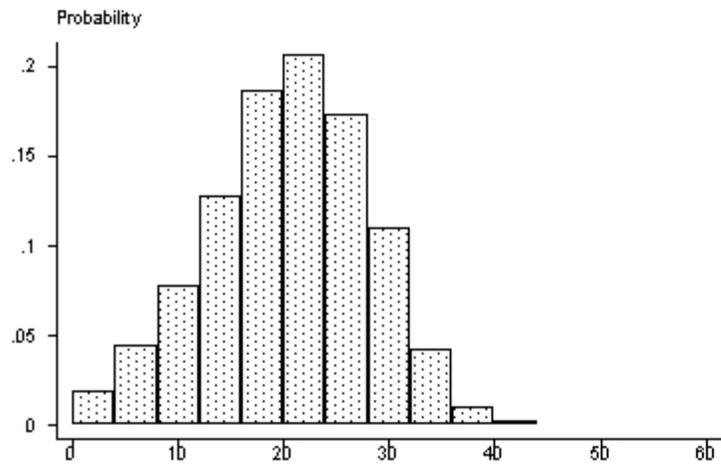
Figure 1: Probability of Acceptance if No Prediction Error



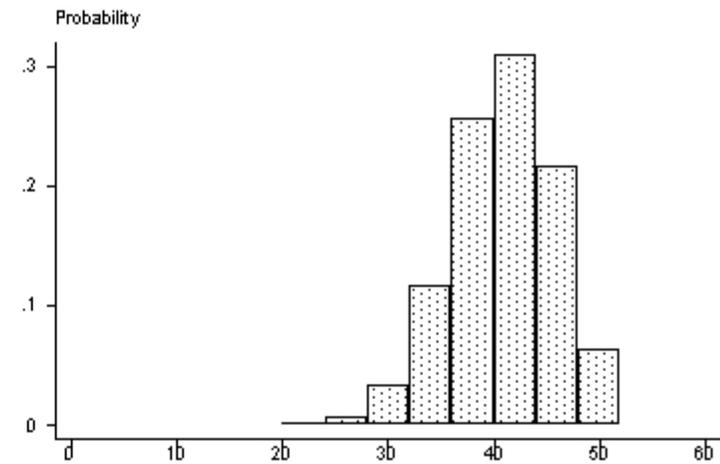
After-Tax Discount Rate Offered to Separating Officers



After-Tax Discount Rate Offered to Separating Enlisted Personnel



Estimated Individual Discount Rate for Separating Officers



Estimated Individual Discount Rate for Separating Enlisted Personnel

Figure 2: Offered and Estimated Discount Rates

The second point to see from Figure 2 is that the distribution of estimated rates for officers is generally much lower than those for enlisted personnel. The third point to see is that the variance of estimated discount rates for officers is larger than it is for enlisted personnel.

2. Extensions

2.1 Time Horizon

The time horizon of the annuity offered to individuals in the field varied directly with the years of military service completed. For every year of service (YOS) the horizon on the annuity was 2 years longer. As a result, the annuities being considered by individuals were between 14 and 30 years in length. With roughly 10% of the sample at each horizon, the average was around 22 years.

This feature of the field experiment implies that we need to control for the horizon length when inferring discount rates, since there is some (controversial) evidence that discount rates vary by time horizon. Specifically, the most popular assumption is currently the “hyperbolic discount function,” which implies that discount rates for shorter time horizons will be higher than those for longer time horizons (see Angeletos, et al. [2001] and Frederick, et al. [2002] for reviews).

These data also provide a qualified opportunity to test the assumption of constancy of discount rates with respect to these time horizons. The opportunity is qualified by the perfect co-linearity of horizon with YOS, the high co-linearity of age with YOS, and the high co-linearity of the stakes with YOS.⁶ In the regression analyses of WP, all three variables appear as explanatory variables. They recognized clearly (p.40) that YOS would also serve as a proxy for time horizon, and noted that an association of YOS with elicited discount rate could be taken as a test of the assumption of constancy of discount rates with respect to horizon if one controlled for age and the after-tax threshold discount rate. In other words, they implicitly claimed that a test of the constancy of discount rates can be undertaken by evaluating the marginal effect of YOS on elicited discount rates, holding constant age and the threshold discount rate.

WP (Tables 4 and 5) find no statistically significant association of YOS on the elicited discount rate, but do find a significant effect from age and the size of the lump-sum amount. Hence when they report average elicited discount rates by YOS (their Table 6), the rates vary significantly with YOS. They appropriately caution (p.48) that this is the result of a joint effect, but do not report the marginal effect in question.⁷ Since YOS, age, and the lump-sum amount are so strongly correlated, and are jointly significant,

⁶ As explained by WP (p.38), the lump-sum amount was 15% times the years of service times the annual final basic pay, and the annuity amount was 2½% times the years of service times the annual final basic pay. Hence one need only include the “stake size” from one or the other financial option.

⁷ The size and statistical significance of the marginal effect cannot be directly inferred from the size and significance of the coefficient in the estimated model. This is particularly true when there is a sample selection equation that includes years of service as well, and the coefficient on years of service is statistically significant in that equation; see WP (Appendix B).

there is a chance that the marginal effect of YOS on discount rates is being masked by the effect on age and lump-sum amount. Their results, if accepted, indicate striking evidence against the hyperbolic discounting model, at least for this sample and over these time horizons.

The correct calculation is to estimate the marginal effect of YOS on the probability of accepting the after-tax discount rate offered, controlling for the size of the lump-sum amount and age, since they varied directly with YOS. This calculation can be undertaken for different levels of YOS, since there is a chance that it is not constant for all levels.

We find that there is no statistically significant marginal effect of YOS on the probability of acceptance of the lump-sum. Evaluated at the mean of all explanatory variables, each additional YOS is associated with a decrease in the probability of acceptance by 0.8 of a percentage point for officers, and an increase in the probability of acceptance by 1.0 of a percentage point for enlisted personnel. However, the 95% confidence intervals span a zero effect: for officers (enlisted personnel) the confidence interval is between minus 4.0 (2.3) and plus 2.4 (4.4) percentage points.⁸ The same qualitative conclusion applies if the marginal effect is calculated at any of the discrete YOS values.

The conclusion is that there is no evidence to support the hyperbolic discount function over the horizons considered here: observed discount rates are effectively constant. The last qualification is important: we do not have any evidence from these data for time horizons shorter than 14 years, or beyond 30 years. It is possible that the discount function is hyperbolic over the time period prior to 14 years, and effectively constant beyond that. Indeed, “quasi-hyperbolic discount functions” that allow for such possibilities have been adopted for reasons of analytic tractability by many studies (e.g., Phelps and Pollack [1968], Laibson [1997], O’Donoghue and Rabin [1999] and Angeletos, et al. [2001]).

2.2 Uncertainty

The main conclusion of WP is contained in their Table 6, which lists estimates of the average discount rates for various groups of their subjects. Using their linear model, they report that the average discount rate for officers was 10.4% and that it was 35.4% for enlisted personnel. What are the standard errors on these means? There is reason to expect that they could be quite large, due to constraints on the

⁸ These results are obtained with the single-equation probit model. Essentially the same qualitative results are obtained with the sample-selection model. In that case the marginal effect for officers (enlisted personnel) is 2.8 (0.4) of a percentage point, with 95% confidence intervals between -1.0 and 6.6 (-0.2 and 1.0), respectively. The significance levels for these marginal effects are 13.8% and 15.2%, respectively, so they are more reliably estimated than those obtained from using the single-equation probit model; however, the sign of the effect is the opposite of the prediction of the hyperbolic models.

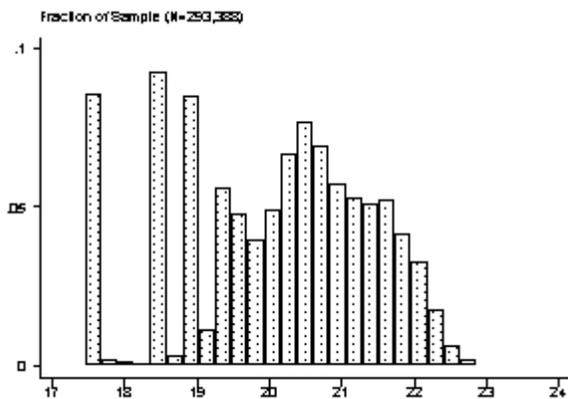


Figure 3: Percent After-Tax Discount Rates Offered

scope of the natural experiment.

Individuals were offered a choice between a lump-sum and an annuity. The *before-tax* discount rate that just equated the present value of the two instruments ranged between 17.5% and 19.8%, which is a very narrow range of discount rates. The *after-tax* equivalent rates range from a low of 14.5% up to 23.5% for those offered the separation option, but over 99% of

the after-tax rates were between 17.6% and 20.4% as shown in Figure 3. Thus the above inferences about average discount rates are “out of sample” in the sense that they do not reflect direct observation of responses at those rates of 10.4% or 35.4%, or *any* rates outside the interval [14.5%, 23.5%]. Figure 2 illustrates this point as well. These averages therefore reflect, and rely on, the predictive power of the parametric functional forms fitted to the observed data.

Even if one accepted the parametric functional forms (probit), the standard errors of predictions *outside* of the sample range of break-even discount rates will be much larger than those *within* the sample range.⁹ The standard errors of the predicted response can be calculated directly from the estimated model. Note that this is not the same as the distributions shown in Figure 1, which are distributions over the sample of individuals at each simulated discount rate that *assume that the model provides a perfect prediction for each individual*. In other words, those predictions just use the average prediction for each individual as the truth, so the sampling error reflected in the distributions only reflects sampling over the individuals. One can generate standard errors that also capture the uncertainty in the probit model coefficients as well.

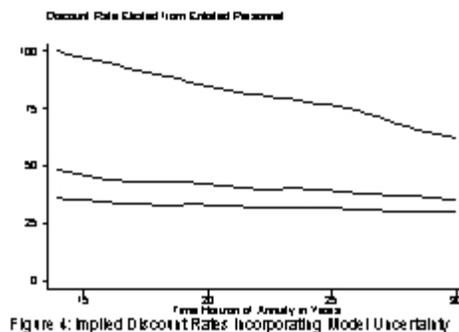


Figure 4: Implied Discount Rates Incorporating Model Uncertainty

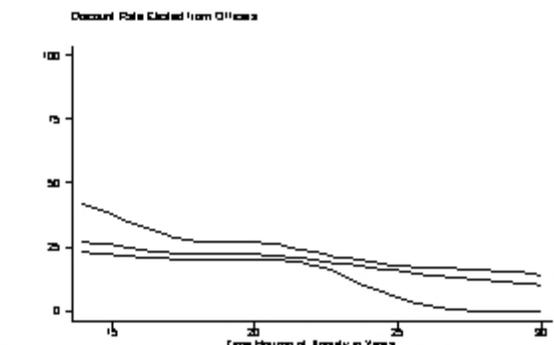


Figure 5: Implied Discount Rates Incorporating Model Uncertainty

Figures 4 and 5 display the results of taking into account the uncertainty about the coefficients of the estimated model used by WP. We use here the single-equation probit model for illustration, and adopt the same vertical scales for ease of comparison. The middle line shows a cubic spline through the predicted average discount rate. The top (bottom) line shows a cubic spline through the upper (lower) bound of the 95% confidence interval allowing for uncertainty in the individual

⁹ Relaxing the functional form also allows some additional uncertainty into the estimation of individual discount rates. The procedures proposed by Klein and Spady [1993] and Manski [1975] provide promising non-parametric possibilities for future research.

predictions due to reliance on an estimated statistical model to infer discount rates. Thus, in Figure 4 we see that there is considerable uncertainty about the discount rates for enlisted personnel, and that it is asymmetric. On balance, the model implies a considerable skewness in the distribution of rates, with some individuals having extremely high implied discount rates. Turning to the results for officers, in Figure 5, we find much less of an effect from model uncertainty. In this case the rates are relatively precisely inferred, particularly around the range of rates spanning the effective rates offered as one would expect.

We conclude that the results for enlisted personnel are too imprecisely estimated for them to be used to draw reliable inferences about the constancy of discount rates of time horizons. However, the results for officers are relatively tightly estimated, and can be used to draw more reliable inferences.

However, there is one additional source of information that should be taken into account when inferring discount rates by individuals in this field setting: the possibility that their responses would have been censored if in fact the higher rates were offered to them rather than counter-factually assumed to be offered to them. This issue may be of considerable significance for the enlisted personnel, given the extremely high discount rates implied for many of them.

2.3 Censoring

To explain the censoring problem in a laboratory setting, assume that you value a cold beer at \$3, which is to say that if you had to pay \$3 for one beer you would. If I ask you whether or not you are willing to pay \$2.50 for a *lab* beer, your response to me will depend on whether or not there is a market price of *field* beer¹⁰ lower than \$2.50. If the market price of the field beer is \$2.00, and you know that you can buy a beer outside the lab at this price, then you would never rationally reveal to me that you would pay \$2.50 for my lab beer. In this case we say that your response is censored by the market price (Harrison [1992; p.1432]). Fortunately, there are simple statistical procedures for allowing for this possibility, and we employ those in our statistical analysis.

It is easy to see how this censoring problem applies to a laboratory setting in which discount rates are being elicited. Consider a subject with a true individual discount rate (IDR) of 30%. In the absence of field substitutes for lab incentives, we would expect this subject to choose to save in the lab when the lab instrument provides a rate of return of 30% or higher. Now assume that this subject can *borrow* in the field at a rate of 14%. Although she demands at least 30% interest to delay consumption and save in the lab, at rates between 14% and 30% she is better off borrowing in the field at 14% and not delaying consumption in the field, leaving the money in the lab earning 14% or more, and repaying the field debt at the time she collects from the experimenter. In this case, the subject should rationally choose to invest in the lab when the lab

¹⁰ Assume further that a beer in the lab is the same product as a beer in the field.

instrument provides a rate of return of 14% or more. Hence, censoring would imply that the true IDR *could* actually be greater than or equal to the observed borrowing rate *when we observe lab investment responses that suggest that the IDR is equal to the borrowing rate.*¹¹ In other words, if we ignored the possibility of censoring of lab responses we would incorrectly infer that this subject had an IDR of 14%. Instead, we can only infer from these lab responses that the subject has a true IDR between 14% and ∞ . The problem is symmetric for censoring with respect to savings rates, although less significant empirically.¹²

The implication of allowing for censoring is that we cannot presume that the “raw” responses in the lab are unbiased indicators of the true IDR of the subject. Moreover, if we ignored field censoring then we could easily be led to think that we were measuring responses with more precision than would be warranted.

The implications for the natural field experiment examined by WP are immediate. TO BE WRITTEN.

2.4 Predicting the U.S. Population

WP (p.49) conclude that they “... find evidence of high personal discount rates in our sample of military separatees. Readers might remain skeptical that the rates we estimate are representative of the general population. But the anecdotal evidence points to high discount rates in subjects of the general populace.” If one can correct for the fact that the sample consists of military separatees, one should be able to generate estimates for the general U.S. adult population. The idea of using behavioral responses from a convenience sample to generate estimates for the general population has been applied to environmental damage assessment surveys by Harrison and Lesley [1996][, and consists of little more than a re-weighting of the sample predictions to match the characteristics of the population.

What is novel here is that we need to recognize that the convenience sample for the experiment is not randomly drawn from the population. Hence there is a sample selection process at work before the natural experiment is applied to the sample. MORE TO BE WRITTEN.

¹¹ When the subject reports an IDR interval that exceeds the borrowing rate that we calculate for the subject, we assume that there are subjective and unobserved transactions costs such that the true (unobserved) market rate for the subject is equal to the lower bound of the reported interval. The subject’s responses are then treated statistically as being censored at that inferred borrowing rate.

¹² Consider, for example, a subject with a true IDR of 3%. In the absence of field substitutes for lab incentives, we would again expect this subject to choose to invest in the lab instrument as long as it provides a return of 3% or higher. Now suppose that this subject can *save* in the field at a rate of 10%. Although she would be willing to save at 3%, at rates between 3% and 10% she is better off investing in the field and refusing to invest in the lab. Hence censoring would imply that the true IDR could actually be less than or equal to the observed savings rate when we observe lab investment responses that suggest that the IDR is close to the savings rate.

3. Conclusions

TO BE WRITTEN.

References

- Angeletos, George-Marios; Laibson, David; Repetto, Andrea; Tobacman, Jeremy, and Weinberg, Stephen, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15(3), Summer 2001, 47–68.
- Coller, Maribeth, and Williams, Melonie B., "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 1999, 107-127.
- Greene, William H., *Econometric Analysis* (New York: Macmillan, Second Edition, 1993).
- Greene, William H., *LIMDEP Version 7.0 User's Manual* (Bellport, NY: Econometric Software, Inc., 1995).
- Frederick, Shane; Loewenstein, George; and O'Donoghue, Ted, "Time Discounting: A Critical Review," *Journal of Economic Literature*, forthcoming 2002.
- Gould, William, "ssi4: Confidence intervals in logit and probit models," in S. Becketti (ed.), *The Stata Technical Bulletin Reprints* (College Station, TX: Stata Corporation, v.3, 1994, 182-184).
- Harrison, Glenn W.; Lau, Morten Igel, and Williams, Melonie B., "Estimating Individual Discount Rates for Denmark: A Field Experiment," *American Economic Review*, 92, forthcoming 2002.
- Harrison, Glenn W., and Lesley, James C., "Must Contingent Valuation Surveys Cost So Much?" *Journal of Environmental Economics and Management*, 31, June 1996, 79-95.
- Klein, Roger W., and Spady, Richard H., "An Efficient Semiparametric Estimator For Binary Response Models," *Econometrica*, 61(2), March 1993, 387-421.
- Manski, Charles F., "Maximum Score Estimation of the Stochastic Utility Model of Choice," *Journal of Econometrics*, 3, 1975, 205-228.
- O'Donoghue, Ted, and Rabin, Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 89(1), 1999, 103–124.
- Phelps, Edmund S., and Pollak, Robert A., "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 1968, 185–199.
- StataCorp, *Stata Statistical Software: Release 7.0* (College Station, TX: Stata Corporation, 2000).
- Warmer, John T., and Pleeter, Saul, "The Personal Discount Rate: Evidence from Military Downsizing Programs," *American Economic Review*, 91(1), March 2001, 33-53.