

The Stability of Measured Time Preferences

Erin L. Krupka^{a,b}, Melvin Stephens Jr.^{c,d}

October 12, 2012

Abstract:

We use a panel dataset to test the stability of measured discount rates over time in response to changes in both macroeconomic events and household-level labor market outcomes. While discount rate measures are constructed to capture a rate of time preference, our evidence is inconsistent with such an interpretation. Our results more closely align with the interpretation that standard methods to elicit discount rates reveal the market interest rate faced by an individual rather than their pure rate of time preference. It follows directly from such an interpretation that factors which influence the interest rate at which a household can borrow and lend, such as the inflation rate and household income, ought to be correlated with the elicited discount rate – a prediction supported by our data.

Key Words: Intertemporal choice, Household behavior, Impatience

JEL Classification: D1, D80, D90, J22

^a The School of Information, University of Michigan

^b Institute for the Study of Labor (IZA), Bonn Germany

^c Department of Economics, University of Michigan

^d National Bureau of Economic Research

*Corresponding author. Erin Krupka, ekrupka@umich.edu

This research was supported, in part, by the Carnegie Mellon Berkman Faculty Development Fund. We thank Robyn Dawes, George Loewenstein, Ted O'Donoghue, Roberto Weber, Howard Kunreuther, Dan Hamermesh, David Laibson, Glenn Harrision and Nikos Askitas.

1. Introduction

At the core of economic decisions involving the allocation of resources over time is the notion of a discount rate. A growing body of economic literature suggests that understanding variation in discount rates across individuals is critical to understanding variation across outcomes such as job selection and wage dynamics, job search, household income, savings and investment behavior (Lawrance 1991; Browning and Lusardi 1996; Eckel et al. 2005; Dellavigna and Paserman 2005; Munasinghe and Sicherman 2005; Munasinghe and Sicherman 2006; Ashraf et al. 2006). A parallel literature has developed methods of eliciting discount rates from respondents in laboratory and field experiments as well as in large-scale surveys with the intent of revealing the respondent's true rate of time preference. (e.g., see Frederick et al. 2002 for a survey of this literature). In fact, a wide range of discount rates have been found when using these methods.

There is some question, however, as to whether the empirical (i.e. measured) discount rate indeed captures the underlying pure time preferences (Fuchs 1982; Loewenstein 1987; Pender 1996; Frederick et al. 2002). If capital markets are perfect then the choice between values at different dates (e.g., "Do you prefer \$X today or \$Y tomorrow?") will depend upon the market interest rate since households can, theoretically, arbitrage the choices that are presented to them at this rate regardless of their true rate of time preference. When contrasted with the assumption that measured discount rates reveal the respondent's true rate of time preference, this alternative view provides a sharp test of the empirical content of discount rate survey questions. If measured discount rates reflect nominal interest rates, then factors that influence an individual's market rate of interest will also influence their measured discount rate. However, under the presumption that measured discount rates capture the respondent's true rate of time

preference, these factors should not affect the measured discount rate (e.g., see Frederick et al. 2002).¹

In this paper, we use panel data from the Seattle and Denver Income Maintenance Experiment (SIME/DIME) to test whether measured discount rates change when factors that influence the respondent's nominal discount rate change. Since the nominal interest rate is the sum of the real interest rate, r , and the inflation rate, π , the factors that we examine in our analysis are based on these two terms. Discount rates were elicited twice from SIME/DIME respondents, first late in 1972 and then again late in 1973-early 1974. The U.S. inflation rate coincidentally increased threefold between these survey waves, providing a unique opportunity to test the response of discount rates to a rapid change in the rate of inflation. In addition, the SIME/DIME collected numerous labor market outcomes for household members including income and hours worked. Given that events in the labor market, e.g., the loss of a job, can influence the real interest rate at which a household can borrow and/or lend, we test whether changes in these labor market measures are associated with changes in elicited discount rates.

The measured discount rates in the SIME/DIME that we examine are elicited using a choice titration technique that presents subjects with a series of hypothetical choices between a smaller sooner reward and a larger later reward. This hypothetical format has several disadvantages relative to recently developed elicitation techniques (Eckel et al. 2005; Andersen et al. 2008a; Chabris et al. 2008). Studies investigating incentive effects on discounting suggest

¹ Among households that face constraints on their ability to borrow and/or lend (i.e., they face imperfect capital markets), measured discount rates will not necessarily reveal true rates of time preference. For constrained households that face finite interest rates on borrowing, differences in measured discount rates may reflect variation across households in the rates that are available to them. For households that are unable to borrow in the market (i.e., they face an infinite interest rate on borrowing), tradeoffs between current and future values will reflect rates of time preference.

that a hypothetical bias is present.² Coller and Williams (1999), Johnson and Bickel (2002), and Epper et al. (2011) show that participants facing hypothetical rewards in the time discounting task exhibit higher impatience compared with participants facing real monetary incentives. Reported differences range widely from (11.10 percentage points in Epper et al. to 53.4 percentage points in Coller and Williams). However, Epper et al. find that the effect is predominantly a level effect. While these results suggest that using hypothetical methods may produce biased discount rates, a level effect would simply shift up or down the size of the discount rate but would not affect changes in measured discount rates over time. In our analysis we do not focus on the levels of the discount rates, but on how responses change between two periods and there is no evidence which indicates that this format induces changes in these measures over time.

Although elicited discount rates have been thoroughly investigated in the lab and in the field, few papers have directly measured the stability of these measures using changes in labor market outcomes. Using cross-sectional data, there is ample evidence that numerous factors such as unemployment (Harrison et al. 2002) and wealth (Hausman 1979, Lawrance 1991) are correlated with measured discount rates. We are only aware of two papers that examine whether discount rates change over time, both of which find that measured discount rates are non-responsive to changes in household-level circumstances (Harrison et al. 2006; Meier and Sprenger 2010). Using a sample of Danish individuals, Harrison et al. find that changes in elicited discount rates are uncorrelated with changes in the household's expectations of both their own and economy-wide economic outcomes. Meier and Sprenger (2010) examine a sample of low-income tax filers who use the services of a tax preparation provider. They find that changes

² Outside of discount rate elicitation techniques, there is evidence on the effects of financial incentives on judgment and decision making which demonstrate that incentives change responses (cf. Camerer and Hogarth, 1999).

in elicited discount rates are unaffected by changes in numerous covariates corresponding to the prior calendar year including total household income and household income from unemployment benefits.

By contrast, using the SIME/DIME data, we find that elicited discount rates systematically vary over time. First, our results show that measured discount rates are responsive to changes in current household-level labor market outcomes. We find that measured discount rates fall when either household income or hours worked increase, both of which are associated with the household's real interest rate to the extent that lending decisions are based on current income. Second, we find that measured discount rates increased by roughly 35 percentage points between the two discount rate elicitation periods, a period during which the nationwide inflation rate increased by a factor of three. Our results provide strong evidence that measured discount rates are responsive to changes in factors that are related to the respondent's nominal rate of interest.

Two important differences between our study and previous tests of the stability of measured discount rates over time may explain the discrepancies in findings. First, our use of a larger sample size, which yields more precise parameter estimates, is one likely reason that we find a significant relationship between measured discount rates and economic factors. Our sample contains 1,194 individuals who provide discount rates at two points in time whereas the corresponding sample sizes for Harrison et al. and Meier and Sprenger are 97 and 250, respectively. Second, we are able to measure household-level economic outcomes in the same month that discount rates are elicited whereas Meier and Sprenger must rely on income measures from the previous calendar year. To the extent that measured discount rates are influenced by the household's immediate economic situation, outcomes measured over an extended period of

time provide a noisy indicator of the relevant household situation and will attenuate parameter estimates when these noisy measures are used as regressors.

We proceed as follows. The next section describes the SIME/DIME data that we use in our analysis. We then test for a relationship between changes in the measured discount rate between the two survey waves and a corresponding increase in the inflation rate. Our subsequent econometric analysis tests for a relationship between changes in measured discount rates and the changes in household-level economic measures. The final section concludes.

2. Data

The Seattle and Denver Income Maintenance Experiments began in 1970 as the last of four major experiments to examine the impact of a negative income tax on a wide range of household outcomes for both dual and single-headed households. The SIME/DIME experimentally manipulated cash welfare benefit levels, tax rates on earned income, access to job training, and treatment duration (three or five years).³ Households were randomly assigned to treatment status with the probability of treatment assignment depending upon race, household status (one or two household heads), and pre-experimental income (Conlisk 1973).

After enrollment in the experiment, all households were subject to periodic interviews that occurred roughly every four to five months.⁴ These periodic interviews contained a set of three core modules that were asked at each periodic interview as well as one or more topical modules which varied in their frequency of being asked. The first core module ascertained employment histories of each adult household member (ages sixteen and over) since the previous periodic interview, gathering information on earnings and hours worked on each job held as well

³ Spiegelman and Yaeger (1980) provide a very useful description of the SIME/DIME.

⁴ A pre-enrollment interview collected a wide range of household information during the twelve months prior to the onset of the experimental period.

as information on time not worked including unemployment durations. The second core module collected benefits and expenses between interviews including all sources of non-earned income such as welfare benefits, food stamps, pensions, interest, etc. The third core module included information on housing, including homeownership and housing subsidies. The discount rate questions, which were collected only a few times, appeared as a topical module.⁵

Although the wealth of information contained in the SIME/DIME data provides numerous measures that can be included in our analysis, the longitudinal empirical specifications that we discuss below include a fixed effect for each sample respondent. As such, fixed household characteristics such as the variables used to determine treatment assignment probabilities (e.g., race, household composition, etc.), the actual treatment assignment, and pre-experimental household measures (e.g., earnings, government transfer income receipt, etc.) are captured by the respondent fixed effect and are not included in our econometric models. As such, we limit our discussion here to the time-varying measures that we use in our analysis.

In order to examine whether measured discount rates respond to individual labor market outcomes, our main labor market measures are total household income and the respondent's average weekly hours worked during the same month in which the discount rates are elicited. Restricting the time period of labor market measures to a single month has two benefits.⁶ First, to the extent that the elicited discount rate may reflect current household circumstances, using labor market conditions measured during the same month rather than, say, during the previous year (as per Meier and Sprenger) provides a stronger test of whether measured discount rates

⁵ In addition, all households receiving negative income tax benefits, as well as a random sample of control households, were required to complete a monthly income report form. While only the data from the periodic interviews are available for our analysis, presumably the monthly reporting of income helped reduce the error in reported income on the periodic interviews.

⁶ The sample correlation coefficient between these two measures is 0.42. Although these two measures are positively correlated, alternative measures provide little additional variation. For example, the estimated correlation coefficient between the fraction of days unemployed and average weekly hours is -0.94.

reflect contemporaneous household conditions. Second, to the extent that past work effort is reported with error, focusing on work effort in the most recent month greatly minimizes concerns about errors in measuring these variables. However, we also present results using these labor market outcomes measured over the past six months to demonstrate that our findings are not sensitive to the time period over which these regressors are measured.⁷

Discount rates were elicited using the titration technique during the third (late 1972) and seventh (late 1973 to early 1974) periodic interviews but only in Seattle.⁸ One person per household, typically a household head, answered discount rate questions with male household heads designated to respond in dual-headed households. Thus, we limit our sample to original male heads of dual-headed households (married males) and original female heads of single-headed households (single females) to ensure that we capture the same respondent in both waves.

The titration technique elicits discount rates by posing a series of choices between a current and a varying future payoff. Households are asked “Suppose now that you have the choice between a cash bonus today and a different cash bonus next year. If you were given the choice of \$100 today or \$200 a year from now, which would you choose?” Choosing \$100 today ends the series while choosing \$200 leads to another question where respondents choose between \$100 today or \$175 a year from now. The later payment continues to decline until the

⁷ Although we could alternatively determine labor market outcomes over longer periods, e.g. twelve months, ten percent of our sample has less than one year between the dates at which discount rates are elicited.

⁸ The ninth periodic interview in Seattle contains an alternative titration question which differs from the measure asked in the third and seventh interviews both in terms of the magnitude of the prospective payment and the time period over which households compare values. As such, we focus our analysis on the titration measure which is consistent over time although we discuss results involving the alternative titration question at the end of the next section. Furthermore, in Denver each of the two versions of the titration measure is asked only once which precludes us from testing the time stability of measured discount rates using respondents at that experimental site.

respondent either chooses \$100 instead of the listed dollar amount or the final dollar amount of \$100 is reached.⁹

We use the midpoint of the range of values to determine the implied discount rate. Thus, if a respondent is willing to wait a year for \$200 but will not wait for \$175, then the midpoint is \$187.50 and the discount rate is 87.5%. Our analysis accounts for the right censoring of the discount rate at 100% for respondents choosing \$100 now over \$200 in a year in response to the first question in the series. In addition, we present results using econometric methods which explicitly treat the discount rate as a categorical measure and yield nearly identical findings.

Table 1 about here.
(Summary Statistics)

Our final sample excludes households with changes in the household head, (rare) changes in treatment status, or missing assignment variables or discount rate responses. Table 1 reports summary statistics for the resulting sample of 1,194 respondents and also displays these measures separately for the 697 married males and 497 single females. In our empirical analysis, we present results in which we pool men and women as well as separate results for each group. Although married men reside in households with higher income and wealth in the year prior to the experiment as shown in Table 1, we have no *a priori* reason to believe that there is a differential response of measured discount rates to labor market outcomes for men and women, especially once we account for fixed individual differences. However, pre-experimental work effort differs sharply between these groups with nearly 90 percent of married males working at some point in the year prior to the experiment while less than 60 percent of single women doing so. To the extent that the labor supply changes reflect differential movements in and out of the

⁹ The later payment amounts are \$200, \$175, \$150, \$140, \$135, \$130, \$127, \$124, \$121, \$118, \$115, \$112, \$109, \$106, \$103, and \$100.

labor force for men and women, the estimated impact of labor market outcomes on the measured discount rate may differ between these groups.

Since our sample is not representative of the population at large, we acknowledge that care must be taken in comparing our findings to the prior literature. To the extent that discount rate elicitation has been performed on samples of college undergraduates, it is not clear that our results necessarily suffer from greater external validity concerns than the relevant literature. Moreover, the population of tax filers used by Meier and Sprenger (2010) to examine temporal discount rate stability is primarily composed of low income families. Finally, our use of a fixed effects estimator accounts for permanent characteristics among households that may influence the *level* of the measured discount rate.

3. Elicited Discount Rates and Economic Conditions

Figure 1 plots the discount rate response in the first wave of the survey on the x-axis and the response in the second wave of the survey on the y-axis. Because responses are categorical there are multiple points in the x-y space that contain multiple respondents thus, the size of each dot represents the number of observations that fall onto each point. Figure 1 clearly shows that there is considerable variation in individual responses between the first and second wave of the survey in our data.

Figure 1 about here.
(Changes in Measured Discount Rates Between Wave 1 and 2)

We examine the temporal stability of elicited discount rates using two approaches. First, we examine whether the measured discount rate responds to a rapid increase in the national inflation rate. Second, we examine whether individual-level labor market changes affect the measured discount rate.

3a. Rising Inflation

The two SIME/DIME interviews eliciting discount rates coincided with a large inflation rate increase. The first interview was fielded between August and November of 1972 and the second was fielded between September 1973 and March 1974. Over this period, the national inflation rate rose from 3.2% to 10.4%.¹⁰ The University of Michigan's Survey of Consumers indicates that inflation *expectations* rose from 3.9% in the third quarter of 1972 to 10.1% in the first quarter of 1974.¹¹ The unemployment rate *fell* slightly over this period from 5.5% to 5.1% before rising between late 1974 and early 1975.¹²

Figure 2 about here.
(Measured Discount Rates, Inflation, and Unemployment)

Figure 2 shows changes in the mean elicited discount rates and in the national unemployment and inflation rates over this period. The average discount rates range from 43% to 49% for married men and 56% to 63% for single women in the fall of 1972. In late 1973 and early 1974, these ranges are between 71% and 80% for married men and 72% and 85% for single women. Thus, while the inflation rate increased by 7 percentage points, average discount rate responses rose by 26 percentage points for married men and 21 percentage points for single women. Furthermore, the fraction of men reporting the maximum discount rate rose from 24% to 59% between the two interviews. Among women, 67% gave the maximum response to the later interview while only 37% did so in the earlier interview. Since the discount rate is censored above at 100%, these results understate the correlation between the measured discount rate and

¹⁰ The CPI-U figures for the Seattle area show an even sharper increase, rising from 1.8 percent in August 1972 to 10.1 percent in February 1974. National and Seattle area CPI data are available at <http://www.bls.gov/cpi/data.htm>.

¹¹ Data from the Survey of Consumers is available at <http://www.sca.isr.umich.edu>.

¹² National unemployment rate data are available at <http://www.bls.gov/cps/>.

the inflation rate. The econometric methods that we employ below explicitly account for the censoring of the discount rate measure.

The results in Figure 2 are broadly consistent with the interpretation of elicited discount rates as measuring market rates of interest since these measures increase with the inflation rate (e.g., Fuchs 1982).¹³ Our finding that these measures increase more than three times as much as the inflation rate is substantially larger than the one-for-one relationship predicted by this interpretation of measured discount rates. Perhaps the hypothetical nature of the measures induces the rather large response although, as we noted earlier, prior research indicates that the hypothetical bias only generates a level effect in measured discount rates (e.g., Epper et al. 2011). Since we only observe households at two points in time, our results are only suggestive of a relationship between the measured discount rate and the national inflation rate and cannot rule out the possibility that unobserved aggregate factors caused the increase in the measured discount.¹⁴ While our finding, that the average measured discount rate increases so rapidly in one year among the same set of households, casts doubt on interpreting such measures as capturing true time-invariant rates of time preference, we next turn to the relationship between measured discount rates and household-level labor market outcomes. By testing for a relationship between measured discount rates and household-level labor market outcomes we can examine whether idiosyncratic changes at the household level lead to changes in the measured discount rate.

¹³ An alternative interpretation of the nominal interest rate is that households use their *expected* rate of inflation, π^* , rather than using the actual rate of inflation, π , in computing their nominal interest rate (i.e., $i=r+\pi^*$). As we noted above in the text, the expected inflation rate rose by roughly the same amount as the actual inflation rate during this period so that the implications for the nominal interest rate are the same based on either the actual or the expected rate of inflation.

¹⁴ We note, however, that the unemployment rate is not one such factor as it remains relatively stable during this period (see Figure 2).

3b. Regression Results

We next test whether the elicited discount rate is correlated with individual labor market outcomes. An individual's true unobserved rate of time preference, $r_{i,t}^*$, is modeled as

$$r_{i,t}^* = \delta(LaborMkt_{i,t}) + \alpha_i + \gamma_t + \varepsilon_{i,t} \quad (1)$$

where $LaborMkt_{i,t}$ is a measure of labor market conditions faced in period t , α_i is an unobserved individual effect, γ_t is a second interview effect, and $\varepsilon_{i,t}$ is an idiosyncratic error term at time t . We treat α_i as a fixed effect in our analysis which allows this time-invariant unobserved effect to be correlated with household characteristics. As a consequence of this assumption, we exclude time-invariant variables from (1) since such individual characteristics are completely captured by α_i .

As we discussed in the data section, our two primary measures of labor market conditions are total household income and the respondent's average weekly hours worked during the month in which the discount rate is measured. However, there is no *a priori* guidance for the length of time over which to measure individual economic conditions that best captures their influence on the elicited discount rate. It is unclear whether the most relevant measure is from a single month which reflects the immediacy of economic well-being or from a longer period of time which reflects a more permanent measure of household circumstances. Therefore we also present results using labor market conditions measured in the six months up to and including the month in which discount rates are collected.¹⁵

In order to estimate equation (1), we must account for both the right censoring of the measured discount rate, since we only observe whether someone has a discount rate that equals

¹⁵ We choose six months since the shortest interval of time between discount rate measurements is nine months for individual respondents. Roughly ten percent of respondents have less than one year between discount rate questions.

or exceeds 100%, and the individual fixed effect. We use Honoré's (1992) fixed effect trimmed least squares censored regression estimator which both accounts for the censoring of the dependent variable and allows for the fixed effect, α_i . One concern with this approach is that it requires us to treat the measured discount rate as a censored continuous variable by using the mid-points of the elicited discount rate ranges. In results not shown here, we have estimated both random effect interval regressions, which explicitly account for the categorical nature of the dependent variable, and random effect censored regressions which treat the dependent variable as a censored continuous variable. The corresponding estimates are nearly identical to each other which suggests that treating the categorical dependent variable as continuous likely has no effect on our findings.¹⁶

Our analysis does not explicitly correct for sample attrition. The SIME/DIME attrition rate between the two discount rate eliciting surveys is less than six percent for married men and nine percent for single women.¹⁷ These low rates of attrition are in stark contrast to the other papers examining time variability in subjective preference measures such as Harrison et al. (2006) and Meier and Sprenger (2010) which have attrition rates of 61 percent and 72 percent, respectively, and must employ econometric methods that account for attrition (both studies find that their results are unaffected after applying these econometric corrections). Furthermore, recall that the SIME/DIME respondents are part of an experiment which incentivizes households to remain in the survey.¹⁸ Thus, employing a sample selection correction would add little to our analysis. Moreover, a useful feature of Honoré's fixed effects censored regression estimator is that it is consistent even if the unobserved individual effect, α_i , is correlated with the included

¹⁶ These results are available upon request.

¹⁷ 697 of the 740 married men who responded to the first discount rate question also answered the second discount rate question while 497 of the 541 single women did so.

¹⁸ Not only did the households assigned to the treatment group have an incentive to remain in the survey but households in the control group were paid a nominal fee for completing a monthly income report form.

regressors. Therefore, if non-random attrition exists, but it is due to an unobserved time-invariant respondent characteristic, then this estimator implicitly accounts for it.

Table 2 about here.
(Correlation Between DR and LM Outcomes)

Table 2 presents our estimates for the pooled sample (Panel A) along with the separate estimates for married men (Panel B) and single women (Panel C). Each column of each panel represents a separate estimate of equation (1) with the estimated coefficients on the labor market measures and the second interview effect (γ_t) being reported in the Table. Consistent with the results shown in Figure 2, the second interview effect is highly significant across all of the estimation results in Table 2. Our findings indicate that the subjective discount rate rose by over 35 percentage points between the two periods where the inflation rate tripled. This estimated increase is very similar for both men and women and, as expected, is larger than the raw estimated increases shown in Figure 2 since the econometric methods account for the censoring of the elicited discount rate.

The results in Table 2 also indicate that changes in the discount rate are correlated with changes in individual labor market outcomes. The results using labor market outcomes in the month that discount rates are measured (columns (1) and (2)) indicate that the measured discount rate decreases when the respondent experiences an increase either in total household income or in hours spent working. All of these estimates are statistically significant except for the coefficient on total income for single women. The magnitude of the estimated coefficient on total income is very similar across the pooled sample estimates and the separate male and female sample estimates. However, the estimated coefficient on hours worked is somewhat larger for the sample of single women.

The last two columns of Table 2 report the estimates which use the labor market outcomes measured over the six months up to and including the month that the discount rate is elicited as opposed to measuring these outcomes only in a single month. For the pooled sample in Panel A, we find that moving to a longer time period yields nearly identical estimates as with the single month measures and that the results remain highly significant. However, this finding is a combination of two divergent results for the male and female samples shown in Panels B and C. For the sample of married men, moving to the longer period for measuring labor market outcomes produces estimates that fall slightly in magnitude with the coefficient on hours worked becoming marginally significant. For single women, however, we find that the estimates actually increase substantially in magnitude although the coefficient on total household income remains insignificant.

These differences for the male and female samples rule out more straightforward explanations for the impact of measuring labor market outcomes over a longer period. A straightforward hypothesis might be that measured discount rates are more likely to be influenced by immediate household economic events (to the extent they are influenced at all) rather than events measured over an extended period of time because the latter provide a noisy indicator of the relevant household situation and will attenuate parameter estimates when these noisy measures are used as regressors. However, we do not find systematic evidence of how alternating the period of time over which labor market outcomes are measured affects the estimated responsiveness of the elicited discount rate. This leaves us unable to reconcile our findings with those in Meier and Sprenger (2010) using an explanation that is based on differences in the time window used for measuring outcomes.

3c. Alternative Titration Question

We examine the stability of elicited discount rates over time using longitudinal data and find that these measures are sensitive both to the inflation rate and to individual labor market outcomes. However, one possible concern with using the titration question in multiple interviews is that previous respondent experience with the titration method might induce higher elicited discount rates in subsequent interviews. Respondents who found the process of answering a series of monetary tradeoffs to be tedious and repetitive might recognize that giving a sooner response would expedite their time spent answering survey questions.

We can examine this issue using another wave of data collected in the ninth periodic interview, fielded between March and October of 1974 (i.e., after the seventh periodic interview used as the second survey date for our main analysis). The data in the ninth periodic interview in Seattle also uses a titration question, however, this question differs in four important ways from the question previously asked in the third and seventh interviews: 1) the magnitude of the money trade-off increases by a factor of ten, 2) a front-end delay of three months is added, 3) the number of choices in the sequence is shortened (from sixteen to ten), and 4) respondents are told the annual interest rate on the money if they choose the delayed amount. A direct comparison to the question we use is difficult, but we can examine this data to see if there is evidence that subjects avoid answering questions.

If subjects were avoiding answering the long list of questions, then we would predict that implied discount rates in the ninth interview would not be lower than in the first wave, but higher. However, we find that respondents go further into the sequence and give lower discount rates.¹⁹ The fraction of men and women reporting the maximum discount rate falls to 43% and

¹⁹ These findings are consistent with the prior literature which finds that larger magnitudes, front-end delay, and information on the annual interest rate all lead to lower elicited discount rates (e.g., Frederick et. al 2002).

48%, respectively, down from 59% and 67% as reported in the second of the two interviews. We take this responsiveness to changes in the wording between the seventh and ninth interviews to indicate that respondents are contemplating their answers to these questions rather than simply expediting the process by giving a quick response.

4. Conclusion

While there is an interest in knowing whether pure rate of time preferences are theoretically constant, we focus on the more immediate and instrumental question of whether measured discount rates (using trade-offs between money at two different time periods) systematically vary with factors other than time preferences. Using panel data, we find that elicited discount rates vary over time in response to changes in the inflation rate and changes in household labor market outcomes.

While discount rate measures are constructed to capture a fixed rate of time preference, our evidence is inconsistent with such an interpretation. Our results more closely align with the observation, prominently noted in Fuchs (1982), that standard methods to elicit discount rates will reveal the market interest rate faced by an individual rather than a pure rate of time preference. It follows directly from this observation that factors which influence the interest rate at which a household can borrow and lend, such as the inflation rate and household income, *ought* to be correlated with the elicited discount rate – a prediction which our data support.

The implications of our findings bear some discussion. An event that affects everyone simultaneously, such as a rise in inflation, might affect the level of the measured discount rate but preserve the ordering among respondents such that more impatient individuals have higher discount rates. However, our results indicate that measured discount rates are also sensitive to

idiosyncratic events such as labor market outcomes. These results suggest that even an ordinal comparison may not capture relative impatience among respondents. The implication of this is that, for example, one would not be able to compare the employed and unemployed responses or make inferences that the unemployed are more impatient than the employed. These results make it challenging to translate measured responses back into relevant measures of discount rates across individuals that reflect the time preferences we think of in our theoretical models. For ourselves, we conclude that these measures are perhaps best used to identify individual perceptions of the market interest rate they face.

References

- Andersen, S., Harrison, G., Lau, M., Rutström, E., 2008. Eliciting risk and time preferences. *Econometrica* 76 (3), 583-618.
- Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the mast, Evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics* 121(2), 635-72.
- Browning, M., Lusardi, A., 1996. Household saving, micro theories and micro facts. *Journal of Economic Literature* 34(4), 1797-1855.
- Camerer, C. F., Hogarth, R. M., 1999. The effects of financial incentives in experiments, a review and capital-labor production framework. *Journal of Risk and Uncertainty* 19, 7-42.
- Chabris, C., Laibson, D., Morris, C., Schuldt, J., Taubinsky, D., 2008. Individual laboratory-measured discount rates predict field behavior. *Journal of Risk and Uncertainty* 27(2-3), 237-269.
- Coller, M., Williams, M., 1999. Eliciting individual discount rates. *Experimental Economics* 2, 107-127.
- Conlisk, J. 1973. Choice of response functional form in designing subsidy experiments. *Econometrica* 41 (4), 643-656.
- Dellavigna, S., Paserman, M., 2005. Job search and impatience. *Journal of Labor Economics* 23(3), 527-88.
- Eckel, C., Johnson, C., Montmarquette, C., 2005. Saving decisions of the working poor, short- and long-term horizons. in, J. Carpenter, G. Harrison and J. List, eds., *Field Experiments in Economics, Research in Experimental Economics, Vol. 10* (Greenwich, CT, JAI Press), 219-260.
- Epper, T., Fehr-Duda, H., Schubert, R., 2011. Energy-using durables, the role of time discounting in investment decisions. IED Working Paper #16.
- Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference, a critical review. *Journal of Economic Literature* 40 (2), 351-401.
- Fuchs, V. R., 1982. Time preference and health, an exploratory study. In *Economic Aspects of Health*, edited by V. R. Fuchs, pp. 93-120. Chicago, University of Chicago Press, 1982.
- Harrison, G., Morten, L., Rutström, E., 2006. Dynamic consistency in Denmark, a longitudinal field experiment. UCF Economics Working Paper No. 05-02.

- Harrison, G., Morten, L., Williams, M., 2002. Estimating individual discount rates in Denmark, a field experiment. *The American Economic Review* 92(5),1606-1617.
- Hausman, J., 1979. Individual discount rates and the purchase and utilization of energy-using durables. *Bell Journal of Economics* 10(1),33- 54.
- Honoré, B. E., 1992. Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica* 60(3), 533-565.
- Johnson, M., Bickel, W., 2002. Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior* 77, 129-146.
- Laibson, D., Repetto, A., Tobacman, J., 2007. Estimating discount functions with consumption choices over the lifecycle. Working paper #13314, National Bureau of Economic Research.
- Lawrance, E., 1991. Poverty and the rate of time preference, evidence from panel data. *Journal of Political Economy* 99(1), 54-77.
- Lowenstein, G., 1987. Anticipation and the valuation of delayed consumption. *The Economic Journal* 97(384), 666-684.
- Meier, S., Sprenger, C., 2010. Stability of time preferences. IZA Discussion Paper #4756.
- Munasinghe, L., Sicherman, N., 2005. Wage dynamics and unobserved heterogeneity, time preference of learning ability? National Bureau of Economic Research Working Paper.
- Munasinghe, L., Sicherman, N., 2006. Why do dancers smoke? Time preference, occupational choice, and wage growth. *Eastern Economic Journal* 32(4),595-616.
- Pender, J. L., 1996. Discount rates and credit markets, theory and evidence from rural India. *Journal of Development Economics* 50(2),257-296.
- Spiegelman, R. G., Yaeger, K. E., 1980. Overview, the Seattle and Denver Income Maintenance Experiments. *Journal of Human Resources* 15(4), 463-479.

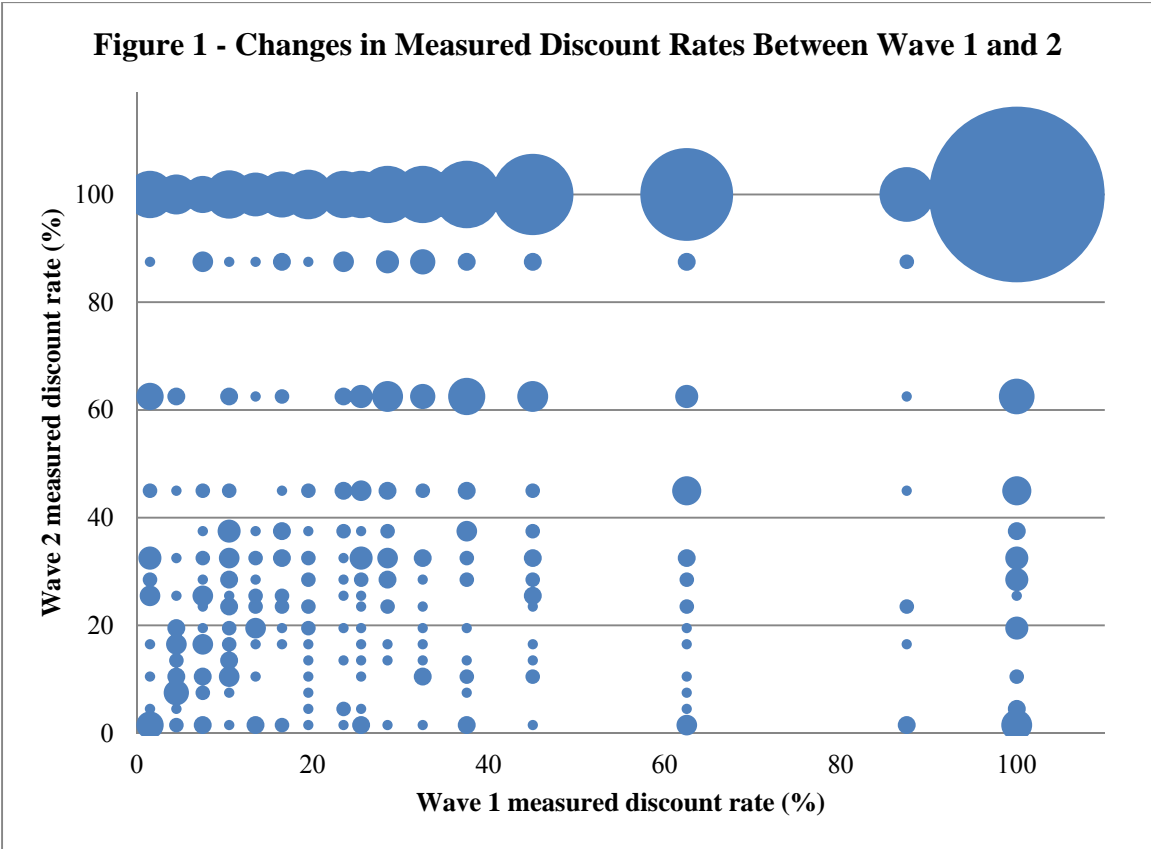
Table 1 - Summary Statistics for the Final Sample

	All Respondents	Men in Dual-Headed Households	Women in Single-Headed Households
Years of Education	11.3	11.4	11.3
Male	0.58	1.00	0.00
Age	35.7	35.8	35.7
Household Size	4.0	4.3	3.5
HH Members Under 16	2.1	2.1	2.1
HH Members Under 5	0.6	0.7	0.5
Black	0.43	0.37	0.53
Five-year Program	0.34	0.35	0.32
NIT Treatment Group	0.56	0.53	0.59
Job Training Treatment Group	0.58	0.57	0.58
Pre-Experimental Total Income	6,330	7,613	4,529
Pre-Experimental Earned Income	4,862	6,655	2,346
Pre-Experimental Transfer Income	1,465	958	2,173
Pre-Experimental Assets	3,147	3,932	2,035
Pre-Experimental Net Worth	3,776	4,773	2,407
Pre-Experimental Employed	0.75	0.88	0.57
Third Interview Discount Rate	52%	47%	59%
Seventh Interview Discount Rate	76%	73%	80%
N	1,194	697	497

Table 2 - Correlation Between the Elicited Discount Rate and Labor Market Outcomes

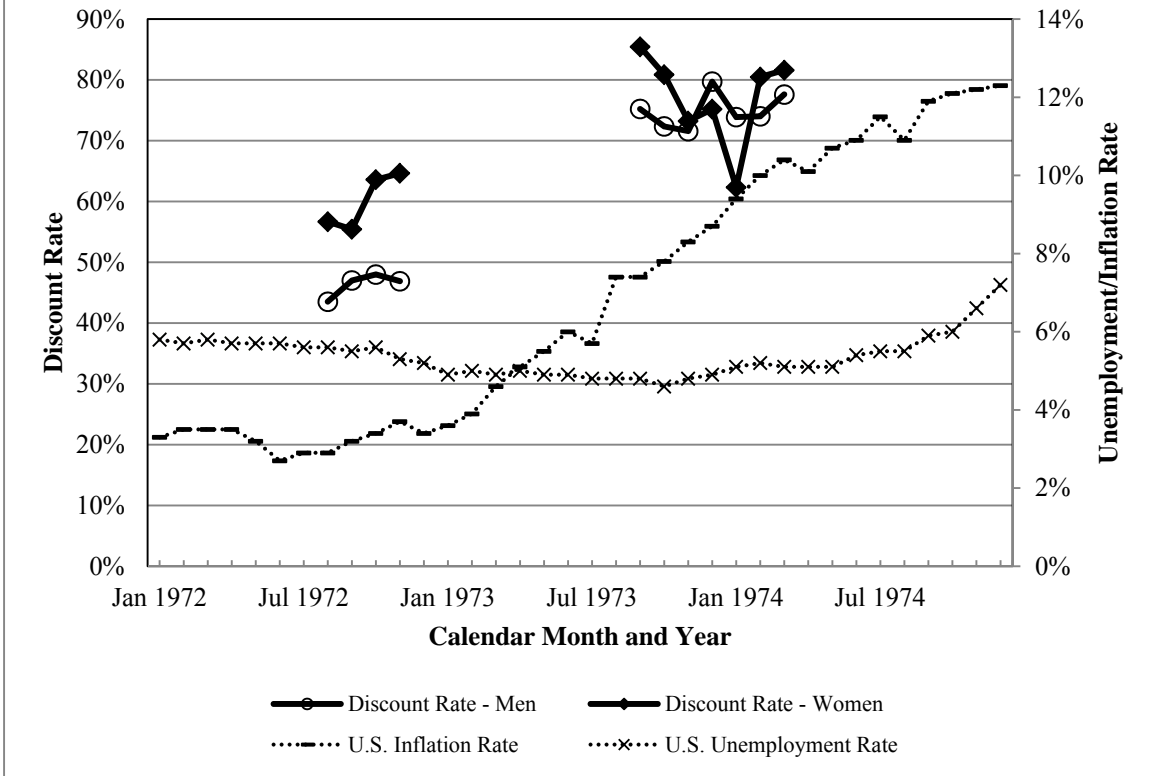
Regressor Period:	Single Month		Six Months	
	(1)	(2)	(3)	(4)
<i>A. All Respondents</i>				
Monthly Household Income/1000	-0.167 (0.060)		-0.170 (0.074)	
Weekly Hours Worked/100		-0.364 (0.101)		-0.352 (0.116)
Second Interview	0.355 (0.015)	0.355 (0.015)	0.356 (0.015)	0.357 (0.015)
<i>B. Men Only</i>				
Monthly Household Income/1000	-0.173 (0.063)		-0.157 (0.072)	
Weekly Hours Worked/100		-0.321 (0.111)		-0.223 (0.120)
Second Interview	0.357 (0.018)	0.355 (0.018)	0.358 (0.018)	0.356 (0.018)
<i>C. Women Only</i>				
Monthly Household Income/1000	-0.128 (0.190)		-0.235 (0.162)	
Weekly Hours Worked/100		-0.466 (0.222)		-0.789 (0.333)
Second Interview	0.352 (0.027)	0.357 (0.027)	0.353 (0.027)	0.362 (0.027)

Note: The results in each column of each panel are from a separate estimate of equation (1) using Honoré's fixed effects censored regression estimator. The labor market regressors in columns (1) and (2) are measured in the same month as the discount rate elicitation while those in columns (3) and (4) are measured in the six months up to and including the month in which the discount rate is elicited. The results in Panel A are based on 1,194 pooled observations comprised of 697 male observations, that are also used in Panel B, and 497 female observations, that are also used in Panel C.



Note: Figure 1 plots the discount rate response in the first wave on the x-axis and in the second wave on the y-axis. Responses are categorical so the size of each dot represents the number of observations that fall onto each point. The relative size of each bubble, along with the clear dispersion of the bubbles across the graph, shows that the measured discount rate is different between waves for many respondents.

Figure 2 - Measured Discount Rates, Inflation and Unemployment



Note: The discount rates for men and women are the respective average values among sample respondents who were surveyed in the relevant month. The sources for the monthly U.S. inflation rate and the monthly U.S. unemployment rate are given in the text.