Loss Aversion and Effort: Evidence from a Field Experiment

Ernst Fehr[*], Lorenz Goette[†], Marion Lienhard[‡]

September 1, 2013

Abstract

This paper contributes new evidence to a recent controversy in labor economics: Is labor supply affected by sensations of gains and losses relative to a narrowly-defined earnings target? We report evidence from a field experiment that exogenously and temporarily manipulates the progress of workers towards a possible earnings target, by randomly delaying subjects in their progress. We find that individuals who were delayed put in more effort. We also find that this effect is specific to the treatments in which we paid the subjects on a piece rate. In a control treatment with fixed wages, we find no effect that the subjects effort choices were affected by the productivity manipulations. This is inconsistent with a wide class of neoclassical utility functions. We argue that this is evidence that the individuals in the piece rate treatment were trying to catch up with the loss in output due to low productivity. There are two models to explain this behavior: Reference-dependent preferences, or concave utility over narrowly defined outcomes. To further probe into the pattern, we also measure risk preferences of the individuals in an experiment involving risky payoffs. We find that the pattern of catching up after a delays is only present for individuals who behave in a risk-averse fashion in simple choice experiments. Furthermore, we find that the effect is concentrated among strongly risk-averse individuals, and not present, for moderately risk-averse individuals. This is evidence against concave utility, thus favoring reference-dependent preferences to explain our result.

Keywords: Labor supply, reference-dependent preferences, field experiment

JEL classification: C93, D64, H41, I18

[*]University of Zurich, Department of Economics, Bluemlisalpstr. 10, 8006 Zurich, Email: ernst.fehr@econ.uzh.ch.

[†]University of Lausanne, Faculty of Business and Economics,Batîment Internef, 1015 Lausanne-Dorigny Email: lorenz.goette@unil.ch.

[‡]Federal Competition Commission of Switzerland

I. Introduction

Models of intertemporal optimization by individuals are widely used in economics. They have been applied to consumption choices (Deaton, 1992), labor supply (Blundell and Macurdy, 1999), or human capital investment (Becker, 1993). The key implication from models of intertemporal maximization is that individuals equalize the expected marginal benefits of any activity over time. In some cases, this has very strong implications: in the context of labor supply, it implies that consumption is chosen such that the marginal utility is equalized over periods. But consider what this implies for labor supply if wages are unexpectedly low in one period: Because of the low wage, the individual will be earning less money than expected. However, because the individual equalizes the marginal utility of consumption over time, this is only going to lead to a very small change in consumption in every period, thus leaving the marginal utility of consumption virtually unchanged. With the marginal utility of income approximately constant, there is a first-oder effect of the wage, leading the individual to cut labor supply in that period (See, e.g. Browning et al., 1985; Fehr and Goette, 2007, for a formal proof). Loosely speaking, the general intuition behind this is to take advantage of the intertemporal relative prices, while smoothing out potential small gains and losses over time.

Yet, anecdotes and evidence exist to document the opposite. There are stories of strict rules about the number of pages a Ph.D. candidate wrote per day to finish her dissertation, or individuals putting in double effort when they feel they are behind relative to a narrowly defined target (Heath et al., 1999). The underlying psychological intuition is that falling short of a perceived target creates a sensation of a loss, which motivates extra effort. It is consonant with a large literature on reference-dependent preferences (Kahneman and Tversky, 1979; Tversky and Kahneman, 2000). This mechanism could also distort labor supply over time: If wages are low, individuals may find it difficult to reach an perceived income target and therefore exert more effort than usual, even though the economic incentives dictate the opposite response. Thus, psychologically, it may be difficult for individuals to smooth out small changes over time, because of these sensations of gain or loss relative to what they expected to happen.

A new literature, using high-frequency data on labor supply and wages, has tested intertemporal substitution of labor. This evidence comes from studies using day-to-day variation in wages of cab drivers (Camerer et al., 1997; Chou, 2002; Farber, 2005), or field experiments with bicycle messengers (Fehr and Goette, 2007). All studies find that daily effort is lower when wages are high, but there is complete disagreement over the interpretation of these results. The results in all studies are consistent with a model where individuals care about reaching a daily income target. The evidence in Fehr and Goette (2007) is particularly suggestive: They find that only bike messengers behaving in a loss-averse fashion in an unrelated choice experiment reduce effort during the exogenously implemented wage increase.

However, it is not always possible to rule out all other plausible alternative explanations, such as fatigue effects. For example, in Fehr and Goette (2007), there is also a large increase in the number of shifts the bike messengers work, raising the possibility that the reduction in effort per shift is a consequence of that increase in the number of shifts.[1] Furthermore, it is hard to directly test for reference-dependent preferences, as these models are notoriously difficult to distinguish from a model in which individuals have concave utility, but evaluate utility over narrowly defined events.[2] In such a model, daily income effects can also occur, leading to a lower effort because the perceived marginal utility of consumption decreases, not because of a perceived loss relative to a goal.

In this paper, we present the most direct test of reference-dependent preferences in the context of labor supply. Our study has three new features. First, we conduct a field experiment in which we manipulate output of our subjects exogenously. We hire high-school students to do data entry work, using a computer interface. The subjects are unaware that they are participating in an experiment. We slow down the subjects in a controlled fashion by causing response delays in the computer software. In the treatment of interest, the subjects are paid on a piece rate per data entered. This treatment allows us to examine directly how the subjects respond after they had been slowed down and earn less than they may have expected. Second, we implement a fixed-wage treatment that allows us to rule out potential confounds related to possible effects of the delays on effort costs.

Third, our setup also allows us to explicitly test for reference-dependent preferences in our setup. In a separate experiment, we measure risk preferences over small-stakes lotteries in our subjects. This allows us test whether risk preferences are related to the response to the delays. It also allows us to test more subtle predictions: As we show, a simple model of reference-dependent preferences (Kőszegi and Rabin, 2006) makes a clear prediction: If an individual is not very loss averse, she will

[1] There is no disagreement that higher wages lead to higher participation of workers (see also Oettinger, 1999). This leads to considerable methodological problems in interpreting the cab driver results, as they could be driven by selection. See **?** and Farber (2005) for a discussion.

[2] See Cox and Sadiraj (2005). Wakker (2005) shows, however, that this criticism is often poorly grounded in formal theory.

not respond to the delays at all. However if an individual is sufficiently loss-averse, she will fully compensate for the income loss due to the delay condition. At the same time, the same loss aversion parameter governs preferences over risky lotteries: The more loss averse an individual, the less risk she is willing to take. Thus, reference-dependent preferences predict a discontinuous relationship between risk preferences and the response to the delays. On the other hand, if individuals have concave valuations, but simply evaluate utility over narrowly defined intervals, risk aversion should have a more gradual impact on the response to delays.

Our results provide clear support to the notion that individuals have a narrowly-defined income target in mind. We find that individuals respond to having been in the delay condition by working harder if they are paid on a piece rate. The effect is statistically significant, and non-negligible. We find no such effect present when the subjects were paid a fixed wage. In the fixed-wage treatment, the only reason to change effort in response to the delay is if they affected effort costs. Thus, the exogenously implemented delays do not affect effort costs per se. We also find that the response to the delays is modulated by risk preferences as measured in the separate experiment. The more loss-averse an individual behaves in the risky-choice experiment, the more strongly he responds by increasing effort subsequent to being in the delay condition. This modulation of the response to the delays is highly specific to the piece-rate treatment, and completely absent in the fixed-wage treatment.

Further, our data point towards reference-dependent preferences as the right explanation for this phenomenon. We find a highly non-linear relationship between risk preferences and the response to the delays. This non-linearity is predicted by reference-dependent preferences, but inconsistent with concave utility. For example, we can explicitly that moderate risk aversion (log utility over narrowly defined events) generate our data. Overall, our results provide the most complete evidence of reference-dependent preferences in labor supply. Though many of the predictions of the reference-dependent model are subtle, they are borne out in the data in every detail.

The remainder of this paper is structured as follows: Section 2 describes the experimental setup and derives the behavioral predictions for the experiment. Section 3 explains the empirical strategy and discusses the results. Section 4 concludes.

II. The Experimental Setup

A. The Task

We recruited high-school students to perform a short-term work task. The task lasted two hours and consisted of entering data using a computer interface. The subjects were told that the data they were entering were used for a reserch project. It is not uncommon to recreuit high-school students for thise kinds of short-term tasks.

The subjects were entering data from two different boxes, containing tags from a bicycle messenger firm. The subjects had to enter various pieces of information from each tag, with specific instructions what to do when, e.g., information was missing or illegible. The boxes differed by the serial number. The subjects were instructed to enter only tags from the box designated on the screen of the computer interface. No explanation was provided why this was the case, and none of the subjects asked. The computer interface would not let the subjects enter the data if it came from the wrong box.

B. The Treatments

We manipulated the progress of the subjects by by sometimes delaying the appearance of the data entry screen. The subjects could not work, or even pick the right receipt, without the screen (because only the screen informed them about from which pile to pick a receipt). In the delay condition, the appearance of the screen was delayed by 6 seconds with probability 5/6, and delayed by 1 second with probability 1/6. In the baseline condition, we switched the probabilities, so that all subjects experienced some delays at some point.

Table 1 gives an overview over the different treatments. In our main treatment, subjects were paid a piece rate of CHF 0.2 per receipt entered. As Table 1 shows, half of the subjects in sessions 1 and 2 were delayed during the first hour of work, while half of the subjects were delayed in the second hour of work. In sessions 3 and 4, all subjects were delayed for both hours.

We also conducted a control treatment in which the subjects were paid a fixed wage of CHF 50. In these two sessions, half of the subjects were delayed in the first hour of work, and half of the subjects were delayed in the second hour of work.

C. Procedures

The experiments were conducted in the June 2005. We recruited subjects from several high schools in Zurich. These high schools were identical in terms of their composition of students, e.g., from different regions of the Canton of Zurich. We organized work sessions on different afternoons. Subjects could sign up for work on an afternoon if they did not have classes on that particular afternoon. Upon signing up for the work episode, the subjects completed a short survey asking them about previous job experience and computer skills. Subjects were only given a vague range for their potential earnings (CHF 40 - 60), and were not told how they would be paid. Therefore, the assignment to the particular treatment was random.

At the beginning of a work session, the subjects were seated in a computer lab at the University of Zurich. The subjects then received written instructions on how to enter the data. On the first page of these instructions, they were told how they would be paid (piece rate or fixed wage). They were walked through one example, and could ask questions if something was unclear. The subjects were then reminded again how they would be paid. After this brief instruction, the subjects started entering the data using the interface. The interface was programmed using *z-tree* (Fischbacher, 2007). The clock was not started to enter the data was not started until the instruction was over.

After the two hours were over, the subjects were asked whether they were willing to participate in a brief study, since they were already here. All subjects agreed to participate. As part of that study, subjects were offered six gambles, in which they could either win CHF 6, or lose CHF x, where x was varied from 2 to 7 in increments of 1 (Gaechter et al., 2007, this procedure is identical to the one used in). For each of the six gambles, the subjects had to indicate whether they would accept it or not. If they rejected the gamble, their earnings were unchanged. Then, one of the gambles was drawn at random, and a random device determined the payoff. Due to a technical lapse, the risk experiment was not conducted in one of the piece-rate treatment sessions.

D. Descriptive Statistics

Figure 1 gives an overview of the output per 15-minute interval during the two work hours. The figure clearly shows that, holding the delay condition constant, output was higher in the piece-rate treatment, as would be predicted by the economic model. Output is about 15 percent higher under piece rates than under fixed wages. This difference is very similar to results found in other studies

(See, e.g., Lazear, 2000; Shearer, 2004, for a recent study).

The figure also shows the impact of the delay condition on output. Comparing the dashed lines to the solid lines in Figure 1, we see that output was clearly reduced by the delays. The effect on output is about the same in both conditions and highly statistically significant ($p < 0.001$, corrected for clustering on individuals). Therefore, the manipulation of the software interface generates an income loss in the piece-rate treatment. Average earnings from the data entry task were CHF 49.2 in the piece-rate treatment, and CHF 50 in the fixed-wage treatment.

Figure 2 provides an overview over the behavior in the risk experiment. The figure shows the acceptance rates for each gamble, with the gambles ordered decreasing in their expected payoff. As can be seen in the figure, the vast majority of subjects behave in a risk averse fashion,as is common in these types of experiments (See, e.g., Holt and Laury, 2002, for a recent study). For example, 55 percent of the subjects reject a gamble in which they could lose CHF 4, or win CHF 6 with probabilty 0.5.

Table 2 provides a check of randomization. In this table, we report the correlations and p-values of various characteristics that may potentially be relevant for productivity in our setting. There are no measurable differences in age, gender, data entry experience or skills between the different treatments. There appears also no correlation between risk preferences and the treatment conditions. For some of the results to follow, we will exclude individuals who indicated a strategy of accepting and rejecting gambles that was non-monotonic in the expected values of the gambles (i.e., they rejected one gamble, but accepted the next *worst*). Approximately 10 percent of the subjects behaved in such a way. Such behavior is difficult to interpret. Importantly, as the table shows, the prevalence of this behavior is uncorrelated with any of the treatments or delay sequences.

E. Behavioral Predictions

This section derives predictions for the experiment to clarify how we identify different behavioral effects. The discussion will be kept informal, with all the analytical derivations relegated to the appendix. We discuss how individuals respond to being subjected to the delay condition in our experiment. The prime effect of the condition is to reduce the individuals' output (and we verify in section III. that it does not affect behavior in any other way). Different theories of behavior predict different responses to such a delay.

E..1 Predictions based on Intertemporal Maximization

The basic intuition from intertemporal maximization is that a small negative income shock to an individual's income should not change her labor supply choice. Because the individual is optimizing over a large number of periods, the small decrease in the income will be spread over many periods, leading to a very small reduction in consumption in every period (Browning et al., 1985). Because this change is small, the marginal utility of consumption is virtually unchanged. Consequently, the marginal return to effort is unchanged, and therefore, effort should stay constant. As shown in Fehr and Goette (2007), in each period, individuals should behave as if they were maximizing

$$u_t = \theta(w_t e_t - z_t) - g(e_t, \theta) \tag{1}$$

where θ is the marginal utility of income, w_t is the (discounted) wage per unit of effort e_t, z_t is a shock to income and $g()$ is a utility cost of effort that can be derived from the overall intertemporal maximization problem together with θ.

As can be seen in this formulation, the delay condition, best thought of as a small change in z_t, can not change effort through the channel of the marginal utility of money. It could, however, potentially change the marginal cost of effort. If, for example, past effort increases current marginal effort costs, then exposure to the delay condition could lower future effort costs, and boost effort, because the individual is prevented from working while delayed. However, this effect also has to be present in the fixed-wage treatment, in which the individuals simply set effort such that the marginal costs of effort are zero. In the appendix, we examine two empirically plausible forms for fatigue for repetitive work tasks in which fatigue can start to develop (See Foster et al., 2005, 2003; Koning et al., 1999, for evidence on pacing in sports.). As we show, if such an effect occurs because the delay condition affects effort costs, it also has to be present in the fixed-wage condition.

Thus, our empirical strategy to test for possible effects of the delay condition on effort costs is to check whether it affected effort choices in the fixed-wage treatment.

E..2 Predictions based on Reference-Dependent Preferences

We formulate a specific version of a model of reference-dependent preferences based on Kőszegi and Rabin (2006). Their model distinguishes between consumption utility and a reference-dependent component to utility. We assume that the individual's consumption utility conforms to the properties dictated by intertemporal optimization as specified in (1). We then follow Kőszegi and Rabin (2006) in how we add a reference-dependent part to the utility function. Loosely speaking, individuals gain utility if they do better than they could have expected, and lose utility if they do worse. These valuations are relative to the utility they could have gotten in every possible outcome, and weighted with the probability of this occurring. The shape of the valuation conforms to the Kahneman and Tversky (1979) value function.

We then ask the following question: Suppose the individual beliefs that she may have an adverse income shock of $-z$ with probability p. What is the optimal response to the shock, given the reference-dependent preferences? We show in the appendix (proposition 2) that optimal behavior must fall into one of two cases:[3] In the first case, the individual ignores the income shocks completely, and chooses the effort level as a neoclassical individual with the same consumption utility. In the second case, the individual *fully* eliminates the income shocks by increasing labor supply by z/w if the shock occurs. Which of the two behaviors is optimal depends on the weight of reference-dependent utility (and the degree of loss aversion) in the individual's objective function. For moderate loss aversion, case 1 is optimal: the best strategy is to choose the effort level that maximizes consumption utility and to live with the disutility from fluctuating income. However, if an individual is sufficiently loss averse, case 2 becomes optimal: it becomes attractive to distort effort in order to minimize the effect of the delays on utility: The reference-dependent disutility now stems from the fluctuation in effort costs on good days and bad days. However, by lowering effort, the individual can lower these costs. The optimal effort level in this case trades off the reductions in reference-dependent utility to the loss in consumption utility.

Thus, the model makes a clear prediction that one of two behaviors should prevail. No response to the delays, if not sufficiently loss averse, or full compensation of the income loss due to the delay condition, if loss aversion has a sufficiently high weight. Our strategy in the experiment will be to obtain a separate measure of loss aversion by having the individuals also participate in a risk experiment suitable to measure loss aversion, as discussed in section C.. We can then examine whether we find the predicted difference in the response to the delays based on this measure of loss aversion.

[3]In order to constrain our predictions, we only consider what Kőszegi and Rabin (2006) call preferred equilibria.

E..3 Predictions based on Narrowly-Defined Concave Utility

In this context,an alternative explanation is that individuals evaluate utility more narrowly, but otherwise conform to the standard assumptions about preferences. In this case, the individual maximizes

$$u_t = v(w_t e_t - z_t) - g(e_t) \tag{2}$$

where $v()$ is a concave function, and g is the effort cost function as before. In such a model, it is straightforward to show (as we do in proposition 8) that an individual increases effort in response to a negative shock z_t. The model also makes the prediction that the more curved $v()$ is, the stronger the response to z_t. The reason is that with more curvature, a negative income shock increases the marginal utility of income by more. Therefore, labor supply has to respond by more.

In our experiment, narrowly-defined concave utility can be distinguished from reference-dependent preferences by using the choices in the risk experiment. Under the null hypothesis of concave utility, more risk aversion in the experiment is associated with more curvature. Thus, there should be a monotonic increase in the reaction to the delays as risk aversion increases. On the other hand, reference-dependent preferences take the all-or-nothing form of responding to the delays, as we discussed earlier. This implication can be tested empirically.

III. Results

A. Empirical Strategy

All of the formal statistical tests will be based on regressions of the following form

$$T_{it} = \gamma D_{it} + x_{it} b + e_{it} \tag{3}$$

where T_{it} is the mean time for entering a receipt, measured in seconds after the screen returns for data entry in the 15-minute episode t. The time it takes to enter a receipt is the natural measure of effort in this setting. The variable D_{it} measures the total time individual i has been subjected to the delay condition up to episode t. This is our key variable measuring the impact of the delay on subsequent effort. The control variables x_{it} contain period effects as well as controls for individual computer skills as measured in the application forms that the subjects filled in. In some specifications, x_{it} will contain an individual fixed effect instead of the controls measured in the application. In all regressions, we allow the residual e_{it} to be heteroskedastic and possibly correlated within a subject, and only report robust standard errors adjusted for clustering on individuals. We estimate (3) separately for the piece-rate and fixed-wage treatment.

In order to test how risk aversion measured in the choice experiment affects the response to the delays, we run the regression

$$T_{it} = \gamma_0 D_{it} + \gamma_1 D_{it} RA_i + \delta RA_i + x_{it} b + e_{it} \tag{4}$$

where RA_i measures the risk aversion of individual i in the following way: RA is a variable counting the number of lotteries the individual rejects. Thus, the higher the index, the more risk averse the individual is. Thus, coefficient γ_1 measures whether more risk-averse individuals respond differently to the delay condition. Again, we estimate (4) separately for the piece-rate and fixed-wage treatment. In running this regression, we also exclude all individuals who display non-monotonic strategies in accepting the lotteries (e.g., reject a lottery with expected gain CHF 1, but accepting again a lottery with expected gain CHF 0.5), as such behavior is difficult to interpret. However, none of our conclusions depend on this exclusion.

B. Baseline Results

Figure 3 provides a first descriptive look at the overall evidence. The figure plots the average entry times as a function of the time during which the individual had been exposed to delays. We subtract overall period means from each observation in order to purge the data of the learning effects evident in Figure 1. The figure also displays a standard error band around the mean. Turning to the results for the fixed wage treatment, Figure 3 shows that there is essentially no relationship between delays and effort in the fixed-wage treatment. When paid on fixed wages, individuals always choose approximately the same work pace, irrespective of the delay condition. This is initial evidence that the delays condition did not affect the marginal costs of effort. Turning to the results for the piece-rate treatment in Figure 3, we see that the delay condition leads individuals to work harder subsequently, as can be seen by the lower entry times as the delay progresses in this treatment.

Table 3 provides a formal statistical test for the qualitative pattern seen in Figure 3. The upper panel of Table 3 displays the results for the piece-rate treatment, the lower panel for the fixed-wage treatment. The table confirms the impression from the figure: In any specification, we find a significant effect of the cumulative delays on effort in the piece rate condition, but never in the fixed-wage condition. In the first three columns in the top panel, we estimate the impact of a one-hour exposure to the delays to be a reduction in the entry time of 2.6 to 2.7 seconds. In the second column, we also include a dummy variable indicating whether the individual is slowed down in the current period t or not. As can be seen, the point estimate is small, and far from significant. Thus, it is the cumulative delay, rather than the current delay condition that changes labor supply. In the third column, we also include individual fixed effects. The point estimate of the coefficient is unchanged, but estimated with more precision than before. Except for the specification in which we also control for the current delay condition, cumulative delays are always significant despite the conservative correction of the standard errors. In order to get a better sense of the quantitative importance of these effects, we re-estimate the equation with $\log(T_{it})$ as the dependent variable. The point estimate now indicates that a one-hour exposure to the delay condition reduces the entry times by 0.1 log points, or approximately 10 percent. Thus, these results confirm that individuals work significantly harder as a result of being exposed to the delay condition in the piece-rate treatment.

The lower panel of Table 3 displays the results for the fixed-wage treatment. Consistent with the figure, we never find a significant effect of the delays on entry times. The point estimate is very unstable and bounces around depending on the specification. The results confirm that there is no significant effect of the delay condition on effort in the fixed-wage condition.

In sum, the baseline results show a considerable effect of the delay condition on effort in the piece-rate condition, but no effect in the fixed-wage condition. This rules out explanations based on non-separable effort costs discussed in section E.. We therefore now turn to exploring the treatment effect in the piece-rate condition in more detail.

C. Interactions with Risk Aversion

Figure 4 displays entry times as a function of exposure to the delay condition in Figure 3. However, in this figure, we divide the sample into two groups based on the behavior in the risk experiment, by splitting them at the median: The group labeled as "less loss-averse than median" is the group of individuals accepting the lottery win CHF 6 / lose CHF 4, the group labeled "more loss averse than median" is the group rejecting this lottery (see also Figure 2). In the piece-rate treatment, the figure reveals a striking separation in behavior between the two groups of individuals: The more risk-averse individuals strongly react to the delay condition by increasing effort substantially as exposure to the delay increase. Conversely, the less risk-averse individuals hardly respond to the delay exposure at all. The shaded areas around the two curves indicate standard error bands. As can be seen, they are clearly separate, indicating a significant difference in the response to the delay condition between the two groups. In the fixed-wage treatment, however, we see no such separation. If paid a fixed wage, risk aversion does not modulate the response to the delay condition anymore: The two curves are very close together, and the standard error bands thoroughly overlap.

Table 4 displays the corresponding formal statistical test by estimating equation (4). The first column in the table shows that the impression that risk-averse individuals respond to the delay condition more strongly is confirmed in the regression. The interaction effect is highly significant. In the second column, we add a dummy variable indicating whether or not the individual had ever been subjected to delays, and also interact this variable with RA_i. Most importantly, the inclusion of the control variable does not diminish the significance of the interaction term between RA_i and exposure to the delays. The third column includes individual fixed effects. They do not change the result and leave the interaction effect significant. The fourth through sixth column in Table 4 repeat this exercise for the fixed-wage treatment. In line with the result from Figure 2, there is no significant interaction between risk aversion and the response to delay exposure.

Putting the results to an even more stringent test, we ask whether the response to the delay condition and the modulation by risk aversion differs between the piece-rate and fixed-wage condition. We test this hypothesis for each of the three specifications. The p-values are reported in the first three columns. The show that there is a significantly different response to the delay condition in the piece-rate treatment, modulated by risk aversion as measured in the risk experiment.

To summarize, the response to the delay condition in the piece rate treatment is clearly influence by an individual's risk aversion. This effect is highly specific to the piece rate condition: It is not significant in the fixed-wage condition. Moreover, we find that it is different from the fixed-wage condition, even by comparing the estimates from the piece-rate treatment to the noise estimates from the fixed-wage treatment. Therefore, this is evidence of, either, reference-dependent preferences or concave utility over narrowly-defined events. The next subsection proposes an explicit test to distinguish between the two.

D. Reference-Dependent Preferences vs. Concave Utility

In order to distinguish between reference-dependent preferences and concave utility, we estimate an augmented version of equation (3). We estimate

$$T_{it} = \gamma_1 RA_{1i} D_{it} + \gamma_2 RA_{2i} D_{it} + ... + \gamma_6 RA_{6i} D_{it} + x_{it} b + e_{it} \tag{5}$$

where RA_{ki} is equal to one if individual i rejected k of the lotteries, and zero otherwise. All other variables are defined as before. Thus, the coefficients $\gamma_1, ..., \gamma_6$ indicate, for each possible degree of risk aversion, how much the individual responded to a one-hour exposure to the delay condition. The pattern then allows us to examine whether it qualitatively supports the predictions from the reference-dependent model or from the model with concave utility. If we see a gradual increase in the response to the delays, this would be evidence against the reference-dependent model, but in line with the predictions of the model with concave utility. Conversely, the reference-dependent model predicts that individuals with moderate loss aversion should not respond to the delays. With sufficiently strong loss aversion, individuals switch to a strategy in which they fully compensate for the income loss due to the delays. Because the distinction between the two models depends on the pattern estimated for the different γ_k, we examine the robustness of the result with respect to how we control for risk aversion. In our baseline specification, we control for RA_i, to preserve comparability with the estimates of equation (4). However, we estimate (5) also with fixed effects $\delta_k RA_{ki}$ to control for risk aversion in the most flexible way.

Figure 5 displays the estimated coefficients γ_k for the two specifications discussed above, with standard errors adjusted for clustering on individuals. The figure shows a striking pattern: For moderate levels of risk aversion (all individuals who accept a coin flip of win 6 / lose 4), there is no reaction at all to the delay condition. The point estimates are very close to zero for each of the three categories, and a joint test reveals that none is significantly different from zero ($p = 0.97$ in the baseline specification, $p = 0.74$ in the fixed effects specification). However, the individuals rejecting the win 6 / lose 4 lottery or better lotteries, respond very strongly to the delay condition, and approximately of the same order of magnitude. Each of these three coefficients is individually significant, and a joint test also rejects that they are zero at conventional significance levels ($p = 0.06$ in the baseline specification, and $p < 0.01$ in the fixed effects specification).

As we argued, such a pattern is consistent only with reference-dependent preferences, and inconsistent with narrowly-defined concave utility. Consider, in particular, the response of the individuals who reject the win 6 / lose 5 coin flip, but accept all better lotteries. More than 30 percent of our subjects fall into this group. Assuming iso-elastic utility, and using the money earned in the experiment as the relevant wealth, we find that the coefficient of relative risk aversion for this category must lie between 1.67 and 4.15. Thus, these individuals should be highly risk-averse, as far more curvature in the utility function is required than, e.g., log utility. With this curvature, individuals should also strongly react to the delay condition, as this amount of curvature can generate potent income effects. However, we find no evidence that they respond to the delay condition at all. Further, we find that the response to the delay does not vary for the three groups that do respond. This, again, is in line with the reference-dependent model, but inconsistent with the concave-utility model.

Overall, we take this pattern as very strong evidence in favor of reference-dependent preferences. The data confirm the predictions from the reference-dependent model in every detail, while they reject important predictions of the model with concave utility.

IV. Concluding Remarks

References

Becker, Gary S., *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, Chicago, Il: University of Chicago Press, 1993.

Blundell, Richard and Thomas Macurdy, "Labor Supply: A Review of Alternative Approaches," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 3A, New York: Elsevier Science, 1999, pp. 1559–1695.

Browning, Martin, Angus Deaton, and Margaret Irish, "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, May 1985, *53* (3), 503–544.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics*, May 1997, *112* (2), 407–441.

Chou, Yuan K., "Testing Alternative Models of Labour Supply: Evidence from Taxi Drivers in Singapore," *Singapore Economic Review*, April 2002, *47* (1), 17–47.

Cox, James C. and Vjollca Sadiraj, "Implications of Small- and Large-Stakes Risk Aversion for Decision Theory," 2005.

Deaton, Angus, *Understanding Consumption*, Oxford, UK: Oxford University Press, 1992.

Farber, Henry S., "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy*, February 2005, *113* (1), 46–82.

Fehr, Ernst and Lorenz Goette, "Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment," *American Economic Review*, March 2007, *97* (1), 298–317.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, June 2007, *10* (2), 171–178.

Foster, Carl, Jesus Hoyos, Conrad Earnst, and Alejandro Lucia, "Regulation of Energy Expenditure during Prolonged Athletic Competition," *Medicine and Science in Sports and Exercise*, April 2005, *37* (4), 670–675.

—, Jos De Koning, Floor Hettinga, Joanne Lampen, Christopher Dodge, Maarten Bobbert, and Jon Porcari, "Effect of Competitive Distance on Energy Expenditure during Simulated Competition," *International Journal of Sports Medicine*, 2003, *25*, 198–204.

Gaechter, Simon, Andreas Herrmann, and Eric Johnson, "Individual-Level Loss Aversion in Risky and Riskless Choice," 2007. Unpublished.

Heath, Chip, Richard Larrick, and George Wu, "Goals as Reference Points," *Cognitive Psychology*, 1999, *38*, 79–107.

Holt, Charles A. and Susan K. Laury, "Risk Aversion and Incentive Effects," *The American Economic Review*, December 2002, *92* (5), 1644–1655.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decisions Under Risk," *Econometrica*, March 1979, *47* (2), 263–291.

Koning, Jos De, Maarten F. Bobbert, and Carl Foster, "Determination of Optimal Pacing Strategy in Track Cycling with an Energy Flow Model," *Journal of Science and Medicine in Sports*, 1999, *2* (3), 266–277.

Kőszegi, Botond and Matthew Rabin, "A Model Of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 2006, *124* (1), 1133–1166.

Lazear, Edward P., "Performance Pay and Productivity," *The American Economic Review*, December 2000, *90* (5), 1346–1361.

Oettinger, Gerald, "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors," *Journal of Political Economy*, 1999, *107* (2), 360 – 392.

Shearer, Bruce, "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," *Review of Economic Studies*, April 2004, *71* (247), 513–534.

Tversky, Amos and Daniel Kahneman, *Choices, Values, and Frames*, Cambridge, MA: Cambridge University Press, 2000.

Wakker, Peter P., "Formalizing Reference Dependence and Initial Wealth in Rabin's Calibration Theorem," 2005.

V. Appendix

Proposition 1 (The Response to Delays with Fatigue with Recovery)**.** *In the case of fatigue with recovery, a reduction in earlier work has the same qualitative impact in the fixed-wage and piece-rate treatment: Effort should go up in the period following the delay.*

Proof. The individual maximizes

$$V_t(k_t) = \max_{e_t} we_t - c(e_t + \alpha k_t) + V_{t+1}(\alpha k_t + e_t) \tag{6}$$

in each work episode $t = 1, ..., M$ on a work day. The first-order condition to this is

$$\begin{aligned} w - c_t' + V_{t+1}' &= 0 \\ \Leftrightarrow c_t' - V_{t+1}' &= w \end{aligned} \tag{7}$$

Now take the derivative of (6) with respect to k_t to obtain

$$\begin{aligned} V_t' &= -\alpha c_t' + \alpha V_{t+1}' \\ &= -\alpha(c_t' - V_{t+1}') \end{aligned} \tag{8}$$

Now shift (7) and (8) one period forward to find $V_{t+1}' = -\alpha w$ for all $t < M$. Substitute this into (7) to obtain

$$c_t' = (1 - \alpha)w \tag{9}$$

for all $t < M$, and

$$c_M' = w \tag{10}$$

because $V_{M+1}' = 0$.

□

Proposition 2 (The Characterization of the Equilibria)**.** *There are two types of equilibria: In the first type of equilibrium, the individual does not react to windfall losses at all and sets effort* *according to $g'(e^*) = w$. In the second type of equilibrium, the individual fully offsets the windfall loss, i.e., $e_B = e_G + z/w$ if*

$$w > g'(e_B)\frac{1+\eta(p+(1-p)\lambda)}{1+\eta\lambda}$$

and

$$w < g'(e_G)\frac{1+\eta(p+(1-p)\lambda)}{1+\eta}$$

Effort on the good day is given by

$$(1-p)g'(e_G) + pg'(e_G + z/w) + p(1-p)\eta(\lambda-1)(g'(e_G+z/w) - g'(e_G)) = w$$

and z is sufficiently small. The second type of equilibrium is the preferred equilibrium if

$$p\eta(\lambda-1) > 1$$

Proof. The proof consists of four steps.

Step 1: Lemma 3 - 5 imply that any personal equilibrium either has to be of the type where $e_B = e_G$ or $e_B > e_G$, and $we_B - z = we_G$ by ruling out all other possibilities.

Step 2: Lemma 6 shows that an equilibrium, in which the individual does not react to z always exists. Since the reference-dependent utility in this case does not depend on e, the best among this class of equilibria is the one that maximizes the material payoff.

Step 3: Lemma 7 shows that an equilibrium, in which the individual fully compensates for z exists if z is small enough. The lemma characterizes the best in this class of equilibria.

Step 4: Thus, finding the preferred equilibrium amounts to determining which of the two equilibria identified in steps 2 and 3 gives the higher utility. Notice that for $z = 0$, both equilibria trivially give the same payoff. Using the equilibrium utilities, we get for the "neoclassical" equilibrium

$$\frac{\partial U^{neo}}{\partial z} = -p(1+(1-p)\eta(\lambda-1))$$

by applying the envelope theorem. Similarly, we obtain for the "income targeting" equilibrium

$$\frac{\partial U^{IT}}{\partial z} = -p(1 + (1 - p)\eta(\lambda - 1))g'(e_G + z/w)/w$$

Taking the difference between the two, we get

$$\frac{\partial U^{IT}}{\partial z} - \frac{\partial U^{neo}}{\partial z} = -p(1 + (1 - p)\eta(\lambda - 1))[g'(e_G + z/w)/w - 1] \tag{11}$$

Thus, if $g'(e_G+z/w) < w$, the last term in (11) will be negative, and therefore, the income targeting equilibrium become better as z increases. In order to ensure that this is the case, recall that e_G solves

$$pg'(e_G + z/w) + (1 - p)g'(e_G) + p(1 - p)\eta(\lambda - 1)(g'(e_G + z/w) - g'(e_G)) = w$$

Thus, if we can show that

$$(1 - p)g'(e_G) + p(1 - p)\eta(\lambda - 1)(g'(e_G + z/w) - g'(e_G)) > (1 - p)g'(e_G + z/w) \tag{12}$$

it is impossible that $g'(e_G + z/w) \geq g'(e^*)$, and therefore $g'(e_G + z/w)/w < 1$. Simple algebra shows that (12) holds, if $p\eta(\lambda - 1) > 1$, as claimed. This completes the proof. □

Lemma 3. *Any situation with $e_B < e_G$ cannot be a personal equilibrium.*

Proof. In any situation with $e_B < e_G$, we have

$$U_B = we_B - z - g(e_B) + (1 - p)\left[\eta\lambda(we_B - z - we_G) + \eta(g(e_G) - g(e_B))\right]$$
$$U_G = we_G - g(e_G) + p\left[\eta(we_G - we_B + z) + \eta\lambda(g(e_B) - g(e_G))\right]$$

Notice that if $g'(e_B) < w$, U_B is increasing in e. Therefore in this case, there is no such equilibrium. If $g'(e_G) > w$, then U_G is decreasing in e. Therefore in this case, there is no such equilibrium. But since the two cases exhaust all possibilities, there is no such equilibrium. □

Lemma 4. *Any situation with $e_B > e_G$, but $we_B - z < we_G$ cannot be a personal equilibrium.*

Proof. In any situation with $e_B > e_G$, but $we_B - z < we_G$ we have

$$U_B = we_B - z - g(e_B) + (1-p)\left[\eta\lambda(we_B - z - we_G) + \eta\lambda(g(e_G) - g(e_B))\right]$$
$$U_G = we_G - g(e_G) + p\left[\eta(we_G - we_B + z) + \eta(g(e_B) - g(e_G))\right]$$

Notice that if $g'(e_B) > w$, U_B is decreasing in e. Therefore in this case, there is no such equilibrium. If $g'(e_G) < w$, then U_G is increasing in e. Therefore in this case, there is no such equilibrium. But since the two cases exhaust all possibilities, there is no such equilibrium. □

Lemma 5. *Any situation with $e_B > e_G$, but $we_B - z > we_G$ cannot be a personal equilibrium.*

Proof. In any situation with $e_B > e_G$, but $we_B - z > we_G$ we have

$$U_B = we_B - z - g(e_B) + (1-p)\left[\eta(we_B - z - we_G) + \eta\lambda(g(e_G) - g(e_B))\right]$$
$$U_G = we_G - g(e_G) + p\left[\eta\lambda(we_G - we_B + z) + \eta(g(e_B) - g(e_G))\right]$$

Notice that if $g'(e_B) > w$, U_B is decreasing in e. Therefore in this case, there is no such equilibrium. If $g'(e_G) < w$, then U_G is increasing in e. Therefore in this case, there is no such equilibrium. But since the two cases exhaust all possibilities, there is no such equilibrium. □

Lemma 6. *There always exists an equilibrium in which $e_B = e_G = e$, where $g'(e^*) = w$. This is the equilibrium giving the highest payoff of all the equilibria with $e_B = e_G$.*

Proof. Consider the case where e solves $g'(e^*) = w$. Then,

$$U_B = we - g(e) - z - (1-p)\eta\lambda z$$
$$U_G = we - g(e) + p\eta z$$

In order to check whether it pays to deviate from this strategy, note that nothing can be gained by either increasing or decreasing e_B. Nothing can be gained for the material payoff by changing e_B by the definition of e. With regard to the reference-dependent part of utility, increasing e causes a loss in effort costs and a reduces the income loss. However, since $g'(e) = w$ these effects offset each other. Decreasing e_B causes a loss in in income, but only a gain in effort costs. Hence, the overall effect would be negative. A similar argument applies to changing e_G. This establishes the personal equilibrium.

Finally, notice that, in equilibrium, reference-dependent utility does not depend on effort. Hence, since the equilibrium maximizes the material payoff, it is the equilibrium in this class with the highest utility. □

Lemma 7. *There exists an equilibrium with $we_B - z = we_G$ for a sufficiently small z if*

and

Proof. Under the proposed equilibrium, the utility on a good and bad day, respectively, is given by

$$U_G = we_G - g(e_G) + (1-p)\eta(g(e_B) - g(e_G))$$
$$U_B = we_B - z - g(e_B) + p\eta\lambda(g(e_G) - g(e_B))$$

In order to check whether an individual has an incentive to deviate, we evaluate the marginal utility of effort at different points. Consider first U_B for $e < e_B$. We have

$$U_B(e) = we - z - g(e) + \eta\lambda(we - z - we_G) + p\eta(g(e_B) - g(e)) + (1-p)\eta\lambda(g(e_G) - g(e))$$

Taking the derivative w.r.t. e, we get

$$\frac{\partial U_B}{\partial e} = w - g'(e) + \eta\lambda w - \eta(p + (1-p)\lambda)g'(e)$$
$$\Rightarrow \frac{\partial U_B}{\partial e} > 0 \text{ if } w > g'(e_B)\frac{1 + \eta(p + (1-p)\lambda)}{1 + \eta\lambda} \tag{13}$$

Thus if condition (13) is satisfied, the individual does not want to reduce effort from e_B. Further, if z is small enough, this condition is necessarily satisfied. Similarly, we can show that for $e < e_G$,

$$\frac{\partial U_G}{\partial e} > 0 \text{ if } w > g'(e_G)\frac{1 + \eta}{1 + \eta\lambda}$$

which is implied by (13). Now, consider $e > e_G$. In this case,

$$U_G(e) = we - g(e) + \eta(w - we_G) + p\eta(g(e_B) - g(e)) + (1-p)\eta\lambda(g(e_G) - g(e))$$

Taking the derivative w.r.t. e, we obtain

$$\frac{\partial U_G}{\partial e} < 0 \text{ if } w < g'(e_G)\frac{1+\eta(p+(1-p)\lambda)}{1+\eta} \tag{14}$$

Again, for small enough z, condition (14) is necessarily satisfied. Similarly, we can show that for $e > e_B$,

$$\frac{\partial U_B}{\partial e} < 0 \text{ if } w < g'(e_B)\frac{1+\eta\lambda}{1+\eta}$$

which is implied by (14). To find the best in this class of equilibria, notice that the equilibrium utility is given by

$$U = we_G - (1-p)g(e_G) - pg(e_G + z/w) + p(1-p)\eta(1-\lambda)(g(e_G + z/w) - g(e_G)) \tag{15}$$

Maximizing (15) w.r.t. e_G yields the first-order condition

$$(1-p)g'(e_G) + pg'(e_G + z/w) + p(1-p)\eta(\lambda-1)(g'(e_G + z/w) - g'(e_G)) = w$$

as claimed in the proposition. □

Proposition 8 (Response to income shocks with narrowly-defined concave utility)**.** *Suppose an individual has utility function*

$$u(e) = v(we - z) - g(e)$$

where $v()$ is a concave utility function and all other parameters are defined as in the text. Then,

$$\frac{\partial e}{\partial z} = \frac{-v''(ew)w}{-v''(ew)w^2 + c''(e)} > 0$$

Proof. The first-oder condition for effort in this case is

$$u(e)' = v'(we - z)w - g'(e) = 0$$

Totally differentiating the first-oder condition yields the desired result. □

Figure 1: Output over Time

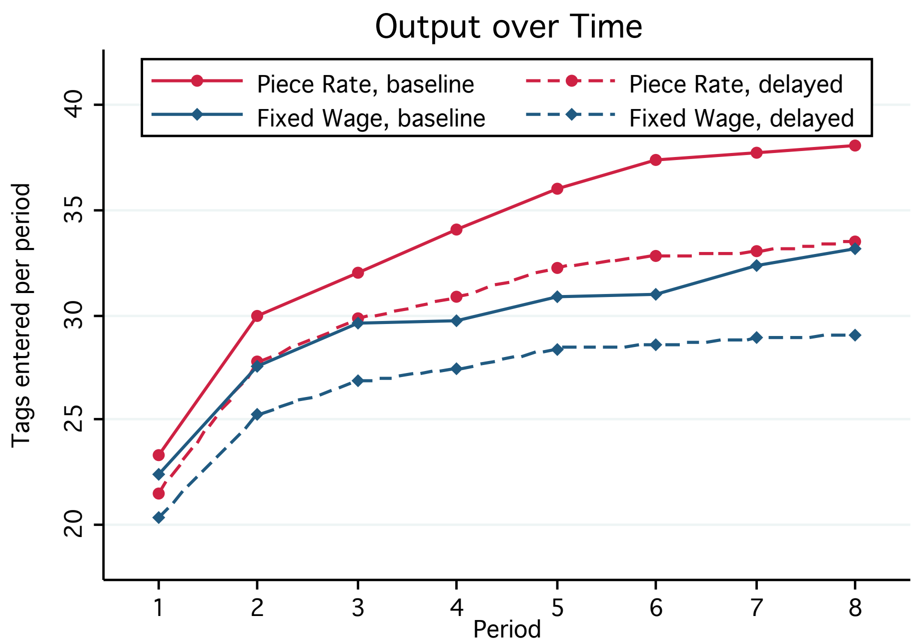

Figure 2: Behavior in Loss Aversion Experiment

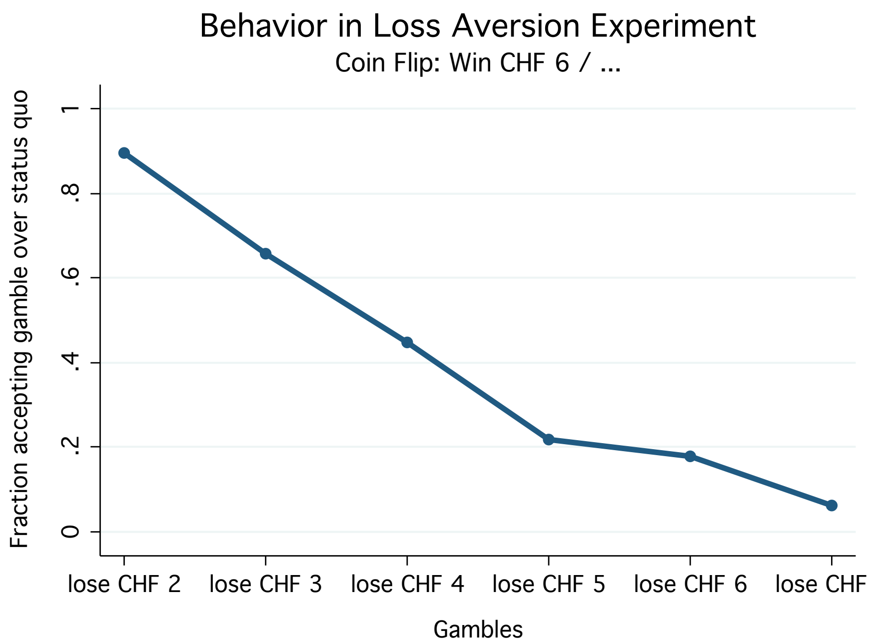

Table 1: The Design of the Study

Appearance of screen after a tag has been entered

	Piece-Rate Treatment				Fixed-Wage Treatment	
	$S1$ and $S2$		$S3$ and $S4$		$S5$ and $S6$	
Observations	$N = 16$	$N = 19$	$N = 20$	$N = 21$	$N = 18$	$N = 18$
Periods 1 - 4	Delayed	Normal	Delayed	Delayed	Delayed	Normal
Periods 5 - 8	Normal	Delayed	Delayed	Delayed	Normal	Delayed

Figure 3: Behavior in Loss Aversion Experiment

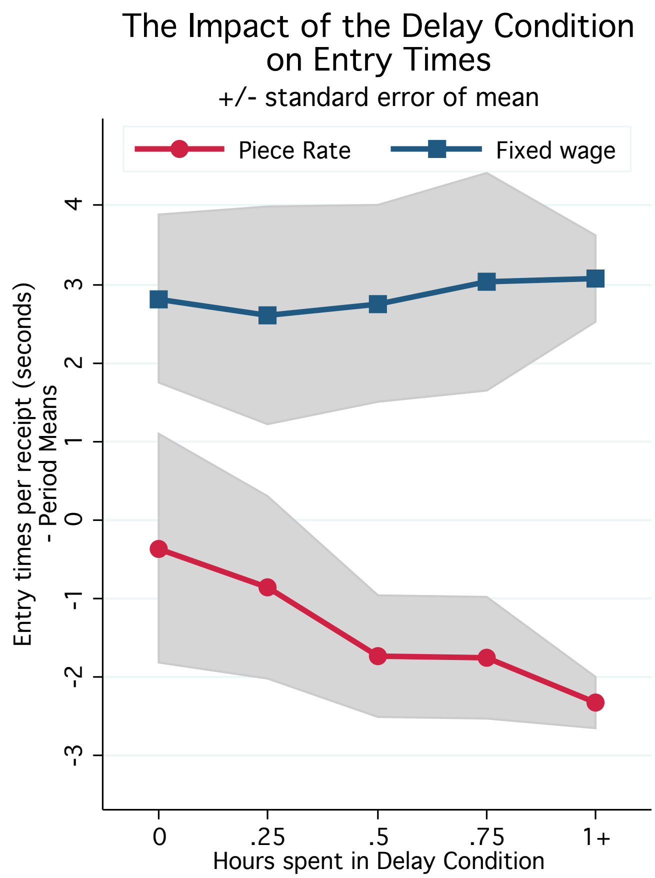

Table 2: Checks of Randomization

Correlations of individual characteristics with treatment condition in period 5

	Loss averse	one switching point	Age	Gender	Data entry experience	typing skills
Delayed in period 5	-0.0888	-0.1101	-0.0067	-0.1407	-0.0618	0.0809
	(0.390)	(0.248)	(0.945)	(0.139)	(0.529)	(0.401)
delayed in periods 1 - 4	0.025	-0.0078	-0.0108	-0.1254	0.1635	-0.1279
	(0.809)	(0.935)	(0.911)	(0.188)	(0.094)	(0.183)
Piece-rate pay	0.0677	0.0387	0.0013	-0.1104	0.0822	-0.1423
	(0.512)	(0.685)	(0.989)	(0.247)	(0.403)	(0.138)

Notes: p-values, unadjusted for multiple hypotheses tests, are in parentheses.

Figure 4: Risk Preferences and Entry Times

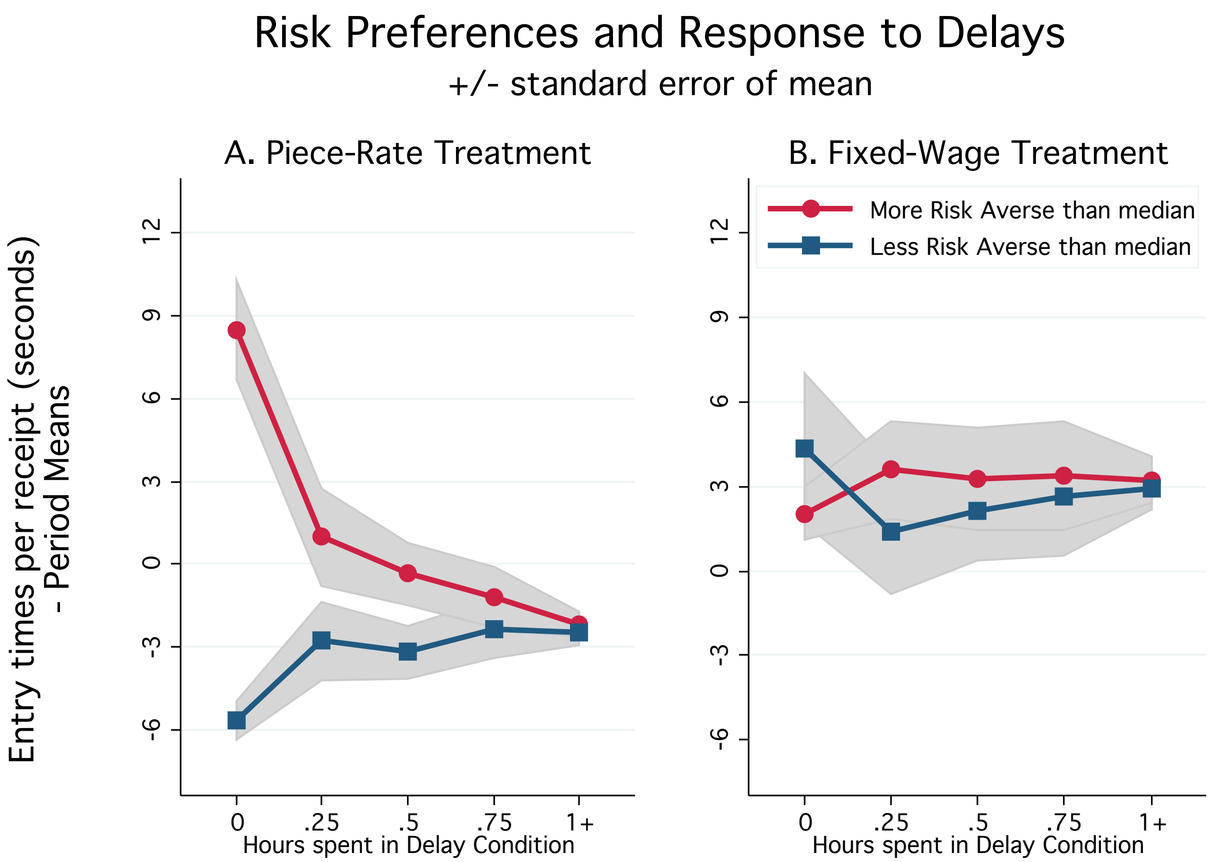

Figure 5: Loss Aversion vs. Concavity

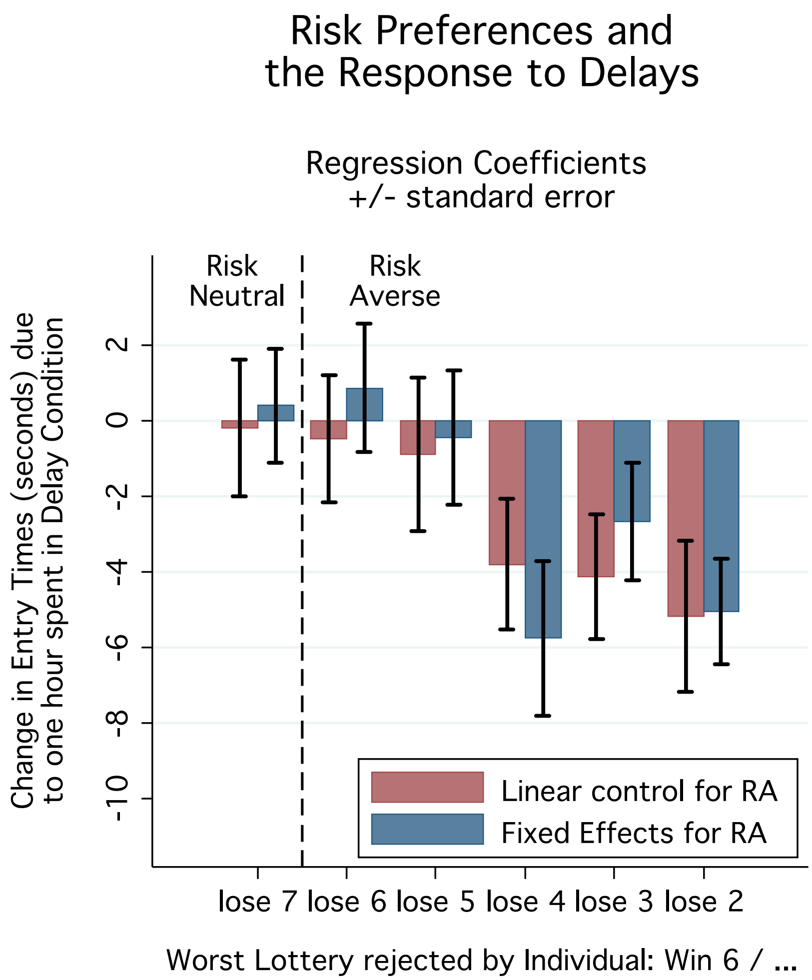

Table 3: The Overall Experimental Results
OLS Regressions

	Piece-Rate Treatment				
Dependent Variable:	T_{it}			$\log(T_{it})$	
Hours spent in delay condition	– 2.767*	– 2.610	– 2.734**	– 0.100*	– 0.095***
	(1.548)	(1.606)	(1.121)	(0.058)	(0.030)
slowed down in t		– 0.244			
		(0.900)			
R^2	0.472	0.472	0.875	0.442	0.918
N	568	568	568	568	568

	Fixed-Wage Treatment				
Dependent Variable:	T_{it}			$\log(T_{it})$	
Hours spent in delay condition	– 1.532	– 1.658	0.446	– 0.042	0.004
	(2.830)	(3.017)	(1.058)	(0.091)	(0.034)
slowed down in t		0.318			
		(0.947)			
R^2	0.275	0.275	0.814	0.286	0.859
N	280	280	280	280	280

Notes: T_{it} is the mean entry time per receipt of individual i in fifteen-minute intervall t. All specifications include period effects and controls for basic demographics and computer skills. Standard errors adjusted for clustering on individuals. *, **, *** denote significance at the 10, 5, 1 percent level, respectively.

Table 4: Interactions with Risk Aversion Index
Dependent Variable: T_{it}
OLS Regressions

	Piece-Rate Treatment			Fixed-Wage Treatment		
Hours spent in delay condition	1.739 (1.835)	0.518 (1.523)	– 0.483 (1.821)	– 4.639 (1.821)	4.101 (4.008)	– 4.639 (5.136)
Has ever been slowed down		12.068** (5.921)			– 11.250 (7.629)	
Hours in delay × RA_i	– 1.178*** (0.262)	– 0.815*** (0.231)	– 0.593** (0.248)	0.616 (1.144)	– 1.229 (0.786)	1.304 (1.277)
ever slowed down × RA_i		– 3.631* (2.077)			2.887* (1.647)	
RA_i Index	2.248*** (1.390)	5.354*** (1.925)		0.822 (0.040)	– 0.392 (1.830)	
Response to delay differs from specification in fixed wage treatment	$p = 0.083$	$p = 0.077$	$p = 0.087$			
Individual Fixed Effects?	No	No	Yes	No	No	Yes
R^2	0.538	0.549	0.892	0.303	0.314	0.797
N	392	392	392	232	232	232

Notes: T_{it} is the mean entry time per receipt of individual i in fifteen-minute intervall t. RA_i is the risk aversion index, counting the number of lotteries that individual i rejected. All specifications include period effects and controls for basic demographics and computer skills. Standard errors adjusted for clustering on individuals. *, **, *** denote significance at the 10, 5, 1 percent level, respectively.