

Institute for Empirical Research in Economics University of Zurich

> Working Paper Series ISSN 1424-0459

Working Paper No. 125

Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment

Ernst Fehr and Lorenz Götte

September 2005

Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment^{*}

Ernst Fehr^{**} University of Zürich and CEPR

and

Lorenz Goette^{***} University of Zürich and CEPR

September 2005

Abstract: Most previous studies on intertemporal labor supply found very small or insignificant substitution effects. It is not clear, however, whether these results are due to institutional constraints on workers' labor supply choices or whether the behavioral assumptions of the standard life cycle model with time separable preferences are empirically invalid. We conducted a randomized field experiment in a setting in which workers were free to choose their working times and their efforts during working time. We document a large positive wage elasticity of overall labor supply and an even larger wage elasticity of labor hours, which implies that the wage elasticity of effort per hour is negative.

While the standard life cycle model cannot explain the negative effort elasticity, we show that a modified neoclassical model with preference spillovers across periods and a model with reference dependent, loss averse preferences are consistent with the evidence. With the help of a further experiment we can show that only loss averse individuals exhibit a significantly negative effort response to the wage increase and that the degree of loss aversion predicts the size of the negative effort response.

^{*} The authors acknowledge support from the Swiss National Science Foundation under project number 1214-051000.97. This paper greatly benefited from the comments of David Card and two excellent referees. In addition, we thank George Akerlof, Henry Farber, David Huffman, Reto Jegen, Rafael Lalive, George Loewenstein, Jennifer Lerner, Stephan Meier, Matt Rabin, Jason Riis, Alois Stutzer, Richard Thaler, and George Wu for their helpful comments.

^{**} University of Zurich, Institute for Empirical Economic Research, Blümlisalpstrasse 10, CH-8006 Zurich. Phone: 0041-44-6343709. Fax: 0041-44-6344907. E-mail: <u>efehr@iew.unizh.ch</u>.

^{***} University of Zurich, Institute for Empirical Economic Research, Blümlisalpstrasse 10, CH-8006 Zurich. Phone: 0041-44-6343704. Fax: 0041-44-6344907. E-mail: <u>lorenz@iew.unizh.ch</u>.

The intertemporal substitution of labor supply has far-reaching implications for the interpretation of important phenomena. If, for example, the intertemporal substitution of labor supply is high, one may interpret the large variations in employment during business cycles as voluntary choices by the workers rather than involuntary layoffs. Intertemporal substitution also plays a crucial role in the propagation of shocks across periods (Romer, 1996; King and Rebelo, 1999). Previous studies have found little evidence for intertemporal substitution of labor, however; the estimated elasticities are often small and statistically insignificant, and sometimes even negative (see, e.g., Mankiw, Rotemberg and Summers 1985; Pencavel 1986; Altonji 1986; Blundell, 1994; Card 1994 and Blundell and MaCurdy 1999).¹

However, the low estimates of intertemporal substitution are difficult to interpret because of serious limitations in the available data. The life cycle model of labor supply predicts intertemporal substitution with regard to *transitory* wage changes or wage changes the workers *anticipate*. Yet, the typical wage changes are not transitory; hence they are associated with significant income effects. In addition, it seems almost impossible to infer reliably from existing data whether the workers anticipated the wage change. Furthermore, serious endogeneity problems arise, as both supply and demand conditions determine wages.² Thus, the typically available data require many auxiliary assumptions when testing the life cycle model of labor supply.

Another issue arises if labor markets are characterized by a significant amount of job rationing or other constraints on workers' labor supply. In fact, there is strong evidence suggesting that workers are not free to set their working hours (Ham 1982, Kahn and Lang 1991, Dickens and Lundberg 1993), rendering the identification of the source of small intertemporal substitution effects difficult, even if the above mentioned problems could be solved. A small intertemporal substitution effect could be due to these constraints or it could be that the behavioral assumptions behind the life cycle model are wrong. Indeed, Camerer, Loewenstein, Babcock and Thaler (1997) put forward the view that New York City cab drivers' daily labor supply is driven by nonstandard, reference dependent, preferences that exhibit loss aversion

¹ After reviewing a sizeable part of the literature, Card (1994) concludes, for instance, that the "very small magnitude of the estimated intertemporal substitution elasticities" can only account for a tiny fraction of the large person-specific year-to-year changes in labor supply.

²Oettinger (1999) shows that if one neglects the endogeneity of wage changes, estimates of labor supply elasticities are severely downward-biased.

around a target income level. This view has recently been called into question by Farber (2004, 2005).

In this paper, we use an ideal data set to study workers' responses to transitory wage changes. We conducted a randomized field experiment at a bicycle messenger service in Zürich, Switzerland. The bicycle messengers receive no base wage that is independent of effort, and are paid solely on commission. We have precise information for all the workers on the number of shifts they work, and the revenues they generate per shift. A shift always comprises five hours and workers in our sample worked at most one shift per day. A key feature of our experiment is the implementation of an *exogenous* and *transitory* increase in the commission rate by 25 percent. Therefore, we can be sure that unobserved supply or demand variations did not induce the change in the commission rate (i.e., the "wage" change). Each participant in the experiment knew ex-ante the precise duration and size of the wage increase. Since the wage was only increased during four weeks, its impact on the workers' lifetime wealth is negligible.

In the firm under study, the messengers can freely choose how many shifts (hours) they work, and how much effort they exert (to generate revenues). This means that our setting also provides an ideal environment for studying the behavioral foundations of labor supply: in our context, the absence of intertemporal substitution effects cannot be attributed to institutional constraints on labor supply. The exogenous change in the commission rate raises the returns from both the number of shifts and effort per shift. Therefore, we have the unique opportunity of studying how hours *and* effort respond to the wage increase *and* how *overall* labor supply (i.e., the number of hours times the effort per hour) is affected.

Our experimental results show that the wage increase caused a large increase in overall labor supply. Our estimate of the intertemporal elasticity of substitution with regard to overall labor supply is between 1.12 and 1.25. This large effect is exclusively driven by the increase in the number of hours worked. In fact, the elasticity of hours worked with regard to the wage is even higher than the elasticity of overall labor supply; it lies between 1.34 and 1.50, considerably in excess of that found in previous studies (e.g., Oettinger 1999). The fact that the elasticity of hours worked is larger than the overall labor supply elasticity suggests that the effort per hour decreased in response to the wage increase. And indeed, a detailed analysis indicates that effort

per shift decreased by roughly 6% in response to the wage increase, which implies a wage elasticity of effort per shift of -0.24.

These results confirm the non-experimental evidence in previous studies of intertemporal labor substitution based on samples where workers were largely unconstrained in choosing hours and effort. Camerer et al. (1997) and Chou (2002) examined how cab drivers, after having decided to work on a given day, vary their daily working time (which is a good proxy for daily effort) in response to wage variations. Both studies report that workers work fewer hours (provide less effort) on high-wage days, indicating a negative effort elasticity. Oettinger (1999) investigated how stadium vendors adjust their probability of working in response to transitory wage variations across different baseball games. He develops a good set of ex-ante predictors of game attendance, which can be used to instrument for the wage: temperature, day of the week, the ranking of the home team, the quality of the opposing team, etc. In his IV estimates, Oettinger finds a positive and significant wage elasticity of participation, ranging from 0.53 to 0.64. The data in all these studies, however, do not allow making inferences on the overall labor supply elasticity: the data by Camerer et al. (1997) and Chou (2002) do not reflect the participation decision and Oettinger's data do not allow reliable inferences about effort per game.

When viewed through the lens of a standard neoclassical model with separable time preferences, the reduction in effort seems puzzling. After all, the rise in the commission rate provides strong economic incentives for working more hours and for working harder during those hours. Our results are, however, immune against the criticisms that have recently been raised against the study of Camerer et al. (Farber 2004, 2005). One problem is that the source of the variation behind cab drivers' wages is not completely clear. If, for example, there are common supply side shocks (e.g. most drivers don't work on the 4th of July), then the supply of cabdriver hours will be small on these days and the wage will be high. As a result, there will be a negative correlation between wages and hours that has nothing to do with loss averse preferences. This criticism does not apply to our study because we vary workers' wages experimentally. A second concern is a possible selection effect: higher wages may induce cab drivers to work a few hours on days when they otherwise would not have worked. Such an effect may generate a negative correlation between daily wages and daily hours even though all individuals behave exactly as the standard model with time separable preferences predicts. Our data enable us to solve this problem as well.

In addition, we would like to point out that a reasonable extension of the standard model can, in principle, explain a negative effort elasticity. In the theory part of our paper, we show that a neoclassical model, in which last period's effort raises this period's marginal disutility of effort, is consistent with our evidence: workers who work in more periods may rationally decide to reduce effort per period. However, we also show that a rational choice model, with reference dependent preferences exhibiting loss aversion around the reference point, is also able to explain the evidence. The intuition behind this model is that workers with loss averse preferences have a daily reference income level.³ Daily incomes below the reference level are experienced as a "loss", and the marginal utility of income is large in the loss domain. In contrast, the marginal utility of income at and above the reference level decreases discontinuously to a lower level. Workers who temporarily earn higher wages are more likely to exceed the reference income level, hence reducing their marginal utility of income and ultimately inducing them to provide less effort. At the same time, however, workers with higher wages have a higher overall utility from working a shift so that they can more easily cover the fixed costs of getting to work. Hence, they are more likely to work.

There are thus two competing theories which are consistent with the facts. In order to discriminate between the two theories, we conducted another experiment based on the idea that loss aversion is a personality trait which affects behavior across several domains (Kahneman and Tversky 2000; Gaechter, Hermann and Johnson 2005). In this experiment, we measured the individual worker's loss aversion in lottery choices. We then used these measures to examine whether the negative response of effort per shift is due to the existence of loss averse workers. We indeed find that the degree of a worker's loss aversion contributes significantly to the negative effort elasticity. Moreover, it turns out that workers who do not show loss aversion in the lottery choices also do not have a significantly negative elasticity. Only workers with loss aversion reduce effort per shift significantly when paid a high wage.

Thus, the result of our second experiment favors the model with reference dependent preferences over the neoclassical model with "disutility spillovers" across periods. Of course, the evidence from the second experiment is not the ultimate arbitrator, but it suggests that future work should not disregard the loss aversion model because it could contribute to a deeper

³ Heath, Larrick and Wu (1999) provide evidence that goals often serve the function of a reference point.

understanding of effort choices. At the same time, we should also point out that one third of the workers in our sample did not exhibit loss aversion and a negative effort elasticity. Thus future work should take the possibility of heterogeneous preferences more seriously. In addition, the results of our first experiment unambiguously show that whatever behavioral forces worked against the intertemporal substitution of labor, they were apparently not capable of generating a negative elasticity of the overall labor supply. The behavioral forces that worked in favor of intertemporal substitution far outweighed any opposing forces.

The remainder of this paper is structured as follows. Section I describes the institutional environment and the details of the field experiment. Section II discusses the implications of different models of labor supply. Section III reports the results from the field experiment. We first report the impact of the wage increase on overall labor supply and then discuss how shifts responded; we finally present the evidence on how the wage increase affected the effort per shift. This section also describes the follow-up experiment and discusses the link between individual loss aversion and workers effort responses. Section IV concludes the paper.

I. Experimental Set-up

Our study is based on the delivery records of two relatively large Swiss messenger services – Veloblitz and Flash Delivery Services (henceforth "Flash") – which are located in Zurich. Each firm employs between 50 and 60 bicycle messengers. The available records contain information about when a messenger worked a shift, all deliveries he conducted during a shift, and the price of each delivery. Thus, we know which messengers worked a shift and how much revenue they generated during the shift for each day in the observation period. We first describe the organization of work at a bicycle messenger service below and then present our experiment in more detail.

A. Work at a Messenger Service

Unless pointed out below explicitly, the arrangements are the same for the two messenger services, Veloblitz and Flash. When a potential worker applies for a job with one of the messenger services, an experienced messenger evaluates him or her with respect to fitness, knowledge of locations, names of streets, courtesy, and skills regarding handling the CB radio.

Once accepted as an employee, messengers can freely choose how many five-hour shifts they will work during a week. There are about 30 shifts available at Veloblitz, and about 22 at Flash on each workday from Monday to Friday. In principle, messengers could work more than one shift per day, but none of them chose to do so during the experiment and the months prior to the experiment. The shifts are displayed on a shift plan for every calendar week at the messenger service's office. There are two types of shifts, called "fixed" and "variable". A "variable" shift simply means that a shift is vacant at a particular time. Any messenger can sign up to work that shift, e.g., on Wednesday from 8 am to 1 pm. If a messenger commits to a "fixed" shift, he has to work that shift every week. For example, if a messenger chooses Wednesday, 8 am - 1 pm as a fixed shift, he will have to fill that shift on every Wednesday for at least six months. Thus, fixed shifts represent a commitment for several months and can only be cancelled with at least four weeks notice. Roughly two-thirds of the shifts are fixed. It is also import to note for our examination that the number and the allocation of fixed shifts across messengers remained the same during the whole experiment; the company refused to change the fixed shifts just because of the experiment. All shifts that are not fixed are variable shifts; they are available for any messenger to sign up for. All workers participating in our study worked both fixed and variable shifts.

Two further items are worth mentioning. First, there is no minimum number of shifts that the messengers have to work at either messenger service. Second, both messenger services found filling the available shifts difficult. There is almost always at least one unfilled shift and, on average, almost 3 shifts per day remain unfilled. For example, during the period before the experiment, from September 1999 – August 2000, approximately 60 shifts remained unfilled every month. This implies that messengers are unlikely to be rationed in the choice of shifts.

Messengers receive no fixed wage. Their earnings are given *solely* as a fixed percentage w of their daily revenues. Hence, if a messenger carries out deliveries that generate revenues r during his shift, his earnings on that day will be wr. An important feature of the work environment concerns the fact that messengers have substantial discretion on how much effort to provide during a shift. They only stay in contact with the dispatcher at the messenger service's office through CB radio. In order to allocate a delivery, say, from location A to location B, the dispatcher will contact the messenger whom he thinks is closest to A to pick up the delivery. All messengers can listen in on the radio. If they believe that they are closer to A than the messenger

who was originally contacted, they can get back to the dispatcher and say so and will then be allocated to that delivery. Conversely, if the messenger does not want to carry out the delivery from A to B, he may just not respond to the call. Messengers have, therefore, several means of increasing the number of deliveries they complete. They can ride at higher speed, follow the radio more actively, or find the shortest possible ways to carry out a delivery.

Thus, work at a bicycle messenger service closely approximates a model where individuals are unconstrained in choosing how many shifts (hours) to work, and how hard to work (i.e., how many deliveries to complete during a shift).

B. The Experimental Design

In order to evaluate the labor supply effect of a temporary wage increase, we randomly assigned those Veloblitz messengers who were willing to participate in the experiment to a treatment and a control group and we implemented a fully anticipated temporary increase in the commission rate by (roughly) 25% for the treatment group. The commission rate for men in the treatment group was temporarily increased from w = 0.39 to w = 0.49 and the rate for women was temporarily increased from w = 0.54. The additional earnings for the messengers were financed by the Swiss National Science Foundation.

In order to participate in the experiment, all messengers had to complete a questionnaire at the beginning and at the end of each experimental period. The messengers were informed that a failure to complete all questionnaires meant that they would not receive the additional earnings from the experiment. All messengers who finished the first questionnaire also filled in the remaining questionnaires.⁴ Thus, the group of messengers who participated in the experiment was constant during the whole experiment, i.e., there was no attrition. Randomization into a treatment and a control group was achieved by randomly allocating the *participating* messengers into a group A and a group B. The randomization was based on the administrative codes that the messenger service uses to identify a messenger in its accounting system. All messengers at Veloblitz are assigned a number depending on the date when they started working for the company. The first messenger who worked at Veloblitz was assigned the number 1, the second 2,

⁴ The messengers at Veloblitz who did not participate in the experiment were almost exclusively workers who were already quite detached from the company or who where on probationary shifts. The "detached" workers typically worked roughly one shift per week during the experiment and the months prior to the experiment.

and so forth. The participating messengers with odd numbers were assigned to group A, participating messengers with even numbers to group B.

The messengers did not know that the purpose of the experiment was the study of labor supply behavior, nor did they realize that we received the full (anonymous) records of each messenger about the number of shifts and the number of deliveries completed. If pressed, we told the participants that we wanted to study the relation between wages and job satisfaction. The purpose of our study was credible because the questionnaires contained several questions related to job satisfaction.⁵

We implemented a 25 percent increase in the commission rate during four weeks in September 2000 for group A. The messengers in group B were paid their normal commission rate during this time period so that they can be used as a control group. In contrast, only the individuals in group B received a 25 percent increase in the commission rate during four weeks in November 2000, while the members of group A received their normal commission rate and therefore served as a control group. Thus, a key feature of our experiment is that there are two experimental periods that lasted for four weeks and that both group A and B served as a treatment and a control group in one of the two experimental periods. This characteristic of our experiment enables us to provide a very clean isolation of the impact of the temporary wage increase. If, for example, the implemented wage change increases labor supply, then we should observe this increase both in the first and the second experimental period. In the first experimental period, the members of group A (who receive the higher wage in this period) should exhibit a larger labor supply than the members of group B while the reverse should be true in the second experimental period – members of group B (who receive the higher wage in this period) should supply more labor. Moreover, since members of both groups are in the treatment and the control group we can identify the treatment effect within subjects by controlling for individual fixed effects.

Our experimental design also enables us to control for the income effect of the wage increase, i.e. we can identify the pure substitution effect for the participating messengers. We announced the experiment in the last week of August 2000, and all additional earnings from the

⁵ These features of the experiment ensure that our results cannot be affected by the Hawthorne effect. The Hawthorne effect means that subjects behave differently just because they know that the experimenters observe their behavior. Yet, our subjects did not know that we could observe their behavior during the wage increase. Moreover, since both the treatment group and the control group are part of the overall experiment, and since our key results rely on the comparison between these groups we control for a potential Hawthorne effect.

experiment – regardless of whether subjects were members of group A or group B – were paid out after the end of the second experimental period in December 2000.⁶ Thus, the budget constraint for both groups of participating messengers was affected in the same way. Due to the randomization of the participating messengers into groups A and B, the income effect cancels out if we identify the treatment effect by comparing the labor supply of control and treatment group.

As demand for delivery services varies from day to day and from month to month, it is useful to control for time effects. The available information about Flash enables us to identify possible time effects across treatment periods because both Veloblitz and Flash operate in the same market. There is a strong correlation between the total daily revenues at Veloblitz and Flash. When we compute the raw correlation between total revenues at the two firms over the two experimental periods plus the four weeks prior to the experiment, we find a correlation of 0.56 (Breusch-Pagan $\chi^2(1) = 18.93$, p < 0.01, N = 60 days). Even after removing daily effects from both series, the correlation is still 0.46 (Breusch-Pagan $\chi^2(1) = 13.16$, p < 0.01, N = 60 days). This shows that the revenues at the two firms are highly correlated, even over quite short a time horizon.⁷

We believe that our experiment represents a useful innovation to the existing literature for several reasons. First, it implements a fully anticipated, temporary and exogenous variation in the (output based) wage rates of the messengers, which is key for studying the intertemporal substitution of labor. The experimental wage increase was large and provides a clear incentive for increasing labor supply. Moreover, the participating messengers are experienced, and daily fluctuations in their earnings are common. Hence, we experimentally implement a wage change in an otherwise familiar environment. Second, the data we obtained from Veloblitz allows us to study two dimensions of labor supply: Hours as measured by the number of shifts and effort as measured by the revenues generated per shift or the number of deliveries per shift. No other study

⁶ In the time period between the announcement of the experiment and the beginning of the first treatment period no new regular workers arrived at Veloblitz. Only workers who worked on probationary shifts arrived during this time period and they were not allowed to participate in the experiment because they often leave the firm after a short time and lack the necessary skills. Including them in the experiment would have created the risk of attrition bias.

⁷ If we add the 8 months prior to the experiment we find a correlation of about 0.75. To check the robustness of our results we also include - in some of our regressions - the non-participating messengers at Veloblitz in the non-experimental comparison group that is used to identify time effects.

that we are aware of can look at these two dimensions simultaneously. Third, we can combine the data set with the full records from a second messenger service operating in the same market. This will prove useful for investigating any effect that the experiment might have had on the non-participating messengers at Veloblitz, and helps to control for demand variations over time.

II. Predictions

In this subsection, we derive predictions about labor supply behavior in our experiment. We use two types of models –neoclassical models and a model of reference dependent utility with loss averse workers. In view of our results, we are particularly interested in the question of which kind of model is capable of predicting an increase in shifts (hours) worked *and* a decrease in effort per shift.

A. Neoclassical Model with Time-Separable Utility

In this subsection, we integrate the institutional setting at our messenger service into a canonical model of intertemporal utility maximization with time-separable utility. We define the relevant time period to be one day. Consider an individual who maximizes lifetime utility

$$U_o = \sum_{t=0}^T \delta^t u(c_t, e_t, x_t) \tag{1}$$

where $\delta < 1$ denotes the discount factor, u() represents the one-period utility function, c_t denotes consumption, e_t is effort in period t and x_t denotes a variable that affects the preference for working on particular days. For example, a student who works a few shifts per week at Veloblitz may have higher opportunity costs for working on Fridays because he attends important lectures on Fridays. The utility function obeys $u_c > 0$, $u_e < 0$ and is strictly concave in c_t and e_t . The lifetime budget constraint for the individual is given by

$$\sum_{t=0}^{T} \hat{p}_{t} c_{t} (1+r)^{-t} = \sum_{t=0}^{T} (\hat{w}_{t} e_{t} + y_{t}) (1+r)^{-t}$$
(2)

where \hat{p}_t denotes the price of the consumption good, \hat{w}_t the period *t* wage per unit of e_t and y_t non-labor income. For convenience we assume that the interest rate *r* is constant and that there is no uncertainty regarding the time path of prices and wages. The sign of the comparative static predictions is not affected by these simplifying assumptions.

In appendix A, we show that along the optimal path, the within-period decisions of a rational individual maximizing a time-separable concave utility function like (1), subject to constraint (2), can be equivalently represented in terms of the maximization of a static one-period utility function that is linear in income.⁸ This static utility function can be written as

$$w(e_t, x_t) = \lambda w_t e_t - g(e_t, x_t), \tag{3}$$

where $g(e_t, x_t)$ is strictly convex in e_t , and measures the discounted disutility of effort and x_t captures exogenous shifts in the disutility of effort. λ measures the marginal utility of life-time wealth and w_t represents the discounted wage in period *t*. Thus, $\lambda w_t e_t$ can be interpreted as the discounted utility of income arising from effort in period *t*.⁹

Workers who choose effort according to (3) respond to an anticipated temporary increase in w_t with a higher effort e_t . A rise in w_t increases the marginal utility returns of effort, λw_t , which increases the effort level e_t^* that maximizes $v(e_t, x_t)$. The situation is, a bit more complicated in our experiment, however, because the messengers can choose the number of shifts and the effort during the shift. Theoretically the existence of shifts can be captured by the existence of a minimal effort level \tilde{e} that has to be met by the worker or by the existence of fixed costs of working a shift. Intuitively, if there is a fixed cost of working a shift, an employee will only work on a given day if the utility of e_t^* , $v(e_t^*, x_t)$, is higher than the utility of not going to work at all. As a wage increase raises $v(e_t^*, x_t)$ workers are more likely to work on a given day, i.e., the number of shifts worked will increase. ¹⁰

⁸ Our characterization is inspired by the results in Browning, Deaton and Irish (1985) who show that the within period decisions can be characterized in terms of the maximization of a static profit function. However, the present exposition is more convenient for our purposes.

⁹ λ is constant along the optimal path of c_t and e_t . This has the important consequence that an *anticipated* temporary wage variation does not affect the marginal utility of life-time wealth. Thus, anticipated temporary variations in wages (or prices) have no income effects. Yet, if there is a non-anticipated temporary increase in the wage, λ changes immediately after the new information about the wage increase becomes available and remains constant at this changed level afterwards. For our experiment, this means that the income effect stemming from the temporary wage increase has to occur immediately after the announcement of the experiment in August 2000. Thereafter, the marginal utility of life-time wealth again remains constant so that there are no further changes in λ during the experiment. The difference in behavior between the treatment group and the control group during the two treatments can thus not be due to changes in λ . Note also that (3) does not only describe the optimal effort choice in period *t* but is also based on the optimal consumption decision in period *t*. For any change in effort, the consumption decision also changes in an optimal manner (see appendix).

¹⁰ More formally, the wage increase raises the utility of going to work for all x. Hence the participation condition will be met for more states x.

B. Neoclassical Model with Non-Separable Utility

The prediction of the previous subsection is, however, not robust to the introduction of nonseparable utility functions. To illustrate this, consider a simple example where

$$v(e_t, e_{t-1}) = \lambda e_t w - g(e_t(1 + \alpha e_{t-1})).$$
(4)

This example captures the intuition that if a worker worked yesterday, he has higher marginal cost of effort today. We assume, for simplicity, that $e_0 = 0$, that there are only two further time periods, period 1 and period 2 and that the wage is constant across time. If we ignore discounting the two-period utility is given by $U = v(e_1, 0) + v(e_2, e_1)$. Therefore, if the wage is high enough to induce the worker to go to work in both periods the worker chooses effort e_1^{**} and e_2^{**} according to

$$\lambda w = g'(e_1) + \alpha e_2 g'(e_2(1 + \alpha e_1))$$
(5)

$$\lambda w = g'(e_2(1 + \alpha e_1))(1 + \alpha e_1).$$
(6)

If work is supplied in both periods, an increase in e_1 causes a higher disutility of labor in period 2 which lowers e_2 . Of course, rational workers take this effect into account when they decide on e_1 which means that the overall marginal disutility of e_1 is higher if e_2 is positive compared to when it is zero. In particular, if wages are low enough so that it is no longer worthwhile to work in period 2 ($e_2 = 0$), the first order conditions are given by

$$\lambda w = g'(e_1) \tag{5'}$$

$$\lambda w < g'(0)(1 + \alpha e_l). \tag{6'}$$

A comparison of conditions (5) and (6) with conditions (5') and (6') shows that it is possible that the optimal effort e_1 according to (5') is higher than e_1^{**} and e_2^{**} . In appendix B we provide an explicit example that proves this point. This possibility arises because the marginal disutility of working in each of the two periods, which is indicated by the right hand side of (5) and (6), is higher than the marginal disutility of working only in period 1 which is given by $g'(e_1)$. In the context of our experiment, this means that messengers who work more shifts when the wage is high may rationally decide to reduce the effort per shift.

The simple model above does not predict that workers who work more shifts (days) will necessarily reduce their effort per shift. It only allows for this possibility. If the wage increase is large enough, it is also possible that workers who behave according this model raise their effort per shift. There is, however, one prediction that follows unambiguously from a neoclassical approach regardless of whether utility is time separable or not. Browning, Deaton, and Irish (1985) have shown that a general neoclassical model predicts that overall labor supply $\sum e_i$ increases in the high wage periods in response to a temporary increase in wages. Applied to our context, this means that during the four-week period where the wage is higher for the treatment group, the total revenue (or the total number of deliveries) of the treatment group should exceed the total revenue (or the total number of deliveries) of the control group.

C. Reference Dependent Utility

Another potential explanation for why effort per shift might decrease in response to a temporary wage increase is that individuals could have preferences that include a daily income target \tilde{y} that serves as a reference point. The crucial element in this approach is that if a person falls short of his or her target, he or she is assumed to experience an additional psychological cost which is not present if income varies above the reference point. This explanation is suggested by the large number of studies indicating reference dependent behavior (for a selection of papers on this see Kahneman and Tversky 2000). Evidence from psychology (Heath, Larrick, and Wu, 1999) suggests that the marginal utility of a dollar below the target is strictly higher than the marginal utility of a dollar above the target.¹¹ A *daily* income target seems plausible for bike messengers in our sample because their daily incomes are a salient feature of their work environment. The

¹¹ See Goette and Huffman (2003) for survey evidence on this point. They present bike messengers with direct survey scenarios to elicit whether messengers care more about making money in the afternoon if they had good luck in the morning than after a bad morning. In their scenarios good luck means that messengers had the opportunity of making particularly profitable deliveries in the morning. For example, good luck means that a delivery just crosses an additional district boundary; such deliveries command a substantially higher price without much additional effort. About 70% of the messengers respond in a fashion consistent with daily income targeting.

messengers keep receipts from each delivery they did on a shift. This makes them acutely aware of how much money they earn from each completed delivery. The messengers also turn in the receipts at the end of the shift, making it difficult for them to keep track of how much money they earned over several shifts. A daily income target may also serve the messengers as a commitment device for the provision of effort during the shift. Zurich is rather hilly and riding up the hills several times during a shift requires quite some effort – in particular if the weather is bad or towards the end of a shift. A daily income target may thus help the messengers overcome a natural tendency to "shirk" that arises from a high marginal disutility of effort.

Formally, the existence of reference dependent behavior can be captured by the following one-period utility function.

$$v(e_t) = \frac{\lambda(w_t e_t - \tilde{y}) - g(e_t, x_t)}{\gamma \lambda(w_t e_t - \tilde{y}) - g(e_t, x_t)} \quad \text{if } w_t e_t \ge \tilde{y}$$

$$(7)$$

where $\gamma > 1$ measures the degree of loss aversion, i.e., the increase in the marginal utility of income if the individual is below the income target. Previous evidence (see Kahneman and Tversky 2000) suggests that $\gamma \approx 2$ for many individuals. Loss aversion at this level creates powerful incentives to exert more effort below the income target. However, once individuals attain the target \tilde{y} , the marginal utility of income drops discretely (from $\gamma \lambda$ to λ), causing a substantial reduction in the incentive to supply effort.

The preferences described in (7) imply that workers increase the number of shifts when they are temporarily paid a higher wage: a rise in wages increases the utility of working on a given day. Thus, at higher wages it is more likely that the utility of working $v(e_t)$ exceeds the fixed costs of working. At the same time, however, the increase in wages makes it more likely that the income target is already met or exceeded at relatively low levels of effort. Therefore, compared to the control group, the workers in the treatment group are more likely to face a situation where the marginal utility of income is λ instead of $\gamma \lambda$, i.e., they face lower incentives to work during the shift.¹² As a consequence, members of the treatment group will provide less effort than members of the control group.

The previous discussion shows that reference dependent preferences and a neoclassical model with non-separable preferences may make similar predictions. In particular, both models are consistent with a reduction in effort per shift during the wage increase. However, the reduction in effort in the income target model should be related to the degree of loss aversion γ , as explained above. Evidence suggests that there is substantial heterogeneity in the degree of loss aversion between individuals, and that individuals who are loss averse in one type of decisions are also loss averse in other domains of life (see Gaechter, Herrmann, and Johnson, 2005). Thus, in principle, the two explanations can be distinguished if one obtains an individual level measure of γ .

III. Results

This section reports the results from our field experiment. Our analysis is based on the four weeks prior to the first experimental period and the two subsequent experimental periods in which first group A and then group B received a wage increase. The data contain the day of each delivery, the messenger's identification number, and the price for each delivery. Thus, we have, in principle, two measures of labor supply – the amount of revenue generated and the number of deliveries completed. Since longer deliveries command a higher price and require more effort, the revenue is our preferred measure of labor supply. However, our estimates of the treatment effect are almost identical for either choice of the labor supply measure.

A. The Impact of the Wage Increase on Total Revenue per Messenger

The first important question is whether there is a treatment effect on total revenue per messenger during the first and the second experimental period. Figure 1a and 1b as well as Tables I and II

¹² If γ is sufficiently high relative to the wage increase, one may obtain the extreme result that the worker provides effort to obtain exactly \tilde{y} before and after the increase. In this case, the worker's effort obviously decreases in response to the wage increase because at higher wages \tilde{y} is obtained at lower effort levels. In general, the larger γ , the sharper the kink in the objective function and the more likely the worker's optimal effort choice e^* will be at the kink, i.e., the more likely $\gamma \lambda \hat{w}_t > g'(e^*) > \lambda \hat{w}_t$ holds. Note, however, that even if the worker is not a "perfect" income targeter, i.e., even if before or after the wage increase he does not earn exactly \tilde{y} , negative effort responses may occur.

present the relevant data. The figures depict the revenue data for groups A and B at Veloblitz; the tables show also the data of messengers at Flash and those messengers at Veloblitz who did not participate in the experiment. Figure 1a and Table I show the "raw" revenue per messenger – uncontrolled for individual fixed effects. Figure 1b and Table II control for individual fixed effects by showing how – on average – the messengers' revenues deviate from their personspecific mean revenues. Thus, a positive number here indicates a positive deviation from the person-specific mean, a negative number indicates a negative deviation.

These figures and tables show that group A and B generate very similar revenues per messenger during the four weeks prior to the experiment. If we control for individual fixed effects, we find that the revenues per messenger are almost identical across groups and close to zero. For example, the difference in revenues between group A and B is only CHF 48.9 if we control for person specific effects with a standard error of CHF 366.6 (see Figure 1b and Table II). This difference is negligible in comparison to the average revenue of roughly CHF 3400 that was generated by a messenger during the pre-experimental period. Thus, in the absence of an experimental treatment the messengers in group A and B behave in the same way.

However during the first experimental period (henceforth, "treatment period 1"), in which group A received the higher wage, the total revenue generated by group A is much larger than the revenue of group B, indicating a large treatment effect. On average, messengers in group A generated roughly CHF 4131 while messengers in group B only generated revenues of CHF 3006 during this period (see Table I and Figure 1a). Then, in the second experimental period, this pattern is reversed and group B, which receives now the higher wage, generates much more revenue. In treatment period 2 group B generates revenues per messenger of CHF 3676 while group A only produces revenues of CHF 2734. If we control for individual fixed effects (see Figure 1b) we can see that the standard errors are relatively small, suggesting that the differences across groups are significant.

Insert Figures 1a and 1b about here

To control more tightly for statistical differences across groups, we performed regressions (1) - (4) in Table III. All regressions are of the form

$$r_{it} = \alpha_i + \delta T_{it} + d_t + e_{it} \tag{8}$$

where r_{it} measures the revenue generated by messenger *i* during a four-week period *t*, α_i is a fixed effect for messenger *i*, T_{it} is a dummy variable that is equal to 1 if the messenger is on the increased commission rate, d_t is a time dummy estimated for treatment period 1 and for treatment period 2, and e_{it} is the error term.

Regression (1) is based only on the data of groups A and B at Veloblitz. Due to the random assignment of the participating messengers across groups and due to the fact that both groups served once as a control and once as a treatment group, this regression allows for a clean isolation of the treatment effect. The regression indicates that the treatment effect is highly significant and that the messengers on a high wage generate roughly CHF 1000 additional revenue compared to the experimental control group.

The other three regressions show that the measured impact of the experimental wage increase on the treated group remains almost the same if we include in the comparison group messengers of Flash and non-participants of Veloblitz. Regression (2) compares the treatment group at Veloblitz with all other messengers at Veloblitz and finds again a large and significant treatment effect of roughly CHF 1000. Regression (3) uses observations from all messengers at Veloblitz and the messengers at Flash. The inclusion of the messengers at Flash is suggested by the strong correlation in revenues between Flash and Veloblitz. Regression (3) also includes a dummy for the whole non-treated group at Veloblitz, i.e. the messengers in the control group and those who did not participate in the experiment. Therefore, this dummy measures whether the non-treated group at Veloblitz behaved differently relative to the messengers at Flash and the treatment dummy measures whether the treated group at Veloblitz behaved differently relative to the messengers at Flash. In this regression, the coefficient of the treatment dummy indicates again a treatment effect of roughly CHF 1000. In addition, the dummy for the whole non-treated group at Veloblitz is small and insignificant, indicating that the non-treated group was not affected by the wage increase for the treated group. This result suggests that the wage increase for the treated group did not constrain the opportunities for working for the non-treated group at Veloblitz. The result is also consistent with the permanent existence of unfilled shifts and with survey evidence; the overwhelming majority of the messengers stated that they could work the number of shifts they wanted to work. Finally, regression (4) uses the data from all Veloblitz and Flash messengers but does not include the dummy for the whole nontreated group at Veloblitz. Therefore, the treatment dummy in this regression measures whether the treated group at Veloblitz generates a different revenue from all the other messengers at Veloblitz and Flash. Again, the treatment effect is of similar size and significance as in the previous regressions.¹³

In summary, the above results indicate a large and highly significant effect of a temporary wage increase on the total effort of the treated group. In contrast to many previous studies, our results imply a large intertemporal elasticity of substitution. The standard way to calculate this elasticity is to estimate (8) in logs. However, because some participants of the experiment did not work at all during a four-week period (because they went on vacation) we cannot use this method. If we did, we would have to drop these observations although the decision to not work during a whole four week period also represents a labor supply decision. For this reason we include all participating messengers in our measure and compare the percentage increase in the revenue per messenger (which is our proxy for overall labor supply per messenger) that is due to the wage increase with the 25% increase in wages. We have seen that the treatment effect is roughly CHF 1000. The average revenue across group A and B is CHF 3568 in treatment period 1; in treatment period 2 it is 3205. Thus, the intertemporal elasticity of substitution is between (1000/3568)/0.25 = 1.12 and (1000/3205)/0.25 = 1.25, which is substantially larger compared to what previous studies have found (see e.g. Oettinger 1999).¹⁴

¹³ It is also noteworthy that we find a negative effect of time on revenues per messenger in all four regressions. However, while the time effect is never significant for the first treatment period, it is higher for the second treatment period and reaches significance at the 5% level in some of the regressions. These time effects suggest that a comparison of the revenues of the same group over time is problematic because revenue is likely to be "polluted" by monthly variations in demand. It is thus not possible to identify the treatment effect by comparing how a group behaved in treatment period 1 relative to treatment period 2.

¹⁴ It is even possible that our measure of the elasticity of labor supply with regard to a temporary wage increase underestimates the true elasticity because we use revenues per messenger as a proxy for labor supply per messenger. If wages w affect effort e and effort affects revenue r the elasticity of e with respect to w, which we denote by ε_{ew} , is given by $\varepsilon_{rw}/\varepsilon_{re}$ where ε_{rw} is the elasticity of r with respect to w (which is observable to us) and ε_{re} is the elasticity of r with respect to e (which is not observable to us). Thus, our measure ε_{rw} implicitly assumes that the elasticity ε_{re} is equal to one. If ε_{re} is less than one our measure even underestimates the true labor supply elasticity. ε_{re} is less than one if the production function r = f(e) is strictly concave and f(0) = 0 holds.

B. The Impact of the Wage Increase on Shifts worked

After we documented the strong impact of the wage increase on total labor supply, the natural question is whether both the number of shifts and the effort per shift increased. In this section we examine, therefore, the impact of the wage increase on the number of shifts worked while in the next section we have a closer look at effort per shift. Figure 2a and 2b provide a first indication of how the wage increase affected shifts. Figure 2a shows the absolute number of shifts per worker in group A and B during the four-week period prior to the experiment and the two treatment periods. Figure 2b controls for person specific effects by showing the average deviation of the number of shifts from the person specific means.

Figure 2a shows that group A worked roughly 12 shifts in the pre-experimental period and group B worked roughly 11 shifts (see also Table I for the precise numbers). However, the standard errors are very large due to large differences between the workers, suggesting insignificant differences across groups. If we control for person specific effects (see Figure 2b) we find that the average deviation from person specific means is very small in both groups and close to zero during the pre-experimental period. Table IV, which shows the concrete numbers relating to Figure 2b, indicates that in group A the deviation from person specific means is 0.22 (with a standard error of 1.29) and in group B it is -0.35 (with a standard error of 0.98). Thus, there are almost no differences in shifts across groups before the experiment.

During the first treatment period, however, the messengers in group A, who are paid the high wage, worked almost 4 shifts more than did the messengers in group B. Likewise, in the second treatment period the messengers in group B, who receive now the high wage, work substantially more shifts than the messengers in group B. Moreover, if we control for person specific effects (see Figure 2b and Table IV), the standard errors become very small, suggesting that the differences across groups are significant.

Insert Figures 2a and 2b about here

To test more rigorously for significant differences, we performed regressions (5) – (8) in Table III. The independent variable in these regressions is s_{it} , the number of shifts that messenger

i worked during the four week period *t*. The right hand side of these regressions is the same as in equation (8), i.e., we included a treatment dummy, individual fixed effects and time dummies for treatment periods 1 and 2. Regression (5) estimates the impact of the treatment by using only data from group A and group B. It shows a large and highly significant treatment effect; the treated group works on average four shifts more than the control group. Regression (6) uses data from all messengers at Veloblitz; the treatment dummy thus compares the treated with the whole group of untreated messengers at Veloblitz. This regression basically replicates the results of regression (5). In regression (7), we use data from all messengers at Veloblitz and at Flash. In addition, we include a dummy variable that takes on a value of 1 if a messenger belongs to the whole nontreated group at Veloblitz (which comprises the experimental control group and the messengers who did not participate in the experiment). As in regression (3), this dummy measures whether the experiment had an effect on the whole non-treated group at Veloblitz by comparing this group with Flash messengers. The coefficient of this dummy is highly insignificant, suggesting that the experiment had no effect on the non-treated group at Veloblitz. The treatment dummy in regression (7) compares the treated group with the Flash messengers and again indicates a significant treatment effect of similar size as in the previous regressions. Finally, regression (8) compares the treated group to all untreated messengers at Veloblitz and Flash; again the treatment effect is roughly 4 shifts per treatment period and significant.

In summary, Figure 2a and 2b, Table IV, and regression (5) - (8) indicate a clear positive treatment effect of the wage increase on shifts. On average, workers supplied about four shifts more if they receive a high commission rate. Since the average number of shifts worked during the two treatment periods is 11.925 and 10.64, respectively, the wage elasticity of shifts is between (4/11.925)/0.25 = 1.34 and (4/10.64)/0.25 = 1.50. Thus, the shift choices are even more responsive to the wage increase than total revenue per messenger. By definition, the wage elasticity of total revenue is equal to the elasticity of shifts plus the elasticity of total revenue per shift. Therefore, the higher wage elasticity of shifts compared to the elasticity of total revenues is a first indication that the elasticity of effort per shift is negative.

C. The Impact of the Wage Increase on Effort per Shift

When examining the revenue per shift, it is useful to restrict attention to behavior during fixed shifts. Recall that the management at Veloblitz did not allow workers to change their fixed shifts after the announcement of the experiment and during the experiment. The increase in the supply of shifts is fully borne by the variable shifts. Therefore, our experiment could not induce any kind of selection effect with regard to the fixed shifts and the revenue change during the *fixed* shifts identifies the impact of the treatment on effort per shift.¹⁵

In Figure 3a, we show the log of revenue per shift in group A and B during the four weeks prior to the experiment and in the two treatment periods. We control for person effects in Figure 3b by showing the deviation from person-specific means. If we control for person-specific effects, we find that both groups generated almost identical revenues per shift during the four weeks prior to the experiment. However, group B, which receives the *lower* wage, generates roughly 5 percent more revenue per shift than group A during the first treatment period. Likewise, in the second treatment period, group A, which receives now the *lower* wage, exhibits roughly a 6 percent higher revenue per shift than group B. Thus, Figures 3 suggests that the wage increase caused a *reduction* in revenue per shift.

Insert Figure 3a and 3b about here

The impression conveyed by Figures 3 is further supported by the two regressions presented in Table V, which are based on observations from group A and B during fixed shifts. The dependent variable is log revenue of messenger i at day t. We include a treatment dummy in both regressions that takes on a value of 1 if messenger i at day t is in the treatment group and we further control for daily fixed effects and i's tenure. Daily fixed effects are important because of demand variations across days; tenure is important because experienced messengers usually have higher productivity. We do not control for individual fixed effects in regression (1), but for a messenger's gender. This regression shows that the wage increase leads to a reduction in revenue per shift by roughly 6 percent. We control for individual fixed effects in regression (2). The

¹⁵ We should, however, mention that the results remain the same when we examine revenue per shift over *all* (fixed and variable) shifts.

treatment effect in this regression is again significantly negative and indicates a reduction in revenues by roughly 6 percent.

Thus, the temporary wage increase indeed reduced revenue per shift. The implied wage elasticity of revenue per shift is -0.06/0.25 = -0.24, which is consistent with our neoclassical model with preference spillovers across periods and the target income model based on loss aversion. It is also worthwhile to point out that this estimate neatly fills the gap between the elasticity of total revenue and the elasticity of shifts. The intermediate value (between the lower and the upper bound) of the elasticity of total revenue is 1.18; the intermediate value for the elasticity of shifts is 1.42. Thus, according to this difference the elasticity of effort per shift should be -0.24. Our estimates in Table V precisely match this value.

D. Does Loss Aversion Explain the Negative Impact on Effort per Shift?

We provide additional evidence in this section that helps us understand the forces behind the negative impact of the wage increase on effort per shift. Our strategy is to measure individual level loss aversion and to examine whether these measures have predictive value with regard to individuals' response of effort per shift. In other words, we ask the question whether the loss averse messengers drive the negative effect of the wage increase on effort per shift or whether the messengers who are not loss averse drive this effect? If mainly the loss averse messengers show a negative effort response, the loss aversion model is supported; if the negative effect on effort is not related to individual's loss aversion, the neoclassical model provides the more plausible explanation.

Loss aversion and reference dependent behavior have implications in a variety of domains. Loss averse choices have been documented, in particular, in the realm of decision-making under uncertainty (Kahneman and Tversky 1979). Therefore, we measured the messengers' loss aversion by observing choices under uncertainty in an experiment that took place eight months after the experimental wage increase. As part of this study, we presented the messengers with the opportunity to participate in the following two lotteries:

Lottery A: Win CHF 8 with probability 1/2, lose CHF 5 with probability 1/2. If subjects reject lottery A they receive CHF 0.

Subjects could participate in both lotteries, or only in one lottery, or they could reject both lotteries.

The above lotteries enable us to construct individual measures of loss aversion. In particular, the observed behavior in these lotteries enables us to classify subjects with regard to their degree of loss aversion γ . If subjects' preferences are given by (7), subjects who reject lottery A have a higher level of γ than subjects who accept lottery A and subjects who reject lottery A and B have a higher level of γ than subjects who reject only lottery A. In addition, if subjects' loss aversion is consistent across the two lotteries, then any individual who rejects lottery B should also reject lottery A because a rejection of lottery B implies a higher level of loss aversion than a rejection of only lottery A. We derive these implications of (7) explicitly in appendix C.

Among the 42 messengers who belong either to group A or B, 19 messengers rejected both lotteries, 8 messengers rejected only lottery A, 1 messenger rejected only lottery B and 14 messengers accepted both lotteries. Thus, with the exception of the one messenger who rejects only lottery B, all messengers who rejected lottery B also rejected lottery A. These results can be neatly captured by a simple loss averse utility function that obeys equation (7).¹⁶

In principle, one might think that the rejection of A and/or B is also compatible with risk aversion arising from diminishing marginal utility of lifetime income. Rabin's calibration theorem (Rabin 2000) rules out this interpretation, however. Rabin showed that a theory of risk averse behavior based on the assumption of diminishing marginal utility of *life-time* income implies that people essentially *must be* risk neutral for low stake gambles like our lotteries. Intuitively, this follows from the fact that risk averse behavior for low stake gambles implies ridiculously high levels of risk aversion for slightly higher, but still moderate, stake levels. Yet, such unreasonably high levels of risk aversion can be safely ruled out. For example, we show in appendix D that if one assumes that the rejection of lottery A is driven by diminishing marginal utility of life time income, then the subject will also reject a lottery where one can lose \$32 with

¹⁶ These results are qualitatively similar to the results obtained in a many other studies (e.g., Read, Loewenstein, and Rabin, 1999; Cubbit, Starmer and Sudgen, 1998; Hogarth and Einhorn, 1992; Keren and Wagenaar, 1987).

probability ¹/₂ and win *any* positive prize with probability ¹/₂. Thus, there is no finite prize that induces this subject to accept a 50 percent chance of loosing \$32. Similar results are implied by a rejection of lottery B.

We illustrate the behavior of messengers with and without loss averse preferences in Figure 4. The figure controls for person specific effects by comparing individual log revenues to the mean of the individual's log revenues. We show that the messengers who did not display loss averse preferences do not change their effort per shift across the treatment and the control period. However, the messengers who displayed loss aversion in the lottery choices exhibit a lower effort per shift in the treatment period compared to the control period. This pattern suggests that the negative effect of wages on effort per shift may only be driven by the loss averse messengers.

Insert Figure 4 here

To examine this possibility in more depth, we conducted the regressions in Table VI. In these regressions, log daily revenue of messenger *i* at day *t* is again the dependent variable and we control for the messengers' tenure and for daily fixed effects in all four regressions. In the first two regressions, we generate a loss aversion dummy *L* that is based on the rejection of lottery A. If a messenger rejects this lottery L = 1, if lottery A is accepted L = 0. In regression (1) and (2) we estimate the treatment effect for the loss averse messengers (by interacting the treatment dummy with *L*) and for the messengers who did not exhibit loss aversion (by interacting the treatment dummy with (1-*L*)). Regression (1), which does not control for individual fixed effects, shows that loss averse messengers generated a roughly 10 percent lower revenue per shift when they received the high wage. In contrast, the treatment effect is much lower and insignificant for the messengers without loss aversion. Regression (2), which controls for individual fixed effects, shows the same pattern. There is no significant decrease in revenue per shift for messengers without loss aversion, whereas the messengers with loss aversion exhibit a significant 10 percent reduction in revenue per shift during the treatment period.

Regressions (3) and (4) provide a further robustness check for these results. We use a finer scale to indicate a messenger's loss aversion in these regressions. Here, we capture the absence of loss aversion, which is indicated by the acceptance of both lotteries, by L' = 0. If a messenger

rejects only one of the lotteries we assign L' = 1 and if both lotteries are rejected we assign L' = 2 to this messenger. The variable "treatment dummy × loss averse" is now defined as the interaction between the treatment dummy and L'. Thus, the interaction term measures how the *degree* of loss aversion affects the messengers' effort responses, whereas the treatment dummy alone measures the effort response of those who did not display loss aversion.¹⁷ Both regression (3) and (4) indicate that the messengers without loss aversion did not show a significant effort reduction in response to the wage increase. In contrast, the interaction term is relatively large and significant; an increase in L' by one integer unit decreases revenue per shift by roughly 8 percentage points. Thus, messengers who rejected both lotteries generated a 16 percentage point lower revenue per shift when they received the high wage. This result suggests that the negative impact of the wage increase on revenue per shift is associated with the messengers' degree of loss aversion H.C.

V. Summary

This paper reports the results of a randomized field experiment examining how workers, who can freely choose their working time and their effort during working time, respond to a fully anticipated temporary wage increase. We find a strong positive impact of the wage increase on total labor supply during the two four-week periods in which the experiment took place. The associated intertemporal elasticity of substitution is between 1.12 and 1.25. The large increase in total labor supply is exclusively driven by the increase in the number of shifts worked. On average, messengers increase their working time during the four weeks in which they receive a higher wage by four shifts (20 hours), which implies a wage elasticity of shifts between 1.34 and 1.50. This is a considerably larger elasticity than what has previously been found on the basis of daily labor supply data (Camerer et al. 1997, Chou 2000, Oettinger 1999). We also find that the wage increase a decrease in revenue (effort) per shift by roughly 6 percent. However, the increase in the number of shifts dominates the negative impact on effort per shift by a large margin such that overall labor supply strongly increases.

¹⁷ Thus, in regressions (1) and (2), the variable "treatment dummy × not loss averse" is constructed as "treatment dummy × (1 - L) whereas the variable "treatment dummy × not loss averse" is given by the treatment dummy alone in regressions (3) and (4). But the variable "treatment dummy × not loss averse" measures the effort response of the messengers who did not display loss aversion in the lotteries in all four regressions.

The standard neoclassical model with separable intertemporal utility is not consistent with the evidence because this model predicts that both the number of shifts and the effort per shift increase in response to the wage increase. However, we show that a neoclassical model with preference spillovers across periods as well as a target income model with loss averse preferences is consistent with the observed decrease in effort per shift. In order to discriminate between these two models, we measured the messengers' loss aversion at the individual level in the domain of choices under uncertainty. We use these measures to examine whether the negative impact of the wage increase on effort per shift is mediated by the degree to which messengers' are loss averse. We find that the degree of loss aversion is indeed related to the response of effort per shift: higher degrees of loss aversion are associated with a stronger negative impact of the wage increase on effort per shift and workers who do not display loss aversion in choices under uncertainty also do not show a significant effort reduction. Thus, it seems that loss aversion drives the negative effect of wages on effort.

We believe that these results contribute to a deeper understanding of the behavioral foundations of labor supply. Our results certainly do not rule out a role for "neoclassical" preferences in labor supply decisions. One third of the workers in our sample did not exhibit loss aversion and the large intertemporal substitution effects on overall labor supply and the supply of shifts document the power of behavioral forces that have always been emphasized in the standard life cycle model. Our results also contrast sharply with the small and insignificant substitution effects that have been found in many previous studies. Therefore, the small effects in these studies may reflect the constraints workers face in their labor supply decisions and – in view of our results – may be less likely due to workers' unwillingness to substitute labor *hours* over time. However, our results on the behavioral sources of the negative wage elasticity of effort per shift also suggest that disregarding reference dependent preferences in *effort* decisions is not wise.

References

- Altonji, Joseph G. (1986), Intertemporal Substitution in Labor Supply: Evidence from Micro Data, Journal of Political Economy, 94(3), Part 2, S176-S215.
- Blundell, Richard (1994), Evaluating Structural Microeconometric Models of Labor Supply. In: Christopher Sims (Ed.), Advances in Econometrics: Sixth World Congress of the Econometric Society, Vol. II, 3 – 48.
- Blundell, Richard; MaCurdy, Thomas E. (1999), Labor Supply: A Review of Alternative Approaches, in Ashenfelter, Orley and David Card (eds.): Handbook of Labor Economics vol. 5.
- Browning, Martin, Angus Deaton and Margaret Irish (1985), A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle, Econometrica 53(3), 503-543.
- Camerer, Colin, Linda Babcock, George Loewenstein and Richard Thaler (1997), Labor Supply of New York City Cabdrivers: One Day at a Time, Quarterly Journal of Economics; 112(2), 407-41.
- Card, David (1994), Intertemporal Labor Supply: An Assessment. In: Christopher Sims (Ed.), Advances in Econometrics: Sixth World Congress of the Econometric Society, Vol. II, 49 – 78.
- Chou, Yuan K. (2000), Testing Alternative Models of Labor Supply: Evidence from Taxi Drivers in Singapore, Research Paper No. 768, Dept. of Economics, University of Melbourne.
- Cubbit, Robin, Christ Starmer and Robert Sugden (1998), On the Validity of the Random Lottery Incentive System, Experimental Economics 1(2), 115 131.
- Dickens, William T. and Shelley J. Lundberg (1993), Hours Restriction and Labor Supply, International Economic Review 34, 169-192.
- Farber, Henry S. (2004), Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers, Working Paper No. 497, Princeton University, Industrial Relations Section.
- Farber, Henry S. (2005), Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers, Journal of Political Economy 113, 46-82.
- Gaechter, Simon, Herrmann, Andreas and Eric J. Johnson (2005), Individual-level Loss Aversion in Risky and Riskless Choice. Working Paper, University of Nottingham.
- Goette, Lorenz and David Huffman (2003), Reference Dependent Preferences and the Allocation of Effort over Time: Evidence from Natural Experiments, Working Paper, Institute for Empirical Research in Economics, University of Zurich.
- Heath, Chip, Larrick, Richard and George Wu (1999). Goals as Reference Points, Cognitive Psychology 38, 79-109.
- Hogarth, Robin and Hillel Einhorn (1990), Venture Theory: A Model of Decision Weights, Management Science 36, 780 803.
- Kahn, Shulamit and Kevin Lang (1991), The Effect of Hours Constraints on Labor Supply Estimates, Review of Economics and Statistics 73, 605-611.

Kahneman, Daniel and Amos Tversky (1979), Prospect Theory, Econometrica 47, 263-291.

- Kahneman, Daniel and Amos Tversky (2000) Eds., Choices, Values, and Frames. Cambridge University Press.
- Keren, Gideon and Willem Wagenaar (1987), Violation of Utility Theory in unique and repeated Gambles, Journal of Experimental Psychology: Learning, Memory, & Cognition. Vol 13(3), 387-391.
- King, Robert G., Rebelo, Sergio (1999), Resuscitating Real Business Cycles. in: Handbook of Macroeconomics.
- Mankiw, Gregory, Julio Rotemberg and Lawrence Summers (1985), Intertemporal Substitution in Macroeconomics, Quarterly Journal of Economics, 100 February, 225 251.
- Oettinger, Gerald S. (1999), An Empirical Analysis of the Daily Labor Supply of Stadium Vendors, Journal of Political Economy; 107(2), 360-92.
- Pencavel, John (1986), Labor Supply of Men: A Survey. in: Ashenfelter, Orley and Richard Layard (Eds.), Handbook of Labor Economics, Amsterdam: North Holland.
- Rabin, Matthew (2000), Risk Aversion and Expected Utility Theory: A Calibration Theorem, Econometrica; 68(5), 1281-92.
- Read, Daniel; Loewenstein, George; Rabin, Matthew (1999), Choice Bracketing, Journal of Risk and Uncertainty; 19(1 3), 171-97.
- Romer, David (1996), Advanced Macroeconomics, Upper Saddle River: McMillan.

Appendix A

In this appendix, we derive the quasi linear objective function in equation (3) of the paper from the underlying intertemporal maximization problem. The intertemporal optimization problem is

max
$$U = \sum \delta^{t} u(c_{t}, e_{t}; x_{t})$$
 subject to $\sum (\hat{w}_{t}e_{t} + y_{t} - \hat{p}_{t}c_{t})(1+r)^{-1} = 0$

where *u* is strictly concave and twice differentiable in *e* and *c*, *e* is labor supply in period *t*, *c* is consumption in period *t*, *x* is a taste shift variable to allow for periods without work, \hat{w}_t is the wage, \hat{p}_t is the price of consumption goods, δ is the discount rate, and *r* is the interest rate. We assume that there are no liquidity constraints, and that the path of wages, prices, and the taste shifter are known, and that the interest rate is constant.

The first order conditions to this problem are

$$\delta^t u_c(c_t, e_t, x_t) = (1+r)^{-t} \lambda \hat{p}_t$$
$$-\delta^t u_e(c_t, e_t, x_t) = (1+r)^{-t} \lambda \hat{w}_t$$

where λ is the Langrange multiplier on the life-time budget constraint. Thus, it can be interpreted as the marginal utility of lifetime wealth. Define the discounted price as $p_t = (1+r)^{-t} \delta^{-t} \hat{p}_t$ and the discounted wage w_t analogously. The first order conditions then have the easily interpretable form

$$u_c(c_t, e_t, x_t) = \lambda p_t \tag{A1}$$

$$-u_e(c_t, e_t; x_t) = \lambda w_t \tag{A2}$$

Equation (A1) implies that, at every date *t*, the individual equates the marginal utility of consumption to the marginal utility of *lifetime* income λ times the discounted price of the consumption good. Similarly, when choosing how hard to work, the individual chooses effort such that the marginal disutility of effort is equal to the marginal utility of lifetime income times the discounted wage per unit of effort w_t . The model also allows for non-participation. If $-u_e(c_t, 0, x_t) < \lambda w_t$ it is optimal to choose e = 0.

It is possible to represent within-period preferences in terms of a static objective function. This is essentially a reformulation of the results in Browing, Deaton, and Irish (1985). Consider equation (A1) again. Since u(.) is strictly concave, u_c is strictly decreasing in c. Thus, (A1) can be solved for c_t

$$c_t = u_c^{-1}(\lambda p_t, e_t, x_t) \tag{A3}$$

Substitute this into (A2) to obtain

$$-u_e(u_c^{-1}(\lambda p_t, e_t, x_t), e_t, x_t) = \lambda w_t$$
(A4)

Now consider the static one-period objective function

$$v(e_t) = \lambda w_t e_t - g(e_t, \lambda p_t, x_t)$$
(A5)

where λ is the lifetime marginal utility of income along the optimal path. Next we show that maximizing this static objective function is equivalent to solving the intertemporal maximization problem, that g() is convex and can be interpreted as the monetary equivalent of the disutility of effort. To see that this, define

$$g(e_t, \lambda p_t, x_t) = -\int_0^{e_t} u_e(u_c^{-1}(\lambda p_t, e_t, x_t), e_t, x_t)dx$$
(A6)

From the construction of g() in (A6), it is obvious that the first order condition (FOC) that results from the static one-period objective function is equivalent to the FOC (A4). To show that g() is convex in e, we need to show that the second derivative w.r.t. e is positive. We proceed in two steps: First, consider how the individual adjusts consumption to a small perturbation in labor supply along the optimal path, i.e., λ remains constant. Differentiation of (A1) yields:

$$\frac{dc_t}{de_t} = -\frac{u_{ce}}{u_{cc}}$$

Now, take the second derivative of g() to obtain

$$g_{ee}(e_t, \lambda p_t, x_t) = -u_{ee} - u_{ce} \frac{dc_t}{de_t} = -u_{ee} + \frac{u_{ce}^2}{u_{cc}} = \frac{-1}{u_{cc}} \left(u_{cc} u_{ee} - u_{ce}^2 \right)$$

To determine the signs of the terms, observe that the conditions for concavity of u() are $u_{cc} < 0$, $u_{ee} < 0$ and $u_{cc}u_{ee} - u_{ce}^2 > 0$. But this establishes the convexity of g(), as claimed. Thus, in the canonical life-cycle model, a rational, forward looking individual behaves as if she maximized the one-period objective function (A5).

Appendix B

In this appendix, we provide a specific example that shows how non-separable time preferences can induce workers to increase the number of shifts but decrease the effort per shift in response to a wage increase. We consider a two-period model in which the workers one-period objective function is given by

$$v(e_t, e_{t-1}) = \lambda w e_t - g(e_t, e_{t-1})$$

We assume that if a worker does not work during a period she has a utility from leisure time of L_0 and that the effort cost function g() is given by

$$g(e_t, e_{t-1}) = e_t(1 + \alpha e_{t-1}) + 0.5ge_t^2$$

If we ignore discounting and set $e_0 = 0$, total utility is given by

$$U = v(e_1, 0) + v(e_2, e_1)$$
.

(a) If an individual works only one period, the first order condition for effort in this period is

$$\frac{\partial U}{\partial e_{t}} = \lambda w - 1 - g e_{t} = 0 \Leftrightarrow e_{t}^{*}(w) = \frac{\lambda w - 1}{g}$$

Substituting this into the utility function, we get the overall utility of working one shift

$$U(\text{one shift}) = \frac{\left(\lambda w - 1\right)^2}{2g} + L_0.$$

(b) If an individual works two shifts, the first order conditions for effort in the two periods are given by

$$\frac{\partial U}{\partial e_1} = \lambda w - 1 - ge_1 - \alpha e_2 = 0$$
$$\frac{\partial U}{\partial e_2} = \lambda w - 1 - ge_2 - \alpha e_1 = 0$$

The two first order conditions imply $e_1 = e_2$. Therefore,

$$e_1^{**}(w) = e_2^{**}(w) = \frac{\lambda w - 1}{g + \alpha}.$$

Substituting this into the objective function, we get

$$U(\text{two shifts}) = \frac{\left(\lambda w - 1\right)^2}{g + \alpha}$$

We can now examine the implications of this model for the number of shifts worked and effort exerted on a shift as a function of the wage *w*.

(i) Shifts: A rational individual works two shifts if U(two shifts) > U(one shift). This implies

$$\frac{\left(\lambda w-1\right)^2}{g+\alpha} > \frac{\left(\lambda w-1\right)^2}{2g} + L_0 \tag{B1}$$

Notice that, in this model, if $\alpha > g$, it is never optimal to work two shifts. The condition $\alpha > g$ has a straightforward interpretation: In this case, yesterday's effort raises today's marginal costs of effort by more than today's effort raises today's marginal costs of effort. Simplifying this inequality, we get

$$\lambda w - 1 > \sqrt{L_0 \frac{2g(g + \alpha)}{g - \alpha}}.$$
(B2)

Denote the wage that satisfies (B2) with equality by w'. As intuition suggests, the individual's participation is increasing in the wage: If w is large enough such that (B2) is satisfied, she will work two shifts.

(*ii*) *Effort*: To examine how effort responds to a change in wages, we choose two wage levels $w_L < w_H$ and set $w_L = w'$, i.e. the low wage is equal to the highest wage at which it is still optimal to work only one shift. If the wage is low, the individual works one shift, and effort is equal to

$$e_1^*(w_L) = \frac{\lambda w - 1}{g} = \frac{\sqrt{L_0 \frac{2g(g + \alpha)}{g - \alpha}}}{g} = \sqrt{L_0 \frac{2(g + \alpha)}{g(g - \alpha)}}$$

Effort on the high wage is equal to

$$e_1^{**}(w_H) = e_2^{**}(w_H) = \frac{\lambda w_H - 1}{g + \alpha}$$
.

(*iii*) The response to a change from w_L to w_H : In this example, $e_1^*(w_L)$ exceeds $e_1^{**}(w_H) = e_2^{**}(w_H)$ if $w_H \in \left(w_L, \frac{g+\alpha}{g}w_L - \frac{\alpha}{\lambda g}\right]$. Thus, changing the wage from w_L to w_H may

decrease effort per shift if the wage increase is not too high. Notice also that the effect crucially

depends on L_0 , the value of leisure. If $L_0 = 0$, the effect cannot occur, because the wage cancels from the participation condition (A7). Then the individual always works the same number of shifts, irrespective of the wage, and effort responds positively to the wage, *irrespective of the strength of the intertemporal spillover* α . By continuity, this also holds for some $L_0 > 0$. Thus, in our example intertemporal spillovers alone can produce the described response of shifts and effort to the wage only if the value of leisure is large enough.

Appendix C

In this appendix, we derive the conditions under which a loss averse individual whose preferences obey (7) in the text will reject lotteries A and B. For the purpose of lottery choices the disutility of effort does not matter so that we can simplify preferences to

$$v(x-r) = \begin{cases} \lambda(x-r) & \text{if } x \ge r \\ \gamma \lambda(x-r) & \text{if } x < r \end{cases}$$

where x is the lottery outcome, and r is the reference point. We take the reference point to be the status quo. The individual will reject gamble A if

$$0.5v(-5) + 0.5v(8) \le v(0)$$

which simplifies to

$$0.5(-5\gamma\lambda) + 0.5(8)\lambda \le 0$$

This condition is satisfied if

$$\gamma \geq \frac{8}{5}.$$

The individual will reject gamble B if

$$\frac{1}{64}v(-30) + \frac{6}{64}v(-17) + \frac{15}{64}v(-4) + \frac{20}{64}v(9) + \frac{15}{64}v(22) + \frac{6}{64}v(35) + \frac{1}{64}v(48) \le v(0)$$

Plugging in our functional form and simplifying, we find that the individual will reject the gamble if

$$\gamma \ge \frac{793}{192}$$

As claimed in the text, the degree of loss aversion required to reject gamble B is greater than the degree of loss aversion needed to reject A.

Appendix D

In this appendix we prove the following: If an individual has a concave utility function u() and rejects a coin flip where she can either win CHF 8 or lose CHF 5 for all wealth levels $[m,\infty)$, then she will reject *any* coin flip in which she could lose CHF 32 no matter how large the positive prize that is associated with the coin flip.

Proof: We proceed in four steps

(i) We adopt the convention that, if indifferent, the individual rejects the coin flip. Rejecting the coin flip implies

$$0.5u(m+8) + 0.5u(m-5) \le u(m)$$
$$\Leftrightarrow u(m+8) - u(m) \le u(m) - u(m-5).$$

It follows from concavity that $8[u(m+8)-u(m+7)] \le u(m+8)-u(m)$ and.

Define MU(x) = u(x) - u(x-1) as the marginal utility of the xth dollar. Putting the last three inequalities together, we obtain

$$MU(m+8) \le \frac{5}{8}MU(m-5)$$
 (D1)

and, because of the premise, it is true that $MU(x+12) \le \frac{5}{8}MU(x)$ for all x > m - 4.

(ii) We now derive an upper bound on $u(\infty)$. The concavity of u() implies

$$u(m+12) \le u(m) + 12MU(m).$$

Using the same logic,

$$u(m+24) \le u(m) + 12MU(m) + 12MU(m+12)$$
$$\le u(m) + 12MU(m) \left[1 + \frac{5}{8} \right]$$
$$u(m+36) \le u(m) + 12MU(m) \left[1 + \frac{5}{8} + \left(\frac{5}{8}\right)^2 \right]$$

and so on. Thus, we can develop a geometric series starting from m. Taking the limit, we obtain

$$u(\infty) \le u(m) + 12MU\frac{8}{3} = u(m) + 32MU(m)$$

(iii) Concavity implies $u(m-32) \le u(m) - 32MU(m)$.

(iv) Using the results from step (ii) and (iii), we get an upper bound on the value of a coin flip where the individual would either lose CHF 32 or win an infinite amount:

 $0.5u(m-32)+0.5u(\infty)\leq u(m)$

This implies that the individual would reject the gamble. This concludes the proof.

		Participating Messengers		Non- Dantinitating	Messengers, Elash	
		Group A	Group B	Veloblitz	1 1(451)	
	Mean Revenues	3500.67 (2703.25)	3269.94 (2330.41)	1461.70 (1231.95)	1637.49 (1838.61)	
Period Prior to Experiment	Mean Shifts	12.14 (8.06)	10.95 (7.58)	5.19 (4.45)	6.76 (6.11)	
	Ν	21	19	21	59	
Treatment Period 1	Mean Revenues	4131.33 (2669.21)	3005.75 (2054.20)	844.21 (1189.53)	1408.23 (1664.39)	
	Mean Shifts	14.00 (7.25)	9.85 (6.76)	3.14 (4.63)	6.32 (6.21)	
	Ν	22	20	21	65	
Treatment Period 2	Mean Revenues	2734.03 (2571.58)	3675.57 (2109.19)	851.23 (1150.31)	921.58 (1076.47)	
	Mean Shifts	8.73 (7.61)	12.55 (7.49)	3.29 (4.15)	4.46 (4.74)	
	Ν	22	20	24	72	

TABLE I: DESCRIPTIVE STATISTICS

Notes

a) Standard deviations in parenthesis

b) Group A received the high commission rate in Experimental Period 1, group B in Experimental Period 2.

TABLE II: REVENUES PER FOUR-WEEK PERIOD

		Participating Messengers		Non-	Messengers,	
		Group A	Group B	Parutipaling Messengers, Veloblitz	1 744577	
Four-Week	Mean Revenues	-48.88 (366.61)	-119.91 (302.61)	456.72 (179.92)	305.08 (131.42)	
Period Prior to Experiment	Difference: Group A - Group B	71.03 (475.37)				
Treatment Period 1	Mean Revenues	721.98 (192.90)	-277.95 (240.62)	-160.77 (173.89)	102.85 (105.76)	
	Difference: Group A - Group B	999.93 (308.40)				
Treatment Period 2	Mean Revenues	-675.32 (288.62)	391.87 (250.55)	-258.95 (137.61)	-342.84 (129.50)	
	Difference: Group B - Group A	1067.19 (382.20)				
Notes	a) Standard error of the means in parenthesis b) Same number of observations as in Table I					

AVERAGE DEVIATIONS FROM INDIVIDUAL MEANS

c) Group A received the high commission rate in Experimental Period 1, group B in Experimental Period 2.

TABLE III: MAIN EXPERIMENTAL RESULTS

OLS REGRESSIONS

	Dependent Variable: Revenues per four-week period			Dependent Variable: Shifts per four-week period				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Observations are restricted to	Messengers Participating in Experiment	All Messengers at Veloblitz	All Messengers at Flash and Veloblitz	All Messengers at Flash and Veloblitz	Messengers Participating in Experiment	All Messengers at Veloblitz	All Messengers at Flash and Veloblitz	All Messengers at Flash and Veloblitz
Treatment Dummy	1033.6*** (326.9)	1094.5*** (297.8)	1035.8** (444.7)	1076.2*** (290.6)	3.99*** (1.030)	4.08*** (0.942)	3.44** (1.610)	3.9*** (0.930)
Dummy for Non- treated at Veloblitz			-54.4 (407.4)				-0.772 (1.520)	
Treatment Period 1	-211 (497.3)	-370.6 (334.1)	-264.8 (239.9)	-290 (200.6)	-1.28 (1.720)	-1.57 (1.210)	-0.74 (0.996)	-1.01 (0.781)
Treatment Period 2	-574.7 (545.7)	-656.2 (357.9)	-650.5** (284.9)	-675.9** (238.0)	-2.56 (1.860)	-2.63** (1.260)	-2.19** (1.090)	-2.51** (0.859)
Individual Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R squared	0.74	0.786	0.753	0.753	0.694	0.74	0.695	0.695
Ν	124	190	386	386	124	190	386	386

Notes:

a) Robust standard errors, adjusted for clustering on messenger, are in parentheses

b) ***, **, * indicate significance at the 1, 5, and 10 percent level, respectively.

		Participating Messengers		Non-	Messengers,		
		Group A	Group B	Participating Messengers, Veloblitz	F' <i>l</i> \$\$!)		
Four-Week Period Prior to Experiment	Mean Revenues	0.22 (1.29)	-0.35 (0.98)	1.57 (0.75)	0.98 (0.53)		
	Difference: Group A - Group B	0.57 (1.62)					
Treatment Period 1	Mean Revenues	2.53 (0.65)	-1.18 (0.79)	-0.48 (0.75)	0.52 (0.42)		
	Difference: Group A - Group B	3.71 (1.02)					
Treatment Period 2	Mean Revenues	-2.74 (0.98)	1.52 (0.77)	-0.96 (0.57)	-1.27 (0.45)		
	Difference: Group B - Group A	4.26 (1.24)					
Notes	a) Standard er b) Same numl	ror of the me per of observa	ans in parent ations as in T	thesis Fable I			
	c) Group A received the high commission rate in Experimental Period 1, group B in Experimental Period 2.						

TABLE IV: SHIFTS PER FOUR-WEEK PERIOD

AVERAGE DEVIATIONS FROM INDIVIDUAL MEANS

TABLE V: THE IMPACT OF THE EXPERIMENT ON LOG REVENUES PER SHIFT

DEPENDENT VARIABLE: LOG(REVENUES PER SHIFT) DURING FIXED SHIFTS , OLS REGRESSIONS

	(1)	(2)
Treatment Dummy	-0.0642** (0.030)	-0.0601** (0.030)
Gender (female = 1)	-0.0545 (0.052)	
Log(tenure)	0.105*** (0.016)	0.015 (0.062)
Day Fixed Effects	Yes	Yes
Individual Fixed Effects	No	Yes
R Squared	0.149	0.258
Ν	1137	1137

Notes: a) Observations are taken from group A and group B while working on fixed shifts.

b) Robust standard errors, adjusted for clustering on messenger, are in parentheses

c) ***, **, * indicate significance at the 1, 5, and 10 percent level, respectively.

TABLE VI: DOES LOSS AVERSION EXPLAIN THE REDUCTION IN EFFORT PER SHIFT?

	Loss Aversio	on Measure 1	Loss Aversion Measure 2		
	(1)	(2)	(3)	(4)	
Treatment Effect × not loss averse	-0.0374 (0.034)	-0.0273 (0.033)	-0.0353 (0.035)	-0.0369 (0.033)	
Treatment Effect × loss averse	-0.0983** (0.040)	-0.105** (0.046)	-0.0827*** (0.032)	-0.0755** (0.034)	
Gender (female = 1)	-0.0485 (0.052)		-0.0457 (0.052)		
Log(tenure)	0.104*** (0.016)	0.00152 (0.061)	0.102*** (0.017)	-0.00131 (0.061)	
Day Fixed Effects	Yes	Yes	Yes	Yes	
Individual Fixed Effects	No	Yes	No	Yes	
R Squared	0.14	0.243	0.14	0.243	
N	1137	1137	1137	1137	

DEPENDENT VARIABLE: LOG(REVENUES PER SHIFT) DURING FIXED SHIFTS , OLS REGRESSIONS

Notes: a) Robust standard errors, adjusted for clustering on messenger, are in parentheses.

b) ***, **, * indicate significance at the 1, 5, and 10 percent level, respectively.

c) Loss aversion measure 1 (denoted by L): L=1 if subject rejects lottery A, L=0 if subject accepts lottery A.

d) Loss aversion measure 2 (denoted by L'): L'=2 if subject rejects both lotteries, L'=1 if subject rejects one of the lotteries, L'=0 if subject accepts both lotteries.

Figure 1: Revenue per four week period



(a) Amount of Revenue







(a) Number of Shifts





Figure 3: Log of Daily Revenues on Fixed Shifts

(b) Deviation of Log (Daily Revenues) from





