

How Many Pears Would a Pear Packer Pack if a
Pear Packer Could Pack Pears at Quasi-exogenously
Varying Piece Rates?
An Empirical Investigation of Inter-temporal Labor Supply¹

Tom Chang²

Tal Gross³

August 2011

¹We would like to thank Josh Angrist, David Autor, Harry DeAngelo, Glenn Ellison, Sara Ellison, Kevin Murphy, Nancy Rose and Mark Westerfield for many helpful comments. We also thank seminar participants at Cornell, MIT, RAND, USC Marshall, and Vanderbilt University and conference participants at the All California Labor Conference and the Trans-Pacific Labor Seminar for their valuable comments. We are grateful for financial support from the George and Obie Schultz Fund. Janice Crew provided superb research assistance. Finally, we are extremely grateful to the CEO, personnel department and workers of the firm this paper studies.

²Marshall School of Business, University of Southern California

³Mailman School of Public Health, Columbia University

Abstract

We examine labor supply using a unique dataset collected from a large pear-packing factory. Pear packers face both expected and unexpected shocks to their wages, and we use this to evaluate different models of inter-temporal labor supply. In contrast to most previous research, the data allow us to examine the workers' effort when hours are constrained. In response to unpredictable, positive wage shocks, workers decrease their effort. In response to predictable, positive wage shocks, workers respond differently depending on how frequent the shocks are. Inter-day shocks lead workers to increase effort, while high-frequency, intra-day shocks lead workers to exert less effort. The results support models of reference-dependent preferences, but suggest that workers fail to generate rational targets for high-frequency wage changes.

1 Introduction

Economists have long sought to understand how workers respond to a change in their wage rate. The topic has implications for many areas of economics. In macroeconomics, it is a critical ingredient in models of the business cycle (Lucas and Rapping, 1969). In public finance, it determines, in part, how changes to the income tax will affect tax revenue (Feldstein, 1995). In personnel economics, it determines how pay-for-performance will affect workers (Lazear, 2000; Prendergast, 1999).

Historically, researchers studying this question have faced several empirical challenges. Few data sets exist that contain accurate measures of both labor supply and wages. In addition, the available data typically involve wage changes that are not transitory, but rather correlated with lifetime income. Moreover, the available data typically describe workers who are not free to set their own hours. These challenges have made it difficult to interpret the (generally low) estimates of inter-temporal labor supply (Card, 1991; Blundell and MaCurdy, 1999). In response, a more recent strand of literature examines the behavior of workers in highly flexible occupations, such as workers who have a very high level of discretion in setting their own working hours.

This paper analyzes how workers respond to changes in their wage rate using a dataset of workers at a California firm that packages pears and ships them to retailers. Specifically, we use a combination of worker-day-level payroll data as well as worker-minute-level data from a field collection effort. The workers in this firm regularly face both expected and unexpected shocks to their wages, and we examine the workers' response to these shocks to directly test several models of labor supply.

We exploit high-frequency wage shocks to disentangle the impact of wage changes from the income effect. In contrast to much of the recent literature, we study an environment in which workers cannot choose their own hours, but can adjust their

effort. To our knowledge, this paper is the first to examine wage changes in the field in an hours-constrained environment. Such an environment may be more representative of the general workforce, as few professions allow workers to freely set their own hours. In addition, to the best of our knowledge we are the first to document the importance of the *frequency* of wage changes in determining worker response.

The workers in our sample face three types of wage shocks. The first two wage shocks are driven by state laws regarding overtime. Workers are required to stay for an entire shift, but the length of a shift varies. Once a worker has worked eight hours in a given day or forty hours in a given week, the firm must increase compensation by 50 percent. This overtime pay leads to both predictable and unpredictable shocks to wages. Workers do not know the exact length of a shift on any given day, thus daily overtime represents a partially unexpected shock to wages. But workers also receive overtime pay once they have worked forty hours in a week. Many weeks thus last longer than forty hours, overtime at the end of a week represents a predictable increase in wage.

The third type of wage shock is driven by the way work is organized on the factory floor. As workers package pears for shipping, they rotate across bins filled with pears of different sizes. Workers must pack a different number of pears into each box, depending on the size of the pear at their current bin. But workers are paid the same piece rate for each box packed. As such, when packing large pears, a worker's implicit wage rises, because fewer pears need to be packed per box. Conversely, when assigned to pack small pears, their implicit wage falls. For the sake of fairness, the workers rotate across bins, spending roughly fifteen minutes at each bin. As a result, variation in pear size generates high-frequency, predictable variation in wages throughout the course of a workday.

Our empirical analysis leads to three main findings. First, workers respond to

unexpected overtime by exerting *less* effort. Second, workers respond to expected overtime by exerting *more* effort. Third, when rotating across bins, workers exert *less* effort when facing a higher implicit wage.

These findings paint a complex picture of how workers respond to wage changes. Workers respond to both expected and unexpected overtime in a manner that is consistent with rational-expectations-based daily targets. But workers do not generate rational-expectations-based targets when rotating across packing stations. Taken together, these findings suggest that different theoretical models may be required to model labor supply in different contexts.

The paper is organized as follows. The subsequent section provides a brief review of the literature on labor supply and describes our contribution to that literature. Section 3 describes the pear-packing factory and the wage shocks experienced by the workers. Section 4 reviews the different labor-supply models and their implications for worker effort in this context. Section 5 describes our data and empirical framework. Section 6 presents empirical estimates of how workers react to wage changes driven by overtime, and 7 presents estimates of how workers react to wage changes driven by pear size. Section 8 concludes.

2 Previous Research on Labor Supply

This paper contributes to a growing literature on the importance of reference-dependent preferences in determining labor supply. We focus on workers who do not control their schedules but do control their effort. Previous studies have focused instead on workers in occupations with flexible hours. Camerer et al. (1997), for instance, study the behavior of New York City taxicab drivers. The authors use the average daily wage of other cab drivers as an instrument for each driver's wage and find that drivers

work fewer hours when average wages are higher; a result consistent with reference-dependent preferences. Chou (2000) finds similar evidence among taxi drivers in Singapore.

Several studies have found evidence in favor of a more complicated model of labor supply. Farber (2005, 2008), for instance, finds mixed evidence of reference-dependent preferences among cab drivers. He concludes that drivers may rely on income targets, but that such targets change each day, are unstable, and are imprecisely estimated. Crawford and Meng (2008), in contrast, use Farber's data to estimate a model of reference-dependent preferences based on the work of Köszegi and Rabin (2006). They model drivers' income and hours targets as being set via rational expectations and find parameter estimates that they argue are plausible and precisely estimated. Similarly, Giné et al. (2010) find support for a model of reference-dependent preferences with rational-expectations-based targets among fishermen. Finally, Fehr and Götte (2007) run a field experiment on bicycle messengers and find that higher wages have a positive effect on the number of hours worked, but a negative effect on effort per hour. The authors conclude that workers have reference-dependent preferences, but find that the total amount of work provided is positively related to the wage.

In contrast, other studies have found results consistent with the neoclassical model of labor supply. Oettinger (1999) finds that stadium vendors are more likely to work on days when they can expect a higher average wage, while Paarsch and Shearer (1999) measure a positive relationship between piece rate and productivity among workers at a tree-planting firm. Similarly, Lazear (2000) finds that workers respond to a switch from fixed wages to piece rates by increasing effort.

A common characteristic of these studies is their focus on the extensive margin of labor (hours worked) and not on the intensive margin (effort).¹ In this paper, in

¹An exception is a paper by Dickinson (1999). In a laboratory experiment, the author finds that when hours are constrained, worker effort is positively correlated with wage. When his subjects

contrast, we focus on the intensive margin rather than the extensive margin and we do so with data from the field. Our data allow us to measure worker response to both expected and unexpected shocks to wages. This allows us to test not only if workers have hourly targets, but whether those targets are set via rational expectations.

3 Institutional Background

We examine the behavior of workers who package pears to be shipped to market at one of the largest pear-packing facilities in California. Due to the nature of the work, pear packers at the factory face both expected and unexpected shocks to implicit wages. This section describes how those shocks arise.

Because of their thin skin, pears do not respond well to bulk shipping. As a result, packers must wrap each pear individually in tissue paper and carefully arrange the pears in boxes for shipment. The procedure is labor intensive and must be done by hand. But such packaging allows the firm to ship pears directly to retail outlets and preserve the value of the fruit in transit.

Pears arrive at the factory each morning from farms throughout Northern California. Once washed, the pears pass through a quality control process in which damaged pears are removed. The remaining pears then travel along a conveyor system that sorts the pears according to size into large rotating bins. Following industry standards, the pears are sorted into the following sizes: 60, 70, 80, 90, 100, 110, 120, 135, 150. The size of each pear indicates the number of pears that will fit into a standard, four-fifth-bushel box. For example, 100 “size 100” pears would be packed into a standard box.

Packers individually wrap and then carefully arrange pears in large cardboard

are allowed to set their own hours, however, some respond to higher wages by working longer but exerting less effort.

boxes according to a pre-determined pattern based on pear size. Once a box is finished, the worker places the box on a conveyor belt that takes the box to a station where it is placed onto a pallet for shipping. A random subset of boxes is inspected to ensure that workers are taking sufficient care in packing each box. If a box is found deficient, the worker receives a lower piece rate for all boxes packed that day. Such fines are extremely uncommon in our data, with the median number of fines per worker equal to zero.

Packers are paid the same piece rate for each box, regardless of the number of pears it contains. Workers thus face a higher implicit wage when packing larger pears. For instance, while a worker must exert fifty percent more effort to fill a box with 120 small pears than with 80 large ones, the worker will be paid the same piece rate for both boxes.² For that reason, the firm requires packers to rotate across packing stations, so that each worker spends the same amount of time with each pear size. Typically, workers spend fifteen minutes at each station and then rotate to the next-larger size. When they reach the end of the line, the workers rotate back to the station with the smallest pears. If a worker is part way through a box when it is time to rotate, the worker finishes the box before moving to the next bin. The bins are large enough that two workers can have simultaneous, unrestricted access to a single bin. Standard pear sizes vary from 70 to 150, so the number of pears that must be packed in a standard box varies by more than a factor of two. Thus a worker will experience a change in piece rate at regular intervals (every 15 minutes) as he or she rotates across the bins containing the different sized pears throughout each shift.

From the firm's founding in 1919 until 2003, workers were paid a single piece rate for each standard box packed, irrespective of the per-box pear count. If a worker's piece rate earnings per day implied a wage lower than the California minimum wage,

²The pears weigh very little, approximately 4–6 ounces. Thus the effort required to pack a single pear does not vary by size.

then the worker was paid minimum wage.³ This system was called “piece or hourly.”⁴

After 2003, the firm switched from the “piece *or* hourly” payment scheme to a “piece *and* hourly” scheme. Under this new scheme packers continued to earn a flat piece rate per box packed, but were also paid an hourly wage. The new arrangement was meant to insure packers against unexpected shocks that could reduce their output. Workers under the piece and hourly compensation scheme were occasionally asked to help with tasks other than packing boxes, such as restocking materials and cleaning the factory floor.

Under both regimes, the firm provided packers with overtime pay in accordance with California law. Specifically, packers received overtime wages when they had been working for more than 8 hours in any given day or for more than 40 hours in any given week. Wages increased by 50% when the overtime limits were reached. For example in 2005, overtime increased a worker’s hourly wage from \$6.25 an hour to \$9.37 an hour, and piece rates from 20 cents a box to 30 cents a box. Packers generally work 6 days per week, Monday through Saturday, and are expected to work for an entire shift, regardless of the work day’s duration.

The “just-in-time” nature of pear packing leads to significant daily variation in shift length. Given limited cold storage space, pears at the plant are generally packed within 24 hours of being picked. Figure 1 plots the daily shift lengths over the 2002 season, and demonstrates that the variance in hours is large. This subjects the plant’s daily production to fluctuations both due to supply and demand. Demand from individual retailers can vary in response to idiosyncratic sales conditions, while the supply of pears to be packed on a given day is affected by the productivity of farms currently being harvested.

³In general, the minimum wage constraint applied only to new workers as they learned to pack, and affects only a small fraction of our sample.

⁴After the 2003 season, the firm changed its compensation system and were unable to provide us with accurate record of worker hours after this change.

In summary, pear packers are thus subject to several wage changes. The first is caused by variation in pear size; the packers' implicit wage changes every 15 minutes. Second, packers' implicit wage changes when overtime begins. In particular, overtime leads to two types of wage shocks. Early in the week, a worker will receive overtime pay for any work done after the 8-hour mark. Workers arrive each day with no knowledge of the length of that day's shift.⁵ As such, daily overtime represents a largely unexpected, 50 percent increase in wages. Towards the end of the week, the 40 hour a week overtime constraint will bind. Weekly overtime is highly predictable. For example, if workers have put in a total of 38 hours from Monday through Thursday, workers know when they come in Friday morning that after two hours of work they will start receiving overtime pay. Thus weekly overtime represents a largely expected positive increase in wage.

4 Models of Labor Supply

This section reviews three different models of labor supply and describes each model's implications for the workers in our sample.

4.1 Basic Model

Suppose that an agent can work for up to two periods, and must choose their effort at the start of each period. The worker's expected utility at the start of period 1 is:

$$E[u(e_1, e_2)] = e_1 - \frac{1}{2}e_1^2 + \rho \cdot \left[\omega e_2 - \frac{1}{2}(e_2 + \gamma e_1)^2 \right]. \quad (1)$$

⁵During lunch, some workers may ask the shift supervisor whether the work day will last into overtime. But often the supervisor may not know.

In this utility function, e_t represents the effort exerted in period t and ω represents the ratio of wages in period one and period two. The parameter ρ is the probability that work continues into period 2 and is designed to capture the extent to which overtime is expected.

The parameter $\gamma \in [0, 1]$ captures fatigue, whereby effort exerted in period 1 increases the dis-utility of effort in period 2. Due to fatigue, packers may face an incentive to pack more slowly early in the day in order to be less fatigued later in the day. As described below in section 6, our empirical results suggest that fatigue is not a significant factor for pear packers ($\gamma \approx 0$), but is included here for completeness. Note that if γ is indeed close to zero, then the effort decisions in each period are independent of one another and thus independent of ρ .

Consider the three types of wage shocks experienced by pear packers. First, consider the case of a worker rotating across bins containing different sizes of pears. Each period in equation (1) then represents the 15 minutes spent at each particular bin. In that case, packers expect that period 2 will occur ($\rho \approx 1$), and given the short time frame fatigue should matter little between periods ($\gamma = 0$). Thus ω is equal to the ratio of implicit wages driven by the two pear sizes.

Next, consider the effect of overtime. Here we treat regular hours as period 1 and overtime hours as period 2. Early in the week, workers are uncertain as to whether overtime will occur at the end of the day. As such, we assume that $\rho < 0.50$ and $\omega = 1.5$ for daily overtime. In contrast, weekly overtime is largely expected, and so we assume that $\rho \approx 1$ and, again, $\omega = 1.5$.

4.2 Neoclassical Workers

At the start of period 1, the worker's expected utility is given by equation (1) above and worker utility at the start of period 2 is

$$E[u(e_1, e_2)] = \omega e_2 - \frac{1}{2}(e_2 + \gamma e_1)^2. \quad (2)$$

Neoclassical workers choose effort by maximizing equations (1) and (2). Their optimal effort is thus

$$(e_1^*, e_2^*) = (1 - \rho\gamma, \omega - \gamma e_1^*) = (1 - \rho\gamma, \omega - \gamma(1 - \rho\gamma\omega)). \quad (3)$$

First, consider the case of a packer rotating from one bin to another ($\rho = 1, \gamma = 0$). The optimal effort in each period, (3), implies that $e_2^*/e_1^* = \omega$, and thus effort simply varies directly with wage. Intuitively, since both the returns and costs of effort are independent across time periods, the packer's optimal effort must vary positively with her wage. The neoclassical model thus predicts that packers should work faster when assigned to larger pears.

Second, consider the worker's response to overtime when it is completely unexpected, such that $\rho = 0$ and $\omega = 1.5$. Those parameters imply that $e_2^*/e_1^* = 1.5 - \gamma$. Thus packers will exert more effort during overtime so long as $\gamma < 0.5$. A neoclassical packer's effort will increase with their overtime wage unless the dis-utility of fatigue is so large as to overwhelm the pecuniary incentives.

Finally, notice that e_2^*/e_1^* is increasing in ρ . Thus, the more packers expect overtime to occur, the more their effort will likely increase when overtime begins. Consequently, packers are more likely to increase effort when facing expected overtime at the end of the week than unexpected overtime at the beginning of the week. But in both cases this neoclassical model suggests that packers will decrease their effort

in response to overtime only if fatigue is sufficiently large relative to the increase in wage. For the 50% wage increases associated with overtime, such a decrease in effort would require the effect of fatigue to be extremely large.

4.3 Reference-Dependent Preferences

Models of targeting assume that workers choose a level as a goal, beyond which they become less responsive to changes in their wage. For example, according to models of daily income targeting, workers hold a specific dollar value as a daily target income level. Under this model, workers may exhibit a negative wage elasticity, because a higher wage allows them to reach their income targets with less labor.

Targeting can be incorporated into the model above by adding an additional term to the agent’s utility function, so that equation (1) becomes:

$$v(e_1, e_2; y^R) = u(e_1, e_2) - I\{y_t > y_t^R\} \cdot \lambda \cdot [e_1 + \omega e_2]. \quad (4)$$

Here, y is the value of a target variable, y^R is its realized value, and $\lambda \geq 0$ represents the relative importance of the “gain-loss utility” that captures the behavioral bias brought about by exceeding the target level (Kőszegi and Rabin, 2006). We focus on a single-target model with loss-averse agents (i.e. $\lambda > 0$). Specifically, we assume that workers have daily hours targets, and will therefore be less responsive to wage changes once their daily hours target has been reached.

Following Kőszegi and Rabin (2006), we assume targets are set equal to the agents’ rational expectations: $y^R = E(y)$. Thus changes in the target variable must be unanticipated to generate behavior that differs from the neoclassical model. In terms of our application, this implies that packers only deviate from the neoclassical model when they work longer than expected. When packers perfectly anticipate (or overes-

timate) the length of their shift, $y \leq y^R$, and the gain-loss utility does not affect their effort.

Consider then the case of a packer rotating from one bin to another ($\rho = 1$ and $\gamma = 0$). Workers should be able to perfectly anticipate the move to the next bin. As such, the gain-loss utility is multiplied by zero in equation (4), and the implications of the neoclassical model follow.

Now consider the workers' response to overtime. When workers fully anticipate overtime, their response will be identical to the neoclassical case above. But when overtime is unanticipated, the hours target is exceeded ($y_t > E(y_t)$), and so workers are affected by two countervailing forces. First, workers are induced to exert more effort during overtime, because their wage increases. But when overtime is unexpected, worker dis-utility of effort increases, since they have already met their target hours. Given a sufficiently large λ relative to ω , workers may actually decrease their effort during overtime.

As a simple example, consider the case where a worker expects to work N hours in a day. If a change in wage happens during those N hours, the gain-loss term is multiplied by zero and the worker responds as the neoclassical model would predict. If, instead, workers are required to work for $N + 1$ hours, during that last hour workers experience an additional dis-utility of effort. The response of reference-dependent workers with rational-expectations based hourly targets to overtime will then depend on whether or not the overtime is anticipated.

Sections 6 and 7 describe empirically how workers react to expected and unexpected shocks. In section 8 we summarize which of these two models—the neoclassical framework or reference-dependent preferences—is supported by the empirical results.

5 Data

Our analysis is based on two sources of data. First, the firm's personnel department provided payroll records for the 2001, 2002, and 2003 packing seasons. For each season, the payroll data list each worker's total boxes packed during regular hours, total boxes packed during overtime, regular hours worked, and overtime hours worked.⁶ Turnover among workers is high; roughly 69% of workers in one year did not return the next year.

The payroll data do not indicate the size of the pears packed in each box, because such information is not relevant to the workers' pay. As such, we cannot use the payroll data to examine how workers respond to the high-frequency wage changes associated with rotating from one bin to another.

For that reason, we collected data on output for a small set of workers on the factory floor during the summer of 2006. A research assistant monitored workers along one production line (row of bins) and recorded the time each worker spent packing each box. Over the course of four weeks, the research assistant measured the output of workers for 20 shifts; this generated data on 3,967 boxes packed by 70 packers.⁷

6 The Reaction of Workers to Overtime

This section studies the workers' response to overtime, first in the on-site data and second in the payroll data.

⁶Workers were given unique employee identifiers for each calendar year, but those identifiers are not unique across years. For that reason, we linked workers across years based on their full names.

⁷We had hoped to collect payroll data on the same workers as in the on-site sample. Unfortunately, the factory has closed, and we were unable to procure that data.

6.1 Overtime in the On-Site Data

Using the on-site data, we examine how productivity varies throughout the work day. To do so, we regress the time it takes a worker to pack a single pear (measured in seconds per pear) on indicators for the time of day, worker fixed effects, and date fixed effects. All shifts in the on-site data started at the same time, thus time-of-day is equivalent to hours into shift.

Table 1 presents the results of this regression. Each column of Table 1 includes a different set of fixed effects. For all regressions the first hour of work is the omitted category. The point estimates from column 2 of that regression are plotted in Figure 2. The figure and all four columns of Table 1 clearly show the same striking pattern: the coefficients are small and statistically insignificant until (or just before) 8 hours into the day, after which they become both positive and statistically significant at conventional levels. The figure suggests that workers pack pears at a remarkably constant rate throughout the day, and start slowing down (increasing the time to pack each pear) after eight hours. For instance, at 8 hours and 30 minutes into the day (and hence well into overtime) the packers take an additional 0.19 to 0.52 seconds longer to pack each pear. Such a slowdown corresponds to a drop in effort of approximately 10–20% compared to an hour before overtime begins.

This basic pattern of coefficients is the same across specifications. Packers, however, worked more than 8 hours on only one of the 18 weekday shifts in our on-site sample. As such, these results should be interpreted with caution.⁸ We, instead, focus on the effect of overtime in the payroll data.

⁸We estimate a similar pattern of coefficients if we limit the sample solely to the one day that exceeded 8 hours, though the estimates become less precise.

6.2 Overtime in the Payroll Data

The payroll data does not contain measures of productivity within each shift, but it does contain daily production data for all packers at the plant for the 2001, 2002, and 2003 seasons. In particular, for each worker-day we observe the total regular hours, overtime hours, boxes packed during regular hours, and boxes packed during overtime hours. This data consists of a total of 1,346,770 boxes packed by 191 distinct workers over 275 shifts.

Table 2 presents sample statistics for the payroll data. Overtime composes 10 percent of the sample. But Table 2 makes clear that this fraction varies across the days of the week. Nearly 32 percent of hours worked on Saturday are overtime, whereas roughly 7 percent of hours during the week are overtime.

To explore how workers react to that overtime pay, we rely on the following two regressions:

$$\text{Total Boxes}_{id} = \alpha_0 + \gamma_1 \text{OT Hours}_{id} + \gamma_2 \text{RT Hours}_{id} + \alpha_i + \alpha_d + \nu_{id} \quad (5)$$

and

$$\frac{\text{Total Boxes}_{id}}{\text{Total Hours}_{id}} = \alpha_0 + \beta_1 \frac{\text{OT Hours}_{id}}{\text{Total Hours}_{id}} + \beta_2 \text{Total Hours}_{id} + \alpha_i + \alpha_d + \varepsilon_{id}, \quad (6)$$

where α_i and α_t are individual and date fixed effects, respectively. Equation (5) regresses the total boxes packed each day on the number of regular hours and overtime hours. This specification measures effort during both regular and overtime hours. Equation (6) regresses each worker's daily rate (total boxes divided by total hours) on the worker's total hours and the share of total hours that are overtime. Both regressions are meant to test how effort responds to overtime. The neoclassical model suggests that for equation (5) $\gamma_1 > \gamma_2$ and that for equation (6), $\beta_1 > 0$.

Table 3 presents ordinary least squares (OLS) estimates of these two equations. The first three columns report the estimates of equation (5) for the entire sample, only weekdays (Monday through Thursday), and only weekends (Friday and Saturday). We combine observations measuring output on Friday and Saturday because both days tend to entail a higher probability of overtime.⁹ The first column shows that overtime and regular time hours lead to roughly the same production. But columns two and three present starkly different patterns for weekdays and weekends, respectively. During the week, when overtime is relatively rare and unpredictable, regular time hours lead to roughly 5 boxes on average, versus only 1.8 for overtime hours. This difference is statistically significant at conventional levels (p -value is 3.8%). In contrast, on weekends, when overtime is both common and largely expected, overtime hours are more productive. The remaining columns of Table 3 present estimates of equation (6). The estimates are much less precise, and neither reject the implication of columns 1–3 nor confirm them.

Nevertheless, the ordinary least squares results in Table 3 may be biased. Hours are set by the firm, and the firm may be forced to offer overtime based on omitted variables that are correlated with output. Neither total hours nor overtime are explicitly assigned at random, so we must interpret the results of Table 3 with caution.

As a complementary approach, we pursue an instrumental variables (IV) strategy. Following Angrist (1991), we use time fixed effects—indicator variables for day of week and for week—as IV’s for overtime hours. This strategy assumes that the fixed effects are not correlated with the cost of effort on a particular day of week or week of the year. Instead, the strategy assumes that on particular days or weeks, the firm has more pears to pack or more orders to fill and thus is more likely to demand overtime from its workers.

⁹All results are fully robust to dropping Fridays or classifying Friday as a weekday (occasionally less precise in the latter case).

Note that this assumption is much less restrictive than it would be in a setting where workers could set their own hours. When workers can choose their own hours, this IV strategy would require that the *opportunity cost* of working for an hour on Saturday is the same as working an hour on Wednesday. In contrast, workers in our setting are required to come to work, and so can only adjust their effort level. Therefore our IV strategy requires only that, conditional on already being at work, the act of packing a pear requires the same degree of effort on Saturday as on Wednesday.

Another concern regarding this IV strategy is that overtime may be correlated with fatigue. Late in the day, when overtime begins, workers may be fatigued. Similarly, workers may be especially fatigued during weeks in which they work more than forty hours. We believe, however, that fatigue is likely not a source for bias in the following IV results. First, equation (6) controls for total hours in a given day, thus controlling for the main source of fatigue at the daily level. Second, we examined worker performance for weeks during which workers did not receive any overtime. For those weeks, we did not find any relationship between cumulative hours and packing rates. Finally, according to the on-site data, other than a sharp discontinuity at the 8 hour mark, worker productivity is extremely consistent throughout a shift. This suggests that fatigue cannot affect the daily overtime results unless fatigue itself is discontinuous at 8 hours.

Table 4 presents the first-stage regressions. The four columns present results for the four endogenous variables in equations (5) and (6): share overtime, total hours, hours overtime, and hours regular time. In all equations, the F -statistics associated with a test that the coefficients on all instruments is zero are large with associated p -values of less than 1%.

Note also that the day-of-week coefficients are small and insignificant during weekdays (with the possible exception of Friday), but large and significant for Saturday.

This pattern is consistent with the assumptions above, namely that overtime is difficult to predict during the week, but easier to predict on Saturdays.

Table 5 presents IV estimates of equations (5) and (6). In general, the table exhibits the same pattern as the OLS results in Table 3: during the week, workers pack more slowly during overtime. For instance, the second column reports that an hour of overtime leads to 13.2 boxes packed, while an hour of regular time leads to 27.5 boxes packed. This difference is statistically significant (p -value less than 1%). Column 5 confirms this via equation (6): a ten percentage point increase in the share of a weekday spent overtime leads to 1.4 fewer boxes packed per hour.¹⁰

In stark contrast, columns 3 and 6 suggest the opposite pattern for weekends. On weekends, workers exert more effort during overtime hours than during regular hours.¹¹ That is, while workers decrease their effort in response to unexpected (i.e. weekday) overtime, they increase their effort in response to the expected (i.e. Saturday) overtime. Similar to the OLS regressions, our IV results show that the significant impact of overtime on worker behavior is masked when the sample is not divided into weekday and weekend days.

The IV strategy may be invalid if unobserved productivity shocks are correlated with the instruments. For instance, the IV estimates would be biased if on certain weeks the factory receives a distribution of pears that affects both overtime hours and the packing rate of workers. We cannot rule out this possibility. Nevertheless, both the OLS and IV results, which rely on different identifying assumptions, generate the same result: workers pack more slowly during overtime when it occurs during the

¹⁰Roughly a third of workers from one season work at the firm the next season. The results of Table 5 do not change substantially when we interact the effect of overtime with a control for experience.

¹¹Given the concern that Friday may be more similar to Saturday than the other days of the week (e.g. overtime is more predictable), we have run both the OLS and IV regressions dropping Fridays from our sample. Doing so does not have a significant effect on either the magnitude or precision of the results.

week but pack more quickly during overtime when it occurs on a Saturday.

7 The Reaction of Workers to Pear Size

This section studies how workers change their behavior as they rotate across packing stations. Every fifteen minutes, workers rotate across different bins into which pears of different sizes have been mechanically sorted. Since the number of pears that must be packed into each box varies with the size of the pear while the remuneration does not, a packer's real piece rate varies along with pear size. At each fifteen-minute mark, workers rotate to the next larger pear size until the largest size is reached. If a worker is at the last station when it is time to rotate, she moves back to the station carrying the smallest pear. In general workers experience an increase in their piece rate of 12.5% each time they move down the line and a decrease of 50% when they rotate up from the back to the front of the line.

Table 6 presents sample statistics for the on-site data. Overall, we observe 3,960 boxes being packed by 70 workers. Workers take an average of 2.4 seconds for each pear, with a standard deviation 0.856. The remainder of Table 6 suggests that workers increase their effort as their implicit wage decreases. For instance, workers take an average of 2.6 seconds to pack each of the largest type of pear, but only 2.0 seconds for each of the smallest pears. That difference alone suggests that workers exert more effort for the smaller pears, when the incentive to pack is weakest.

Table 6, however, does not control for variation that is driven solely by worker- and day-specific shocks. To account for worker- and day-specific variation, we regress the packing speed of packer i when packing box b on day t on the size of the pear.

Specifically, we estimate the regression:

$$\begin{aligned} \text{Seconds per Pear}_{ibt} = & \alpha_0 + \beta_{90} \cdot I\{\text{Pear Size 90}\} \\ & + \beta_{100} \cdot I\{\text{Pear Size 100}\} + \beta_{110} \cdot I\{\text{Pear Size 110}\} \\ & + \beta_{120} \cdot I\{\text{Pear Size 120}\} + \alpha_i + \alpha_t + \varepsilon_{ibt}, \end{aligned} \quad (7)$$

where α_i is a worker fixed effect and α_t are date fixed effects.

Table 7 presents estimates of equation (7) with size-80 pears as the omitted category. Size-80 pears are the largest pears and therefore involve the largest implicit piece rate.¹² The estimates of β_i decrease monotonically as the pears become smaller, meaning that workers take *less* time to pack a pear as their effective piece rate decreases. Compared to the smallest pears (120 count), workers take an additional half second (or 17% more time) to pack the largest pears. The third and fourth columns of Table 7 present similar estimates, but include only a single variable representing pear size. Such a model imposes a linear functional form on size. That specification leads to the same conclusion: workers pack smaller pears faster.

We interpret Table 7 as a test of the models presented in section 4. But that interpretation must address three concerns. First, smaller pears may simply require less effort to pack. In that case, workers may speed up as they pack smaller pears, not because their implicit wage changes but because of the change in the difficulty of packing. We believe that this is unlikely to be the case. The difference in sizes between pears is not large. The casual observer would find it difficult to distinguish pears taken from adjacent bins. Individual pears are also quite light (approximately 4 to 6 ounces), so the effort lies essentially in the act of hand-wrapping each pear in tissue paper, an act likely unaffected by pear size. Finally, when asked, the packers

¹²We exclude size 70, 135, and 140 pears as they are rarely observed in the data.

stated that there was very little difference in the effort required to pack different sized pears.

A second concern is that workers may simply be unable to “solve” the effort optimization problem. Packers face a relatively simple optimization problem, and encounter the wage changes frequently, thus this also seems unlikely. Moreover, Table 7 suggests that the workers not only fail to optimize, but instead exert more effort as incentives to do so decrease, a strategy that actively decreases their wage relative to a strategy of constant effort. The magnitude of the point estimates in Table 7 suggest that packers could earn substantially more money simply by providing effort at a constant rate. Specifically, they would earn 7.3% more money each hour if they packed at such a constant rate.

A final concern is that the workers’ effort may be determined, in part, by peer pressure. A growing literature has emerged showing the importance of social networks in determining how workers respond to incentives (Mas and Moretti, 2009; Bandiera et al., 2010). The concern is that packers pressure each other to exert less effort when packing larger pears. But the nature of pear packing makes it unlikely that workers collude in such a fashion. In most settings in which network effects are discussed, the workers either collaborate (Bandiera et al., 2010) or can easily observe co-workers (Mas and Moretti, 2009). In contrast, pear packers work alone at each station, and it is difficult for them to monitor each other. The distance between packers, and the noisiness of the factory floor means that passive monitoring is difficult and the act of packing requires their full attention.

In summary, Table 6 and Table 7 suggest that workers exert more effort as the incentive to do so diminishes. The transition across pear sizes is an expected wage shock; the packers rotate throughout the day at regular intervals. Thus these results suggest behavior that is inconsistent with both neoclassical preferences and the

rational-expectations-based, reference-dependent preferences of section 4.

8 Discussion

The payroll and on-site data suggest two general patterns in how packers respond to changing wages. First, packers respond to overtime pay differently depending on whether the overtime is expected or unexpected. When overtime pay is largely predictable (during weekends), workers respond to overtime by increasing their effort. In contrast, when overtime is unexpected (during weekdays), worker effort is negatively associated with overtime pay. Fatigue could explain the latter finding only if its effect were both large and discontinuous in time (i.e. at the 8-hour mark). While we cannot definitively rule out such a mechanism, we consider it unlikely. In contrast, this pattern is entirely consistent with rational-expectations-based targets. When overtime is expected, workers behave in a manner consistent with the neoclassical model, but when overtime is unexpected, workers exhibit a negative wage elasticity. Since overtime corresponds to a 50% increase in wage, the fact that we find a decrease in effort when overtime is unexpected implies that the “targeting” effect is quite powerful.

The second general pattern we observe is that packers fail to change their effort as they rotate across bins. In fact, we observe packers monotonically *decrease* their effort as the incentive to exert effort increases. That finding is inconsistent with the neoclassical model. And since rotation across bins is routine and predictable, it is also inconsistent with a rational-expectations-based targeting model of labor supply.

One possible explanation for this behavior comes from the workers themselves: hourly targets. When asked about their packing strategy, many workers say that they try to pack a certain number of boxes per hour. If workers follow such a heuristic, then they would increase effort when packing smaller pears in order to maintain a

constant rate. The behavior we observe is consistent with that hypothesis, although it begs the question of why packers do not adjust their targets.

Our results have normative implications for the incentives given to workers at other firms. The results suggest that workers respond to overtime based on their expectations. When overtime is necessary, managers should thus handle workers' expectations with care. If workers are surprised by overtime, they may exert less effort when it arrives.

This paper is part of a growing literature that presents a more nuanced picture of how workers supply their labor. This literature is far from uniform, with many studies suggesting that the neoclassical labor supply model alone does not adequately describe how workers choose how much labor to supply. The results of this paper support that general conclusion and suggest that different theoretical models may be required to model labor supply in different contexts, even for the same worker.

References

- Angrist, J. D. (1991). Grouped-data estimation and testing in simple labor-supply models. *Journal of Econometrics* 47(2-3), 243–266.
- Bandiera, O., I. Barankay, and I. Rasul (2010, April). Social incentives in the workplace. *Review of Economic Studies* 77(2), 417–458.
- Blundell, R. and T. MaCurdy (1999). Labor supply: A review of alternative approaches. *Handbooks In Economics* 5(3), 1559–1696.
- Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997, May). Labor supply of New York City cabdrivers: One day at a time. *The Quarterly Journal of Economics* 112(2), 407–441.
- Card, D. (1991). Intertemporal labor supply: An assessment. Working Paper 3602, National Bureau of Economic Research.
- Chou, Y. K. (2000). Testing alternative models of labor supply. evidence from taxi-drivers in singapore. Working Paper.
- Crawford, V. P. and J. Meng (2008). New York City cabdrivers’ labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. Working Paper.
- Dickinson, D. L. (1999). An experimental examination of labor supply and work intensities. *Journal of Labor Economics* 17(4), 638–670.
- Farber, H. S. (2005). Is tomorrow another day? The labor supply of New York City cabdrivers. *Journal of Political Economy* 113(1), 46–82.
- Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of New York City taxi drivers. *American Economic Review* 98(3), 1069–82.
- Fehr, E. and L. Götte (2007). Do workers work more if wages are high? Evidence from a randomized field experiment. *American Economic Review* 97(1), 298–317.
- Feldstein, M. (1995). The effect of marginal tax rates on taxable income: A panel study of the 1986 Tax Reform Act. *The Journal of Political Economy* 103(3), 551–572.
- Giné, X., M. Martinez-Bravo, and M. Vidal-Fernandez (2010). Intertemporal substitution, weekly target earnings or both? Evidence from daily labor supply of southern indian fishermen. Working Paper.
- Köszegi, B. and M. Rabin (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics* 121(4), 1133–1165.

- Lazear, E. (2000). Performance pay and productivity. *The American Economic Review* 90(5), 1346–1361.
- Lucas, R. E. and L. A. Rapping (1969). Real wages, employment, and inflation. *The Journal of Political Economy* 77(5), 721–754.
- Mas, A. and E. Moretti (2009, March). Peers at work. *American Economic Review* 99(1), 112–145.
- Oettinger, G. S. (1999, April). An empirical analysis of the daily labor supply of stadium vendors. *The Journal of Political Economy* 107(2), 360–392.
- Paarsch, H. J. and B. S. Shearer (1999). The response of worker effort to piece rates: Evidence from the British Columbia tree-planting industry. *The Journal of Human Resources* 34(4), pp. 643–667.
- Prendergast, C. (1999). The provision of incentives in firms. *Journal of Economic Literature* 37(1), 7–63.

Figure I. Hours and Boxes Packed over the 2002 Season





Figure II. Average Packing Speed by Hours Into Shift




Table I: Packing Speeds by Hours into Shift, On-Site Data
 Dependent Variable: Seconds per pear packed

	(1)	(2)	(3)	(4)
1-2 hours	0.047 (0.054)	0.037 (0.049)	0.089 (0.051)	0.036 (0.045)
2-3	0.135 (0.059)	0.098 (0.056)	0.161 (0.052)	0.090 (0.052)
3-4	0.143 (0.064)	0.110 (0.050)	0.171 (0.059)	0.089 (0.047)
4-5	0.130 (0.078)	0.070 (0.065)	0.091 (0.066)	0.017 (0.058)
5-6	0.052 (0.069)	-0.058 (0.061)	0.000 (0.065)	-0.111 (0.057)
6-7	0.095 (0.069)	0.048 (0.058)	0.071 (0.056)	0.013 (0.052)
7-7.25	-0.030 (0.204)	-0.135 (0.176)	0.048 (0.179)	-0.067 (0.171)
7.25-7.5	0.145 (0.124)	0.041 (0.116)	0.152 (0.105)	0.046 (0.105)
7.5-7.75	-0.023 (0.142)	-0.112 (0.125)	0.013 (0.138)	-0.079 (0.126)
7.75-8	0.480 (0.438)	0.299 (0.330)	0.713 (0.428)	0.312 (0.334)
8-8.25	0.298 (0.194)	0.324 (0.099)	0.588 (0.215)	0.316 (0.104)
8.25-8.5	0.186 (0.191)	0.241 (0.042)	0.520 (0.218)	0.257 (0.061)
8.5-8.75	0.017 (0.175)	0.198 (0.081)	0.279 (0.205)	0.190 (0.051)
Worker FE		X		X
Date FE			X	X

R²

Note: $N = 3,908$. The omitted category indicates the first hour of the day. Standard errors in parenthesis allow for auto-correlation between observations based on the same worker. All specifications include a fixed effect for the type of box packed. The data set contains a record for every box packed by a sample of workers, and the time it took them to pack each box. A research assistant collected the data over the course of one month. See text for details.

Table II: Sample Statistics for Payroll Data

	Percent Overtime	Std. Dev. Percent Overtime)	Packing Rate	Hours Overtime	Total Hours
Monday	7.31	9.03	13.17	0.73	8.39
Tuesday	7.48	8.87	13.33	0.75	8.39
Wednesday	7.78	9.34	13.12	0.79	8.30
Thursday	6.60	9.01	13.39	0.67	8.29
Friday	8.15	10.40	14.74	0.83	8.54
Saturday	31.90	39.15	16.96	2.23	6.80
All	10.58	18.37	14.02	0.94	8.19

Table III: The Effect of Overtime on Output Estimated via OLS, Payroll Data

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	Total Boxes Packed			Boxes Per Hour		
Sample:	All	Weekday	Weekend	All	Weekday	Weekend
Hours Overtime	2.997 (1.350)	1.846 (1.367)	3.729 (1.608)			
Hours Regular Time	2.661 (0.930)	4.909 (0.997)	1.195 (1.548)			
Share Overtime				1.748 (0.849)	2.207 (2.039)	1.889 (1.210)
Total Hours				- 0.937 (0.090)	- 0.907 (0.119)	- 1.090 (0.231)
F-statistic to test that OT = RT	0.09	4.38	6.29			
Constant	108.000 (9.714)	79.825 (10.084)	38.738 (15.852)	23.036 (0.899)	21.846 (1.067)	14.325 (1.906)
R ²	0.680	0.667	0.753	0.583	0.579	0.683
N	5,002	4,375	627	5,002	4,375	627

Note: Standard errors in parenthesis allow for auto-correlation between observations based on the same worker. Week and worker fixed effects not shown. The data set is based on payroll records covering the 2001 through 2003 packing seasons. See text for details.

Table IV: First Stage Regressions, Payroll Data

	(1)	(2)	(3)	(4)
Dep.Var.	Share OT	Total Hours	Hours OT	Hours RT
Tuesday	0.001 (0.002)	0.028 (0.037)	0.026 (0.024)	0.002 (0.027)
Wednesday	0.025 (0.003)	- 0.044 (0.057)	0.200 (0.032)	- 0.244 (0.053)
Thursday	0.006 (0.002)	- 0.015 (0.044)	0.040 (0.024)	- 0.055 (0.043)
Friday	0.007 (0.003)	0.201 (0.043)	0.111 (0.040)	0.090 (0.032)
Saturday	0.245 (0.026)	- 1.333 (0.090)	1.583 (0.170)	- 2.916 (0.203)
Week F-Stat	930.7	21,138.0	556.9	30,727.6
Day F-Stat	38.6	68.4	33.0	97.0
R ²	0.464	0.455	0.434	0.454
N	5,002	5,002	5,002	5,002

Note: Standard errors in paranthesis allow for auto-correlation between observations based on the same worker. Week and worker fixed effects not shown. The data set is based on payroll records covering the 2001 through 2003 packing seasons. See text for details.

Table V: The Effect of Overtime on Output Estimated via IV, Payroll Data

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	Boxes Per Hour			Boxes Packed		
Sample:	All	Weekday	Weekend	All	Weekday	Weekend
Hours Overtime	19.422 (1.496)	13.252 (1.982)	16.929 (1.166)			
Hours Regular Time	19.340 (1.275)	27.502 (2.942)	13.700 (1.749)			
Share Overtime				1.926 (0.820)	- 14.000 (5.006)	2.823 (1.163)
Total Hours				0.807 (0.137)	1.621 (0.369)	0.252 (0.159)
F-statistic to test that OT = RT	0.010	14.010	5.240			
Constant	- 56.500 (12.464)	- 204.000 (23.765)	- 38.800 (6.900)	8.290 (0.981)	- 5.420 (2.571)	8.430 (1.609)
<i>N</i>	5,002	4,375	627	5,002	4,375	627

Note: In all specifications, share overtime and hours variables are predicted using week fixed effects and day-of-week fixed effects as instrumental variables. Standard errors in paranthesis allow for auto-correlation between observations based on the same worker. Week and worker fixed effects not shown. The data set is based on payroll records covering the 2001 through 2003 packing seasons. See text for details.

Table VI: Sample Statistics of the On-Site Data

Pear Size (Pears per Box)	Obs.	Minutes per Box	Seconds per Pear (SPP)	Std. Dev. of SPP
All	3,960	3.94	2.41	0.86
80	586	3.57	2.64	1.01
90	943	3.76	2.49	0.92
100	1391	3.91	2.33	0.81
110	843	4.40	2.36	0.73
120	197	4.21	2.08	0.62

Note: Packing speed is measured in pears packed per minute. The data set contains a record for every box packed by a sample of workers, and the time it took them to pack each box. See text for details.

Table VII: The Effect of Expected Shocks on Effort, On-Site Data
 Dependent Variable: Seconds per Pear

	(1)	(2)	(3)	(4)
90 pears per box	- 0.009 (0.082)	- 0.097 (0.063)		
100 pears per box	- 0.136 (0.099)	- 0.219 (0.074)		
110 pears per box	- 0.303 (0.096)	- 0.387 (0.065)		
120 pears per box (smallest pears)	- 0.539 (0.114)	- 0.438 (0.091)		
Pears per box			- 0.012 (0.003)	- 0.012 (0.002)
Constant	2.747 (0.099)	2.997 (0.175)	3.755 (0.326)	3.987 (0.261)
Worker Fixed Effect		X		X
Date Fixed Effect		X		X
R^2	0.039	0.250	0.036	0.249

Note: $N = 3,908$. The "pears per box" variables equal unity when the box packed has the given number of pears. Standard errors in paranthesis allow for auto-correlation between observations based on the same worker. All specifications include a fixed effect for the type of box packed. The data set contains a record for every box packed by a sample of workers and the time it took them to pack each box. See text for details.