Life Cycle Labor Supply:
Time to Sow and Time to Reap

J.Angrist

MIT 14.661 Fall 2024 (rev. 09-24-24)

A Background: the LCH/PIH

Consider the long run: as male manufacturing wages rose through the 20th Century, hours worked fell. Labor supply seems either inelastic or backward bending for prime-age men.

- Yet, as everyone knows, Uber *raises* driver pay when they want more drivers on the road - riders see this via surge pricing; drivers see it in surge pricing and driver promotions. So Uber (and other rideshare companies and their millions of shareholders) must believe that driver effort is elastic (Apparently, they haven't read Thaler 2015).
- Life-cycle models unpack this puzzle, reconciling flat or even backward-bending long-run labor supply responses with highly-elastic responses to transitory and evolutionary changes in pay.
  - transitory changes = wage *surprises*, like the bonuses now on offer for seasonal service work
  - evolutionary changes = *predictably* higher wages as we age, or on certain days (like holidays)
- The life-cycle labor supply model extends the life-cycle theory of consumption (LCH; attributed to Modigliani) and the permanent income hypothesis (PIH; attributed to Friedman)

At the beginning of adult life, person i chooses consumption each period (c_{i0},..., c_{iT}) to max utility of lifetime consumption:

$$U(c_{i0}, c_{i1}, ..., c_{iT}),$$

assuming known paths for wages and prices (later, we add uncertainty)

- To make this problem tractable, we assume *intertemporally additive* utility:

$$U(c_{i0}, c_{i1}, ..., c_{iT}) = \sum_{t=0}^{T} u_t(c_{it}) = \sum_{t=0}^{T} \left(\frac{1}{1+\rho}\right)^t u(c_{it})$$

  This assumes we discount future consumption at rate ρ. Our personal discount rate, also called the *rate of time preference*, need not equal the interest rate, r (but it might)

- The lifetime budget constraint is

$$\sum_{t=0}^{T}\left(\frac{1}{1+r}\right)^{t}[p_t c_{it} - y_{it}] = A$$

where y_{it} is income in period t (so far, income is just given, no labor supply needed), and A is the value of my (surely unearned) bar-mitzvah bonds

- The FOCs tell us to choose consumption in each period to satisfy:

$$\left(\frac{1}{1+\rho}\right)^{t} u'(c_{it}) = \lambda_i \left(\frac{1}{1+r}\right)^{t} p_t$$

where λ_i is person i's Lagrange multiplier; a function of prices, y_{it}, and A.

- Suppose $\rho = r$ and prices are constant at p:

$$u'(c_{it}) = \lambda_i p$$

  - Consumption here is constant: y_{it} has no bearing on my optimal c_{it}
  - The PIH/LCH suggests consumption is likely to be much smoother than income. Why does this matter for policy?

- With time-varying prices,

$$u'(c_{it}) = \lambda_i p_t$$

  - Note that utility is assumed concave in consumption. Show that this implies we consume when it's cheap to reap (i.e. prices are low), while transitory income shocks matter not
  - Why does the Lagrange multiplier have an i on it?

B The Life-Cycle Labor Supply (LCLS) Setup

Live long and prosper

Utility is a function of the lifetime stream of consumption (c_{i0},..., c_{iT}) and leisure (l_{i0},..., l_{iT}), where $h_{it} = \tau - l_{it}$ is hours worked in period t, capped at τ.

- At the beginning of adult life, we plan ahead to max:

$$U(c_{i0}, c_{i1},..., c_{iT}; l_{i0}, l_{i1}, ..., l_{iT}),$$

assuming known paths for wages and prices

- Intertemporally additive preferences simplify our lives:

$$U(c_{i0}, c_{i1},..., c_{iT}; l_{i0}, l_{i1}, ..., l_{iT}) = \sum_{t=0}^{T}\left(\frac{1}{1+\rho}\right)^{t} U(c_{it}, l_{it}) \tag{1}$$

- We often simplify yet further, invoking *within-period additivity*:

$$U(c_{it}, l_{it}) = u(c_{it}) + v(l_{it})$$

Ashes to ashes

The lifetime budget constraint is

$$\sum_{t=0}^{T}\left(\frac{1}{1+r}\right)^{t}\left[p_t c_{it} - w_{it}(\tau - l_{it})\right] = A_i \tag{2}$$

The lifetime planner maxes the RHS of (1) while constrained to die with a clean slate, as described by (2).

- *Recall the LCH/PIH*: Keep your eyes peeled for similar insights
- Maintaining within-period additivity, FOCs for period t choices go like this:

$$u'(c_{it}) = \left(\frac{1+\rho}{1+r}\right)^{t}\lambda_i p_t \tag{3}$$

$$v'(l_{it}) = \left(\frac{1+\rho}{1+r}\right)^{t}\lambda_i w_{it} \tag{4}$$

- Note that:

$$[c_{it}, h_{it}] = f(r, \rho, \lambda_i, p_t, w_{it})$$

  We also have:

$$\lambda_i = f(r; \rho, ; p_0, ..., p_T; w_{i0}, ..., w_{iT}; A_i),$$

  so the marginal utility of wealth is a function of prices and wages in all periods (where does this come from?)

  - The Lagrange multiplier again gets an i subscript - why?
  - This looks terribly abstract, yet has important concrete implications:
    * Conditional on λ_i, labor supply in period t depends only on contemporaneous wages. Wages and prices in all other periods operate through changes in λ_i, the marginal utility of lifetime wealth
    * Modest, short-lived, and/or perfectly anticipated changes in wages change λ_i little, and so we can ignore these *lifetime wealth effects* when evaluating the response to small, short-term, or anticipated wage changes
    * Modest, short-lived, and/or perfectly anticipated changes in wages must <u>increase</u> labor supply in the period in which they occur. To see this, simplify by assuming $\rho = r$ and differentiate:

$$v''(l_{it})dl_{it} = \lambda_i dw_{it},$$

      so

$$\frac{dl_{it}}{dw_{it}} = \frac{\lambda_i}{v''(l_{it})} < 0,$$

      because of diminishing marginal utility. This establishes $\frac{dh_{it}}{dw_{it}} = -\frac{dl_{it}}{dw_{it}} > 0$ when λ_i is fixed (recall that Lagrange multipliers are positive: they're MU wealth)

    * The wage-derivative of a $\lambda-$*constant* labor-supply function is called an *intertemporal substitution effect*. Like the static substitution effect in the Slutsky equation, it must be positive

  - These conclusions hold even when future wages and prices are uncertain, as they must be, provided we can forecast them reasonably well
  - Altonji (1986) details an uncertainty-inclusive version of the basic LCLS setup pioneered in MaCurdy (1981)

C Heck-MaC and the ISE

- The following "Heck-Mac (1980) utility function" generates a neat life-cycle labor supply equation:

$$U(c_{it}, l_{it}) = c_{it}^{\delta_1} - \gamma(\tau - l_{it})^{\delta_2} = c_{it}^{\delta_1} - \gamma h_{it}^{\delta_2}$$

- From the FOC for leisure (equation 4), we get

$$\gamma \delta_2 h_{it}^{\delta_2 - 1} = \left(\frac{1+\rho}{1+r}\right)^t \lambda_i w_{it}$$

  where second order conditions require $\delta_2 > 1$

- Futzing and putzing, this yields a linear-in-logs labor supply function:

$$\ln h_{it} = \left[\frac{\ln \lambda_i - \ln \gamma - \ln \delta_2}{\delta_2 - 1}\right] + \frac{t}{\delta_2 - 1} \ln\left(\frac{1+\rho}{1+r}\right) + \frac{1}{\delta_2 - 1} \ln w_{it}$$

- Using the approximation $\ln\left(\frac{1+\rho}{1+r}\right) \approx \rho - r$ yields the Heck-Mac labor supply equation:

$$\ln h_{it} = \mu_i + \delta(\rho - r)t + \delta \ln w_{it} + \varepsilon_{it} \tag{5}$$

  where $\delta = \frac{1}{\delta_2 - 1} > 0$ and ε_{it} is a residual representing random unexplained variation in hours worked

Understanding the ISE

The parameter δ, an *intertemporal substitution elasticity* (ISE), interests us greatly

- The ISE must be positive (as a matter of theory) and is (weakly) larger than a traditional static substitution elasticity, which of course, exceeds the uncompensated elasticity
- The ISE describes labor supply responses that hold the marginal utility of wealth (λ_i) fixed. Examples:
  - Consider my lifetime work plan: when my wage profile is known; my marginal utility of wealth is fixed. But I work harder at age 30 than 25. How come and how much? The ISE answer this question, describing how I allocate my hours over my lifetime to best exploit the low-hanging fruit on offer when wages are high, while binge-watching *HotD* when my time is cheap

  - Cab drivers who anticipate trip demand over days of the week and hours of the day make the same calculation: they're on the road when the driving is good (Uber and Lyft driver apps help with this by pinpointing high-wage periods and locations)
  - We don't all drive cabs (at least not yet)
    * The ISE also approximates the response to short-run or small changes that change lifetime wealth little
    * The ISE looms large in macro: cyclical variation that is either anticipated or modest enough to leave λ_i unchanged (perhaps a temporary tax reduction) generates an ISE-mediated supply response

- The ISE/ISH concept isn't unique to Heck-MaC utility; any LCLS model has one
- What the ISE/ISH doesn't explain: the response to changes in wealth. The Heck-MaC model implies that labor supply responds to a (log) wage shock of amount Δ by:

$$\frac{\partial \ln h_{it}}{\partial \Delta} = \delta \frac{d \ln w_{it}}{d\triangle} + \delta \frac{d \ln \lambda_i}{d\triangle}$$

  An wage increase of Δ percent every period, for example, reduces the marginal utility of wealth (because such a shift makes me wealthier, and marginal utility is declining):

$$\frac{d \ln \lambda_i}{d\triangle} < 0,$$

  and therefore adds a negative wealth effect on hours that may dominate the ISE

A sense of smoothness

- The LCLS framework presumes perfect credit markets: workers borrow and lend freely at parametric interest rates, *frictionlessly* exploiting the fact that to every thing there is a season and a time to every purpose under heaven ... when wages are high, it's time to work, but when they're low, time to loaf, using earnings from good times to fund consumption in the bad
- Workers who can't take advantage of transitory wage gains to fund consumption when wages are low are said to be *liquidity constrained*
- Liquidity constraints and simple myopia look similar: both generate strong within-period wealth effects
- How *relevant* is the ISE? On this, reasonable labor economists can disagree. Card (1994) resists claims for the ISH as an important determinant of labor supply, while my work (Angrist 1990, 1991; Angrist, Caldwell, and Hall 2021) leaves me an ISE optimist.

D LCLS 'Metrics

D.1 Life's identification challenges

- The iconic empirical life-cycle labor supply function looks like this:

$$\ln h_{it} = \mu_i + \delta(\rho - r)t + \delta \ln w_{it} + \varepsilon_{it} \tag{6}$$

where $\delta = \frac{1}{\delta_2 - 1}$ and ε_{it} is an error term that's tacked on

- Trouble in mind
  - Control variable μ_i isn't found in the CPS. Being a function of the marginal utility of wealth, this omitted variable is negatively correlated with wages, w_{it}
  - We have limited data on hourly wages; we often work with average hourly earnings, $AHE_{it} \equiv \frac{y_{it}}{h_{it}}$. So we're naively regressing *hours worked* on something involving $(hours\,worked)^{-1}$; The results might not be pretty; rather, they're pretty often negative!

D.2 Division bias details

Suppose the labor supply equation of our dreams is

$$\ln h^*_{it} = \alpha + \delta \ln w^*_{it} + \varepsilon_{it} \tag{7}$$

For the purposes of this discussion, start by assuming we'd be happy to estimate (7) by OLS.

The supply function of interest uses AHE with well-measured hours,

$$w^*_{it} = \frac{y_{it}}{h^*_{it}},$$

where y_{it} is annual earnings. This is the hourly wage for those who are paid hourly, and its a notional time price for others.

Alas, hours are poorly measured:

$$h_{it} = h^*_{it} \upsilon_{it},$$

where υ_{it} is proportional measurement error (log normal, perhaps). Then:

$$\ln h_{it} = \ln h^*_{it} + \eta_{it}, \tag{8}$$

where $\eta_{it} = \ln \upsilon_{it}$. This implies that:

$$\ln w_{it} = \ln y_{it} - \ln h_{it} = \ln y_{it} - \ln h^*_{it} - \eta_{it} = \ln w^*_{it} - \eta_{it}. \tag{9}$$

Substituting for log hours and log wages in (7), we now have:

$$\ln h_{it} = \alpha + \delta(\ln w_{it} + \eta_{it}) + \varepsilon_{it} + \eta_{it} = \alpha + \delta \ln w_{it} + \{\varepsilon_{it} + (1+\delta)\eta_{it}\} \tag{10}$$

The OVB in OLS estimates of (10) is

$$OVB = \frac{Cov(\ln w^*_{it} - \eta_{it}, (1+\delta)\eta_{it})}{\sigma^2_{\ln w}} = -(1+\delta)\left[\frac{\sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{\ln w^*}}\right]$$

which is real bad, even compared to the usual measurement error attenuation bias, which is just one minus reliability. (Note that the term in brackets is one minus the signal-to-noise ratio–also called *reliability*–of log wages.)

Analysis of covariance aggravates division bias

To kill unobserved fixed effect, μ_i, we might difference or deviate from means. Suppose you have a two-period panel, so (7) with fixed effects becomes OLS on first diffs:

$$\Delta \ln h^*_{it} = \delta \Delta \ln w^*_{it} + \Delta \varepsilon_{it}, \tag{11}$$

while the noisy average hourly earnings (wage) variable becomes:

$$\Delta \ln w_{it} = \Delta \ln w^*_{it} - \Delta \eta_{it}.$$

Assuming measurement error is serially uncorrelated, the variance of $\Delta\eta_{it}$ is $2\sigma^2_\eta$.

Actual wages, by contrast, are highly persistent. Suppose, $w^*_{it} = w^*_i$, in which case:

$$\Delta \ln w_{it} = -\Delta \eta_{it}.$$

In other words, *wage changes are pure noise.* You can't estimate (11); you must first-diff (10) instead. The OVB here is:

$$OVB = \frac{Cov(-\Delta\eta_{it}, (1+\delta)\Delta\eta_{it})}{\sigma^2_{\Delta \ln w}} = \frac{-(1+\delta)2\sigma^2_\eta}{2\sigma^2_\eta} = -(1+\delta)$$

so differencing in this context aggravates attenuation bias to the point where our putative ISE is $\delta-(1+\delta) = -1$! Research on measurement error in hours and wages bears this out: measured wage changes are noisy (see, e.g., Bound and Krueger, 1991), while those who've ventured to compute estimates of equations like (11) by estimating the mismeasured analog (10) in first diffs indeed suffer the ignominy of large negative labor supply elasticities.

- For more on how and why covariates and differencing aggravate attenuation bias in regressions with mismeasured regressors, see MM Chapter 6.

2SLS to the rescue

Rewrite (6) as

$$y_{it} = \mu_i + \alpha t + \delta x_{it} + \varepsilon_{it} \tag{12}$$

where $x_{it} = \ln w_{it}$ and $y_{it} = \ln h_{it}$.

Now, average this model by year:

$$\bar{y}_t = \bar{\mu} + \alpha t + \delta \bar{x}_t + \bar{\varepsilon}_t \tag{13}$$

In an asymptotic sequence that increases the number of workers,

$$\begin{aligned} plim\ \bar{\mu} &= E[\mu_i] \\ plim\ \bar{x}_t &= E[\ln w^*_{it}|t] \\ plim\ \bar{\varepsilon}_t &= 0 \end{aligned}$$

From this we conclude that the OLS estimate of δ in (13) is consistent.

But, of course! For OLS on (13) is the same as 2SLS using time dummies as instruments. Assuming, as we have implicitly done, that:

- measurement error has time mean of zero
- correctly measured wages vary over time
- time effects can be omitted from individual labor supply equations,

The vector of $Z = \{D_t; t = 1,...,T-1\}$ can be used to instrument (6). To see why OLS on the grouped version is the same as 2SLS with time dummies, note that first stage fits using Z (and a constant) as instruments are group means, and that the projection operation that creates them is idempotent. In principle, we should weight the grouped equation by sample size, but the PSID is a balanced panel, so weights there are constant.

The goodness-of-fit statistic for the fit of (13) is

$$\sum_t \frac{N(\bar{y}_t - \hat{\mu} - \hat{\alpha}t - \hat{\delta}\bar{x}_t)^2}{\sigma_u^2} \sim \chi^2(k)$$

where k is the difference between the number of periods and the number of parameters to be estimated. This chi-square statistic is algebraically the same as the over-identification test statistic associated with 2SLS using time dummy instruments. Estimates in Angrist (1990, 1991) show a surprisingly good fit:

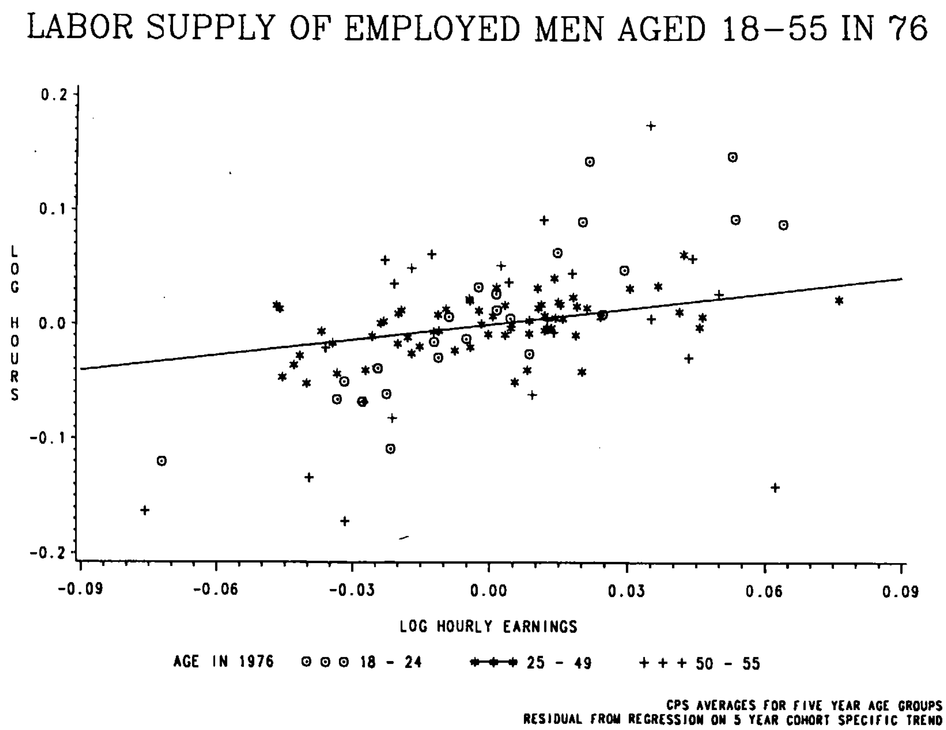

E LCLS Showdown: Intertemporal Substitution vs. Target Earning

Camerer, et al (1997)

> Estimated wage elasticities are significantly negative ... Our interpretation of these findings is that cab drivers (at least inexperienced ones): (i) make labor supply decisions "one day at a time" instead of intertemporally substituting labor and leisure decisions across multiple days (ii) set a loose daily income target and quit once they reach that target.

Farber (2005)

> I am puzzled by these findings ... target earning implies that, on days when it's easy to make money (pick low-hanging fruit, so to speak), drivers quit early, whereas on days when fares are scarce, drivers work longer hours.

Making sense of target earning

Target earning behavior can be explained by two closely-related economic models:

1. Large within-period (daily, for cab drivers) income effects in response to lifetime-wealth-neutral wage changes, due, perhaps, to liquidity constraints
2. *Reference-dependent preferences*: very low or even zero MU(income) above some target earnings level (NYC yellow cab drivers, for example, have been said to target the cost of their medallion lease plus a couple hundred dollars/day.) Transitory wage gains that push me across the target drive marginal utility way down; Fehr and Goette (2007) detail this

Occ rock

Take me out to the ballgame (Oettinger, 1999)

- Stadium vendors show up for work at as many of 81 home games as they like
- The fraction who sell at each game is an increasing function of game-time AHE, instrumented with game demand parameters

Consider the lobster (Stafford, 2015)

- Lobsters come out when the moon is in (because it's darker)
- Lobstermen come out when lobsters come out, and are almost unit elastic

Flash and Veloblitz (Fehr and Goette, 2007)

- An RCT that randomized commission rates for some of Zurich's finest riders, while keeping prices to customers fixed
- This generates ISEs in excess of 1, with small reductions in average daily hours

- Timing:
  - September 2000 (Treatment Period 1), Rider Group A gets 25% more
  - November 2000 (Treatment Period 2), Rider Group B gets 25% more
  - Only Veloblitz riders were treated; Flash riders were not involved.

VOL. 97 NO. 1 FEHR AND GOETTE: DO WORKERS WORK MORE IF WAGES ARE HIGH? 307

TABLE 1—DESCRIPTIVE STATISTICS

		Participating messengers		Difference groups	Nonparticipating messengers,	Messengers,
		Group A	Group B	A and B	Veloblitz	Flash
Four-week period prior to experiment	Mean revenues	3,500.67	3,269.94	241.67	1461.70	1637.49
		(2,703.25)	(2,330.41)	[563.19]	(1,231.95)	(1,838.61)
	Mean shifts	12.14	10.95	1.20	5.19	6.76
		(8.06)	(7.58)	[1.75]	(4.45)	(6.11)
	N	21	19		21	59
Treatment period 1	Mean revenues	4,131.33	3,005.75	1,125.59	844.21	1,408.23
		(2,669.21)	(2,054.20)	[519.72]	(1,189.53)	(1,664.39)
	Mean shifts	14.00	9.85	4.15	3.14	6.32
		(7.25)	(6.76)	[1.53]	(4.63)	(6.21)
	N	22	20		21	65
Treatment period 2	Mean revenues	2,734.03	3,675.57	−941.53	851.23	921.58
		(2,571.58)	(2,109.19)	[513.2]	(1,150.31)	(1,076.47)
	Mean shifts	8.73	12.55	−3.82	3.29	4.46
		(7.61)	(7.49)	[1.65]	(4.15)	(4.74)
	N	22	20		24	72

Notes: Standard deviations in parentheses, standard error of differences in brackets. Group A received the high commission rate in experimental period 1, group B in experimental period 2.
Source: Own calculations.

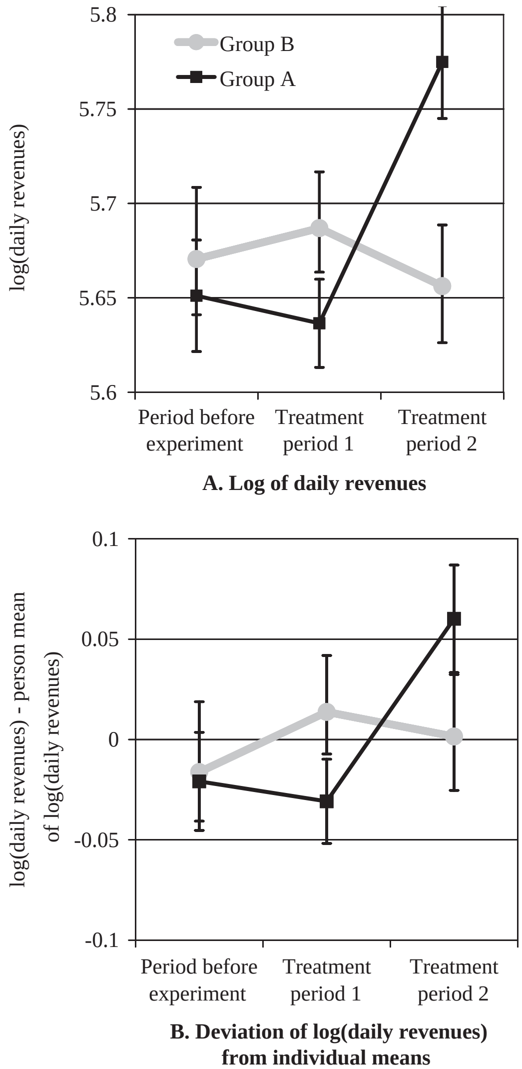

Figure 1. Log of Daily Revenues on Fixed Shifts

Table 5—The Impact of the Experiment on Log Revenues per Day

(Dependent variable: log (revenues per shift) during fixed shifts, OLS regressions)

	(1)	(2)
Treatment dummy	−0.0642**	−0.0601**
	(0.030)	(0.030)
Gender (female = 1)	−0.0545	
	(0.052)	
Log(tenure)	0.105***	0.015
	(0.016)	(0.062)
Day fixed effects	Yes	Yes
Individual fixed effects	No	Yes
R-Squared	0.149	0.258
N	1,137	1,137

Note: Robust standard errors, adjusted for clustering on messengers, are in parentheses.
*** Indicates significance at the 1-percent level.
** Indicates significance at the 5-percent level.
* Indicates significance at the 10-percent level.
Source: Own calculations.

The treatment effect in this regression is virtually unchanged and indicates a reduction in revenues of 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