Willingness to Pay for Workplace Safety

Massimo Anelli[1] and Felix Koenig[2]

[1]Bocconi University, IZA, CESifo

[2]Carnegie Mellon University, CEP, IZA

July 22, 2021

DRAFT – click here for the most recent version.

Abstract

What value do workers attach to workplace safety? This paper develops a novel revealed-preference approach that uses bunching around benefit thresholds to measure the value of non-wage amenities. Using this approach, we find that workers are willing to accept 34% lower incomes to reduce the fatality rate by one standard deviation, or 1% of income to reduce weekly fatality rates by one in a million. Our approach measures weekly labor supply decisions among hourly workers who repeatedly face the decision of how many hours to work at varying levels of Covid-19 risk and leverages 21 state-specific discontinuities in unemployment insurance eligibility criteria to identify the labor supply behavior. Results show a large baseline labor supply response around the eligibility threshold for the $600 Federal Pandemic Unemployment Compensation and rising labor supply responses when health risk increases. These estimates imply the value of a statistical life of $6.9 million, a value consistent with estimates from the literature. We additionally relax the perfect information assumption implicit in such measures and show that failing to account for imperfect risk assessments can produce substantial biases in these estimates.

JEL-Codes: J17, J22, J28
Keywords: Hazard Pay, Workplace Safety, Nonwage Amenities, Partial Unemployment Insurance, COVID19, Labor Supply, Value of Life

1 Introduction

Compensating differentials are a cornerstone of labor economics. Non-wage features of jobs are understood to play a major role in the labor market; they are crucial, for example, in rationalizing wage differences between similar workers (Rosen, 1986, 1974; Lucas, 1977; Masters, 1969). Furthermore, the theory of compensating differentials is important for the purpose of formulating appropriate labor market regulation. Unfortunately, though, quantifying parameters in compensating differential models has proven difficult. In this paper we develop a novel approach for this purpose, and apply our approach to a canonical, policy relevant case—the value of workplace safety.

Every year around 3 million Americans – nearly 2% of the labor force – suffer work-related injuries or illnesses.[1] This fact has raised questions about how these workplace risks should be addressed. Are workers adequately compensated for work risks?[2] Are companies doing enough to protect the health of their employees?[3] The workplace safety challenge has become particularly salient during the current crisis, which has exposed thousands of frontline workers to fatality risks. Major employers such as Amazon, Best Buy, and Target introduced Covid-19 hazard pay and raised wages in response.[4] However, it is unclear how these employers set these hazard bonuses and whether these payments compensated for the increased risk. More generally, how do workers value the benefit of safer workplaces? In this paper, we address these questions and estimate the value of safe workplaces using a revealed preferences approach.

To measure the value of safe workplaces, we develop a revealed-preference estimator for non-wage amenities. This approach builds on the bunching literature and shows that worker preferences over non-wage amenities can be estimated from the amount of excess mass at budget discontinu-

[1] Source of injury and illness rates is the BLS series ISU00000000031100

[2] Health is a major source of lifetime income risk: Dobkin et al. (2018) estimate that worker hospitalizations lead to a 20% decline in earnings three years after the initial event.

[3] Workplace safety is particularly deficient in low-paid occupations and labor market inequalities are therefore bigger than wage differences alone suggest.

[4] For these companies, the hazard pay ranged between \$2 and \$2.50 per hour (Kinder, Stateler, and Du (2020)). By comparison, hourly wages increased on average by 4.5% (about \$0.83) within employment spells in our sample of small businesses.

ities. Work dis-amenities operate like an additional tax on work, reduce returns to work, and thus amplify the effect of budget notches and the magnitude of excess mass around these notches. Consider a job with large non-wage amenities. In this case financial incentives matter less for work decisions, and the impact of budget discontinuity is muted; to the contrary, when workplace safety deteriorates, returns to work depend more strongly on financial incentives, and budget discontinuities have an amplified impact on labor supply. The willingness to pay for workplace safety can be computed based on the differential response of excess mass to changes in workplace safety.

We implement this approach with discontinuities in work incentives from partial unemployment insurance rules, together with variation in workplace safety during Covid-19 outbreaks. The launch of Federal Pandemic Unemployment Compensation (FPUC) creates a jump in worker's budget sets that allows us to separate labor supply behavior from other factors. A worker is entitled to the \$600 FPUC wage supplement if her income falls below an earnings threshold, and by moving across the threshold, the worker loses eligibility. State-specific rules lead to 21 different eligibility thresholds – one for each US state that is part of the analysis. As a result of these different rules, equally paid workers are treated in some states but not in others and we compare two equally paid workers but on different sides of the benefit eligibility threshold to identify labor supply responses. Once we identify labor supply, we study how workplace safety affects such behavior. We study weekly labor supply decisions among hourly workers who repeatedly face the decision of how many hours to work at varying levels of Covid-19 risk.[5]

Our results show a sizable baseline labor supply response around the eligibility threshold for the \$600 Federal Pandemic Unemployment Compensation. Eligible workers reduce earnings to levels below the UI eligibility threshold and we see substantial missing and excess mass in the earnings distribution.[6]

We then show that deteriorating workplace safety leads to additional labor supply responses and magnified excess mass around the UI eligibility thresholds. For this analysis, we use workplace safety shocks during Covid-19 outbreaks. Our measure combines data on local outbreaks

[5] This is similar in spirit to Sorkin (2018) who uses job switch patterns to identify non-pecuniary values of firms.

[6] By contrast, in a placebo test, we find no such effects among similar workers who do not qualify for FPUC.

with tasks specific Covid-19 susceptibility scores. For example, local outbreaks expose restaurant workers to greater increases in risk than workers in automotive repair. Our analysis focuses on front-line hourly workers, mostly in services and retail jobs, who have limited opportunities to work from home and face a choice between risking health and losing income. We study their weekly labor supply under varying risk scenarios and find significant increases in excess mass around the UI thresholds when health risks increase. These estimates imply that individuals are willing to give up around 34% of their income to reduce Covid-19 risk by one standard deviation. In hourly wages, this is equivalent to a $6 wage decrease or equivalent to a 1% decrease in pay for a one in a million lower fatality rate. These findings show substantially higher willingness to pay than a canonical hedonic wage regression. Such a regression shows a WTP of around 0.5% of income for a standard deviation in risk. One reason for the difference is that wages are slow to adjust and do not fully reflect changes in risk. As a result, the hedonic regression results are downward biased.

The variation in workplace safety in this setting has three key advantages. First, the change in workplace safety was unexpected, large, and salient and thus alleviates concerns about downward biased estimates from less prominent variation. Second, we observe the same workers over time and can identify responses while controlling for selection with fixed effects. We can thus rule out that heterogeneities in labor supply elasticities across industries are driving the results. Third, we can compare individuals who live in the same locations and are exposed to the same policies, but face different exposure risks due to the nature of their work tasks. By controlling for location-specific time effects, we can rule out that the findings are driven by shocks that affect the location contemporaneously (e.g., school or transport closures).

We finally show that workers substantially overestimated the fatality risks and acted as if the risk was substantially higher. We show that failing to account for such imperfect information leads to biases in standard Value of Statistical Live (VSL) estimates.[7] Our benchmark estimates imply

[7] Our results echo a sizable behavioral literature on belief formation under uncertainty. See e.g., classic work on prospect theory by Kahneman and Tversky (1979) and recent work on over-emphasize of salient decision features in Bordalo, Gennaioli, and Shleifer (2013). For empirical evidence, see the review by Robinson and Hammitt (2011) and Viscusi (1990)) for an application to health.

the value of a statistical life of $6.9 million, a value consistent with estimates from the literature.[8] Accounting for perceived risk reduces the VSL estimate to $2.6 million. In addition, our findings suggest that the heightened and persistent sense of vulnerability can potentially explain why labor supply was slow to recover, even after risks abated.

An important objective of much labor market research is to inform policies around minimum work standards. A large and influential literature studies minimum wage standards. However, minimum standards of other work dimensions have received far less attention. This limitation comes in part from the difficulty of quantifying the gains from such non-wage regulations.

Our estimates provide a quantitative basis to study the efficacy of workplace safety policies. In a perfectly competitive market, no safety regulation would be necessary, as competition ensures that high-risk employers pay high compensating differentials or go out of business. With imperfect competition, however, firms may not fully internalize the cost of high-risk jobs and may expose workers to excessive risks. Our estimate of the monetary value of improved workplace safety determines the benefits from such policies and is therefore vital for designing such policies. Our results suggest that in the U.S. context, the gains from more stringent safety regulations would be large. In the specific application of hazard pay during the Covid-19 outbreak, our estimates suggest that hazard pay would need to be as high as $6 per hour to fully offset the non-pecuniary costs of added workplace risk, a level that is substantially higher than the increases actually implemented by most employers ($2-2.5). More generally, we find that reducing workplace fatalities to the level observed in the UK and Germany would lead to substantial gains for workers. In construction, for instance, such a reform would be equivalent to a 2.5% wage increase, a gain that's similar to the wage effect of introducing a $15 minimum wage.

Related Literature – The topic of nonwage amenities goes back to classic work in economics (see Rosen (1986, 1974); Lucas (1977); Masters (1969)), but estimating the value of such amenities is difficult in practice.[9]

[8] For recent work on VSL, see for example Guardado and Ziebarth (2019); Lee and Taylor (2019); Lavetti (2020) and for meta-studies on earlier work Viscusi (2018); Viscusi and Aldy (2003).

[9] Compensating differentials can help understand inequality and rationalize wage differences between otherwise similar employers, which play an important role in the labor market (see, e.g. Card, Heining, and Kline (2013)).

The empirical literature on the value of workplace safety typically uses hedonic wage regressions to estimate such values.[10] Hedonic regressions relate occupational wage differences to workplace risk (Brown, 1980; Hwang, Reed, and Hubbard, 1992). A limitation of this approach is that it assumes efficient markets and relies on the idea that the value of workplace safety is priced into wages (Altonji and Paxson, 1992; Bonhomme and Jolivet, 2009; Ruppert, Stancanelli, and Wasmer, 2009). One implication of the efficient labor market assumption is that there is no scope for policy interventions. As a result, hedonic regressions are ill suited to study policy questions like the optimal level of safety regulation and minimum work standards. For these questions we would need to know the workplace risk that is *not* priced into wages. We allow for potential market failures in pricing the cost of workplace safety into wages and our estimates therefore can be used to inform policy decisions on workplace safety.

Another challenge for hedonic wage regressions is to isolate worker preferences from other omitted variables. The canonical hedonic regressions use cross-occupation comparisons: a coal miner, for example, faces greater workplace risk than an administrative assistant and in a competitive market, this leads to compensating wage differences: a coal miner will earn a higher wage which compensates for the added workplace risk. Confounding unobserved productivity differences make it, however, difficult to separate the role of individual labor supply decisions from confounding factors in practice. Several studies address this challenge in the context of other non-wage amenities and leverage policy reforms to estimate their value (Gruber (1994); Gruber and Krueger (1991); Fishback and Kantor (1995)). We use a similar quasi-experimental approach to estimate the value of workplace safety and leverages budget discontinuities to identify labor supply preferences. The theoretical framework is related to the bunching/notch estimation approach in (Kleven, 2016; Kleven and Mazhar, 2013; Chetty, Friedman, and Saez, 2013). We expand their canonical two good labor-leisure approach to a three good economy with workplace safety. Fi-

Recent applications of the compensating differentials approach to inequality and worker sorting, include e.g. Goldin and Katz (2011, 2016); Morchio and Moser (2019); Taber and Vejlin (2020).

[10]For other types of amenities researchers have used stated preference surveys as an alternative to revealed preference estimates (e.g., Flory, Leibbrandt, and List (2015); Wiswall and Zafar (2017); Maestas et al. (2019)). In addition, Mas and Pallais (2017) use a field experiment to estimate the value of schedule autonomy. Similar experiments with workers' health are, however undesirable.

nally, our estimate of a monetary value of avoiding Covid-19 death risks also relates to the large literature on the value of a statistical life (Prominent examples include Ashenfelter and Greenstone (2004); Viscusi and Aldy (2003)).

2 Partial UI Eligibility and the Federal Pandemic Unemployment Compensation

The quasi-experiment at the heart of our empirical strategy is the launch of the Federal Pandemic Unemployment Compensation (FPUC), a lump-sum \$600 expansion of unemployment insurance (UI) benefits, independent of the generosity of the UI payment. This reform was introduced as part of the Coronavirus Aid, Relief, and Economic Security (CARES) Act, enacted on March 27, 2020, and ended on July 31, 2020.[11]

The FPUC payments are also available to individuals who continue to be employed and are paid up to an earnings threshold. The complete withdrawal of benefits above the threshold generates a notch in the budget set of workers. Figure 1A shows a simplified unemployment insurance schedule and the notch that arises with the launch of FPUC. Ordinary UI benefits are gradually withdrawn as earnings increase and benefits decrease at the benefit reduction rate t with each \$ earned. The \$600 FPUC, by contrast, is not phased out and instead is completely withdrawn once the earning ceiling m^* is reached. Workers thus stand to lose the full \$600 if their earnings exceed m^*. This creates an incentive not to exceed m^* and potentially generates excess and missing mass in the earnings distribution, as shown by the light grey area in Figure 1B.

Not all workers are eligible to FPUC, and we limit our sample to those who are eligible and qualify for the maximum weekly benefit amount (MWB).[12] The eligibility is based on contribution to the UI system in prior quarters. And while we do not directly observe eligibility, we apply state-

[11] No FPUC benefits were payable between July 31, 2020, and December 26, 2020. FPUC was re-established by the Continued Assistance Act as a \$300 per week supplement to unemployment benefits from December 26, 2020, to March 14, 2021. Please consult Appendix C for more details on FPUC and subsequent programs.

[12] Some workers are eligible for a fraction of MWB benefits. For these people, different earnings allowances apply, and we exclude them from our analysis.

Figure 1: Bunching at m^*

Panel A

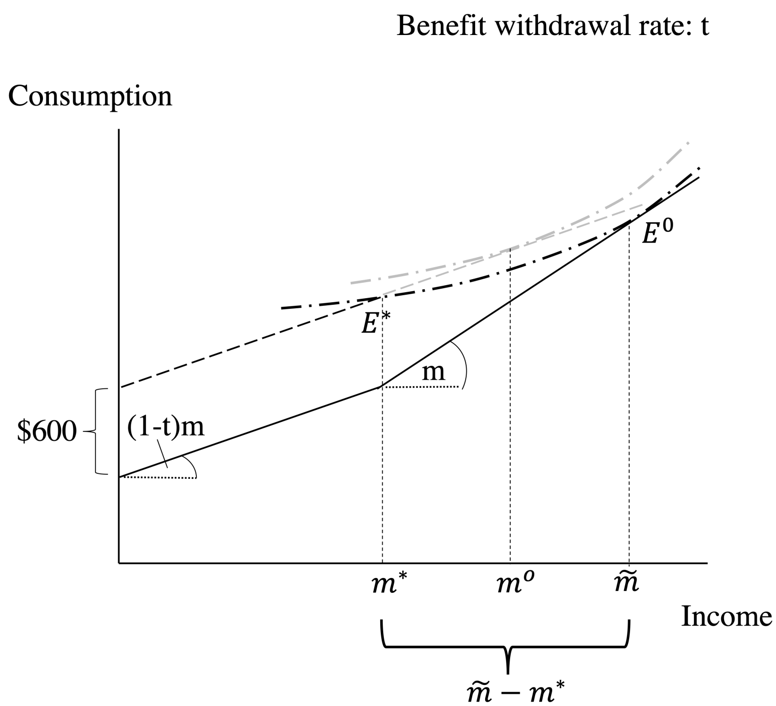

Panel B

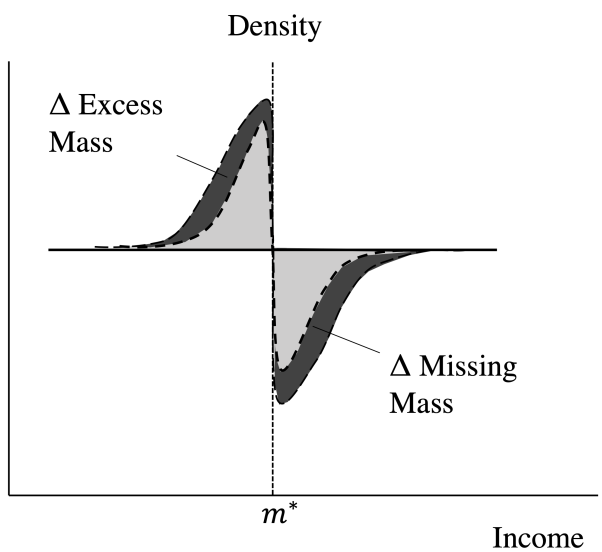

specific eligibility rules and identify workers who are eligible for MWB.[13] For these workers, the earnings threshold that qualifies them for FPUC is the same within state, independently of their earnings.

In what follows we illustrate how we exploit the differential behavioral response to FPUC under changing Covid-19 risk for a revealed preference estimate for the WTP for workplace safety.

3 Willingness To Pay for Non-Wage Amenities

This section presents a revealed preference approach to identify the WTP for workplace safety, which leverages budget discontinuities to identify worker preferences. Such discontinuities are typically used to estimate preferences over leisure and income. We extend this framework to a 3 good economy with leisure, income and health.

For an intuition of our approach, consider this simple logic: when non-wage amenities are the main component of returns to work (e.g. workers enjoy spending time with coworkers and the tasks they do) than discontinuities in financial incentives should play a minor role for labor supply decisions. Conversely, when dis-amenities are large, financial returns are the main driver to work, and labor supply should be more responsive to financial incentives. Figure 1B illustrates the implication for excess/missing mass around budget discontinuities. An increase in dis-amenities (in our case the increase in health risk) magnifies role of financial incentives, and triggers "magnified excess mass" around the FPUC threshold (illustrated in black. Our proposed empirical strategy quantifies this "magnified excess mass" and thereby estimates how the labor supply elasticity changes when work (dis-)amenities fluctuate.

To see this formally, take an individual who obtains utility from after-tax income (or consumption), pre-tax income (cost of effort), and a third good (health in our case). The utility function is $U(m - T(m), m/a, h)$, where m are earnings, $T(m)$ the tax schedule, a the worker ability and h the health of the worker. Heterogeneity in ability is captured by a distribution function $f(a)$. As-

[13] MWB is determined in most states as a function of quarterly earnings of the second-last quarter

sume this distribution and the tax system and preferences are smooth so that the resulting earnings distribution is also smooth. The tax rate is t and benefit Δt are available below earning level m^*. This benefit schedule is illustrated in Figure 1 and is:

$$T(m) = \begin{cases} t * m + \Delta t & m \leq m^* \\ 0 & m > m^* \end{cases}$$

The loss of benefits at m^* generates a notch in the budget that will incentivize movements from the right to the left of the eligibility threshold m^*.

We show how responses at the notch reveal the WTP for workplace safety. We focus on workplace safety, but the same method could be applied to all aspects of work that affect utility (e.g., a sense of purpose, interaction with colleagues, fear of mobbing, etc.). The worker experiences a negative health shock with probability θ during a work hour and utility in the injured state is $U(m, h_i)$. To simplify notation, we assume that risk increases with income m, rather than work hours and for a given workday the risk is: $r = \theta m$ and the expected utility is:

$$E(U(m,h)) = (1-r)U(m,h_0) - rU(m,h_i)$$

Denote the WTP to avoid an injury by W, such that $U(m, h_i) = U(m - W, h_0)$. Analog to the canonical iso-elastic quasi-linear assumption of the two-good economy, we assume that utility is separable and quasi-linear in income.[14] This utility takes the form:

$$U(m,h) = m - T(m) - \frac{a}{1+1/e}(\frac{m}{a})^{(1+1/e)} + h$$

where e is the labor supply elasticity. Using the definition of W, expected utility becomes: $E(U(m,h)) = U(m,h_0) - rW$.

Using the definition of W in this utility function and normalizing $h_0 = 0$, we can express

[14]The assumption of additive value of amenities is common in the literature (e.g., Morchio and Moser (2019)). For a more general utility function, see Appendix D.

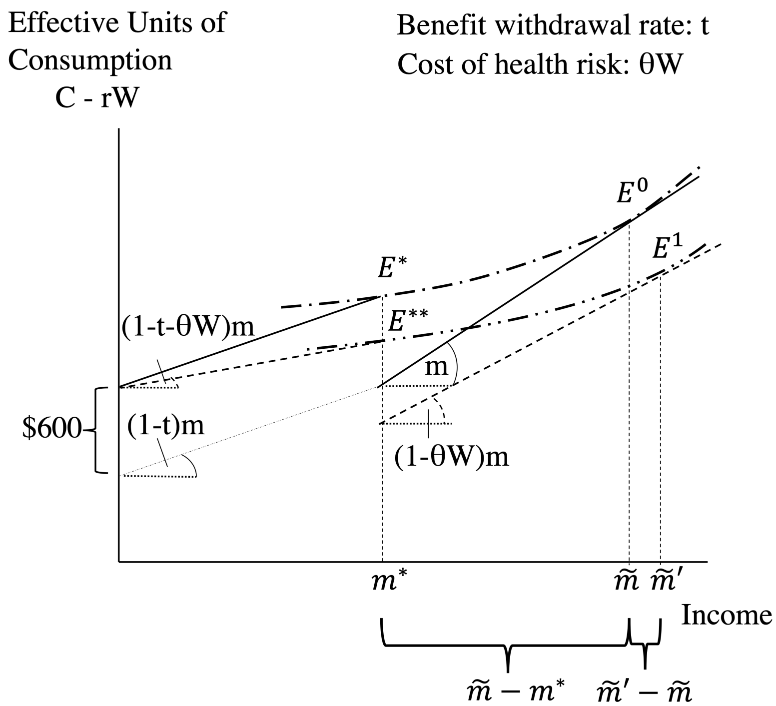

Figure 2: Cost of health risk

expected utility as:

$$E(U(m)) = m - T(m) - \theta W m - \frac{a}{1+1/e}(\frac{m}{a})^{(1+1/e)}$$

The health risk acts like an additional tax with a tax rate θW and reduces the expected return to work. Figure 2 illustrates the impact on labor supply. And we can measure the resulting response in the amount of excess mass left of the eligibility threshold. Variation in this perceived tax on work generates fluctuation in the excess mass that identifies W.

To derive an expression for W we leverage the fact that the marginal worker is indifferent between choosing the notch point m^* and an interior point $\tilde{m}$, $EU^* = E\tilde{U}$. The case is illustrated in Figure 2. At the interior point $\tilde{m}$ the first order condition implies:

$$\tilde{m} = a(1 - \theta W)^e$$

and hence $E\tilde{U}$ is:

$$E\tilde{U} = a(1-\theta W)^{(1+e)} - \frac{a}{1+1/e}(1-\theta W)^{(1+e)}$$
$$= \frac{a}{1+e}(1-t-\theta W)^{(1+e)}$$

The utility at the notch point m^* is given by

$$EU^* = (1-t-\theta W)m^* - \frac{a}{1+1/e}\left[\frac{m^*}{a}\right]^{(1+1/e)} + \Delta t$$

Using that $EU^* = E\tilde{U}$ and that the interior solution at the lower tax rate implies $a = m^o/(1-t-\theta W)^e$, we can obtain the following expression for W:

$$\frac{m^o}{m^*}\frac{1}{1+e}\left[\frac{1}{1-t/(1-\theta W)}\right] - \frac{\Delta t/m^*}{1-t-\theta W} = \left[1 - \frac{e}{1+e}\left(\frac{m^*}{m^o}\right)^{1/e}\right]$$

This expression pins down the WTP for health W in terms of measurable quantities $t, \Delta t, \theta, m^*$ and parameters that we can estimate based on behavioral responses: e, m^o.[15]

Without health risk ($\theta = 0$), the previous expression collapses to the standard bunching formula. In all other cases, W is an additional unknown parameter and we require an additional behavioural equation to solve for W. This additional condition comes from observing workers in high and low risk states. The labor supply elasticity in the low risk state is $e = \frac{(\tilde{m}-m^*)/m^*}{\Delta t/(1-t)}$ and in the high risk state: $e = \frac{(\tilde{m}'-m^*)/m^*}{\Delta \tilde{t}/(1-\tilde{t})}$, where $(\tilde{m}'-m^*)$ is the labor supply response in the high risk state, and $\tilde{t} = t + \theta W$ is the implicit tax rate in this state. We can combine the two elasticity expressions to solve for the willingness to pay for workplace safety ($WTP(r)$). $WTP(r)$ is the fraction of disposable income a workers is willing to forgo for a safer workplace. Deviding the income sacrifice for safety improvement r by the disposable income, $WTP(r)$ is:

$$WTP(r) = \frac{rW}{(1-t)m^*} = \frac{(\tilde{m}'-\tilde{m})}{(\tilde{m}'-m^*)} \tag{1}$$

[15] It can be shown that $(m^o - m^*)$ is closely related to the amount of excess mass created by the budget discontinuity. The link between excess mass (E) and $(m^o - m^*)$ is $E = \int_{m^*}^{m^o} h_0 = (m^o - m^*)h_0$, where h_0 is the baseline income distribution. The last equality assumes h_0 is constant and simplifies the expression, the same approach, however, also works for cases with more flexible functions of h_0.

The final equality uses the two previous elasticity expressions and assumes that health risks are smooth throughout the cut-off ($\Delta\tilde{t} = \Delta t$).[16] The final expression states that the WTP is a function of the labor supply response in the high-risk state ($\tilde{m}' - m^*$) and the additional labor supply response in the high state compared to the low-risk state ($\tilde{m}' - \tilde{m}$). Intuitively, the WTP calculation normalizes the labor supply response to an increase in risk by the labor supply response to a monetary incentive. This allows us to express the WTP in terms of an equivalent $ amount.

Empirically, the additional labor supply response shows up as additional excess left of m^* as illustrated above in Figure 1B. If the excess mass with and without health risk is the same ($\tilde{m}' = \tilde{m}$), then $WTP(r) = 0$. By contrast a large $WTP(r)$ implies that excess mass increases sharply with risk ($\tilde{m}' > \tilde{m}$). In short, the magnitude of additional excess mass identifies the WTP for health in a revealed preference sense.

So far, the analysis assumed that W is a constant. The framework can, however, accommodate more complex risk preferences. A flexible extension lets W vary across demographic groups (g). We can estimate $W(g)$ by allowing for different levels of responses among alternative demographic sub-groups. Such estimates will trace out the WTP profile in a non-parametric fashion and can accommodate preference heterogeneity.

4 Scheduling data

Our analysis leverages data from a private company, Homebase, which provides scheduling and HR services to small businesses relying on hourly waged workers. Typical businesses covered by the data operate in the restaurant, food and beverage, retail, health and beauty, and healthcare sectors. These are exactly the sectors where front-line workers typically work and thus the type of workers who face the decision to reduce their work hours to diminish the risk of contracting Covid-19.

The data has three major advantages. First, it provides third-party reported data on weekly

[16] Note that this result holds independent of the structural assumptions about the utility function. The derivation only uses the definition of earning elasticities, which hold for a general set of utility functions.

work hours and earnings. Obtaining reliable labor supply records has been a key challenge as many survey data sources suffer from rounding and recall errors that make it difficult to measure labor supply changes reliably (see, e.g., classic work by Bollinger (1998); Bound and Krueger (1991)). The Homebase software was created to help companies maintain accurate work hour records – the core feature of the Homebase app is a time-clock app. Workers clock in and out on a mobile phone app when they start and end their shift and the software uses the phone's geo-location to ensure accurate clocking.[17] As a result we obtain one of the cleanest source of work hours data that we know off and can track work hours and compensation to the minute. Importantly, these records are not used to administer UI benefits which alleviates concerns about strategic misreporting and avoidance. The data therefore combines the accurate measurement advantage of administrative records but minimises the pitfall of misreporting, and enables us to measure true behavioral responses with greater precision.

A second advantage of the data is that it covers many states and is available in real-time with daily frequency. Several previous studies of UI benefits use records from the state unemployment administrations. Such data also have a high degree of accuracy, however, they often cover only a single state and become available with multiple years of delay. In our case, the data covers 21 states in 2020, which enables us to analyze current policies and control for nuanced state-specific shocks.

A third advantage is the coverage of the data. The data mainly covers small service sector businesses with hourly workers, which in our application has two further advantages.[18] First, these high-street service sector workers are typically "frontline workers" and directly exposed to Covid-19 risk at the workplace. The experience of sharp changes in workplace risk makes this group an ideal sample to study the responses to such risks. Second, adjustment frictions are smaller among hourly workers as schedules are usually adjusted weekly, thus increasing the power to identify hour

[17] An app on workers' smartphones recognizes when workers get to or leave the workplace and send a check-in/out notification as shown from the app screenshot in Appendix figure A1.

[18] In Appendix B, we systematically compare the Homebase data universe and our analysis sample with a nationally representative survey. That exercise shows that workers in our analytical sample have weekly earnings, hourly wages and hours worked very much in line with the average hourly worker in small firms of the 21 states under analysis.

adjustments.

The Homebase data does not directly contain information on unemployment insurance eligibility. As discussed above, we use state-specific eligibility rules to compute benefit eligibility and restrict our analysis to workers eligible for MWB. For these workers, the earning threshold and the resulting notch in the budget constraint can be easily computed.[19] For some workers where we only observe an incomplete work history during the qualifying period and we predict eligibility based on full-time earnings at the hourly wage rate of the most recent observed weeks. Since such imputations inevitably introduce noise, and we thus down-weight such observations.[20]

A drawback of this type of private-sector data is that we lack data on individuals who exit the sample. We, for instance, do not know whether these individuals left the labor force or changed employers. As a result, we only have a noisy measure of extensive margin responses. Our core analysis focuses on intensive margin responses and is thus unaffected by this problem. For completeness, we provide additional robustness checks on extensive margin responses. In these, We treat exits from the Homebase data as exits from the labor force and find similar results to our baseline estimates.

Our analysis focuses on the period from November 1st, 2019 until the end of the FPUC program on July 31st, 2020. The period covers four months before and four after the start of the Covid-19 pandemic in March 2020. Finally, since UI eligibility and benefits are calculated based on weekly earnings, we aggregate hours worked and earnings by week for each worker.

Table 1 - Panel A shows summary statistics for the workers in our analytical sample. The 6,861 included workers work on average 36 hours and receive weekly earnings of $634. The median hourly wage is $16 and does not vary much (the 25^{th} percentile is $14, and the 75^{th} is $20).

Panel B of Table 1 instead characterizes the 2,771 small businesses included in our analysis sample. On average, they have 1.2 branches and 14.2 employees, of which 97% are hourly waged workers in the median firm. 36% of all firms operate in the Food and Drink sector, with Retail,

[19]To cross-validate the quality of our prediction exercise, in Figure 4B we show that workers who are predicted as not eligible for UI at the notch threshold exhibit no behavioral response.

[20]Our weight variable is the share of qualifying earnings that is observed directly in the data. To avoid biases from changing weights, we treat the weights as fixed over time.

Health Care and Professional Services being the next most represented sectors in the data.

Table 1: Descriptive statistics

	Mean	S.D.	p50	p25	p75
Panel A: Workers					
Weekly earnings	634.08	332.70	600.01	431.25	780.91
Weekly hours	35.97	13.05	38.19	28.29	44.24
Hourly wage	17.86	7.71	16.00	14.00	19.99
Number of weeks in data per worker	23.35	7.22	26.00	18.00	30.00
Observations	119020				
Number of workers	6861				
Panel B: Firms					
Size	14.18	22.11	8.58	4.58	16.68
Share of salaried workers	0.10	0.16	0.03	0.01	0.13
Number of Branches	1.16	0.70	1.00	1.00	1.00
Food and Drink	0.36	0.48	0.00	0.00	1.00
Retail	0.15	0.35	0.00	0.00	0.00
Health Care and Fitness	0.12	0.32	0.00	0.00	0.00
Professional Services	0.04	0.20	0.00	0.00	0.00
Observations	55646				
Number of firms	2771				

5 Estimation Strategy

We implement the WTP approach with the notch created by the US Federal Pandemic Unemployment Compensation (FPUC). A worker is eligible if earnings fall below a state-specific threshold and by moving across the threshold the worker loses \$600. Our identification stacks 21 difference in difference (DiD) analysis across US states. While FPUC was introduced uniformly in all US states, the administration of the benefits was left up to states and states applied different eligibility thresholds. Figure 3 shows the variation across states. A worker earning \$400 would be eligible for benefits in California and South Carolina, but not in Arizona or Florida.

A key identification challenge is to isolate variation that distinguishes labor supply responses from the economic crisis and lockdown restrictions that coincide with the launch of FPUC. A

Figure 3: Notch point by state

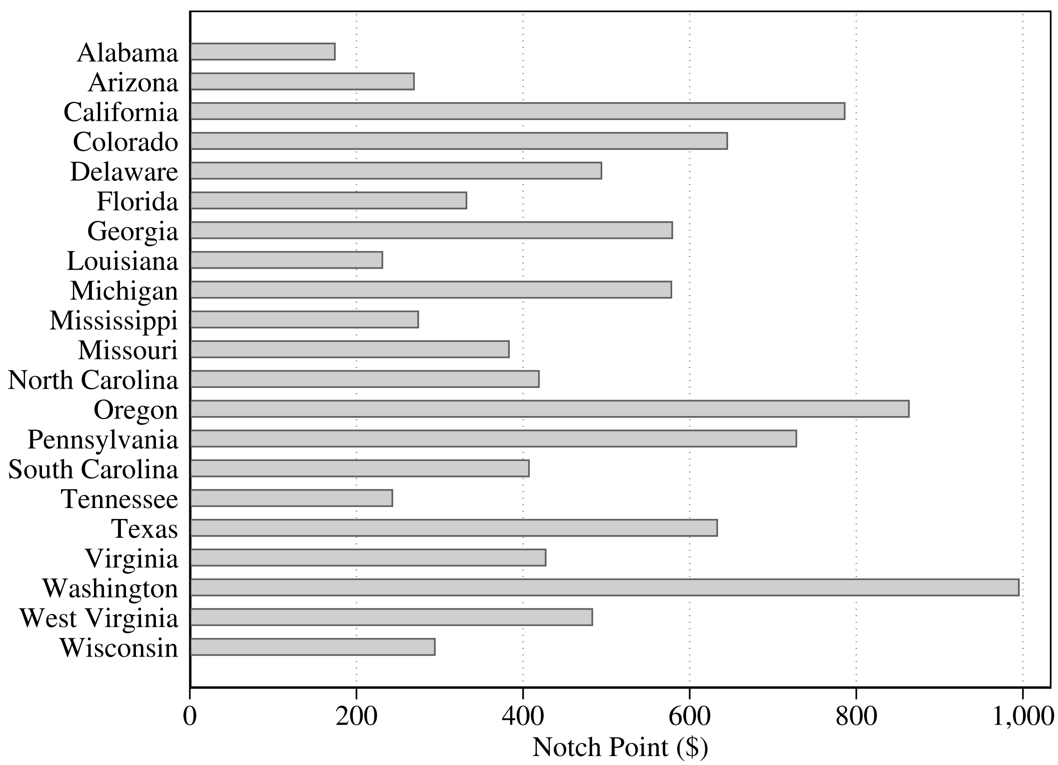

Note: The Figure shows maximum allowable earning while receiving FPUC payments for MWB recipients across US states.

standard DiD regression controls for aggregate fluctuations with time fixed effects and we also use such fixed effects. In addition, one may worry that the recession has different impacts on high- and low-income workers. To capture such effects, we allow time fixed effects to vary by income bins. This is possible here because states have different eligibility thresholds. Two equally paid workers are thus eligible for FPUC and hence treated in some states but not in others. Our analysis compares individuals with identical earnings (held constant with fixed effects for \$100 income bins), who happen to fall on different sides of their state's eligibility threshold. While the two people may hold identical positions, they face very different labor supply incentives, and reducing earnings may be beneficial for the person who is currently ineligible for FPUC.

We estimate the following DiD specification:

$$E_{w,t,m,r} = \pi_{m,t} + \sum_{k=-650}^{1300} \beta_r \cdot I[r = k] + \sum_{k=-650}^{1300} \delta_r \cdot I[r = k] \cdot C_t + \varepsilon_{w,t,m,r} \tag{2}$$

where $E_{i,t,m,r}$ is a dummy with value 100 if a worker w is employed in income range m, in week t, \$$r$ from the UI eligibility threshold, C_t is an indicator with value 1 after the launch of FPUC. $\pi_{m,t}$ are time fixed effects that vary by \$100 income bins and FPUC. Instead of a single eligibility indicator, we use finer dummies that capture the distance to the eligibility threshold, as theory would predict that responses are starkest close to the eligibility threshold and weaker further away from the threshold. $I[r = k]$ is an indicator that takes value 1 if income is \$$k$ from the UI eligibility threshold. β_r captures excess and missing mass around the eligibility threshold *before* FPUC and δ_r captures the same *after* the introduction of FPUC. Given the controls for absolute income levels with $\pi_{m,t}$, δ_r captures differences in behavior of individuals with identical income, say \$300, but on different sides of the eligibility thresholds. Finally, notice that this set-up turns into stacked difference in difference regressions if $\pi_{m,t} = \pi_t$.

Our baseline results focus on intensive margin responses[21] and restricts the sample to obtain a balanced number of work-week observations before and after FPUC for each worker. This "bal-

[21]Cross-sectional studies require this restriction as an identifying assumption. In our difference in difference regression, we could, in principle, relax this assumption and estimate extensive margin responses with the caveat that we cannot distinguish exits from our data-set from exits from the labor force.

anced" sample has two main advantages. First, the spurious fluctuations in the workforce size do not drive our findings and it alleviates the impact of demand shocks since shop closures are not impacting the analysis.[22] Second, restricting the sample to a balanced number of observations before and after FPUC guarantees that any missing mass shows up as excess mass elsewhere in the distribution and excess and missing mass sum to zero by construction. While the sample restriction is not strictly needed, it simplifies both the estimation strategy and the link of the estimates to the theoretical framework.

We strengthen the identification strategy in two ways. First, we run a placebo test with workers ineligible for FPUC. The group faces the same labor market shocks but their work incentives are unchanged by FPUC. We can thus check if there are spurious shocks that generate the observed patterns. Second, we narrow in on counties at state borders. This border design has been used to study the effect of minimum wage and control even more flexibly for demand shocks. Many of these border communities have integrated labor markets and thus again share many of the shocks. Our regressions compare people with identical incomes but work on different sides of a state border and thus have different UI eligibility.

To implement our WTP approach, we additionally require estimates of the change in behavior at different levels of Covid-19 risk. For these results, we simplify equation 3 and summarise the average excess/missing mass with a single coefficient for a \$400 treatment window around the threshold (results with alternative window sizes are reported in Appendix A2). To implement this, we replaces the granular $I[r = k]$ bins in 3 with a categorical variable ($W_{r,m}$) that takes value 0 outside the treatment window, and inside the window takes value 1 to the left of the threshold (excess mass), and value -1 to the right of the threshold (missing mass). For simplicity, we will refer to this as the excess mass, although the coefficient captures both excess and missing mass effects. The resulting simplified version of the baseline model in 2 is given by:

[22] For individuals with temporary absences (e.g. sickness or holidays), the active weeks are included in the sample, and we reduce the number of weeks before and after FPUC to maintain a balanced number of work-week observations before and after FPUC.

$$E_{i,t,m,r,\theta} = \pi_{m,t} + \beta_{r,\theta} \cdot W_{r,m} + \delta_{r,\theta} \cdot W_{r,m} \cdot C_t + \varepsilon_{i,t,m,r,\theta} \tag{3}$$

the estimation coefficients ($\delta_{r,\theta}$ and $\beta_{r,\theta}$) are allowed to vary with the risk level θ and changes in the coefficients capture variation in excess mass with Covid-19 risk.

6 Behavioral response to FPUC

Our first set of results documents the intensive margin labor supply response to the sudden $600 increase in the unemployment benefits available to partially unemployed workers. Figure 4A plots the differences in mass before and after FPUC introduction for each $50 bin around each state-specific notches. All notches are normalized to zero and bins are defined relative to them. The figure shows a sharp response to the increase in UI generosity. Workers move from relative wage bins above the notch to bins below it. The magnitude of these effects is substantial with a missing mass of almost 3 percentage points for the $50 bins with the largest drop (i.e., the bin "notch+$250").[23]. This corresponds to a 33% decrease in frequency relative to a baseline frequency of around 9%.

A number of features of Figure 4A are noteworthy. First, our results show excess and missing mass spread out over broader income ranges around the eligibility threshold instead of a single spike right at the threshold. There are multiple reasons for this. First, adjustment frictions in scheduling prevent workers from freely choosing their earnings, leading to changes to a broader range of income levels. Second, we do not observe the exact incomes used to determine UI eligibility and right at the threshold we may thus miss-classify the treatment status of workers.[24] Third, such behavior is consistent with income effects from the FPUC benefits.

For these reasons, estimates of labor supply elasticities that ignore the excess mass to the left of the threshold may underestimate the behavioral response. Our DiD approach allows us to identify effects over broader income ranges and can thus capture impacts more broadly. The

[23] 9% of all workers in our analytical sample used to work in the bin "notch+250" before the start of the pandemic

[24] Differences in observed earnings and UI relevant earnings arise in some jurisdictions from allowances for families or special circumstances or from multiple jobholders. As described above, we down-weight cases where we cannot measure benefit-relevant incomes well.

Figure 4: Excess and missing mass around the partial UI notch

Panel A: FPUC eligible workers

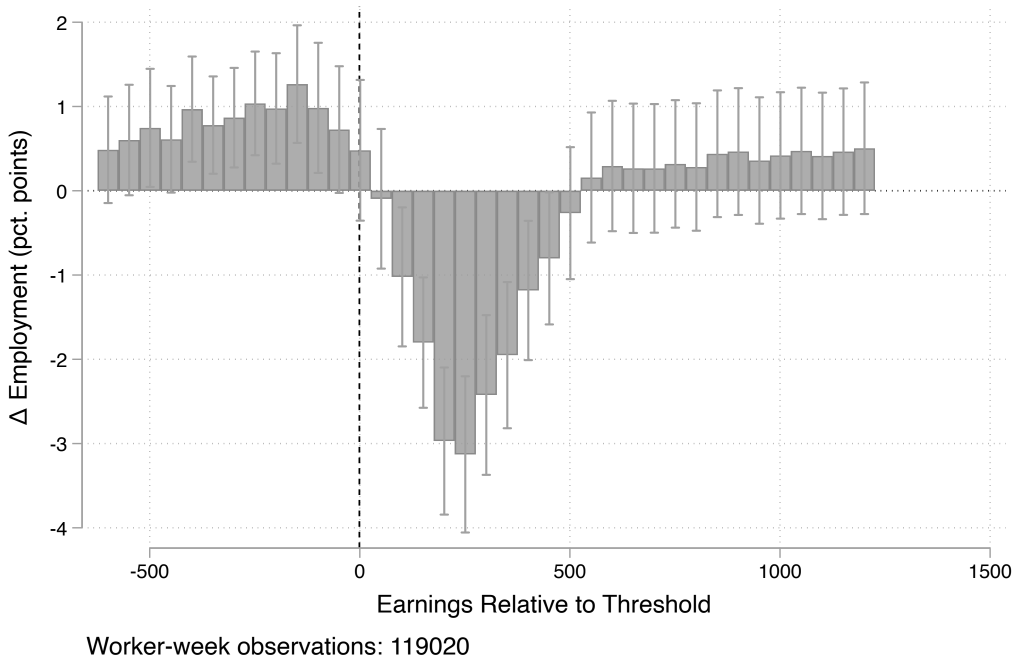

Panel B: FPUC ineligible workers

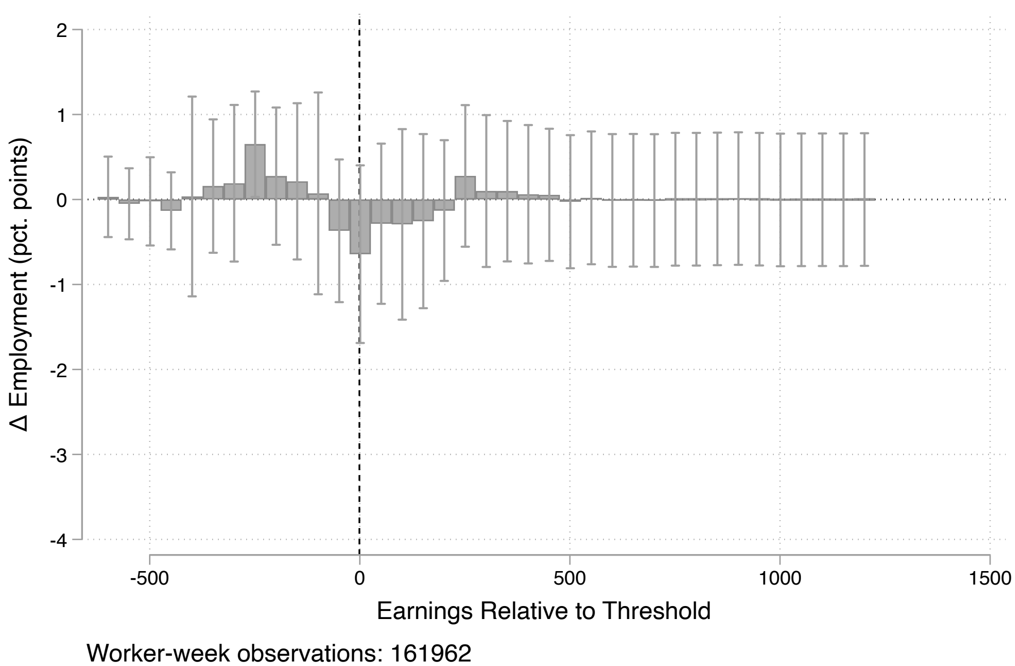

Note: The Figure shows δ_r coefficients from the equation 2. Standard errors are clustered at the state, income bin, week level and 95 percent confidence intervals are reported. In panel a, the sample is hourly workers with sufficient past earnings to qualify for MWB payments in their home state. In panel B the sample is instead hourly workers with insufficient past earnings to qualify for MWB payments in their home state.

variation in benefit eligibility is indeed cleaner slightly away from the threshold and in line with this, the treatment effects are most pronounced as we move slightly away from the threshold.[25] It is also important to note that the exact value of the elasticity is not needed for our WTP approach and the approach can accommodate a wide range of adjustment frictions. Our WTP estimate is based on the ratio of behavioural responses in high vs low-risk states and all factors that impact these responses proportionally will cancel out. For instance, if adjustment frictions attenuate labor supply responses by x%, estimates of labor supply elasticities would be downward biased but our approach would still produce unbiased WTP estimates.

A potential concern with the setup is that we pick up differential labor demand shocks, even after controlling for earning level-specific time effects. We produce a range of robustness checks to analyze this possibility.

First, we repeat the analysis for workers who are ineligible for the benefit. This group has no incentives to respond to the benefit eligibility thresholds and we can thus use this group for a placebo test. The results are shown in Figure 4B, which plots the behavioral response around the eligibility threshold for ineligible workers. The effects are insignificant and small in magnitude, confirming that there are no spurious shocks and that our baseline findings for eligible workers reflect responses to the benefit threshold.

Second, in the first row of Table 2 we estimate our baseline specification from equation 3 and then gradually introduce additional controls to capture potential spurious shocks affecting different income ranges in different states in a way that correlates with the state-specific eligibility rules (e.g. one may worry that our results pick up the general decline in wages). In practice, we interact state dummies with a continuous earning variable in the Covid-19 period. These income trends capture broader state-specific shifts in the earning distribution. At the same time, we can still identify our effect of interest through local shifts in the earning distribution around the FPUC threshold. The estimation results remain close to the baseline (approximately 1 percentage point) and thus

[25]In Figure A2 we test the sensitivity of our DiD estimate to changing treatment windows around the threshold. Our estimate is statistically significant already if we consider a window of $150 around the threshold. The dominated window runs up to $600 above the threshold, however estimates obtained with a $600 window are robust to smaller windows (i.e. for any threshold above $250).

confirm that our findings are orthogonal to state-specific demand shocks (column 2). Next, we allow for even more local policy shocks and repeat the exercise with county-level controls and again find that such controls do not affect our results (column 3). Finally, we control for industry and individual-specific shocks and again find similar results to the baseline (columns 4 and 5).

To provide further evidence of the absence of potential demand-driven confounding factors, we narrow our analysis to adjacent counties sharing a state border. In such a setting, empirical identification relies on comparing equally-paid workers across state borders with different incentives: such workers are likely to face similar demand shocks, however, one might be eligible for UI while the other might not, simply because of differences in the pre-determined exogenous eligibility thresholds. Our data covers border stretches at 17 state-pair boundaries (see Figure A3). And results based on the border counties align with the baseline estimates, adding confidence that we are not picking up spurious demand effects (see Appendix D.2).

7 Willingness to pay for workplace safety

This section applies our magnified excess mass framework to an empirical context with changes to workplace safety from Covid-19.

7.1 Measuring Covid-19 Exposure

To implement this approach, we need a measure of workplace safety that is independent of local worker decisions. We construct such a measure of Covid-19 exposure by combining time-invariant pre-determined task risk information of industries with data on local outbreaks. We denote the time-invariant riskiness of industry i by T_i and the local fatality rate by $F_{c',t}$. Our measure of exposure is:

$$E_{i,c,t} = F_{c',t} \cdot T_i \tag{4}$$

Variable $F_{c',t}$ is the fatality rate in the counties c' adjacent to county c. By using adjacent counties, we can rule out potential reverse causality issues. Also, note that we focus on fatality rates –

rather than infection rates– as our measure of health risk because of the massive underestimation of infection rates due to limited testing capabilities during the first months of the pandemic. The industry risk T_i is obtained by combining task-specific infection risk data from Basso et al. (2020) with American Community Survey data on the distribution of tasks across industries, and then computing a risk score for each 3-digit industry.[26]

This task-based risk variation has the advantage to provide variation *within* local areas across industries and thus allows us to hold potential confounding local policies constant. In practice, our exogenous measure captures three variation dimensions: it compares two workers in a high vs low risk industry, under intense vs mild local outbreaks, to similar workers in non-outbreak areas. Of course, we do not simply compare these groups, but rather specific income ranges around the FPUC threshold, which isolate labor supply responses. In appendix D.3 we present two alternative measures of Covid-19 exposure (the simple cross-industry variation T_i and variation in local Covid-19 outbreaks within each county) and show estimates of the labor supply response to their variation.

Since $E_{i,c,t}$ has no natural units, we normalise this variable to start at 0 and have a standard deviation of 1. This implies that regression coefficients on this variable capture the effect of one standard deviation increase in our exogenous measure of exposure. To provide an interpretable scale for this variation we calculate the relation between our exposure measure $E_{i,c,t}$ and actual fatality rates.[27] One standard deviation of workplace risk $E_{i,c,t}$ is equivalent to an increase in fatality rates by 31.15 cases per million workers. This variation is large, but overall comparable in magnitude to the pre-Covid-19 cross-occupation variation in fatality rates: one standard deviation in fatality rates across US occupations is 17.1 cases per million workers with the highest risk of

[26] Basso et al. (2020) use O*NET data to compute task specific risk measures based on proximity to others at work and the possibility of working remotely. The risk scores are reported at the occupation level and we compute industry averages by taking a employment weighted average of occupational risks in each industry.

[27] Data on weekly local industry specific death rates are not available. We therefore rely on county/week death counts ($F_{c,t}$) and compute the death in each industry based on the industries employment share ($\frac{e_{i,c}}{\sum_i e_{i,c}}$) and by fatality risks in industry i. For example, a worker in an industry with twice the fatality rate gets a weight of 2, and we therefore assign twice as many deaths to the industry relative to the average industry. This exercise requires data on industry specific fatality rates (w_i). Such data are not available at the national level and we instead use data from California, where such rates are published by Chen et al. (2021). Employment counts come from the ACS 2014-2018. Combining all these steps, our proxy for local industry specific fatality rate is $\theta_{i,c,t} = F_{c,t}\frac{e_{i,c} \cdot w_i}{\sum_i e_{i,c} \cdot w_i}$

121 for fishing and hunting workers.[28]

7.2 The value of workplace safety

For our WTP approach we require information on changes in labor supply with Covid-19 risk. We consider the differential response to different levels of risk by re-estimating the baseline DiD regression at different levels of the Covid-19 workplace risk shock. We split our exogenous Covid-19 exposure variable into five quintiles and estimate the responses separately for those five risk levels. Potential spurious effects from labor demand shocks were discussed in detail in the previous section and they are clearly relevant here again. Higher risk industries, such as beauty salons, experienced stronger demand shocks. We again leverage FPUC eligibility rules to separate labor supply behaviour from such demand shocks and measure excess mass around the FPUC eligibility threshold. Unlike the previous section, we now also allow labor supply responses to differ with workplace risk and study whether workplaces with the biggest increase in risk produce most excess mass.

Figure 5 shows the behavioral response and finds that excess mass does indeed increase with rising dis-amenities at work. The excess/missing mass in grey represents the behavioral response to FPUC among worker-week spells in the lowest risk quintile. We compare this baseline response to effects for spells facing higher risk levels. In the top left panel of Figure 5 the black area represents the response among workers in the second lowest risk quintile. Despite the relatively low risk, it is visible that such workers are responding more to the FPUC incentives. Workers appear to move down the earning distribution more than in the low risk settings. As before, we are controlling for demand shocks with income range specific time effects. The remaining three panels show how the response gets magnified at higher fatality rates. The excess mass of workers shifting to the left of the state-specific notches is particularly pronounced for the top risk quintiles, consistent with the model prediction that dis-amenities at work magnify the excess mass at budget notches.

[28]Source: BLS Census of Fatal Occupational Injuries (CFOI) - Current `https://www.bls.gov/iif/oshcfoi1.htm#2019`

Figure 5: Excess and missing mass around the partial UI notch by fatality rates

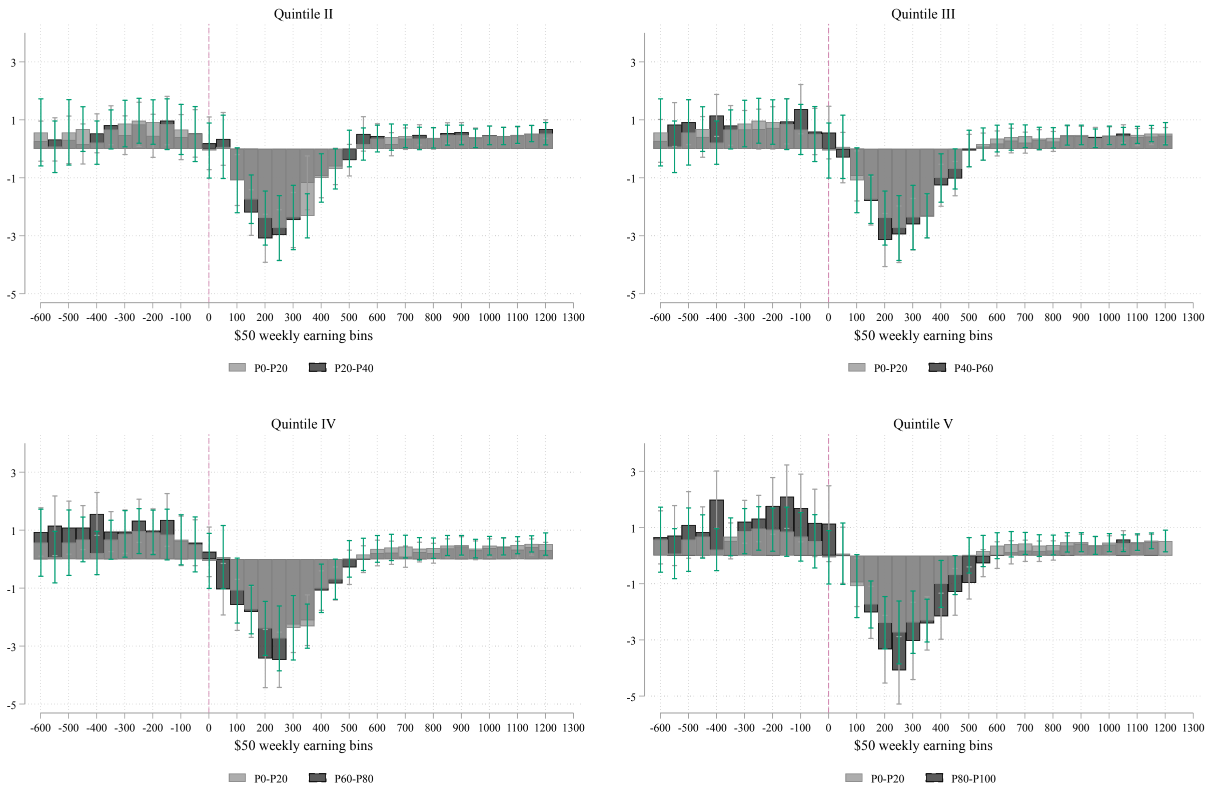

Note: The Figure shows $\delta_{r,\theta}$ coefficients from equation 3. Results for 5 quintiles of Covid-19 risk (θ) are plotted. The gray bars represent the response in the lowest quintile and black bars in the sub-panels respectively show responses in risk quintiles 2 to 5. The sample covers hourly workers with sufficient past earnings to qualify for MWB payments in their home state and is based on 119,020 work-week spells. Source: Homebase.

We next use these results to compute the willingness to pay (WTP) for workplace safety. Equation 1 shows that the WTP is given by the ratio of additional labor supply response created by workplace risk and the baseline response to FPUC. We quantify these responses in Table 2. Panel A shows the excess mass around FPUC at average risk levels. As discussed above, we find that FPUC creates an excess mass of around one percentage point in the income bins surrounding the earnings threshold. Panel B shows how this excess mass changes as workplace risks increase. A standard deviation increase in risk leads to 0.35 percentage points more excess mass and these effects are largely unaffected by the inclusion of controls. The implied willingness to pay is between 31% and 34% of weekly income, or around $216. Expressed in terms of fatality rates, this implies that workers are willing to pay around 1% of their income to cut weekly fatality rates by one per million (Panel C).

We then explore the robustness of these results. In columns 2 to 5 we repeat the robustness checks from above and introduce increasingly granular controls. First, we address worries about contemporaneous changes to local policies and include county or state FE that vary over time and exploit the cross-industry heterogeneity in $E_{i,c,t}$ (columns 2 and 3). Next, we study potential selection effects across industries. To affect our results, one would have to believe that industries with Covid-19 vulnerable tasks employ workers with larger labor supply elasticities. To explore this possibility, we allow for industry specific labor supply elasticities by letting the responses to FPUC vary by industry. Specifically, we interact industry dummies with the launch of FPUC and a continuous income variable (column 4). Next, we can control for individual heterogeneity at an even more granular level by including individual fixed effects. We go beyond standard time-invariant fixed effects and allow for time varying fixed effects to accommodate differences in workers' responsiveness to labor supply incentives. Specifically, we interact individual fixed effects with the launch of FPUC and income. Such regressions then study the behavior of the same individual in high vs low risk weeks. The fixed effects absorb heterogeneity in labor supply responsiveness and allow us to rule out that a rich set of unobserved individual heterogeneity are biasing the results (column 5). Throughout all robustness tests the results remain close to the baseline.

We next compare our estimates to results from a canonical hedonic wage analysis and estimate how hourly wages change with our measure of workplace risk. Individual fixed effects control for time-invariant worker ability and ensure that selection effects do not bias these results. We find that wages are unchanged and the coefficient on workplace risk is insignificant. The point estimate is also quantitatively small and suggests that wages increased by 11 cents, a 0.5% wage increase (results are not reported and are available upon request). The small coefficient is not surprising, considering that wages are slow to adjust and unlikely to fully price in changes in workplace risks. These estimates would lead us to conclude that workers attach next to no value to workplace safety. Our novel approach, by contrast, suggests that these low estimates are biased. Workers do indeed respond substantially to workplace risks and our revealed preference estimate implies a WTP that is an order of magnitude greater than the hedonic result.

Table 2: Willingness To Pay for Workplace Safety

	(1)	(2)	(3)	(4)	(5)
	Panel A: Baseline Excess Mass				
FPUC	1.044	1.044	1.044	1.044	1.044
	(0.102)	(0.104)	(0.104)	(0.102)	(0.104)
	Panel B: Additional Excess Mass				
Workplace Risk std. dev.	0.353	0.330	0.328	0.347	0.325
	(0.0565)	(0.0553)	(0.0554)	(0.0561)	(0.0555)
	Panel C: WTP (% of weekly income)				
Workplace Risk (std. dev.)	33.8	31.6	31.4	33.2	31.1
Workplace Risk (deaths per mio.)	1.1	1.0	1.0	1.1	1.0
	Panel D: Value of Statistical Life (million \$)				
VSL (perfect information)	\$ 6.89	\$ 6.44	\$ 6.40	\$ 6.78	\$6.34
VSL (actual information)	\$ 2.60	\$ 2.43	\$ 2.42	\$ 2.56	\$2.40
FE, interacted with income x time FE		state	county	industry	individual

Note: The Table shows how Covid-19 risk affects excess mass at the FPUC eligibility threshold. Panel A shows excess mass around the FPUC threshold for average risk from estimating equation 3 with $\delta_{r,\theta} = \delta_r$. Panel B, shows how excess mass changes with fatality rates ($\delta_{r,\theta}$). Willingness to pay in Panel C is based on equation 1, and is the ratio of panel B and panel A. Panel D computes $VSL = \frac{WTP}{\Delta fatality} * m$, where m is income. And one standard deviation of workplace risk increases fatality rates by 31.15 cases per million workers and believes about fatality rates by 82.33 cases per million workers. Controls are state, county and two digit NAICS fixed effects, interacted with a dummy for the Covid-19 period and a continuous income variable. The results are based on 119,020 worker-week spells. Source: Homebase, Chen et al. (2021).

A related appealing feature of our approach is that it allows for potential market failures in pricing risks into wages. During the Covid-19 crisis, many companies introduced hazard pay that aimed to compensate frontline workers for the added risks they faced. Critics of hazard rates argue that these rates were too low and did not fully compensate for large risk exposure. Our estimates shed light on this debate and quantify the non-pecuniary value of Covid-19 risk exposure. We do indeed find that hazard rates were lower than the cost of risk exposure. Our results suggest that hazard pay would need to be as high as 34% when risks increase by one standard deviation to fully offset the non-pecuniary costs of added workplace risk. In other words, workers were worse off at work during Covid-19, despite the introduction of hazard pay.

These results relate to the broader policy debate about workplace safety regulations. Addressing this issue has been challenging, in part because it is difficult to quantifying the gains from such non-wage regulations. The rationale for policy interventions is however similar to minimum wage regulation: in imperfect competition firms may not fully internalizing the cost of high-risk jobs and thus may expose workers to excessive risks. In practice, all governments implement some level of worker safety regulation, albeit with large differences in stringency and enforcement. In the design of such policies, the monetary value of improved workplace safety plays a central role and determines the welfare gains from such policies.

Our results suggest that workers value workplace safety highly and that the gains from more stringent safety regulations are substantial. Take for example the construction industry. Weekly fatality rates in this industry in the US are 3 workers per million full-time employees per week.[29] Our estimates suggests that eliminating fatality risks in construction would be equivalent to a 3% increase in wages. Reducing risk thus potentially offers large gains for workers. The reduction of fatality risks to zero is perhaps an unattainable target. However, even reducing fatality rates to the level seen in the UK or Germany would be a substantial improvement and equivalent to a wage increase of 2.5%.[30] The gains for workers would be comparable to an introduction of a

[29] ILO data is converted to weekly deaths per million workers for comparison. Annual fatality rates are 160 per million workers in 2018. Source: ILOSTAT, series "INJ FATL ECO RT A" 2018.

[30] ILO estimates for Germany and the UK are respectively 0.4 and 0.7 weekly deaths per million workers.

\$15 minimum wage.[31] Another useful point of reference are the wage gains implied by switching industries. The gains from greater safety by changing from the construction sector to the safer accommodation and food services sector are worth around 2.5% of income. Moving to the riskier agricultural sector is equivalent to a wage loss of 8%. The magnitude of these gains are comparable to the value of other work amenities analysed by Maestas et al. (2019), who find values ranging from 2% to 16%.[32]

A large related literature uses shocks to fatality risks in a variety of settings to estimate a "value of a statistical life" (VSL). Such estimates require additional assumptions: An important assumption of the VSL literature is that the fear of dying is the sole driver of the observed behaviour. Since higher fatality rates are typically accompanied by unpopular safety measures and imply a greater risk of non-fatal injuries, this assumption amounts to assuming that workers attach zero value to such aspects. If we are willing to make this assumption, we can compute VSL from the observed behaviour. The VSL is the ratio of WTP and the change in fatality risk: $VSL = \frac{WTP}{\Delta fatality}$. Using our estimates, we find $VSL = \frac{\$216}{31.15/1{,}000{,}000} = \$6.9mio$. A value of \$6.9 million broadly aligns with the literature, a recent meta-study by Viscusi (2018) concludes that VSL is somewhere between \$3 and \$13 million (in 2020 USD). Our results fall towards the middle of this range and suggests that our approach yields estimates with sensible magnitudes.

7.3 Imperfect information

VSL calculations have to make assumptions about worker believes and ideally, researchers would want to compute $VSL = \frac{WTP}{E[\Delta fatality]}$, where $E[\Delta fatality]$ is the workers' perception of fatality risk. Since these perceptions are not usually observed, studies instead use the statistical fatality rates as a measure of perception, thereby imposing perfect information and rational expectations. Kahneman and Tversky (1979) famously point out that these assumptions may fail. Individuals, for

[31]The minimum wage calculation computes the wage floor that is equivalent to a 2.5% mean wage increase (assuming no employment loss). The data source is the 2019 and 2020 CPS ASEC data.

[32]Maestas et al. (2019) study the value of schedule autonomy, telecommute, physical activity, sitting, relaxed work environment, work autonomy, PTO, team-work, training, opportunity to serve.

instance, appear to put disproportionate weight on salient issues in their decision-making (Bordalo, Gennaioli, and Shleifer (2013)). Accounting for the correct information set is thus essential to produce reliable VSL estimates. In our context, we have a rare opportunity to observe perceptions about fatality risks.

During the Covid-19 outbreak, beliefs about fatality risks were collected as part of the Understanding America Study (UAS). Individuals were asked about the probability of contracting Covid-19 and conditional on this, the probability of dying. The data covers a representative sample of the US population and uses weekly rounds of interviews. We use this data to compute expectations at the week-state-industry level and then use these to impute expectations for our sample. The expectation measure thus undoubtedly includes measurement error. Our approach can be thought of as an instrumental variable approach that instruments with our exogenous risk measure.

We find that one standard deviation in risk increases expectations of fatality risks by 82.52 deaths per million. The result is highly significant as local Covid-19 outbreaks lead to a sharp increase in the beliefs about the risk of dying. It is noteworthy that the perceived fatality risk increases by more than the actual fatality risk. People thus overreact relative to a perfect information, rational agent setting. This seems highly plausible, since there was extensive news coverage of Covid-19 deaths, making such risks extremely salient. Moreover, little was known about the true fatality rates. Consistent with the behavioral literature on "salience" (Bordalo, Gennaioli, and Shleifer (2013)), we find that workers who faced this uncertainty over-emphasize the very salient fatality rates. Finally, we use the perception results in the VSL computation and find a VSL estimate of $2.57 million (Panel D of Table 2), about a third of the rational expectation estimate. The VSL estimate is reduced because perceived fatality rates enter the VSL calculation as the denominator and replacing rational with actual expectation thus increase the denominator of the VSL calculation. In other words, the same behavioral response is produced by a larger shock to risk perceptions than the perfect information benchmark suggests. A large and growing behavioral literature on belief formation shows similar deviations from the perfect information benchmark in numerous settings. Our result highlights that taking account of such deviations has important implications for VSL estimates. Moreover, persistent high perceptions of risk can potentially explain why workers have been reluctant to return to work and hence why labor supply remained lower when Covid-19 rates fell.

8 Conclusions

This paper uses a revealed preferences approach to measure the value workers attach to safe workplaces. We first present a new method to identify WTP based on excess mass around budget discontinuities. The labor supply response created by these discontinuities increases with injury risk at work. A large WTP is thus reflected in a strong increase in excess mass when risks increase.

We apply this framework to estimate the value attached to fatality risk. The introduction of FPUC introduces eligibility thresholds that produce notches in worker's budget constraints. We first evaluate the labor supply response to these notches and find substantial labor supply responses.

We next estimate how these labor supply responses change with Covid-19 risk. Our results show that Covid-19 risk leads to additional responses. The resulting estimates imply that workers are willing to sacrifice 34% of their weekly disposable income to avoid a standard deviation in risk. This is equivalent to giving up 1% of income to avoid a one-in-a-million risk of dying.

Such WTP estimates have often been used to infer values of a statistical life (VSL). We show that the implied value of VSL depends crucial on how risk is measured. The perceived fatality rate is substantially higher than the observed one and estimates based on rational expectation will thus overstate VSL. This result illustrates that measuring individual information sets accurately is essential for credible VSL estimates.

References

Altonji, Joseph G and Christina H Paxson. 1992. "Labor supply, hours constraints, and job mobility." *Journal of Human Resources* :256–278.

Ashenfelter, Orley and Michael Greenstone. 2004. "Using Mandated Speed Limits to Measure the Value of a Statistical Life." *Journal of Political Economy* 112 (S1):S226–S267. URL `http://www.jstor.org/stable/10.1086/379932`.

Basso, Gaetano, Tito Boeri, Alessandro Caiumi, and Marco Paccagnella. 2020. "The new hazardous jobs and worker reallocation." .

Bollinger, Christopher. 1998. "Measurement Error in the CPS: A Nonparametric Look." *Journal of Labor Economics* 16 (3):576ff.

Bonhomme, Stephane and Gregory Jolivet. 2009. "The pervasive absence of compensating differentials." *Journal of Applied Econometrics* 24 (5):763–795.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121 (5):803–843.

Bound, John and Alan B. Krueger. 1991. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics* 9 (1):1–24.

Brown, Charles. 1980. "Equalizing differences in the labor market." *The Quarterly Journal of Economics* 94 (1):113–134.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128 (August):967–1015.

Chen, Yea Hung, Maria Glymour, Alicia Riley, John Balmes, Kate Duchowny, Robert Harrison, Ellicott Matthay, and Kirsten Bibbins-Domingo. 2021. "Excess mortality associated with the COVID-19 pandemic among Californians 18-65 years of age, by occupational sector and occupation: March through November 2020." *PLoS ONE* 16 (6 June):1–10. URL `http://dx.doi.org/10.1371/journal.pone.0252454`.

Chetty, Raj, John N. Friedman, and Emmanuel Saez. 2013. "Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings." *American Economic Review* 103 (7):2683–2721.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2018. "The Economic Consequences of Hospital Admissions." *American Economic Review* 108 (2):308–352.

Fishback, Price and Shawn Everett Kantor. 1995. "Did Workers Pay for the Passage of Workers' Compensation Laws?" *The Quarterly Journal of Economics* 110 (3):713–742.

Flory, Jeffrey A, Andreas Leibbrandt, and John A List. 2015. "Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions." *The Review of Economic Studies* 82 (1):122–155.

Goldin, Claudia and Lawrence F. Katz. 2011. "The Cost of Workplace Flexibility for High-Powered Professionals." *Annals of the American Academy of Political and Social Science* 638 (1):45–67.

———. 2016. "A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation." *Journal of Labor Economics* 34 (3):705–746.

Gruber, Jonathan. 1994. "The Incidence of Mandated Maternity Benefits." *America Economic Review* 84 (3):622–641.

Gruber, Jonathan and Alan B. Krueger. 1991. "The Incidence of Mandated Employer-Provided Insurance: Lessons from Workers' Compnesation." *Tax Policy and the Economy* 5:111–143.

Guardado, José R. and Nicolas R. Ziebarth. 2019. "Worker Investments in Safety, Workplace Accidents, and Compensating Wage Differentials." *International Economic Review* 60 (1):133–155.

Hwang, Hae-shin, W Robert Reed, and Carlton Hubbard. 1992. "Compensating wage differentials and unobserved productivity." *Journal of Political Economy* 100 (4):835–858.

Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." URL `http://www.jstor.org/stable/1914185{%}5Cnhttp://www.jstor.org/{%}5Cnhttp://www.jstor.org/action/showPublisher?publisherCode=econosoc.{%}5Cnhttp://www.jstor.org`.

Kinder, Molly, Laura Stateler, and Julia Du. 2020. "Windfall Profits and Deadly Risks." *Brookings* .

Kleven, Henrik J. and Waseem Mazhar. 2013. "Using Notches To Uncover Optimization Frictions And Structural Elasticities: Theory and Evidence from Pakistan." *The Quarterly Journal of Economics* :669–723.

Kleven, Henrik Jacobsen. 2016. "Bunching." *Annual Review of Economics* 8 (1):435–464. URL `https://doi.org/10.1146/annurev-economics-080315-015234`.

Lavetti, Kurt. 2020. "The Estimation of Compensating Wage Differentials: Lessons From the Deadliest Catch." *Journal of Business and Economic Statistics* 38 (1):165–182.

Lee, Jonathan M. and Laura O. Taylor. 2019. "Randomized safety inspections and risk exposure on the job: Quasi-experimental estimates of the value of a statistical life." *American Economic Journal: Economic Policy* 11 (4):350–374.

Lucas, Robert E B. 1977. "Hedonic Wage Equations and Psychic Wages in the Returns to Schooling." *American Economic Review* 67 (4):549–558.

Maestas, Nicole, Kathleen Mullen, David Powell, Till von Wachter, and Jeffrey Wenger. 2019. "The Value of Working Conditions in the United States and Implications for the Structure of Wages." .

Mas, Alexandre and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *American Economic Review* (1454476). URL `https://www.ssrn.com/abstract=2837954`.

Masters, Stanley. 1969. "An Interindustry Analysis of Wages and Plant Size." *The Review of Economics and Statistics* 51 (3):341–345.

Morchio, Iacopo and Christian Moser. 2019. "The Gender Gap: Micro Sources and Macro Consequences." *SSRN Electronic Journal* .

Robinson, Lisa A. and James K. Hammitt. 2011. "Behavioral Economics and the Conduct of Benefit-Cost Analysis: Towards Principles and Standards." *Journal of Benefit-Cost Analysis* 2 (2):1–51.

Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1):34–55.

———. 1986. "The theory of equalizing differences." *Handbook of Labor Economics* 1 (C):641–692.

Ruppert, Peter, Elena Stancanelli, and Etienne Wasmer. 2009. "Commuting, Wages and Bargaining Power." *Annals of Economics and Statistics/Annales d'Économie et de Statistique* :201–220.

Sorkin, Isaac. 2018. "Ranking firms using revealed preference." *Quarterly Journal of Economics* 133 (3):1331–1393.

Taber, Christopher and Rune Vejlin. 2020. "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market." *Econometrica* 88 (3):1031–1069.

Viscusi, W. Kip. 1990. "Do Smokers Underestimate Risks?" *Journal of Political Economy* 98 (6):1253–1269.

———. 2018. "Best Estimate Selection Bias in the Value of a Statistical Life." *Journal of Benefit-Cost Analysis* 9 (2):285–304.

Viscusi, W Kip and Joseph E Aldy. 2003. "The value of a statistical life: a critical review of market estimates throughout the world." *Journal of risk and uncertainty* 27 (1):5–76.

Wiswall, Matthew and Basit Zafar. 2017. "Preference for the Workplace, Investment in Human Capital, and Gender*." *The Quarterly Journal of Economics* 133 (1):457–507. URL `https://doi.org/10.1093/qje/qjx035`.

A Appendix Figures and Tables

Figure A1: Scheduling app screenshot

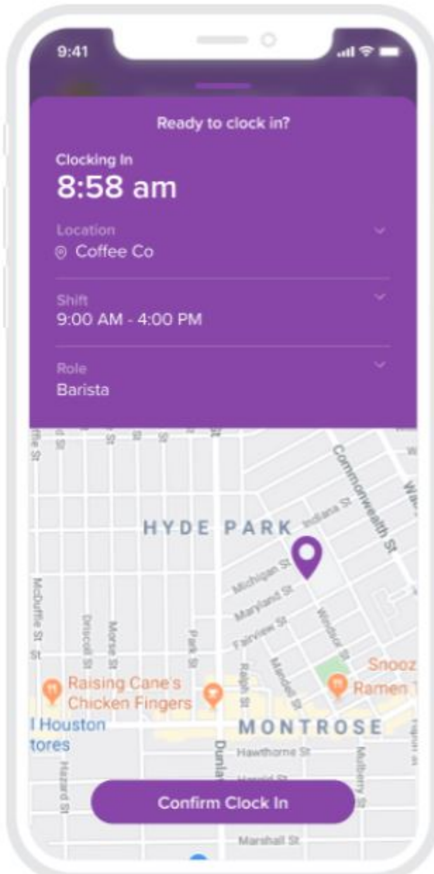

Figure A2: Effect of FPUC with Alternative Treatment Window

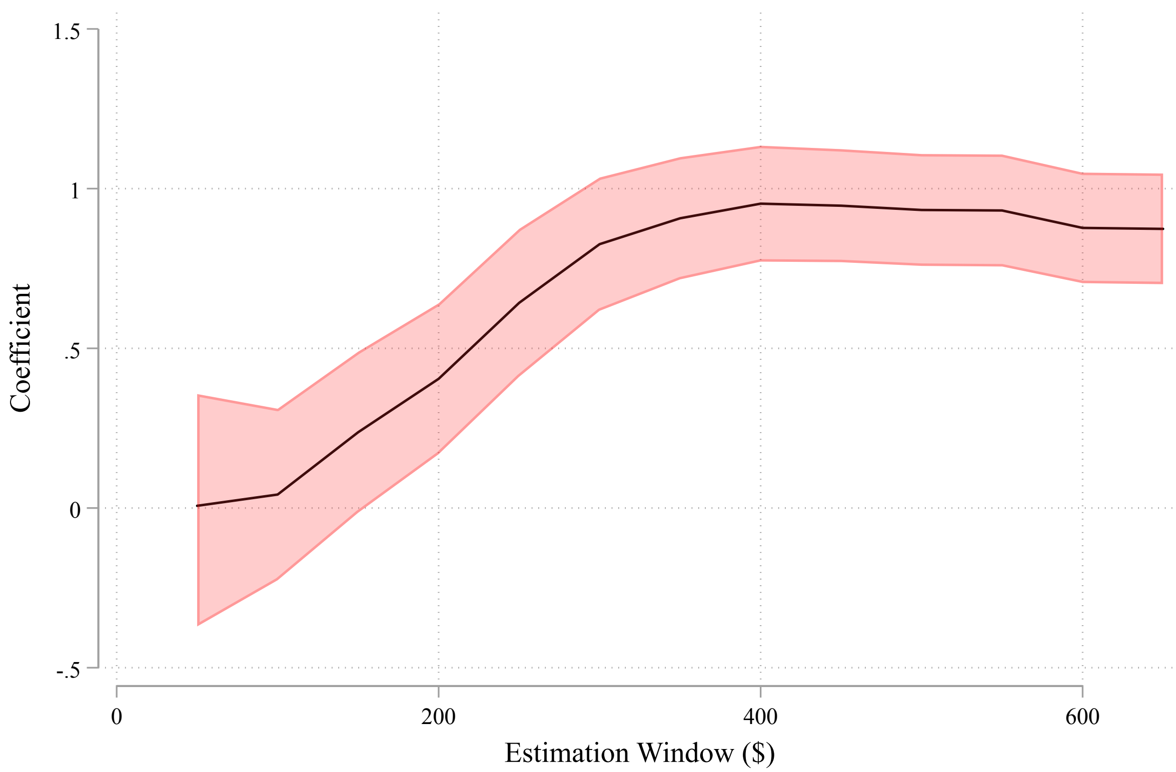

Note:

B Homebase data generability

Table A1 presents summary statistics for wages, weekly earnings, and hours worked across Annual Social and Economic (ASEC) supplement to the Current Population Survey (CPS), Homebase data, and Quarterly Workforce Indicators (QWI). Summary statistics reported for ASEC combine 2019 and 2020 ASEC supplements. Column (1) presents summary statistics without any restrictions; column (3) restricts ASEC data to the 21 Homebase states; and column (4) applies further restrictions that allow for comparability of ASEC to the Homebase sample. Specifically, the restricted ASEC sample comprises of hourly workers, not self-employed, in small-businesses that correspond to Homebase North American Industry Classification System (NAICS) codes in Homebase states.[33]

Homebase provides 6 digit NAICS codes but ASEC does not provide an industry classification that uses NAICS. Therefore, the industry classification in ASEC is first crosswalked to NAICS using the crosswalk provided by IPUMS.[34] Next, the ASEC sample is restricted to Homebase NAICS codes in a step-by-step manner: if an ASEC industry is linked to a 6-digit NAICS code, it is classified as in the Homebase sample only if it matches a 6-digit Homebase code, and it is classified as not in the sample if it does not match any 6-digit Homebase code. Next, if an ASEC industry is linked to a 5-digit NAICS code, it is classified as in the Homebase sample if it matches the first 5 digits of a 6-digit Homebase NAICS code. This process is repeated until all ASEC NAICS codes are classified and the resulting crosswalk is used to restrict ASEC in column (4).

Column (2) presents Homebase summary statistics without any restrictions, inclusive of 2019 and 2020. Column (5) restricts Homebase to the study sample that comprises of individuals eligible for full UI benefits with a balanced number of week spells before and after FPUC. Appropriate survey weights are applied for each ASEC and Homebase summary statistic.

Column (6) presents QWI summary statistics in 2019, restricted to privately owned small firms with fewer than 20 employees in 21 Homebase states.

Table A2 lists the distribution of observations across different 2-digit NAICS sectors. The same restrictions are applied as in Table A1.

[33]Workers are classified as hourly according to the paidhour variable. Firmsize variable is used to restrict the sample to small businesses with fewer than 25 employees. The classwkr variable is used to remove self-employed workers from the sample.

[34]See `https://usa.ipums.org/usa/volii/indtoindnaics18.shtml` and `https://www2.census.gov/programs-surveys/cps/methodology/Industry%20Codes.pdf`.

Table A1: Summary statistics: hourly wages, weekly earnings, and hours worked

	(1) ASEC Full	(2) HB Full	(3) ASEC HB States	(4) ASEC Sample	(5) HB Sample	(6) QWI Sample
Hourly wage	18.69	12.41	18.46	16.68	17.86	
	(10.84)	(4.846)	(10.66)	(9.101)	(7.710)	
Weekly earnings	1016.7	377.6	999.6	631.8	634.1	805.5
	(724.4)	(244.6)	(716.1)	(432.3)	(332.7)	(328.4)
Hours usually worked per week	39.25		39.32	35.78		
at all jobs	(11.30)		(11.11)	(11.00)		
Hours usually worked per week	38.55	29.58	38.66	35.13	35.97	
at main job	(10.84)	(13.27)	(10.69)	(10.60)	(13.05)	
Hours worked last week	38.45		38.49	34.84		
	(12.80)		(12.63)	(12.06)		

Mean coefficients and standard errors in parentheses

ASEC and HB Full data include 2019 and 2020. QWI data is 2019 only.

(3) ASEC is restricted to the 21 HB states.

(4) ASEC sample is restricted to hourly workers, not self-employed, in small businesses (< 25 employees) corresponding to HB NAICS codes in HB states.

(5) HB sample is restricted to individuals eligible for full UI benefits (defined as meeting state specific earning requirements in previous quarters) with balanced number of week spells before and after FPUC.

(6) QWI sample is restricted to privately owned small firms ($<$ 20 employees) in HB states.

(6) Weekly earnings are calculated from beginning-of-quarter employment average monthly earnings. Monthly averages are divided by 4.345 to get weekly averages.

Table A2: Weighted number of observations across NAICS 2

	(1) ASEC Full	(2) HB Full	(3) ASEC HB States	(4) ASEC Sample	(5) HB Sample	(6) QWI Sample
11 Agriculture						
Average obs	2,494,551	16,975	1,627,307	258,873	218	309,331
Percent	1.52	0.33	1.60	2.29	0.25	2.20
21 Mining						
Average obs	811,883	21	565,338	12,832		45,202
Percent	0.50	0.00	0.55	0.11		0.32
22 Utilities						
Average obs	1,376,672	54	880,748	45,739	2	15,133
Percent	0.84	0.00	0.86	0.41	0.00	0.11
23 Construction						
Average obs	11,570,959	74,571	7,542,626	1,738,690	3,345	1,550,910
Percent	7.06	1.45	7.40	15.41	3.89	11.05
31–33 Manufacturing						
Average obs	16,094,836	36,759	9,863,228	323,252	1,137	684,610
Percent	9.82	0.71	9.67	2.86	1.32	4.88
42 Wholesale Trade						
Average obs	3,542,738	139	2,255,483	67,278		643,128
Percent	2.16	0.00	2.21	0.60		4.58
44–45 Retail Trade						
Average obs	17,088,628	697,119	10,869,406	1,211,602	14,354	1,432,060
Percent	10.43	13.53	10.66	10.74	16.68	10.20
48–49 Transportation						
Average obs	7,898,989	53,642	5,011,724	553,461	1,473	364,792
Percent	4.82	1.04	4.92	4.90	1.71	2.60
51 Information						
Average obs	2,981,793	22,408	1,819,702	90,754	249	162,600
Percent	1.82	0.43	1.78	0.80	0.29	1.16
52 Finance&Insurance						
Average obs	7,756,432	9,152	4,427,566	155,808	286	414,341
Percent	4.73	0.18	4.34	1.38	0.33	2.95
53 Real Estate						
Average obs	3,345,098	16,122	2,187,264	166,022	931	444,381
Percent	2.04	0.31	2.15	1.47	1.08	3.17
54 Professional Serv.						

Average obs	13,114,794	113,104	8,199,848	674,572	3,770	1,486,339
Percent	8.00	2.19	8.04	5.98	4.38	10.59
55 Management						
Average obs	149,666	68,697	102,192		2,351	36,842
Percent	0.09	1.33	0.10		2.73	0.26
56 Admin.&Support						
Average obs	7,060,830	52,337	4,704,635	782,275	2,317	772,320
Percent	4.31	1.02	4.61	6.93	2.69	5.50
61 Educational Serv.						
Average obs	15,165,151	77,058	8,954,306	469,870	1,695	216,475
Percent	9.25	1.50	8.78	4.16	1.97	1.54
62 Health Care						
Average obs	22,289,792	243,838	13,349,828	1,630,383	9,837	2,271,235
Percent	13.60	4.73	13.09	14.45	11.43	16.18
71 Arts,Entertain.						
Average obs	3,767,820	194,149	2,311,177	355,276	2,953	292,277
Percent	2.30	3.77	2.27	3.15	3.43	2.08
72 Accomm.&Food						
Average obs	11,943,838	3,209,476	7,647,894	1,820,426	34,866	1,648,446
Percent	7.29	62.28	7.50	16.13	40.52	11.74
81 Other Services						
Average obs	7,839,014	263,762	4,978,586	801,013	6,096	1,249,554
Percent	4.78	5.12	4.88	7.10	7.09	8.90
92 Public Admin.						
Average obs	7,600,328	1,947	4,648,719	126,967	161	0
Percent	4.64	0.04	4.56	1.13	0.19	0.00
76 Misc. Repair						
Average obs		2,233				
Percent		0.04				
Total						
Average obs	163,893,812	5,153,563	101,947,572	11,285,090	86,041	14,039,978
Percent	100.00	100.00	100.00	100.00	100.00	100.00

ASEC and HB Full data include 2019 and 2020. QWI data is 2019 only.

(3) ASEC is restricted to the 21 HB states.

(4) ASEC sample is restricted to hourly workers, not self-employed, in small businesses (< 25 employees) corresponding to HB NAICS codes in HB states.

(5) HB sample is restricted to individuals eligible for full UI benefits (defined as meeting state specific earning requirements in previous quarters) with balanced number of week spells before and after FPUC.

(6) QWI is avg. beginning of quarter employment in privately owned small firms (< 20 employees)

restricted to HB states.

C FPUC

Federal Pandemic Unemployment Compensation (FPUC), the weekly $600 supplement to unemployment benefits, was introduced by the CARES act enacted on March 27, 2020, and ended on July 31, 2020.[35] No FPUC benefits were payable between July 31, 2020 and December 26, 2020. FPUC was re-established by the Continued Assistance Act as a $300 per week supplement to unemployment benefits from December 26, 2020 to March 14, 2021.[36] American Rescue Plan Act extended FPUC through September 6, 2021.[37]

Any individual eligible to receive at least $1 of state unemployment benefits is also eligible to receive federally-funded FPUC for that week. Individuals who are working part-time and who fulfill state eligibility requirements for partial UI benefits are also eligible to receive FPUC payments.[38]

FPUC payments are federally funded; however, states can opt out of participating in the program. As of date, 7 of the 21 study states are planning to terminate some or all federally funded pandemic unemployment compensation programs early, citing labor supply shortages. FPUC will terminate on June 12, 2021 in Mississippi and Missouri; on June 19, 2021 in Alabama; on June 26, 2021 in Georgia; on June 30, 2021 in South Carolina; on July 3, 2021 in Tennessee; and on July 10, 2021 in Arizona.[39]

During the gap in FPUC payments, from August 1, 2020, Lost Wages Assistance (LWA) program was funded through Federal Emergency Management Agency (FEMA). States had the option of choosing between two weekly benefits amounts, $300 or $400, with different cost-sharing requirements. The $400 weekly benefit required the state to contribute $100 (25% of the benefit). The $300 weekly benefit was funded entirely by FEMA and states would satisfy the 25% match, without additional state pay-out, if the state funding for regular state UI benefits at the aggregate level amounted to at least 25% of total LWA benefits paid.[40] All 21 study states were approved for LWA, but only West Virginia picked the $400 weekly benefit option.[41] Individuals receiving

[35] U.S. Department of Labor news release dated April 4, 2020, www.dol.gov/newsroom/releases/eta/eta20200404.

[36] U.S. Department of Labor news release dated January 5, 2021, www.dol.gov/newsroom/releases/eta/eta20210105. U.S. Department of Labor news release dated December 30, 2020, https://www.dol.gov/newsroom/releases/eta/eta20201230-1.

[37] U.S. Department of Labor news release dated March 16, 2021, www.dol.gov/newsroom/releases/eta/eta20210316.

[38] Attachment to Unemployment Insurance Program Letter No.15–20, Change 1, U.S. Department of Labor, dated May 9, 2020.

[39] The governor of Mississippi announced the termination of the programs over social media. Office of Governor Michael Parson (Missouri) press release dated May 11, 2021, https://governor.mo.gov/press-releases The Office of Alabama Governor press release dated May 10, 2021, https://governor.alabama.gov. Georgia Department of Labor press release dated May 13, 2021, https://dol.georgia.gov The Office of the Governor news release dated May 11, 2021, www.tn.gov/governor/news The Office of Governor Doug Ducey news release dated May 13, 2021, https://azgovernor.gov/governor.

[40] U.S. Department of Labor news release dated August 12, 2020, www.dol.gov/newsroom/releases/eta/eta20200812-0. Lost Wages Supplemental Payment Assistance Guidelines, www.fema.gov.

[41] Lost Wages Assistance Approved States, www.fema.gov The Office of the Governor Jim Justice (West Virginia)

at least \$100 of weekly unemployment benefits were eligible for LWA – a stricter eligibility requirement than FPUC's requirement that an individual be eligible to receive at least \$1 in weekly unemployment benefits.[42]

Participating states provided LWA to eligible individuals retroactively, beginning with the week of unemployment ending on August 1, 2020. Due to the fund's early depletion, benefits were paid for at most 6 weeks, until the week ending September 5, 2020.[43] All 21 study states except Florida received 6 weeks of funding. Florida was approved for 4 weeks, until the week ending August 22, 2020.[44]

FPUC and LWA together supplemented weekly unemployment benefits in the following periods depending on eligibility: \$600 (FPUC) from March 28, 2020 through July 31, 2020; \$300 (LWA) or \$400 (LWA, West Virginia) from August 1, 2020 through week ending September 5, 2020 (week ending August 22, 2020 in Florida); gap between September 5, 2020 and December 26, 2020; and \$300 (FPUC) from December 26, 2020 through September 6, 2021, with some states ending the program early.[45]

D Cobb-Douglas

Consider a case where utility is non-separable in health and consumption and take the cobb-douglas case with $g(m,h) = m^{\alpha}h^{1-\alpha}$. The FOC becomes:

$$1 - t - \Delta t = (1-r)\alpha(\frac{h}{m})^{(1-\alpha)} + r\alpha(\frac{h'}{m})^{(1-\alpha)} + \theta[m^{\alpha}h'^{1-\alpha} - m^{\alpha}h^{1-\alpha}]$$

From $U(c,l,h') = U(c - W, l, h)$ we can derive an expression for h':

$$m^{\alpha}h^{1-\alpha} = W(m) + m^{\alpha}h'^{1-\alpha}$$

using this in the FOC simplifies to

$$1 - t - \Delta t - (1+\alpha)\theta W(m) = \alpha(\frac{h}{m})^{(1-\alpha)}$$

Notice that this increases the implicit tax imposed by the health risk by factor α. This additional cost arises from the health effect on the marginal utility of leisure. A second change is that the utility cost of health now depends on the level of income m. Experiencing a health shock is more costly when an individual is working a lot. This increasing cost with m makes health risks operate like a non-linear progressive tax system.

press release dated September 9, 2020, `https://governor.wv.gov`.

[42] Although the size of the benefits are different for eligibility, the same programs qualify for both FPUC and LWA: regular unemployment compensation; Pandemic Emergency Unemployment Compensation (PEUC); Pandemic Unemployment Assistance (PUA); Extended Benefits (EB); Short-Time Compensation (STC); Trade Readjustment Allowances (TRA); Disaster Unemployment Assistance (DUA); and Self-Employment Assistance (SEA) program. U.S. Department of Labor news release dated April 4, 2020, `www.dol.gov/newsroom/releases/eta/eta20200404`. Lost Wages Supplemental Payment Assistance Guidelines, `www.fema.gov`.

[43] See, for example, Lost Wages Assistance, NC Department of Commerce, `https://des.nc.gov`.

[44] Florida Department of Economic Opportunity press release dated Sep 16, 2020, `www.floridajobs.org`.

[45] Unemployment Insurance Program Letter No. 14–21, U.S. Department of Labor, dated March 15, 2021.

D.0.1 Cobb-Douglas case without uncertainty

Without uncertainty, the health cost is a simple function of the time spent at work. An hour of work (1-l) has the health cost κ and $h = \kappa(1-l) = \kappa\frac{m}{a}$.

Denote the substitution elasticity between m and the composite good by e, with the Cobb-Douglas structure In the Cobb-Douglas case:

$$g(\frac{m}{a}, h) = \frac{a}{1+1/e}(\frac{m}{a}^{1-\alpha} h^{\alpha})^{(1+1/e)}$$

If $\alpha = 0$ this model becomes the canonical 2 good leisure-labor economy. From the FOC of the utility maximization, the optimal m^o follows:

$$m^o = \theta(1 - T'(m^o))^e$$

with $\theta = a\kappa^{-\alpha(1+e)}$. The canonical bunching approach for notches identifies e from the marginal buncher. The marginal buncher is the person who is just indifferent between the notch point and a higher income level. This persons' IC is thus tangent to the BC and also touches the notch point. Call the utility at the notch point U^* and the utility at the tangent point U^o, for the marginal buncher $U^* = U^o$. The notch utility U^* for the marginal buncher $\hat{a}$ is:

$$U^* = (1-t)m^* - \frac{\hat{a}}{1+1/e}(\frac{m^*\kappa^{\alpha}}{\hat{a}})^{(1+1/e)}$$

using the FOC result $m^o = \theta(1 - T'(m^o))^e$ we can write U^o as:

$$U^o = (1-t-\Delta t)\theta(1-t-\Delta t)^e - \frac{\hat{a}}{1+1/e}(\frac{\theta(1-t-\Delta t)^e\kappa^{\alpha}}{\hat{a}})^{(1+1/e)}$$

$$U^o = \frac{1}{1+e}\hat{a}(1-t-\Delta t)^{1+e}\kappa^{-\alpha(1+e)}$$

We combine the two utility expression and use the relation $\hat{a} = \tilde{m^o}\kappa^{\alpha(1+e)}/(1-t)^e$, to derive e:

we can write down an implicit solution for e in terms of κ, α, m and t:

D.1 Income Effects

The canonical bunching approach uses quasi-linear utilities and thus assumes that there are no income effects. In many contexts where notches are small, the absence of income effects is plausible. Recent work, however, stresses that small notches may not be salient (Saez and others). Moving to larger notches is thus attractive but leads to the added complication that such notches produce income effects. Structural estimates have previously used utility functions with income effects (Blundell, Meghir etc.) and income effects also feature prominently in most of the early empirical literature on labor supply (Hanoushek). Below we aim to cover a middle ground between the functional form flexibility of structural work and the quasi-experimental approach to identification of the bunching literature. We will show that excess mass does not only appear at m^* but also appears at lower income ranges with income effects.

D.1.1 Estimating Labor Supply Responses

Consider a more general labor supply function that allows for income effects:

$$\tilde{m^o} = \tilde{a} + e\tilde{w} - \gamma\tilde{y} \tag{5}$$

$\tilde{x}$ indicates log values and $\gamma\tilde{y}$ captures the income effect. When $\gamma = 0$ this equation collapses to the canonical quasi-linear utlity case without income effects.

The introduction of a lump sum benefit Δt reduces labor supply if $\gamma < 0$. This effect changes the impact of the non-linear benefit schedule studied above. For a worker just to the right of the eligibility notch at $m^* + \epsilon$, labor supply falls to $m^* + \epsilon - \gamma\Delta$ which is to the left of m^* if ϵ is small (illustrated in Figure XYZ). The labor supply response thus creates excess mass left of m^* and the excess mass at the notch point therefore does not fully capture the labor supply response. A special scenario where all excess mass occurs at the kink point is the case without income effects ($\gamma = 0$). In such a scenario, the worker at $m * +\epsilon$ would move to m^* and since the income effect is smaller for all workers with higher initial incomes, all bunchers will move to m^* and all excess mass occurs in a single point at m^*.[46] With income effects E does not appear at one specific point of the distribution but spreads out across a broader range of incomes, which creates additional identification challenges. We will return below to the question of how to identify excess mass over a wider income range.

The excess mass E is closely linked to the labor supply response of the marginal buncher. Individuals with pre-period income between m^* and the income of the marginal buncher $m^* + \Delta m$ make up the excess mass and E is thus given by:

$$E = \int_{m^*}^{m^*+\Delta m} h_0 dm$$

$$\Delta m = E/h_0 \tag{6}$$

where h_0 is the pre-notch wage distribution between m^* and $m^* + \Delta m$. To keep notation simple, we assume that the pre-period wage distribution is constant over this segment.[47]

To compute Δm we need to estimate h_0 and E. We can directly compute h_0 from the data if data on the pre-notch distribution is available. Such a pre-period distribution provides a valid counterfactual under a parallel trend assumption, similar to the assumption required in a difference in differences regression.[48]

[46] Note that this is a kink point and hence $\tilde{m^o}$ does not hold

[47] This assumption simplifies notation but is not required and richer baseline distributions can be included in the estimation.

[48] Without data on the pre-period, h_0 can still be estimated with "untreated" income ranges away from the notch point. This requires to estiamte h_0 in such untreated income ranges and then extrapolate to incomes in the treatment range. The researchers will need to make an assumption about which income ranges are untreated, and this requirement of an ad-hoc assumption has been controversial (XYZ). The presence of income effects worsens the problem. Bunching is more spread out with income effects and less sharp at the cut-off, making it harder to define untreated income bins.

A second step is to estimate E, the extra mass generated by bunching individuals. E is the difference between the observed post-notch income distribution (h_1) and the distribution of non-bunchers (h'_0):

$$h_1 = E + h'_0, \tag{7}$$

In practice, h'_0 is not directly observed and needs to be estimated. Typically $h'_0 \neq h_0$ and the pre-distribution does not provide a valid counterfactual. To see why, consider workers at m^* in the pre-period, they are to the left of the notch and thus part of the non-bunchers. However, the notch still affects their behavior, with the introduction of benefit labor supply falls to $m^* - \gamma\Delta$. As a result, there is no mass at m^* and $h'_0(m^*) = 0 \neq h_0(m^*)$. Using h_0 as counterfactual will bias the results, $h'_0(m^*) = 0$ implies that *all* individuals at $m = m^*$ are bunchers and the spike in density above neighboring cells ($\hat{E} = h_1(m^*) - \hat{h_0}(m^*) < E$) underestimates the true extend of bunching. Much of the debate about income effects focuses on the difference in compensated and uncompensated labor supply elasticities. It is important to note that the impact is more severe in the context of bunching estimates. Here, income effects not only affect the interpretation of the elasticity as (un)compensated but additionally bias the labor supply response estimate itself.

Unbiased estimates can be obtained without income effects. This is the canonical bunching assumption with $\gamma = 0$. Here $h_0 = h'_0$ as the labor supply of non-bunchers is unaffected by the introduction of the notch and as a result the spike in mass relative to neighboring regions provides an unbiased estimate ($\hat{E} = h_1(m^*) - \hat{h_0}(m^*) = E$). Assuming income effects away is thus an important underlying assumption of the canonical bunching approach.

For the more general case, valid estimates can be obtained with a difference in difference analysis. A first advantage of the difference in difference approach is that it ca detect any deviations from the pre-notch distribution, not just spikes in one specific location. As we saw above, this is important with income effects. Additionally, the difference in difference approach can overcome the identification challenge created by $h'_0 \neq h_0$. When leisure is a normal good, the introduction of benefits reduces labor supply among the non-bunchers. Note, that while the local distribution of m is changed, the total mass of non-bunchers below m^* is unaffected by the notch:

$$\int_0^{m^*} h'_0 = \int_0^{m^*} h_0 \equiv \pi$$

Using this result in 7, we can show that the notch generates total excess mass:

$$\int_0^{m^*} E = \int_0^{m^*} h_1 - \int_0^{m^*} h_0$$

which is the difference in the total density below the notch before and after the notch-reform. $\int_0^{m^*} E$ can be estimated in a difference in difference regression that compares the density below m^* before and after the introduction of the notch. In difference in differences notation:

$$Pr(I = m)_{t,m} = \phi \cdot 1[t > t^*] + \pi \cdot 1[m < m^*] + \bar{E} \cdot 1[t > t^*] \cdot 1[m < m^*] + \varepsilon_{t,m}$$

where t^* is the time of the reform, π is captured by the coefficient on the dummy $1[m < m^*]$. The coefficient $\bar{E}$ captures the average rise in density below m^*. Substituting this estimate into 6 yields the labor supply response of interest Δm.

The setting also yields an identification check in the spirit of a parallel trend check. This test is based on the distribution of the excess mass relative to the notch point. If the notch generates the excess mass, excess mass should peak near the notch and decline as we move away from the notch. To test this, we estimate a specification similar to a dynamic DiD, and let the E coefficient vary across income ranges:

$$Pr(I = m)_{t,m} = \phi \cdot 1[t > t^*] + \pi \cdot 1[m < m^*] + E_m \cdot 1[t > t^*] \cdot 1[m < m^*] + \varepsilon_{t,m}$$

Plotting E_m provides a visual check on the assumption that the notch generates excess mass. The excess mass should peak at m^*, and it's mirror image, missing mass, should peak to the right of m^*. Finally, for m further from m^*, the effects should diminish.

Similar "difference in bunching" approaches have been used in the literature (XYZ), typically as a check on the identification assumption of canonical bunching estimators. In the set-up above we explicitly leverage the additional degrees of freedom to broaden the applicability of bunching methods to preferences with income effects.

These issue can be addressed with a difference in differences set-up. For such an approach both a period before and after the introduction of the notch needs to be observed. Comparing the income distribution below m^* before and after the introduction of the budget notch identifies E.

D.1.2 Compensated Elasticity

The previous section's reduced form labor supply response can be used to estimate structural preference parameters that have validity beyond the specific context. The canonical reduced form approach calculates an upper bound of the compensated labor supply elasticity as $e_c < \frac{(1-\Delta m)^2}{1-\Delta t}$. With income effects the formula yields an upper bound for the uncompensated labor supply elasticity (e_u).

To make further progress, we need to specify a functional form for preferences. The standard approach is to assume a quasi-linear utility function. Such preferences do not have income effects and to allow for a more general case we will thus use a preference structure that includes the possibility of income effects. Preferences can be specified either by assuming a functional form of the utility function, the indirect utility function or the labor supply function.[49] A large empirical literature estimates a linear labor supply function as in 5 (Hausman etc.) and we will follow this literature and use the same labor supply function. Such preferences allow for income effects. Specifically, the substitution and income effects are respectively captured by e and γ. Notice that this assumption nests the quasi-linear case with $\gamma = 0$.

The estimation largely follows the same procedure as the canonical bunching approach, however there is one additional parameter, the income effect γ. To solve for this additional parameter requires one additional moment condition and we can use the dispersion of excess mass for this purpose. Without income effects all excess mass arises at m^*, while the excess mass is more spread out the bigger the income effects.

To derive a solution for γ we leverage the location of bunching. Note that all bunchers below m^* are at an interior solution. There will be one bunching person for whom m^* is an interior solution, call this person the marginal buncher from the left (see Figure XYZ). Before the notch

[49] The can be imposed on any any of the three functions. Roy's identity allows to derive direct and indirect utility functions from labor supply functions (up to a constant, which is meaningless for ordinal utility) and vice-versa.

the income of this person was $h_0 = m^* + p$. And using those two labor supply decisions in 5, we can show that:

$$m^* + p - \tilde{a} - e\tilde{w} + \gamma\tilde{y} = m^* - \tilde{a} - e\tilde{w} + \gamma(\tilde{y} + \Delta T)$$

$$\gamma = p/\Delta T$$

We can thus solve for γ by deriving p. Notice that everyone with $h_0 \leq m^* + p$ is an interior buncher and the total mass of interior bunchers is thus:

$$I = \int_{m^*}^{m^*+p} h_0$$

The excess mass below the notch point (I) thus pins down p, e.g. with h_0 constant $p = I/h_0$. And using p, we can solve for $\gamma = \frac{I}{h_0 \Delta T}$. If all excess mass arises at the notch point then $I = 0$ and consequently $\gamma = 0$ and the analysis collapses to the quasi-linear case. This approach can thus be used to check the validity of canonical bunching estimates. But more powerfully, it can be used to identify labor supply responses from large and salient notches in budget constraints.

D.2 Border Design

In this section we narrow our sample to counties along to border of two states, and thus with similar characteristics but different UI eligibility rules. The border counties are shown in Figure A3. Our sample states have 19 border stretches, and our sample includes 14 of those, excluding places where our sample has no observations in border counties (5 border stretches).

In a first step, we repeat the baseline analysis on the sample of border counties and find very similar effects to the baseline (Column 1 of Table A3). Next, we exploit the idea that neighboring counties experience similar demand shocks and allow all fixed effects to be specific to each border stretch. In practice, this implies that each border stretch is it's own DiD experiment and we stack the 14 border DiDs into a single regression. Since our power in these regressions is limited, we interact local time effects with a continuous measure of income, instead of letting time effects vary non-parametrically by income bins. The results of these regressions are again close to our baseline estimates (column 2).

Figure A3: Border Counties in Sample

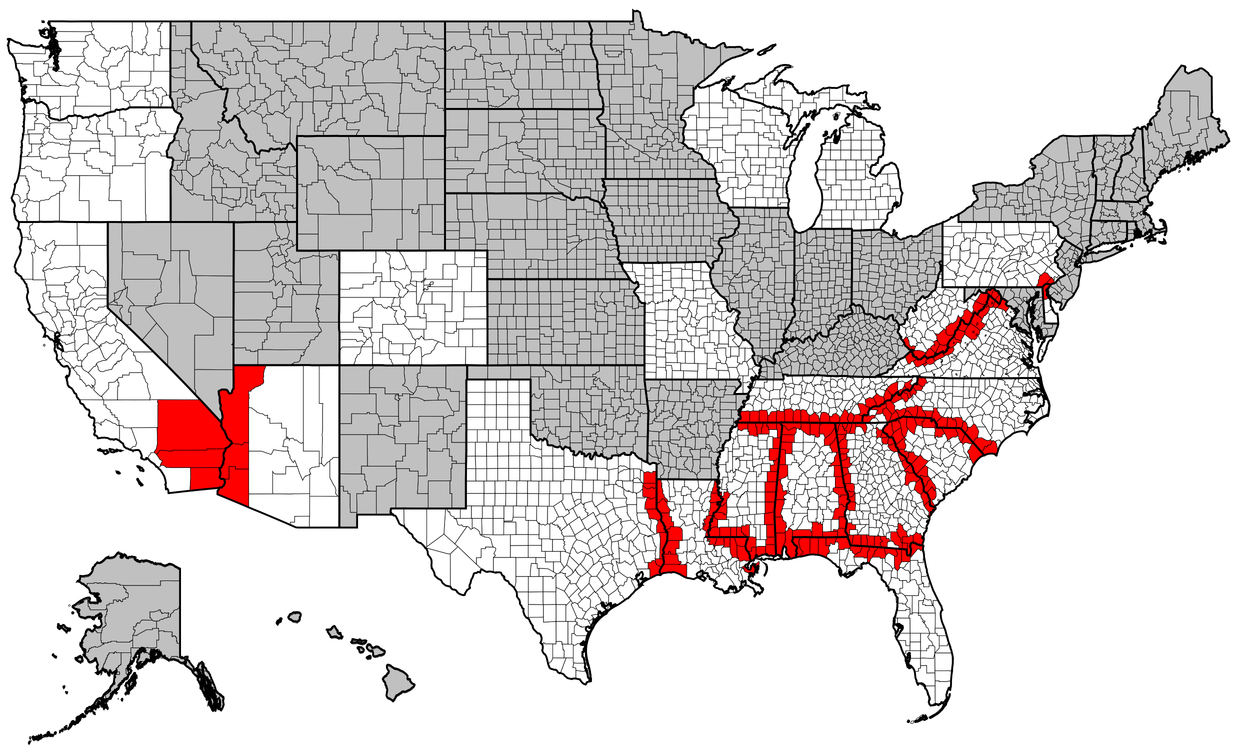

Note: The Figure shows counties along state boundaries that are included in our border sample.

D.3 Alternative measures of Covid-19 exposure

In this section we estimate the labor supply response to increase workplace risk using two alternative measures of Covid-19 exposure: the simple cross-industry task variation (i.e. the T_i component of $E_{i,c,t}$ from equation 4) and local Covid-19 outbreaks *within* the county $D_{i,c,t}$, measured as by the number of deaths per 1 million people.

Table A3: Excess Mass around UI Eligibility Threshold - Border Counties Sample

	(1)	(2)
Excess Mass (ptp)	0.819	0.961
	(0.159)	(0.127)
Interact income x		
time FE with		border
Observations	539932	539932

Note: The Table reports results from equation 3. The border sample is restricted to counties at state borders shown in Figure A3. Source: Homebase.

Figure A4 plots how task risks T_i affect labor supply behaviour. An increase in pre-determined industry risk results in a significantly greater amount of excess mass around the FPUC threshold and hence a reduction in labor supply. These results are highly significant, with excess mass increasing 0.5 percentage point for a standard deviation increase in risk. Workers thus shy away from high risk workplaces, in line with negative compensating differentials. The strong explanatory power of the regression shows that workplace safety is an important driver of labor supply behaviour. The R^2 of the regression is 0.44, workplace risks thus explain almost half of the variation in labor supply behaviour across industries.

In Figure A5 we focus instead on the labor supply response to county/week variation in Covid-19 death risk. We split $D_{i,c,t}$ into five categories and estimate the responses separately for those five risk levels. The excess/missing mass in red –replicated identical in all four panels– represents the behavioral response to FPUC in counties with zero recorded new deaths. In the top left panel of Figure A5, the blue area represents the "excess response" to FPUC in counties with a relatively low observed Covid-19 risk (between 0 and 15 weekly new deaths per million people). Despite the relatively low risk, it is visible that workers in these counties are responding more vigorously to the FPUC incentives relative to counties with zero risk. In these counties, workers appear to choose to move down the earning distribution more than in the zero-risk ones. As before, we are controlling for demand shocks with income range specific time effects. The remaining three panels show how the "excess response" increases with death risk. The excess mass of workers shifting to the left of the state-specific notches is particularly pronounced for very high-risk counties (more than 45 weekly new deaths per million people)

We next use the "excess response" results to compute the implied WTP for workplace safety. Using equation (1) presented in Section 3 we compute. WTP increases in a rather linear pattern and varies between a bit more than 10% of disposable income for counties with relatively low risk (top left panel of Figure A5) to 50% of disposable income in very high-risk counties.

Figure A4: Effect of Workplace Safety on Labor Supply

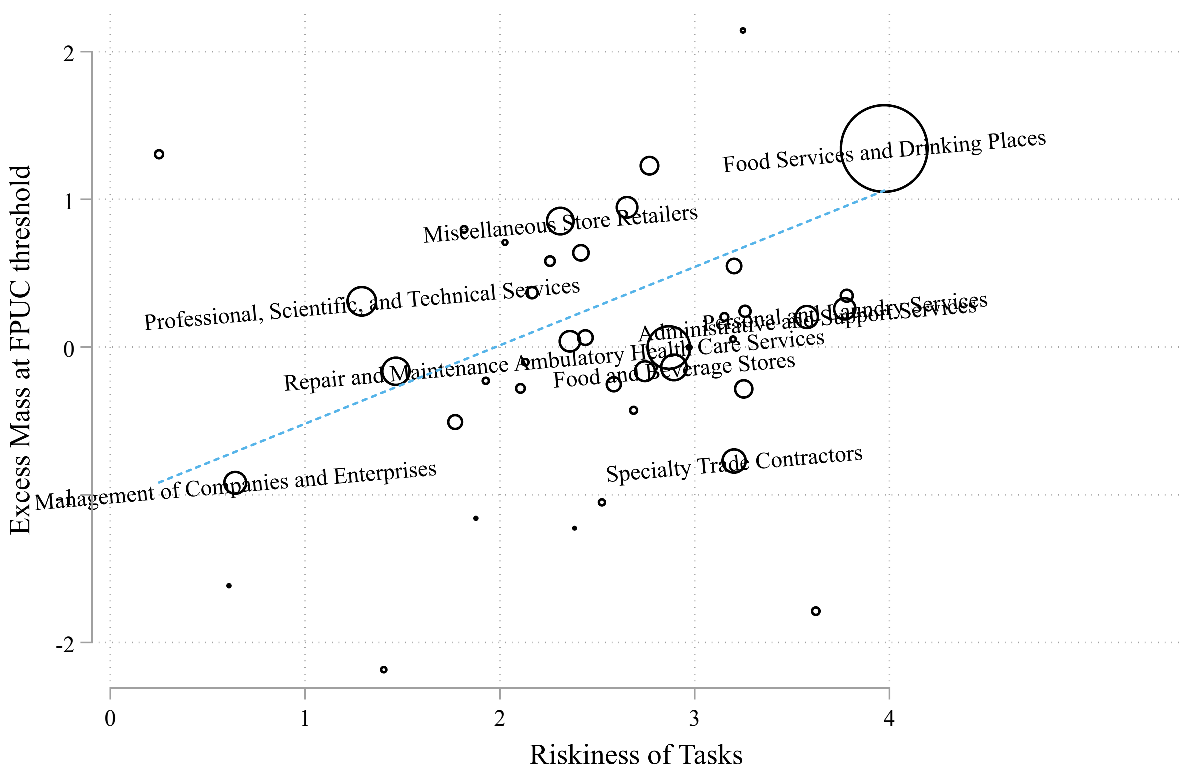

Note: The Figure shows the amount of excess mass at the FPUC threshold for 3-digit NAICS industries. The riskiness of tasks is the average risk of Covid-19 infections among tasks performed in the industry. The task risk data comes from Basso et al. (2020) and risk scores are standardized to have a standard deviation of 1. The y-axis shows the amount of excess mass generated by the FPUC eligibility threshold and is estimated in equation 2. The omitted industry is industry 111 (crop production). Industry titles are shown for the ten largest industries and for display purposes we only show industries with at least 1,000 observations. The size of the markers corresponds to the cell-size and regressions weight by cell-size. The fitted line has a slope coefficient of 0.5 and an $R^2 = 0.44$ Source: Homebase.

Figure A5: Excess and missing mass around the partial UI notch by observed death risk in county

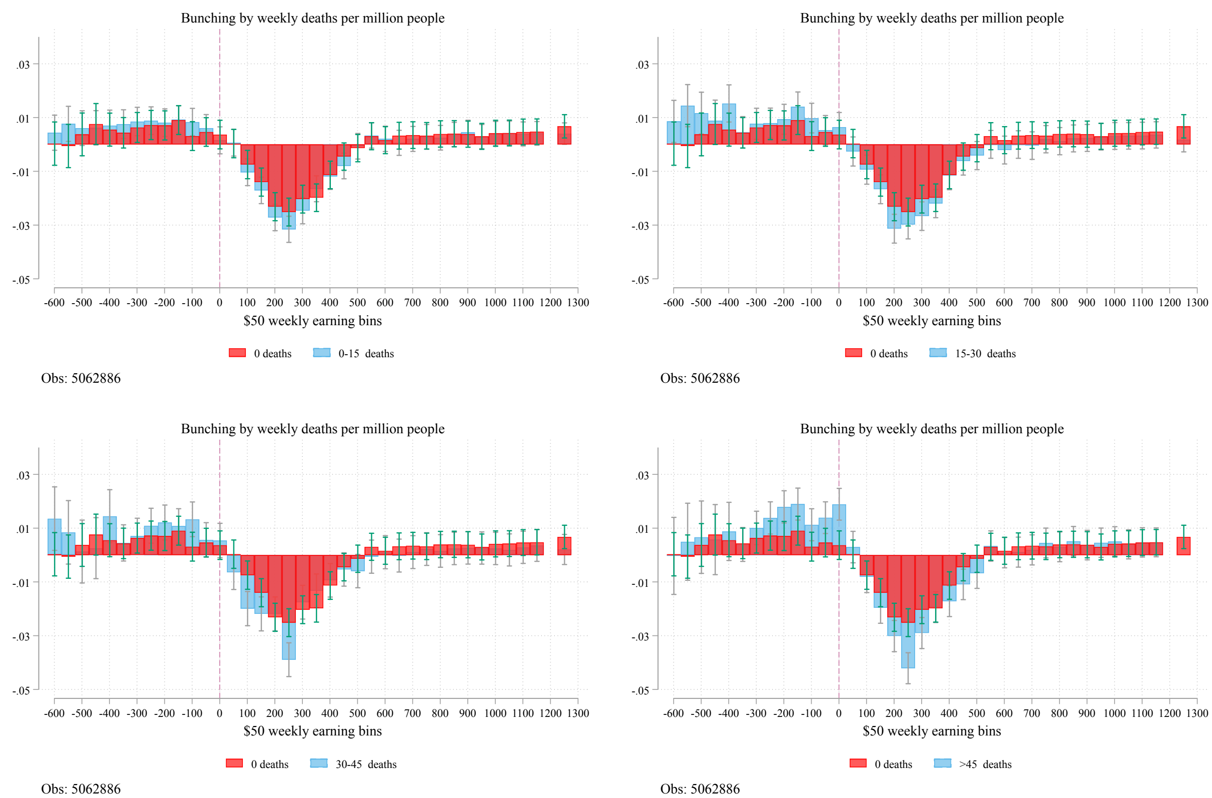

Note: The Figure shows $\delta_{r,\theta}$ coefficients from equation 3. Results for 4 different Covid-19 risk levels (θ) are plotted in each panel in blue. Covid-19 risk is measured as deaths per million in the week in the local area. The red bars are the benchmark response in areas with 0 Covid-19 deaths. The sample covers hourly workers with sufficient past earnings to qualify for MWB payments in their home state. Source: Homebase.