Willingness to Pay for Workplace Safety

Massimo Anelli[1] and Felix Koenig[2]

[1]Bocconi University, IZA, CESifo

[2]Carnegie Mellon University, CEP, IZA

November 17, 2021

DRAFT – click here for the most recent version.

Abstract

This paper develops a novel revealed-preference approach that uses discontinuities in budget constraints to value workplace safety. We measure 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 leverage 21 state-specific discontinuities in unemployment insurance eligibility criteria to identify the labor supply behavior. Results show large baseline excess and missing mass at the eligibility threshold and increasing responses at higher health risks. The observed behaviour implies that workers are willing to accept 34% lower incomes to reduce the fatality rate by one standard deviation, or 1% of income for a one in a million chance of dying. Most of this cost is not priced into wages, and canonical methods to measure the value of workplace safety miss such costs.

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

1 Introduction

The idea that there are large non-monetary returns to work has been a cornerstone of labor economics since the beginning of the field of economics. In “The Wealth of Nations,” Adam Smith conjectures that “Honour makes a great part of the reward of all honourable professions” (p. 117). Influential subsequent work discusses possible ways to measure the value of such amenities. In perfectly competitive labor markets, amenities will create “compensating differentials,” and differences in wages for otherwise similar work can identify the value of such amenities (Rosen, 1986, 1974; Lucas, 1977; Masters, 1969). The empirical implementation of this approach has, however, proven difficult in practice. For one, labor markets may not be perfectly competitive and wages will not necessarily price in the value of amenities. Second, workers typically select jobs with amenities they enjoy, thereby creating additional selection challenges. In this paper, we develop an alternative approach to measure the value of non-wage amenities and apply our approach to the value of workplace safety.

Workplace safety is a canonical case of such non-wage work amenities. Every year around 3 million Americans – nearly 2% of the labor force – suffer work-related injuries or illnesses[1] and two recent trends have further fueled interest in worker safety. For one, a growing number of workers are not directly employed by firms but work on their own account as contractors or “Gig workers.”[2] While wages in these positions may be comparable, there are stark differences in the availability of non-wage amenities; in particular, such workers typically lack access to paid time off and workers’ compensation, and are thus particularly vulnerable to workplace health risks.[3] Second, the outbreak of Covid-19 exposed thousands of frontline workers to health risks at work and sparked a debate about compensation for such risks. Major employers such as Amazon, Best Buy, and Target introduced Covid-19 hazard pay and increased wages for their workers.[4] How-

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

[2] See, e.g., Agrawal et al., 2013

[3] Studies of earnings in the Gig-economy include Cook et al. 2018, Hall & Krueger 2016

[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.

ever, it is unclear how these hazard bonuses were set and whether they fully compensated for the increased risk. More generally, this raises the question about adequate compensation for work risks.[5] How do workers value the benefit of safer workplaces? And are companies doing enough to protect the health of their employees?[6] This paper aims to address these questions.

Our estimation strategy infers the value of workplace safety from worker behavior at budget discontinuities and thus builds on a long tradition of revealed preference studies in economics. We analyze workers who must decide between working at alternative risk levels and losing income.[7] The empirical challenge is to generate exogenous variation in the risk-return ratio that credibly identifies workers' preferences. We build on the quasi-experimental approach developed in the bunching literature and use discontinuities in the US benefit system for this purpose. The intuition is that workers are less responsive to financial incentives when they are mainly motivated to work by non-financial factors such as the enjoyment of the work and the colleagues. Conversely, when work produces substantial dis-amenities like health risk or mobbing, the returns to work depend more strongly on financial incentives, and budget discontinuities have an amplified impact on labor supply. Following this intuition, we show that excess mass at budget "notches" or similar discontinuities in work incentives can identify worker preferences over non-wage amenities. Our estimation approach thus introduces a quasi-experimental identification strategy to estimate the value of non-wage amenities of work, an area that has struggled with clean identification strategies[8], and, additionally, relaxes the friction-less wage setting assumption inherent in classic approaches.

Our application uses discontinuities in work incentives from partial unemployment insurance rules and the variation in workplace safety during the Covid-19 outbreak. For identification, we use

[5] 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.

[6] Workplace safety is particularly deficient in low-paid occupations and labor market inequalities are therefore bigger than wage differences alone suggest. Using Bureau of Labor Statistics data, we calculated that fatality risk for occupations with earnings below the median is 75% higher than for occupations above the median. Sources: TABLE A-5. Fatal occupational injuries by occupation and event or exposure, all United States, 2019, Census of Fatal Occupational Injuries, Bureau of Labor Statistics and May 2019 National Occupational Employment and Wage Estimates, United States, Bureau of Labor Statistics

[7] Our approach is most similar in spirit to Sorkin (2018), who uses job switch patterns to identify non-pecuniary values of firms.

[8] Identification challenges in WTP estimates are for instance discussed in Kahn and Lang (1988); Viscusi (2018); Black and Kniesner (2003)

the launch of the Federal Pandemic Unemployment Compensation (FPUC), which creates a jump in workers' budget sets. 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 eligibility rules, equally paid workers are treated in some states but not in others. We can hence compare two equally paid workers on different sides of the benefit eligibility threshold to identify labor supply responses. Once we identify labor supply behavior, we study how workplace safety affects such behavior.

Our results show a sizable baseline labor supply response to the eligibility threshold for the \$600 FPUC. Eligible workers reduce earnings to levels below the UI eligibility threshold, resulting in substantial missing and excess mass in the earnings distribution.[9] We then show that deteriorating workplace safety leads to additional labor supply responses and magnifies the 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 with task-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 their health and losing their 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.

Our estimates imply that individuals are willing to give up around 34% of their income to reduce Covid-19 risk by one standard deviation, which is equivalent to a decrease of weekly fatalities rates by 31.15 per million workers. Converting the willingness to pay (WTP) into an hourly wage rate, this effect is equivalent to a \$6 wage decrease; or it implies a willingness to pay equal to 1% of income for a one in a million lower fatality rate. The variation in risk is broadly comparable to the change in risk associated with changing occupations. The highest risk occupation in the US,

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

fishing and hunting, has a fatality rate of 28 per million workers, while the lowest risk occupation, educational instruction, has a fatality rate of 0.06 per million workers.[10] In addition, Covid-19 risks were widely publicized and salient to workers.

Our findings show substantially higher willingness to pay for workplace safety than canonical hedonic wage regressions. We replicate the hedonic regression approach in our setting and find that a standard deviation in risk is worth 0.5% of income. This is roughly two orders of magnitude less than our baseline WTP estimate. One plausible reason for the difference is the frictions in wage settings; when wages are slow to adjust, they do not fully reflect changes in risk. As a result, frictions in wage setting lead to downward biased estimates in hedonic wage regressions.

We probe the robustness of our identification strategy to three possible threats. First, we analyze potential spurious local demand shocks. We use three strategies to investigate such shocks: a placebo test, a border design and finally, controlling for proxies for local demand shocks directly. The results show that such demand shocks are orthogonal to our threshold design and do not bias our results. Next, we address possible spurious changes in labor supply incentives from school closures and find again only minor effects on the results. Finally, we address the impact of selection effects on our results. We run specifications with worker-spell fixed effects and thus hold time invariant firm and worker characteristics constant. Such specifications again yield similar results to the baseline.

We additionally decompose the WTP into a selfish component, due to concerns for one's own safety, and a pro-social component, due to concerns about transmitting Covid-19 to others. The selfish component is more directly comparable with the WTP for non-transmittable illnesses, which do not pose secondary risks for others. Assuming that individuals put an equal utility weight on the health of all members of their household and using the Covid-19 intra-househould secondary fatality rate, we find that pro-social concerns account for a very small part of the WTP estimate, as a result of the relatively limited transmission. The vast majority of the observed behaviour is imputable to concerns for one's own safety, and the selfish only WTP is roughly 33.6%.

[10]Source: BLS Census of Fatal Occupational Injuries (CFOI) - Current `https://www.bls.gov/iif/oshcfoi1.htm#2019`, converted to weekly rates.

A central objective of much labor market research is to inform policies around minimum work standards. An influential literature studies the impact of minimum wages; however, minimum standards of other work dimensions have received far less attention. This lack of studies comes in part from the difficulty of measuring 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 high-risk employers pay high compensating differentials or go out of business. However, with imperfect competition, firms may not fully internalize the cost of high-risk jobs and may expose workers to excessive risks. Our approach provides a method to value such amenities, and our estimates yield a benchmark value for the WTP for safe workplaces. A first-order implication of our work is the pricing 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 implemented by most employers (\$2-2.5). Our estimates are informative also for pricing risk in non-Covid contexts: reducing workplace fatalities to the level observed in the UK and Germany would lead to substantial gains for workers. For instance, in construction, 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.

To link our results to prior work on the Value of a Statistical Life (VSL), we convert our WTP estimate into a VSL. This approach requires additional data and assumptions. First, VSL assumes that fatality risks are the sole driver of behavior, and second, VSL requires information on the beliefs about fatality risks—here the literature typically assumes that individuals are perfectly informed about risks. Under this assumptions, our benchmark estimates imply a VSL estimate of \$6.9 million, a value consistent with estimates from the literature.[11] In a further contribution, we use data on fatality beliefs to relax the perfect information assumption and show that workers substantially overestimated the fatality risks from Covid-19 and thus acted as if the risk was markedly higher. Failing to account for such imperfect information leads to biases in standard VSL

[11] 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).

estimates.[12] Accounting for perceived risk reduces the VSL estimate to \$2.6 million.

Related Literature – The topic of nonwage amenities goes back to classic work in economics (this includes Rosen (1986, 1974); Lucas (1977); Masters (1969)), but estimating the value of such amenities is difficult in practice.[13]

The empirical literature on the value of workplace safety typically uses hedonic wage regressions to estimate such values (Lucas, 1977; Brown, 1980; Hwang, Reed, and Hubbard, 1992).[14] Hedonic regressions relate occupational wage differences to workplace risk. 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 can therefore help inform policy decisions on workplace safety.[15]

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

[12]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.

[13]Compensating differentials may help explain 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)). 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).

[14]For other types of amenities researchers have also used hedonic regressions (Summers, 1989; Gruber and Krueger, 1991; Gruber, 1994, 1997; Fishback and Kantor, 1995; Stern, 2004). And recent studies pioneered stated preference surveys as an alternative to revealed preference estimates (e.g., Flory, Leibbrandt, and List (2015); Wiswall and Zafar (2018); Maestas et al. (2018)). 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.

[15]Similar to us, Sorkin (2018) also uses a revealed preference approach based on worker decisions rather than wages. He uses this to infer the aggregate value of non-pecuniar amenities at firms.

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 nonwage amenities and leverage policy reforms to estimate their value (Summers, 1989; Gruber and Krueger, 1991; Gruber, 1994, 1997; Fishback and Kantor, 1995). We use a similar quasi-experimental approach to estimate the value of workplace safety and leverage budget discontinuities to identify labor supply preferences. The theoretical framework of our work is related to the bunching/notch estimation approach (Kleven, 2016; Kleven and Waseem, 2013; Chetty, Friedman, and Saez, 2013). We expand their canonical two good labor-leisure approach to a three good economy with workplace safety. Finally, our estimate of the monetary value of avoiding Covid-19 death risk also relates to the large literature on VSL (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.[16]

Employed individuals can also receive FPUC payments if they meet the eligibility criteria, including an earnings test that requires that their earnings are below an earnings threshold.[17] Above

[16] 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.

[17] More precisely, FPUC is paid to workers on partial UI benefits. The precise qualifying criteria for partial UI varies by state. For our sample states these criteria always include an earnings test. Note that there are typically institutional rules that aim to prevent moral hazard in UI systems. For the most part, UI recipients are not allowed to refuse job offers and a job losses has to be not the fault of the worker. However, these rules are notoriously difficult to enforce

the threshold workers become ineligible for FPUC, which generates a notch in the budget set. 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).[18] The eligibility is based on contribution to the UI system in prior quarters. And while we do not directly observe eligibility, we apply state-specific eligibility rules and identify workers who are eligible for MWB.[19] 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 by leveraging budget discontinuities. 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

and a large literature on UI benefits finds moral hazard problems despite these rules.

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

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

Panel A

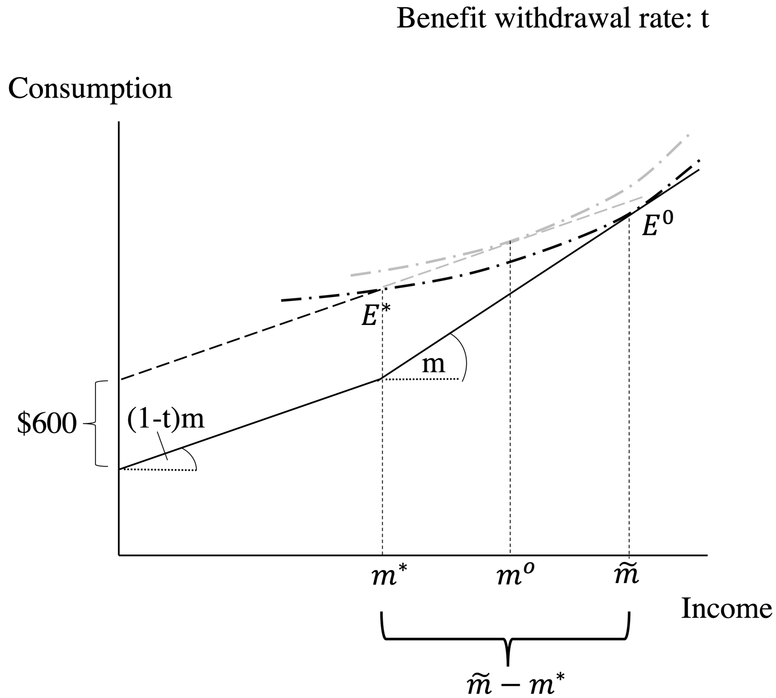

Panel B

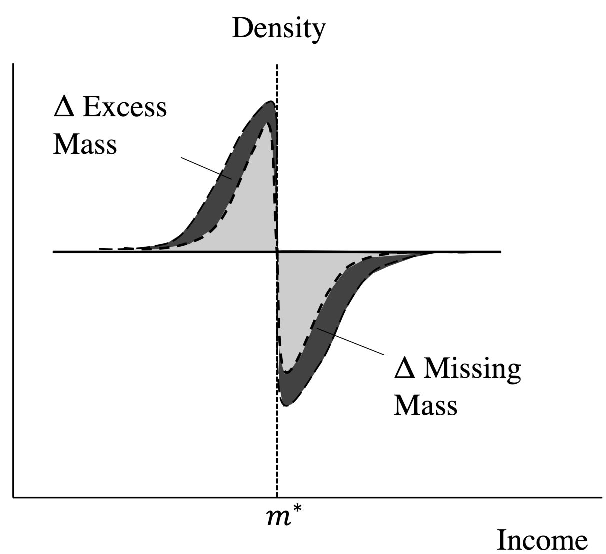

Figure 1: Bunching at m^*

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 $f(a)$. Assume this distribution and the tax system and preferences are smooth so that the resulting earnings distribution is also smooth. The benefit withdrawal 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 value of avoiding 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.[20] 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 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:

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

[20]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.

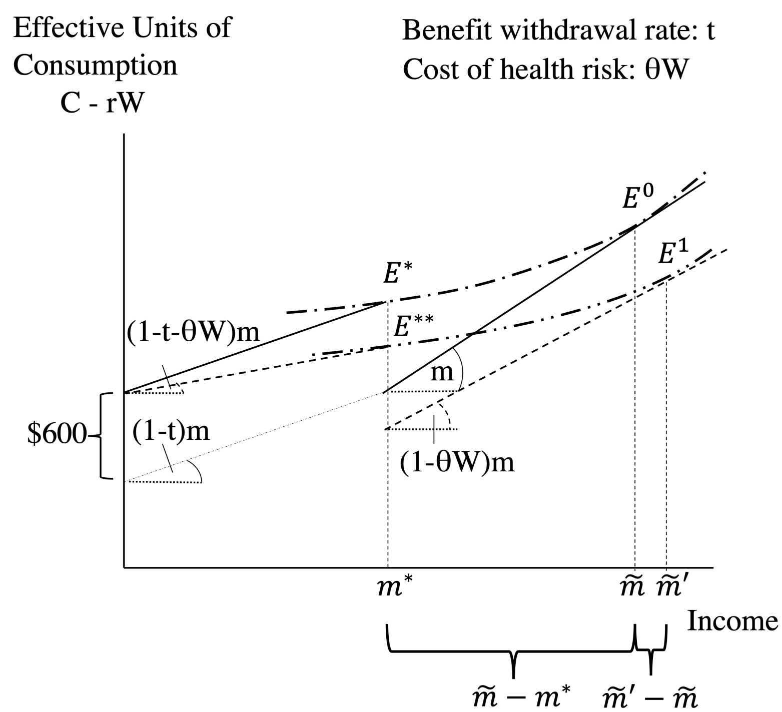

Figure 2: Cost of health risk

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 W in terms of measurable quantities $t, \Delta t, \theta, m^*$ and parameters that we can estimate based on behavioral responses: e, m^o.[21]

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 $(\tilde{m} - m^*)$ is the labor supply response in the low risk state. The implicit tax rate in the high risk state is $\tilde{t} = t + \theta W$. We can combine the two elasticity expressions to solve for the willingness to pay for a risk reduction by r as a share of disposable income, denoted by $WTP(r)$:

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

The final equality uses the two previous elasticity expressions and assumes that health risks are smooth throughout the cut-off ($\Delta \tilde{t} = \Delta t$).[22] 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

[21] 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.

[22] 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.

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.

3.1 Adjustment Frictions

We now consider how adjustment frictions, such as adjustment costs or inattention, affect estimates of WTP. It is well known that such frictions reduce the bunching at thresholds and in such cases the amount of bunching identifies a combination of preferences and the extend of adjustment friction, making it difficult to separate the two.[23]

Our approach, by contrast, can accommodate a wide range of optimisation frictions. First, consider a case where frictions make $1-\alpha$ percent of workers unresponsive. The excess mass (E) is reduced accordingly and now captures both preferences and α: $E = \int_{m^*}^{m^o} h_0 = \alpha(m^o - m^*)h_0$. Multiple combinations of α and labor supply response ($m^o - m^*$) are thus consistent with the observed E. Note however, that this is not the case for WTP, WTP is the ratio of excess mass in high (E_H) and low (E_L) risk settings and hence:

$$WTP = 1 - \frac{E_L}{E_H} = 1 - \frac{\alpha(\tilde{m}^o - m^*)h_0}{\alpha(m^o - m^*)h_0} = 1 - \frac{(\tilde{m}^o - m^*)}{(m^o - m^*)} \tag{2}$$

α cancels out as it affects both responses proportionally. We can thus identify WTP, even if some workers cannot adjust their work hours.

A second and related challenge are cases with hours constraints, for example cases where workers bargain with the employer over hours and therefore have to chose from a limited number of hour options. Such frictions deviate from the ones typically discussed in the literature in two important ways. First, a worker may be able to reach the UI threshold by working one hour less

[23]The canonical bunching literature has developed tools to identify preferences in a frictional labor market (Chetty et al., 2011; Chetty, 2012; Kleven and Waseem, 2013). In principle, any of these methods could also address the impact of frictions on WTP estimates. However, such tools use additional assumptions that are not necessary to identify WTP.

but is only allowed to drop a full shift. As a result she may move to an income further below the threshold and the excess mass in the income distribution therefore occurs over a wider income range and not just right at the threshold. A second issue is that the worker may not respond at all because they would have to move to lower income level. This effect would make workers less responsive to the threshold than in the friction-less benchmark.

The first challenge is relatively easy to address. Excess mass, E, is now spreads over a wider income range and while it may be more difficult to empirically identify the spread out excess mass, such spread out mass does not pose any conceptual challenges for the approach.[24] In other words, the first challenges affects the estimation of E but not the relation to the WTP.

The second challenge can be addressed in a similar fashion as the canonical adjustment friction above. Denote the fraction of individuals who do not respond because of the indivisibility friction by $1-\alpha$. If α is constant equation 2 applies again and implies that the WTP estimate is unaffected by the second challenge.

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 most front-line workers 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 work hours and earnings. Obtaining reliable labor supply records has been a key challenge as many survey data sources suffer from measurement errors that make it difficult to measure labor supply changes accurately (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

[24] Canonical bunching methods focus on excess mass right at the threshold and would fail to fully capture more spread-out excess mass.

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.[25] 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.

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, which also have a high degree of accuracy, however, they often only cover 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. Relatedly, the Homebase records are not used to administer UI benefits which alleviates concerns about strategic misreporting and avoidance.

A third advantage is the coverage of the data. The data mainly covers small service sector businesses with hourly workers, which in our application is helpful for two reasons.[26] First, these high-street service sector workers are typically "frontline workers" and directly exposed to Covid-19 risk at the workplace and 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 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.[27] 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

[25]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.

[26]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.

[27]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.

imputations inevitably introduce noise we down-weight such observations.[28]

A drawback of this type of private-sector data is that we lack data on individuals who exit the sample and therefore do not know whether these individuals left the labor force or changed employers. This is however not a limitation for our study since our methodological contribution is designed to identify intensive margin responses.

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

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.

[28]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.

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				

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 standard DiD regression controls for aggregate fluctuations with time fixed effects. One may worry that the recession has different impacts on high- and low-income workers that aren't captured by time fixed effects. 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:

Figure 3: Notch point by state

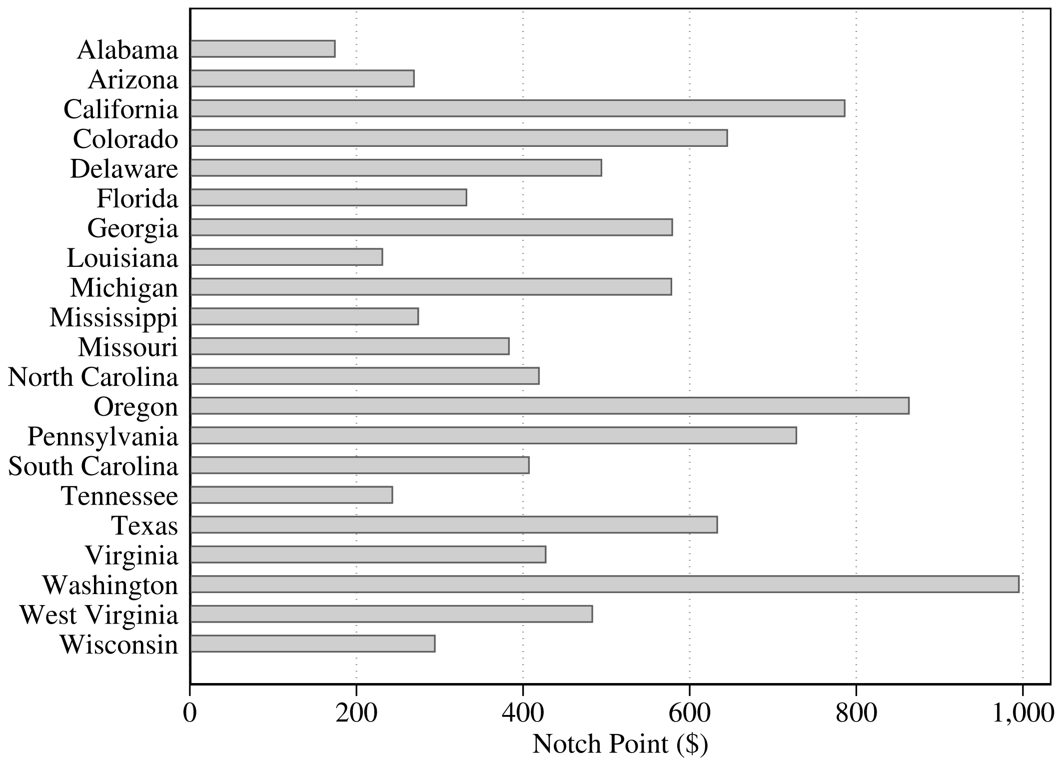

Note: The Figure shows maximum allowable earning while receiving FPUC payments for MWB recipients across US states.

$$E_{w,t,m,r} = \pi_{m,t} + \sum_{k=-650}^{1300} \beta_r \cdot I[r = k] + \sum_{k=-650}^{1300} \delta_r \cdot I[r = k] \cdot C_t + \varepsilon_{w,t,m,r} \tag{3}$$

where $E_{i,t,m,r}$ is a dummy with value 1 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. 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 and β_r captures the associated 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 and restricts the sample to obtain a balanced number of work-week observations before and after FPUC for each worker.[29] This "balanced" 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.[30] 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 4 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 ($T_{r,m}$) that takes value 0

[29] Cross-sectional studies of bunching have to assume that there is no entry or exit. 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.

[30] 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.

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 3 is given by:

$$E_{i,t,m,r,\theta} = \pi_{m,t} + \beta_{r,\theta} \cdot T_{r,m} + \delta_{r,\theta} \cdot T_{r,m} \cdot C_t + \varepsilon_{i,t,m,r,\theta} \quad (4)$$

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”).[31]. 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, the adjustment frictions in scheduling discussed in section 3 may prevent workers from freely choosing their earnings, and could spread earning responses of broader ranges. 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.[32] Third, such behavior is consistent with income effects from the FPUC benefits.

[31] 9% of all workers in our analytical sample used to work in the bin “notch+250” before the start of the pandemic

[32] Differences in observed earnings and UI relevant earnings arise in some jurisdictions from allowances for families

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

Panel A: FPUC eligible workers

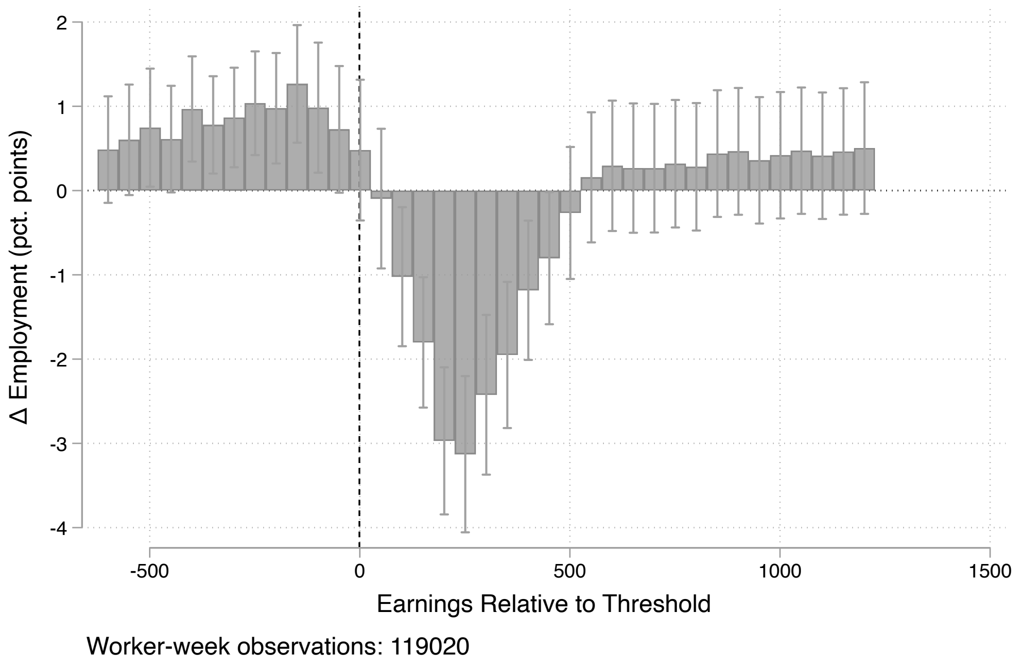

Panel B: FPUC ineligible workers

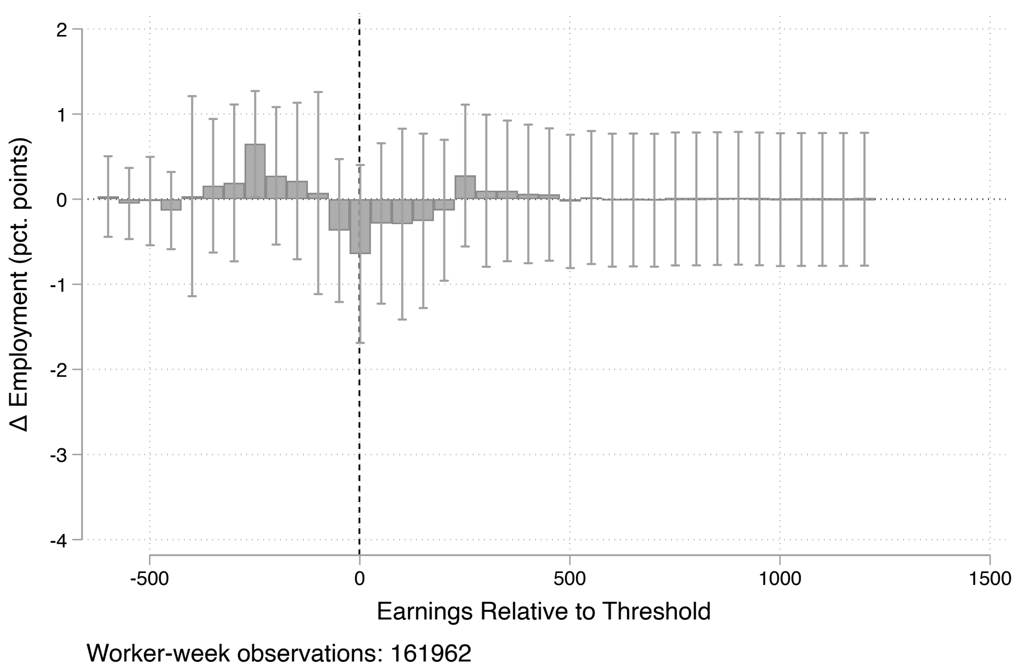

Note: The Figure shows δ_r coefficients from the equation 3. Standard errors are clustered at the state, income bin, week level and 95 percent confidence intervals are reported. In panel a, the sample is hourly workers with sufficient past earnings to qualify for MWB payments in their home state. In panel B the sample is instead hourly workers with insufficient past earnings to qualify for MWB payments in their home state.

or special circumstances or from multiple jobholders. As described above, we down-weight cases where we cannot measure benefit-relevant incomes well.

A potential concern with the setup is that we pick up differential labor demand shocks, even after controlling for earning level-specific time effects. Note, however, that any such spurious shocks would have to be correlated with the FPUC thresholds. In other words, shocks only pose a threat if they affect specific income ranges and different ones in different states in ways that correlates with the state specific thresholds. 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, we introduce additional controls for demand shocks. We allow such shocks to affect different income ranges in different states by interacting state dummies with a continuous earning variable in the Covid-19 period. These state specific 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 confirm that our findings are orthogonal to state-specific demand shocks (Table 2 column 2). Next, we allow for even more local shocks and repeat the exercise with county-level fixed effects and again find that such controls do not affect our results (column 3). Finally, we also 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 A4). 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 E).

7 Willingness to pay for workplace safety

We now use equation 1 to estimate the WTP for workplace safety. The denominator of this equation is the baseline response discussed in the previous section. The numerator is the additional response when fatality rates increase. We thus estimate how excess mass changes with change in fatality rates.

7.1 Measuring Covid-19 Exposure

To implement this approach, we need a measure of workplace safety that is orthogonal to local worker decisions. We construct such a measure of Covid-19 exposure by combining time-invariant 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 $F_{c',t}$.[33] Our measure of exposure is:

$$\theta_{i,c,t} = F_{c',t} \cdot I_i \tag{5}$$

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 that make fatality risks an endogenous outcome. For instance, we rule out that fatality rates are affected by mitigation measures which are likely endogenous to workers' WTP. Also, note that we focus on fatality rates – rather than infection rates – because of the lack of reliable infection data during the first months of the pandemic, caused by limited testing capabilities. The industry risk I_i is obtained by combining task-specific infection risk data from Basso et al. (2020) with American Community Survey data on the distribution of

[33] Appendix F additionally presents results for the two sources of variation separately (the simple cross-industry variation I_i and variation in local Covid-19 outbreaks within each county $F_{c,t}$).

tasks across industries, and then computing a risk score for each industry.[34]

Since $\theta_{i,c,t}$ 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 $\theta_{i,c,t}$ and actual fatality rates.[35] One standard deviation of workplace risk $\theta_{i,c,t}$ is equivalent to an increase in fatality rates by 31.15 cases per million workers. This variation is large, but in the same ballpark as the pre-Covid-19 cross-occupation variation in fatality rates: one standard deviation in fatality rates across US occupations is 4 cases per million workers per week with the highest risk of 28 for fishing and hunting workers.[36]

7.2 The value of workplace safety

We now estimate how the response to FPUC thresholds varies at different levels of risk and estimate equation 4. We split $\theta_{i,c,t}$ into five quintiles and estimate the responses separately for those five risk levels.[37]

We find that excess mass does indeed increase with rising dis-amenities at work. Figure 5 plots the response to FPUC for the five risk quintiles. The benchmark response in the lowest risk quintile is shown in grey and riskier quintiles in black. In the top left the black area represents the

[34] 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 for the lowest digit industries available in the ACS (mostly 3 and 4 digit) by taking an employment weighted average of occupational risks in each industry.

[35] 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 by apportioning the deaths to industries based on time-invariant fatality rates in industries and based on the employment share of the industry ($\frac{e_{i,c}}{\sum_i e_{i,c}}$). 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 $D_{i,c,t} = F_{c,t}\frac{e_{i,c}\cdot w_i}{\sum_i e_{i,c}\cdot w_i}$

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

[37] A more parametric alternative would interact UI thresholds with a continuous measure of $\theta_{i,c,t}$, which yields similar results. Splitting the data into five quintiles allows for more flexible non-linear effects and provides a more transparent look at the data.

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

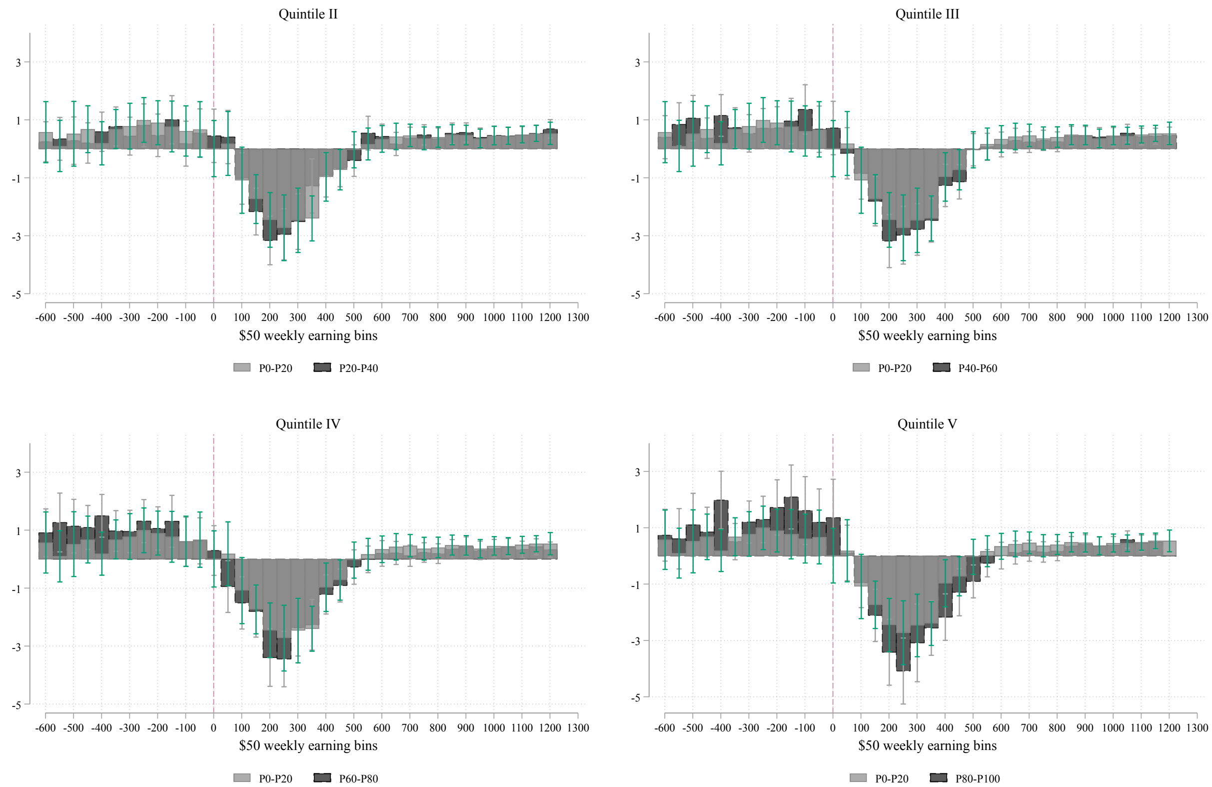

Note: The Figure shows $\delta_{r,\theta}$ coefficients from equation 4. Results for 5 quintiles of Covid-19 risk (θ) are plotted. The gray bars represent the response in the lowest quintile and black bars in the sub-panels respectively show responses in risk quintiles 2 to 5. The sample covers hourly 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.

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 threshold than in the low risk setting. 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 is particularly pronounced for the top risk quintile, consistent with the model prediction that dis-amenities at work magnify the excess mass at budget notches.

It is important to note that $\theta_{i,t}$ is correlated with other shocks to the labor market. However, such correlations only create an omitted variable problem if they are also correlated with the state specific UI thresholds. Formally, omitted variable bias (OVB) only occurs if two conditions are met: First, there are omitted variables which correlate with the treatment variable *and* second, these variables are also correlated with the outcome variable. Our identification assumption is based on the second condition. Our outcome variable is excess mass at UI thresholds and our identification assumption is that such behaviour at the UI thresholds is uncorrelated with demand shocks. Notice that this identifying assumption is the same assumption that is used for canonical bunching estimators. In other words, we can identify WTP when the identifying assumptions for canonical bunching estimates are met. Importantly, we do not require that $\theta_{i,c,t}$ is uncorrelated with demand shocks, all shocks that are orthogonal to the state specific thresholds will not affect the results.

To probe whether this identification assumption holds, we again run a placebo test with ineligible workers. Figure A3 repeats the analysis for workers have no incentives to respond to the benefit eligibility thresholds. If the additional labor supply response for higher risk quintiles depicted by Figure 5 were to capture spurious demand shocks correlated to the increased risk, such additional response would show also for non-eligible workers. The absence of any significant additional mass for non-eligible workers adds further confidence that the research design is valid.

We now combine the estimates to compute the WTP for workplace safety. Equation 1 shows that WTP is given by the ratio of the baseline response to FPUC (seen in Figure 4) and the additional labor supply response created by workplace risk (seen in Figure 5). We now quantify these responses in Table 2, which presents WTP estimates from the $400 treatment window from specification 4.[38] Panel A shows the excess mass around FPUC at average risk levels – the denominator of equation 1. 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 results for the numerator – changes in excess mass as workplace risks increase. A standard deviation increase in risk leads to 0.35 percentage points more excess mass. 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).

[38] In Figure A2 we test the sensitivity of our DiD estimate to changing treatment windows around the threshold. Our estimate is statistically significant if we consider a window of $150 around the threshold. We identify only a subset of the response if we focus on a narrow window, once the window is $250 or bigger the effect is very stable.

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 $\theta_{i,c,t}$ (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 these robustness tests the results remain close to the baseline. Finally, in Table A4 we explicitly control for granular proxies of demand shocks, such as weekly number of active employees by industry and state and weekly percent change in net revenues of small businesses by industry and state. In column 5 we also control for school closures, a factor that might confound labor supply responses.[39]. Estimates are robust to the inclusion of each of these controls and of all these controls simultaneously (column 6), confirming the efficacy of our identification strategy in isolating labor supply responses to changes in health risk.

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

[39]Employment and Small Businesses daily data are obtained from Chetty et al. (2020b), while the share of in class instruction is obtained from Parolin and Lee (2021a)

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 4 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).

is also quantitatively small and suggests that wages increased by 11 cents, a 0.5% wage increase (results are not reported but are available upon request). This small coefficient likely reflects the fact that wages are slow to adjust and, despite several high profile cases of hazard pay, 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 estimate implies a WTP that is an order of magnitude greater than the hedonic result.

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, several 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 implicit cost of risk exposure. To fully offset the non-pecuniary costs of added workplace risk, hazard pay for a standard deviation increase in risk would need to be as high as 34%. In other words, workers were worse off at work during Covid-19, despite the introduction of hazard pay.

8 Discussion

8.1 Workplace Safety Policy

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. To illustrate this point, we perform a back of the envelope calculation for the construction industry, a large and relatively risky industry. Weekly fatality rates in this industry in the US are 3 workers per million full-time employees per week.[40] Our estimates suggests that eliminating fatality risks would be equivalent to a 3% increase in wages. Reducing risk thus potentially offers large gains for workers. Reducing fatality risks to zero is perhaps an unattainable target. However, even reducing fatality rates to the level seen in the UK or

[40]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.

Germany would be a substantial improvement and equivalent to a wage increase of 2.5%.[41] Such gains happen to be similar in magnitude as the worker gains from the introduction of a $15 minimum wage, a popular labor market intervention.[42] Another useful point of reference are the wage gains implied by switching industries. Such an exercise helps evaluate the potential for compensating differentials to explain wage dispersion. 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. (2018), who find values ranging from 2% to 16%.[43]

When generalizing our estimates to non-covid workplace risk, we need to consider that the WTP for a non-transmissible illness or injury might be lower. Indeed, our WTP estimate might in part reflect workers internalizing the risk of covid transmission to others. The higher the weight workers place on others in their utility function, the more likely our estimate represents an upper bound for non-transmissible workplace risk. Conversely, note that in the canonical case of myopic individuals, who only care about their own utility, the WTP for transmittable and non-transmittable health risks coincide.[44] To get a sense of the importance of the pro-social feature in our WTP for lower covid risk, we perform a back of the envelope calculation for a worker who cares about the well-being of other household members. Denote the utility weight of other household members by w, the number of other household members by n and the intra-household secondary fatality rate by r. The relation of WTP for a transmittable (WTP_T) and non-transmittable diseases (WTP_{nT}) is: $WTP_T = (1 + r \cdot w \cdot n) \cdot WTP_{nT}$. For the back of the envelope calculation, assume that the worker cares as much about others' utility as her own ($w = 1$) and note that the intra-househould secondary fatality rate is $r = 0.002$.[45] For household size, consider a four person household, i.e. a

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

[42] 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.

[43] Maestas et al. (2018) study the value of schedule autonomy, telecommute, physical activity, sitting, relaxed work environment, work autonomy, PTO, team-work, training, opportunity to serve.

[44] The prior literature almost exclusively considers myopic agents when interpreting risky behaviour of individuals.

[45] r is obtained multiplying the 30% intra-household transmission rate. (Lewis et al., 2020) times the 0.68% infection fatality rate, that is the fatality rate conditional to being infected (Meyerowitz-Katz and Merone, 2020)

household at the 90th percentile of the US size distribution. In this case, our estimate is 0.6% larger than the WTP_{nT}.[46] In other words, our baseline estimate of a WTP of 33.8% would be reduced to roughly 0.338/1.006=33.6% of weekly income for a non-transmittable disease. Quantitatively, the concern for once own health is thus the main component of the WTP estimate, with a quantitatively small additional contribution from pro-social concerns.[47]

8.2 Value of a Statistical Life

A popular approach for quantifying responses to health risks is to compute a "value of a statistical life" (VSL). Such estimates infer what implicit value of life is consistent with the observed behaviour. Such estimates require additional assumptions: First, they assume 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 non-fatal aspects. Second, VSL estimates typically assume perfect information. In other words, it assumes that the true change in risk is known to decision makers.

If we are willing to make these assumptions, we can compute VSL as 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 are in line with these findings and are at the middle of this range.

Ideally, researchers would relax the perfect information assumption and 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 this assumption often fails. Individuals, for instance, appear

[46]This uses $r \cdot w \cdot n = 0.002 \cdot 1 \cdot 3 = 0.006$.

[47]Pro-social concerns will play a more important role for diseases with more aggressive transmission rates and play a minor role in this setting because r is small.

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 fatality beliefs with our 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. This increase is bigger than the increase in actual fatality risks. People thus overreact relative to a perfect information, rational agent setting. The extensive news coverage of Covid-19 deaths made such risks extremely salient and an in line with the behavioral literature on "salience" (Bordalo, Gennaioli, and Shleifer (2013)) we find that peoples beliefs overshoot the the true changes in risk. We next use the perception data to re-estimate VSL and find a VSL of \$2.57 million (Panel D of Table 2), about a third of the rational expectation estimate. Perceived fatality rates enter the VSL calculation as the denominator and replacing rational with actual expectation thus increase the denominator of the VSL calculation and lowers the VSL estimate. In other words, accounting for imperfect information implies that the same behavioral response is produced by a larger shock to risk perceptions. This illustrates the importance of accounting for belief sets when evaluating worker behaviour. 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.

9 Conclusions

This paper presented a new method to identify WTP based on excess mass around budget discontinuities. We then use this approach to measure the value workers attach to safe workplaces. 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. We show that risk beliefs can deviate substantially from the perfect information benchmark and measuring individual information sets accurately is essential for revealed preference estimates, like VSL.

We find that the cost of risk is not fully priced into wages and hazard bonuses were not quantitatively large enough to offset the non-pecuniary cost of working under greater health risk. We show that this has stark implications for WTP estimates from hedonic wage regressions. First, such estimates understate the true cost of workplace risk by only focusing on the costs that are priced into wages. The cost that are not picked up by hedonic regressions are particularly policy relevant, since policy makers aim to target exactly the costs of workplace safety that are *not* priced into wages. We show that such costs exists and are sizable.

References

Altonji, Joseph G. and Christina H. Paxson. 1992. “Labor Supply, Hours Constraints, and Job Mobility.” *The Journal of Human Resources* 27 (2):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.

Basso, Gaetano, Tito Boeri, Alessandro Caiumi, and Marco Paccagnella. 2020. “The New Hazardous Jobs and Worker Reallocation.” Centre for Economic Policy Research Discussion Paper 15100.

Black, Dan A. and Thomas J. Kniesner. 2003. “On the Measurement of Job Risk in Hedonic Wage Models.” *Journal of Risk and Uncertainty* 27 (3):205–220.

Bollinger, Christopher R. 1998. “Measurement Error in the Current Population Survey: A Nonparametric Look.” *Journal of Labor Economics* 16 (3):576–594.

Bonhomme, Stéphane and Grégory 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.” *The Quarterly Journal of Economics* 128 (3):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):e0252454.

Chetty, Raj. 2012. “Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply.” *Econometrica* 80 (3):969–1018.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team. 2020a. “Data for: The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data.” URL `https://tracktherecovery.org/`. Opportunity Insights Economic Tracker: Employment data series from Paychex, Intuit, Earnin and Kronos; Small business openings and revenue data from Womply.

———. 2020b. “The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data.” NBER Working Paper 27431.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *The Quarterly Journal of Economics* 126 (2):749–804.

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 V. 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." *The 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." *American Economic Review* 84 (3):622–641.

———. 1997. "The Incidence of Payroll Taxation: Evidence from Chile." *Journal of Labor Economics* 15 (S3):S72–S101.

Gruber, Jonathan and Alan B. Krueger. 1991. "The Incidence of Mandated Employer-Provided Insurance: Lessons from Workers' Compensation Insurance." In *Tax Policy and the Economy*, vol. 5, edited by David Bradford. The MIT Press, 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.

Kahn, Shulamit and Kevin Lang. 1988. "Efficient Estimation of Structural Hedonic Systems." *International Economic Review* 29 (1):157–166.

Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2):263–292.

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

Kleven, Henrik J. and Mazhar Waseem. 2013. "Using Notches To Uncover Optimization Frictions And Structural Elasticities: Theory and Evidence from Pakistan." *The Quarterly Journal of Economics* 128 (2):669–723.

Kleven, Henrik Jacobsen. 2016. "Bunching." *Annual Review of Economics* 8 (1):435–464.

Lavetti, Kurt. 2020. "The Estimation of Compensating Wage Differentials: Lessons From the Deadliest Catch." *Journal of Business & 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.

Lewis, Nathaniel M., Victoria T. Chu, Dongni Ye, Erin E. Conners, Radhika Gharpure, Rebecca L. Laws, Hannah E. Reses, Brandi D. Freeman, Mark Fajans, Elizabeth M. Rabold, Patrick Dawson, Sean Buono, Sherry Yin, Daniel Owusu, Ashutosh Wadhwa, Mary Pomeroy, Anna Yousaf, Eric Pevzner, Henry Njuguna, Katherine A. Battey, Cuc H. Tran, Victoria L. Fields, Phillip Salvatore, Michelle O'Hegarty, Jeni Vuong, Rebecca Chancey, Christopher Gregory, Michelle Banks, Jared R. Rispens, Elizabeth Dietrich, Perrine Marcenac, Almea M. Matanock, Lindsey Duca, Allison Binder, Garrett Fox, Sandra Lester, Lisa Mills, Susan I. Gerber, John Watson, Amy Schumacher, Lucia Pawloski, Natalie J. Thornburg, Aron J. Hall, Tair Kiphibane, Sarah Willardson, Kim Christensen, Lindsey Page, Sanjib Bhattacharyya, Trivikram Dasu, Ann Christiansen, Ian W. Pray, Ryan P. Westergaard, Angela C. Dunn, Jacqueline E. Tate, Scott A. Nabity, and Hannah L. Kirking. 2020. "Household Transmission of Severe Acute Respiratory Syndrome Coronavirus-2 in the United States." *Clinical Infectious Diseases* 73 (7):e1805–e1813.

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 J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger. 2018. "The Value of Working Conditions in the United States and Implications for the Structure of Wages." NBER Working Paper 25204.

Mas, Alexandre and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *American Economic Review* 107 (12):3722–3759.

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

Meyerowitz-Katz, Gideon and Lea Merone. 2020. "A Systematic Review and Meta-analysis of Published Research Data on COVID-19 Infection Fatality Rates." *International Journal of Infectious Diseases* 101:138–148.

Morchio, Iacopo and Christian Moser. 2019. "The Gender Gap: Micro Sources and Macro Consequences." Available at SSRN.

Parolin, Zachary and Emma Lee. 2021a. "U.S. School Closure & Distance Learning Database." Data retrieved from OSF.

Parolin, Zachary and Emma K. Lee. 2021b. "Large Socio-economic, Geographic and Demographic Disparities Exist in Exposure to School Closures." *Nature Human Behaviour* 5:522–528.

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." In *Handbook of Labor Economics*, vol. 1, edited by Orley C. Ashenfelter and Richard Layard. Elsevier, 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* (95/96):201–220.

Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *The Quarterly Journal of Economics* 133 (3):1331–1393.

Stern, Scott. 2004. "Do Scientist Pay to Be Scientist?" *Management Science* 50 (6):835–853.

Summers, Lawrence H. 1989. "Some Simple Economics of Mandated Benefits." *American Economic Review, Papers and Proceedings of the Hundred and First Annual Meeting of the American Economic Association* 79 (2):177–183.

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):205–246.

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. 2018. "Preference for the Workplace, Investment in Human Capital, and Gender." *The Quarterly Journal of Economics* 133 (1):457–507.

A Appendix Figures and Tables

Figure A1: Scheduling app screenshot

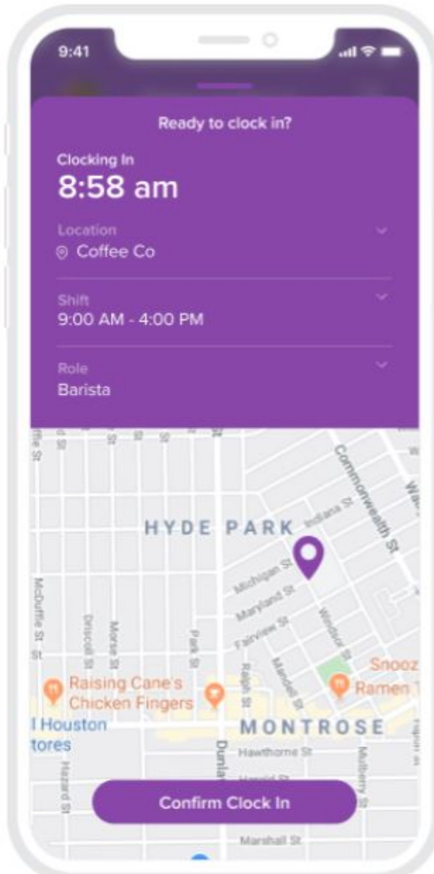

Figure A2: Effect of FPUC with Alternative Treatment Window

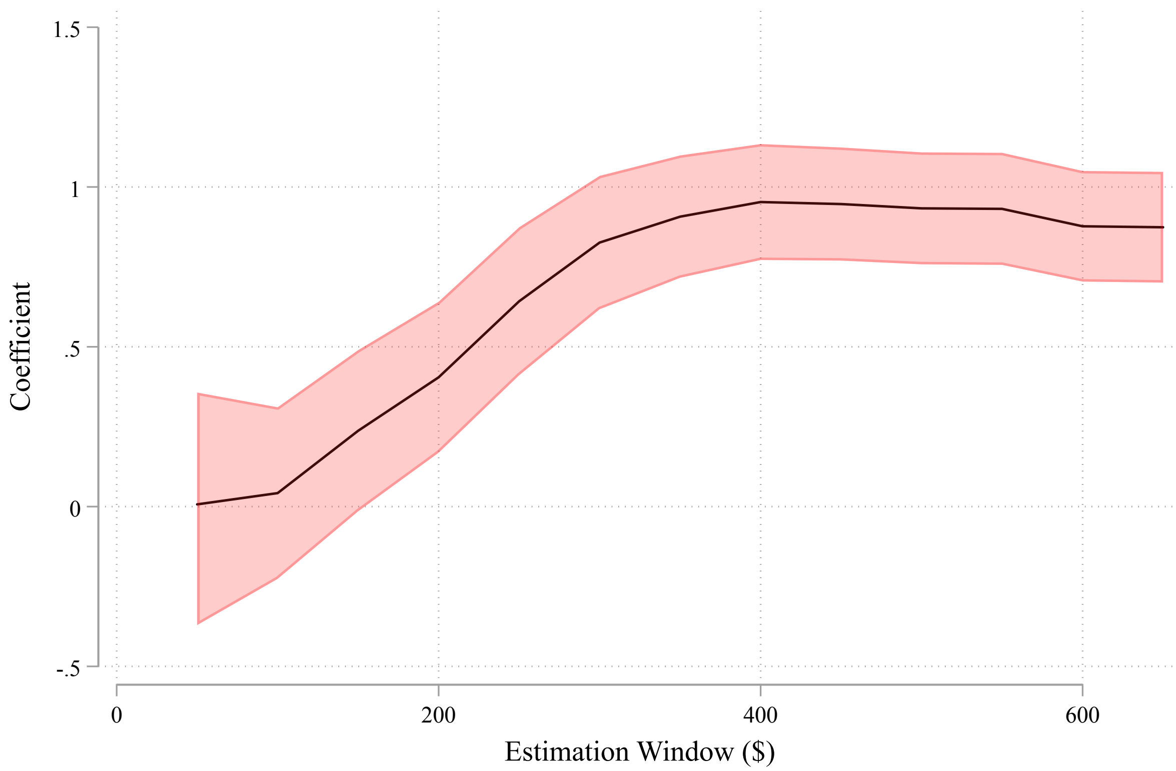

Figure A3: Placebo - Ineligible Workers - Excess and missing mass around the partial UI notch by observed death risk in county

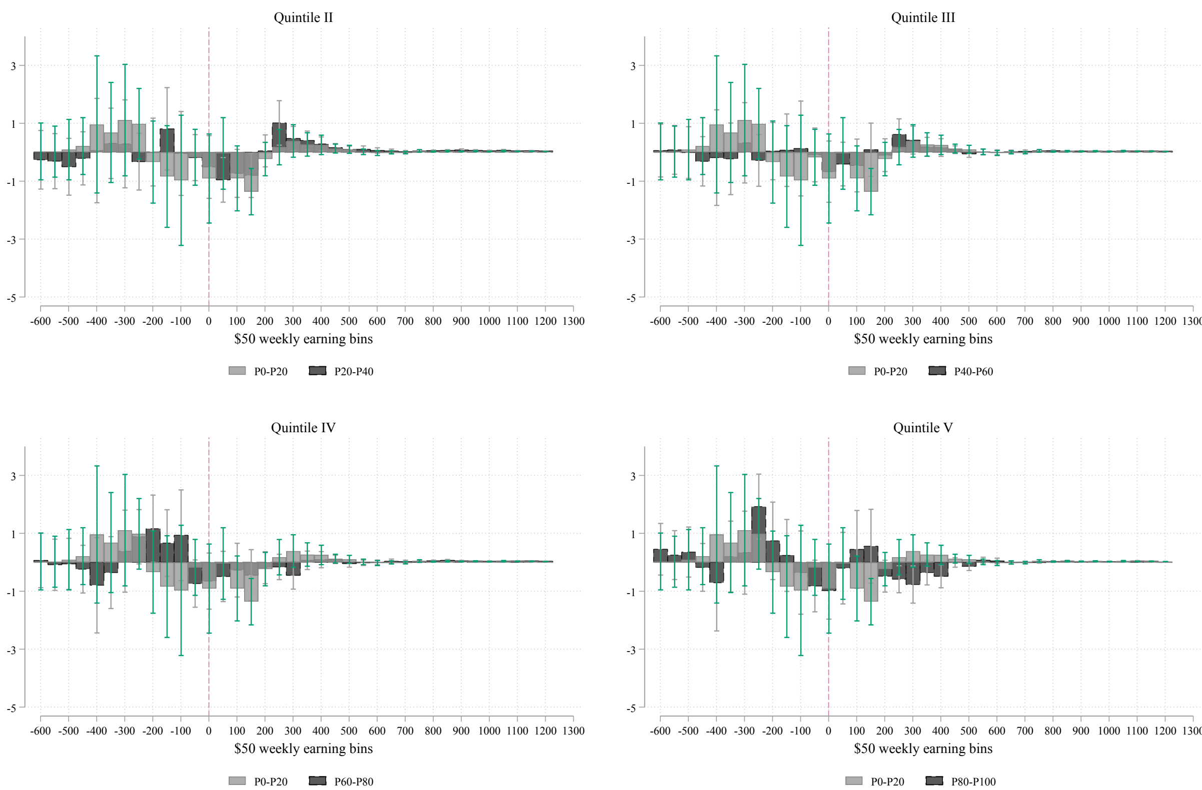

Note: The Figure shows $\delta_{r,\theta}$ coefficients from equation 4. Results for 4 different Covid-19 risk levels (θ) are plotted in each panel in blue. Covid-19 risk is measured as deaths per million in the week in the local area. The red bars are the benchmark response in areas with 0 Covid-19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

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.[48]

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.[49] 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.

[48] 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.

[49] 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)		

Note: Mean coefficients and standard errors in parentheses. ASEC and HB Full data include 2019 and 2020. QWI data is 2019 only. Column (3) ASEC is restricted to the 21 HB states. Column (4) ASEC sample is restricted to hourly workers, not self-employed, in small businesses (< 25 employees) corresponding to HB NAICS codes in HB states. Column (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. Column (6) QWI sample is restricted to privately owned small firms (< 20 employees) in HB states. 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, Forestry, Fishing and Hunting*						
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, Quarrying, and Oil and Gas Extraction*						
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 and Warehousing*						
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 and 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 and Rental and Leasing*						
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, Scientific, and Technical Services*						
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 of Companies and Enterprises*						
Average obs	149,666	68,697	102,192		2,351	36,842
Percent	0.09	1.33	0.10		2.73	0.26
56 *Administrative and Support and Waste Management and Remediation Services*						
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 Services*						
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 and Social Assistance*						
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, Entertainment, and Recreation*						
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 *Accommodation and Food Services*						
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 (except Public Administration)*						
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 Administration*						
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

Note: ASEC and HB Full data include 2019 and 2020. QWI data is 2019 only. Column (3) ASEC is restricted to the 21 HB states. Column (4) ASEC sample is restricted to hourly workers, not self-employed, in small businesses ($<$ 25 employees) corresponding to HB NAICS codes in HB states. Column (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. Column (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.[50] 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.[51] American Rescue Plan Act extended FPUC through September 6, 2021.[52]

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.[53]

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.[54]

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.[55] All 21 study states were approved for LWA, but only West Virginia picked the $400 weekly benefit option.[56] Individuals receiving 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

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

[51]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.

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

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

[54]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.

[55]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.

[56]Lost Wages Assistance Approved States, www.fema.gov The Office of the Governor Jim Justice (West Virginia) press release dated September 9, 2020, https://governor.wv.gov.

unemployment benefits.[57]

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.[58] All 21 study states except Florida received 6 weeks of funding. Florida was approved for 4 weeks, until the week ending August 22, 2020.[59]

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.[60]

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.

[57] 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`.

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

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

[60] 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:

$$\tilde{m^o} = \tilde{a} + e\tilde{w} - \gamma\tilde{y} \tag{6}$$

$\tilde{x}$ 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^*.[61] 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:

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

$$\Delta m = E/h_0 \tag{7}$$

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.[62]

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.[63]

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

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

[63] 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'):

$$h_1 = E + h_0', \tag{8}$$

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 8, 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 7 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.[64] A large empirical literature estimates a linear labor supply function as in 6 (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

[64]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 6, 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.

E 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 A4. 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 A4: Border Counties in Sample

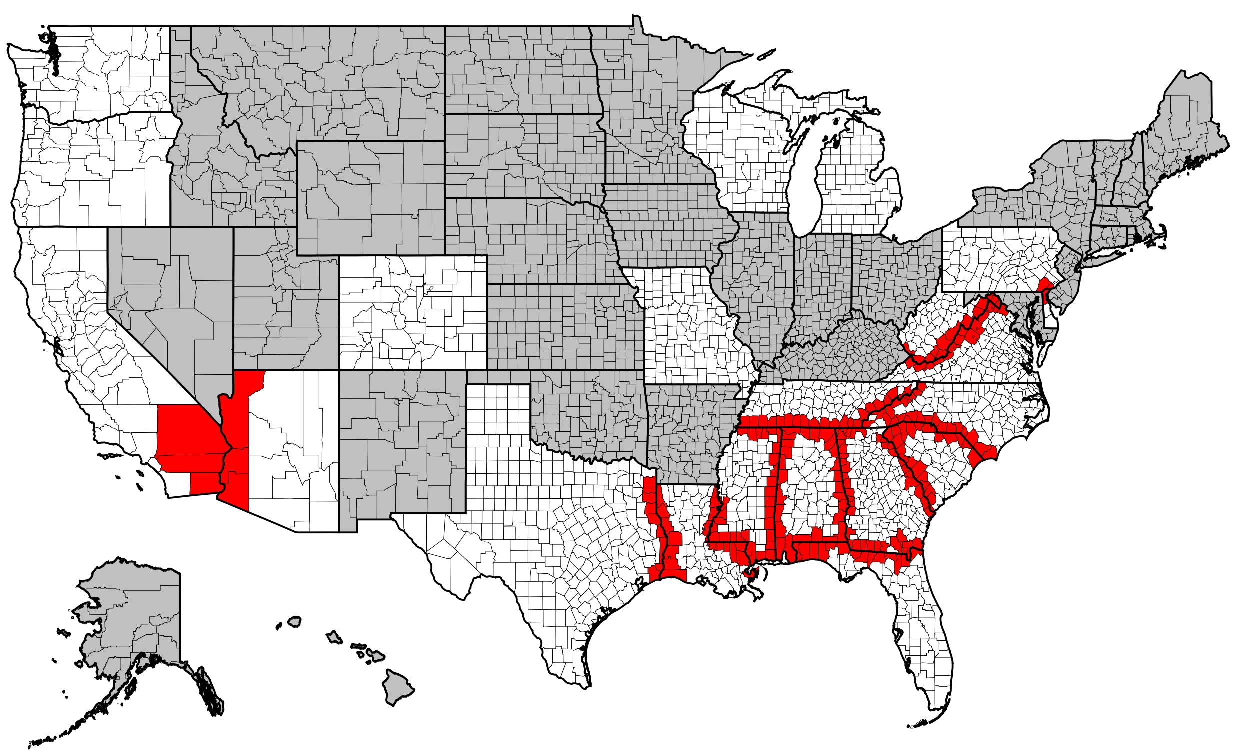

Note: The Figure shows counties along state boundaries that are included in our border sample.

F 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

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 4. The border sample is restricted to counties at state borders shown in Figure A4. Source: Homebase.

of $E_{i,c,t}$ from equation 5) and local Covid-19 outbreaks *within* the county $D_{i,c,t}$, measured as by the number of deaths per 1 million people.

Figure A5 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 A6 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 A6, 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 A6) to 50% of disposable income in very high-risk counties.

Figure A5: Effect of Workplace Safety on Labor Supply – Task Risk Proxy

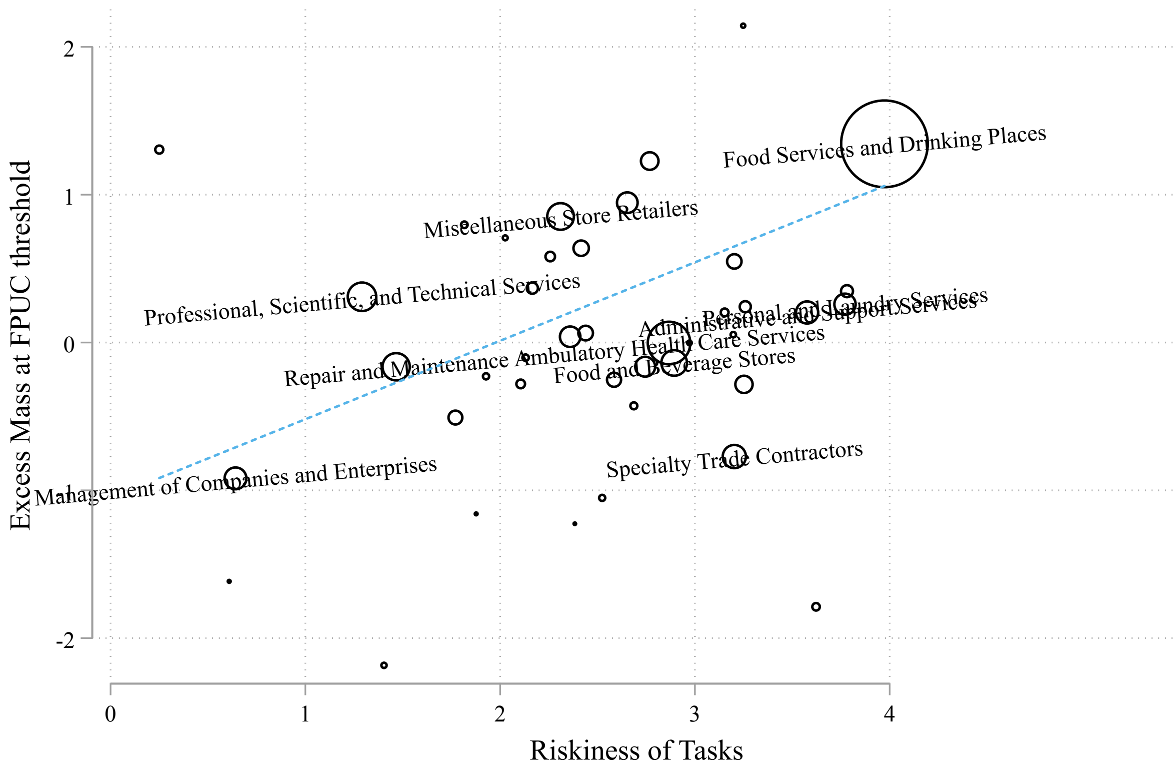

Note: The Figure shows the amount of excess mass at the FPUC threshold for 3-digit NAICS industries. The riskiness of tasks is the average risk of Covid-19 infections among tasks performed in the industry. The task risk data comes from Basso et al. (2020) and risk scores are standardized to have a standard deviation of 1. The y-axis shows the amount of excess mass generated by the FPUC eligibility threshold and is estimated in equation 3. 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 A6: Excess and Missing Mass around the partial UI notch – Fatality Rate in County

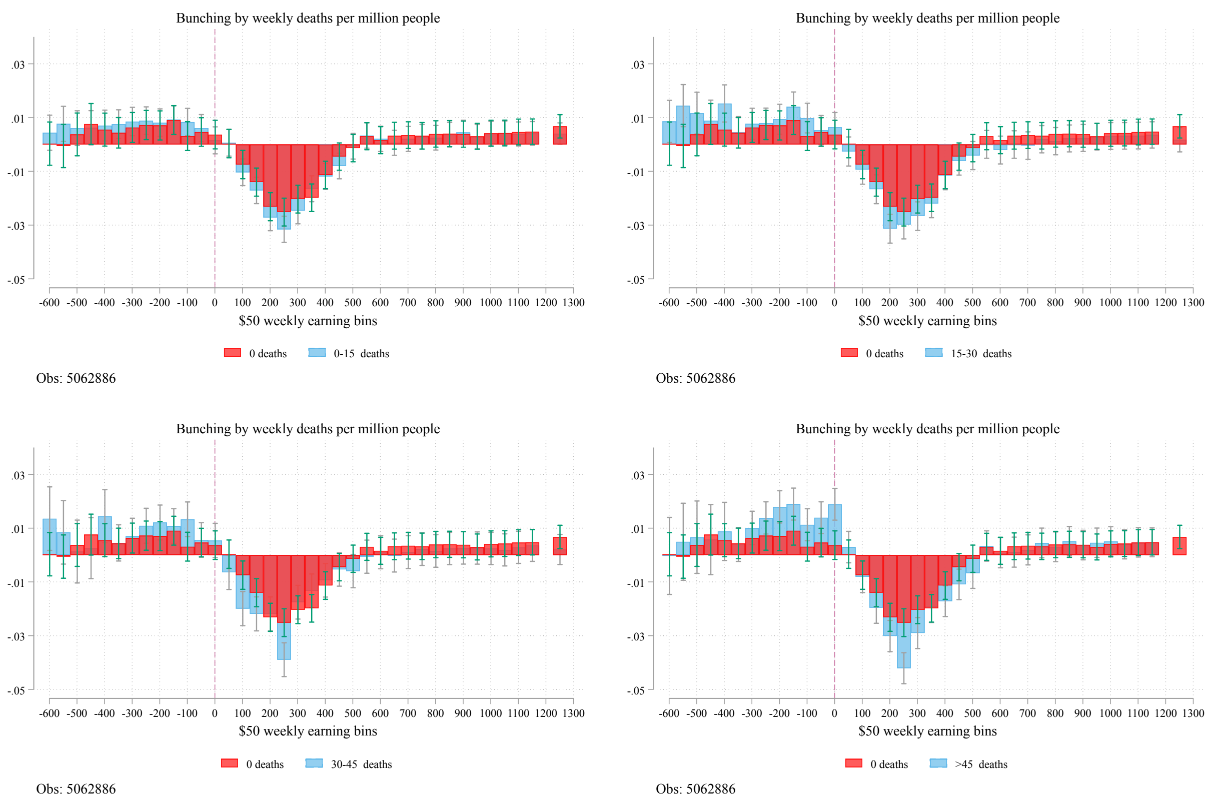

Note: The Figure shows $\delta_{r,\theta}$ coefficients from equation 4. Results for 4 different Covid-19 risk levels (θ) are plotted in each panel in blue. Covid-19 risk is measured as deaths per million in the week in the local area. The red bars are the benchmark response in areas with 0 Covid-19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

G Robustness Checks

Table A4: Robustness to Labor Demand Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Additional Excess Mass			
Workplace Risk	0.353	0.347	0.348	0.347	0.346	0.354	0.348
(std. dev.)	(0.0565)	(0.0561)	(0.0562)	(0.0562)	(0.0561)	(0.0566)	(0.0563)
Controls		# Employees	Small Business Revenues	Change in # merchants	Revenues X Merchants	Share of in-class instruction	All

Note: Columns (2) through (7) supplement the main specification of Panel B of Table 2 (also presented in column (1)) by controlling for demand shock proxies, interacted with a dummy for the Covid-19 period and a continuous income variable. Column (2) controls for the number of active employees from Paychex, Intuit, Earnin and Kronos, varying at state-week-industry level. Column (3) controls for the percent change in net revenue for small businesses from Womply, varying at state-week-industry level. Column (4) controls for the percent change in number of small businesses open from Womply, varying at state-week-industry level. Column (5) interacts the percent change in net revenue with percent change in the number of small businesses from Womply. Employment, revenue and merchants data are downloaded from Opportunity Insights Economic Tracker. Column (6) controls for the share of in class instruction from Parolin and Lee (2021), varying at county-month level. The share of in class instruction is defined as the complement of the share of all schools in an area with at least 50% year-over-year decline in visitors, consistent with the Parolin and Lee definition. Column (7) controls for all demand shock proxies together. Sources: Chetty et al. (2020b); Parolin and Lee (2021b); Parolin and Lee (2021a).