Bend and SNAP? The effect of food availability on workplace injuries

Gina Pagan

Department of Agricultural and Resource Economics, University of California, Davis

gpagan@ucdavis.edu

November 3, 2026

Abstract

Poverty is closely linked to poor health, but causal evidence on how financial hardship affects working-age adults is limited. I study whether running short of resources raises the risk of workplace injury among low-wage workers, who are concentrated in dangerous jobs and among whom hardship and reliance on the Supplemental Nutrition Assistance Program (SNAP) are common. The SNAP cycle provides a way to isolate scarcity from other differences between poor and non-poor workers: recipients spend their benefits quickly, and food spending and calorie intake fall until the next disbursal. I study this question in Oregon, where the last digit of the case head's Social Security number determines the day on which the household receives SNAP benefits, allowing me to compare similar workers at different points in the cycle on any given day. Using administrative workers' compensation claims over 2005–2018 and a predicted measure of SNAP eligibility, I find no detectable cycle when occupations are weighted equally, but injury rates are about 5.7 percent higher in the last ten days of the benefit month than in the first six when occupations are weighted by injury risk. This increase is concentrated among men aged 18 to 24, whose injury rates rise by 28.5 percent in the second half of the benefit month, and, less precisely, among women aged 45 to 54. The composition of injuries also changes: on the day benefits are received, the rate of the most severe injuries rises by about 28.1 percent while less severe injuries are unchanged, a shift that changes in hours worked would not produce. These intent-to-treat estimates likely understate effects for SNAP participants by a factor of 2.1 to 3.2.

1 Introduction

Low-income households in the United States (U.S.) experience worse health outcomes and face additional barriers to healthcare access compared to high-income households (Braveman et al., 2010; Ohlson, 2020). Disparities in food access and quality are significant contributors to this socioeconomic health gradient, as evidenced by the well-established associations between food insecurity and health conditions arising from poor nutrition (Seligman et al., 2007; Eicher-Miller et al., 2009) or heightened stress levels (Black et al., 2012). Workers with comparatively low earnings also experience higher rates of workplace injuries (Steege et al., 2014), higher mortality rates from injuries (Steenland et al., 2003), and are less likely to have employer-provided healthcare coverage (BLS, 2023).

It is important for policymakers to have a complete picture of how acute financial strain affects the health and safety of workers. There are around 3 million nonfatal and 5,000 fatal workplace injuries reported in the U.S. every year, many of which are borne by workers with low socioeconomic status (Dembe, 1999; d'Errico et al., 2007; Guardado and Ziebarth, 2019), and yet little is known about whether such strain impacts worker safety.[1]

This paper aims to fill this gap by examining whether the depletion of financial resources affects the risk of workplace injury for low-wage workers. To answer this question, I leverage an institutional feature of the SNAP program in Oregon where the day that a participating household receives its benefits each month is determined by the last digit of the household head's social security number (SSN). This schedule generates exogenous variation in the number of days since last receiving SNAP benefits among otherwise similar workers, and other researchers have used similar empirical strategies to link resource depletion to reduced test scores (Gassman-Pines and Bellows, 2018; Cotti et al., 2018; Kuhn, 2018; Bond et al., 2022), crime and disciplinary infractions (Gennetian et al., 2016; Carr and Packham, 2019, 2021), health and healthcare utilization (Seligman et al., 2014; Arteaga et al., 2018; Cotti et al., 2020). Since the last digit of one's SSN should otherwise be exogenous from a worker's risk of injury on any given date, any systematic pattern that emerges when comparing similar workers with different digits can only be attributable to SNAP receipt.

I use administrative workers' compensation (WC) claims records from 2005 through 2018 where critically I observe the last SSN digit for each injured worker and the date that the injury happened, allowing

[1] The Bureau of Labor Statistics (BLS) estimates that in 2024, the nonfatal occupational injury rate in the warehouse industry (median wage $\approx$ \$26; 45 percent Black and Hispanic) was 4.2 percent, while the rate in finance and insurance (median wage $\approx$ \$38; 23 percent Black and Hispanic) was only 0.3 percent (BLS, 2026b,a).

me to determine the distance in days between the injury date and the date of last SNAP receipt.[2] However, I do not observe SNAP participation, so my analysis restricts to workers whose observable characteristics - namely their gender, age, occupation, and wage - indicate eligibility for SNAP. Thus all results I present correspond to intent-to-treat (ITT) effects that likely understate effects for true participants.

I find no evidence of an injury cycle over the SNAP benefit month when occupations are equally weighted. However, when I weight occupations by their injury risk I find that injury rates rise by about 5.7 percent in the final ten days of the benefit month relative to the first six, which suggests that resource availability matters most for workers in relatively dangerous jobs. This end-of-month increase is concentrated among young men (aged 18 to 24), whose injury rates rise by 28.5 percent in the second half of the benefit month, and to a lesser extent older women (aged 45 to 54). Unweighted estimates suggest that injury rates for young women move in the opposite direction, though this pattern weakens when occupations are weighted by injury risk, and further investigation is required to determine its mechanism and why effects do not systematically appear for workers between 25 and 44.

I further demonstrate that the benefit cycle affects the *composition* of injury severity, though my results indicate that this effect operates through a different channel. The point estimates place the increase in injuries late in the benefit month among mild to moderate injuries, 75 percent of which are sprains and strains and a third are from from overexertion, which would be consistent with increasing work effort or inattention, though these estimates are imprecise. Conversely, I find the rate of severe injury is about 21.7 percent higher on the day that benefits are received relative to the following 20 days, while less severe injuries are unaffected. A change in hours worked would scale injuries of every severity proportionally, so this shift in the composition of injuries indicates that the timing of food assistance affects safety on the job directly.

It remains an open question why I do not find an effect for workers in the middle age groups. One possibility is that take-up rates may be lower in these subgroups, in which case estimates for this sample would be disproportionately attenuated relative to others where participation rates are high. Assuming this is not the case, another explanation could be that workers age 25 to 44 have significantly different household compositions than the youngest and oldest workers in my sample. They may be more often in two-adult households, which plausibly could reduce acute financial burden or allow households to adjust labor supply more flexibly when resources are scarce. They may also be more likely to have school-age children, and

[2] Ongoing work extends this analysis to other states.

school meals should in theory slacken the household food budget. In an on-going analysis, I try to test this directly by comparing injury profiles for these workers during the summer and school year.

This work contributes to three distinct strands of literature. By linking availability of food assistance to injury risk, I expand our knowledge of the socioeconomic threats to the welfare of low-income households (Bitler et al., 2005; Hoynes et al., 2012; Bitler and Hoynes, 2015) and document that food assistance has health benefits that extend beyond nutrition (Gregory and Deb, 2015). This paper contributes to the body of work linking the cyclical timing of food assistance receipt to reduced-form economic and health outcomes to elucidate how benefit receipt and subsequent depletion affects behavior and cognition. My paper is the first to my knowledge to link the timing of food assistance with workplace health and safety outcomes. Much of the existing work in this area has found similar adverse end-of-benefit month effects consistent with resource depletion (Gennetian et al., 2016; Carr and Packham, 2019; Cotti et al., 2020; Bond et al., 2022), while other papers have studied how receipt affects behavior (Cotti et al., 2016).

Finally, this paper contributes to our understanding of how economic conditions affect worker safety. Recent papers have examined the effect of minimum wage laws on occupational injuries, with mixed findings - both Davies et al. (2026) and Hradil (2018) find a positive relationship between the minimum wage and workplace injuries, while LaFiandra and Schwab (2023) finds a negative effect in the restaurant industry. Other papers have considered how economic conditions affect reporting behavior of workplace injuries (Card and McCall, 1996; Boone and Van Ours, 2006; Boone et al., 2011; Dillender, 2015). My work illustrates that the probability of workplace injury is responsive to short-run economic scarcity.

This paper proceeds as follows. Section 2 provides relevant background about SNAP and WC and how both programs operate in Oregon. In section 3, I discuss my data sources, provide summary statistics, and introduce my empirical strategy. Section 4 presents my main results and explores heterogeneity along several dimensions, and in section 5 I discuss my exploration of potential mechanisms. Finally, section 6 concludes and discusses future directions.

2 Background

2.1 Supplemental Nutrition Assistance Program

SNAP is funded through the Department of Agriculture but the details of program administration are delegated to individual state agencies. Since 2008, benefits have been automatically deposited into participant

accounts linked to debit-like cards that can be used at authorized retailers. SNAP is a means-tested program and eligibility is primarily determined by a household's financial means being below a treshold and family composition. Individual state agencies determine the specific eligibility criteria for participation in that state conditional on meeting federal requirements. Federal guidelines stipulate that households are eligible for SNAP if their gross monthly income is at or below 130 percent of the federal poverty line (FPL) and net income is below the FPL, but most states have set gross income requirements above the 130 percent requirement. SNAP participants must generally either be citizens or legal permanent residents, though states have some flexibility in theory to offer benefits to mixed-status households.

State departments are also responsible for setting the monthly benefit disbursement schedule. Most states stagger disbursement throughout the month to minimize burden for retailers and reduce the probability of stock-outs, though a handful of less populated states distribute all benefits on a single day. Among the states that stagger disbursement, there is significant variation in both the timing and variable(s) used to randomly assign benefit days to participants. For example, Connecticut and Delaware both use the first letter of the client's last name to determine their benefit day, but Connecticut staggers benefit disbursement over the first three days of the calendar month while Delaware staggers issuance from the second through the 22nd. Alternatively, many states use characteristics of individuals' SNAP case numbers, SSNs, or birth dates to assign benefit days.

2.2 Workers' Compensation

Workers' compensation (WC) insurance is a state-regulated system that provides medical benefits and missed wages to employees who are injured at work or incur an illness attributable to work exposure. Importantly, injured workers are entitled to benefits regardless of fault or negligence. There is no federal law requiring employers to have WC insurance, but most state departments have laws mandating universal coverage, Texas being the exception (NASI, 2024).[3] All employers in states with mandated coverage are generally required to purchase WC insurance, but there are typically exemptions for a subset of *employees*. State laws vary in which employees are exempted from coverage, though some of the common exemptions are private household workers, independent contractors, and casual workers. In 2022, the average cost to employers for providing WC insurance was 1 percent of covered wages, and the total cost to employers was

[3]Using workers' compensation data from Texas, Cabral et al. (2022) estimate the demand for WC insurance and argue that mandating coverage may not be surplus-maximizing relative to elective coverage.

approximately $103 billion.

Given the no-fault nature of WC benefits, an extensive body of economic research has attempted to size the degree of moral hazard in WC injury reporting. Krueger (1990b) and Krueger (1990a) show that increases in WC benefit generosity induces more participation and leads to workers' taking more days away from work after sustaining an injury, though Bronchetti and McInerney (2012) argue that the elasticity of participation to changes in benefit generosity is less economically meaningful than previously thought. Smith (1990) provides the first evidence of a 'Monday effect' where sprains and strains are more likely to be reported on Mondays, indicating that workers may be filing WC claims for off-the-job injuries, though Card and McCall (1996) find that uninsured workers are no more likely to report Monday injuries than those likely to have other medical coverage. However, more recently Dillender (2015) documents that reliance on WC for medical care does in part depend on availability of other health insurance coverage.

2.3 SNAP and WC in Oregon

Oregon was one of the early adopters of the EBT system (1996) and of broad-based categorical eligibility (BBCE) for SNAP participation (2000). Under BBCE, households that are eligible for Temporary Assistance for Needy Families (TANF) are automatically considered eligible for SNAP. Oregon households are eligible for SNAP if their gross monthly income is no higher than 185 percent of the FPL and their net monthly income is no higher than 100 percent of the FPL, regardless of household assets. Consequently, Oregon has been consistently ranked in the top 5 states nationwide for SNAP accessibility (FNS, 2025). Almost 18 percent of the Oregon population participated in SNAP on average over my study period (see appendix figure A.2). At the state's peak in 2012, more than 1 in 5 persons received SNAP benefits and the total value of benefits issued exceeded $1.2 billion.

Oregon staggers SNAP benefit issuance over the first 9 days of the calendar month using the last digit of the primary account holder's SSN. Households for whom the account holder's SSN ends in 0 receive benefits on the first of the month, and those where the relevant digit d falls between 1 and 9 are issued benefits on calendar day d. Non-citizens without SSNs obtaining SNAP benefits on behalf of citizen children also receive their benefits on the first of the month. In appendix figure A.1, I show how the SNAP benefit month overlaps with calendar time for a 31-day month.

All employers with at least one subject worker who is not an independent contractor are required to purchase WC insurance in Oregon. Workers are considered subject workers unless they meet one of the

approximately 30 exempted conditions. Most of the exemptions pertain to sole proprietors, though privately employed household workers, volunteers for nonprofit or charitable organizations, and referees for recreational sports are also excluded from the subject worker definition. Employers may purchase coverage from one of the several hundred authorized insurers or they can opt to self-insure, conditional on being approved by Oregon's Department of Consumer and Business Services (DCBS). Premium rates in Oregon generally fall below the national median, which in 2024 meant that Oregon employers paid $0.89 per $100 of payroll while employers in states at the median paid a little over a dollar per $100 of payroll. However, employers and workers also each contribute roughly 1 cent per hour worked to Oregon's Workers Benefit Fund, which was established to supplement financial support for workers whose injuries meet certain extenuating circumstances.

When an injury occurs, the injured person fills out a first report of injury (FROI) form detailing the event and submits it to their employer. The employer then has 5 days to report the claim to the insurer *if* the injury required medical treatment beyond simple first aid. Employers are not required to report injuries that do not require any medical treatment. The insurer then has 60 days to accept or deny the claim and classify accepted claims as either disabling or non-disabling. Accepted disabling claims entitle workers to coverage for medical services and disability or death benefits, while workers with non-disabling claims are only entitled to coverage for medical services. After a decision has been made, the insurer must report all accepted disabling claims and all denied claims to Oregon's DCBS within 14 days. The insurer is not required to report accepted non-disabling claims to DCBS. The injured worker has 60 days to appeal decisions regarding the disability status of their claim or request a hearing to dispute a denial. Workers that are granted indemnity payments for lost wages receive installments equal to two-thirds of their average weekly wage (AWW) paid on a biweekly basis.

Indemnity payments compensate workers for lost wages after a disabling injury, and the exact award owed to the worker depends on whether a doctor has approved *any* work-related activity and on whether disability persists after the worker is deemed medically stationary. Most accepted disabling claims result in a *temporary* disability award, which entitles workers to payments for lost wages at two-thirds of their average wage rate. If reduced or modified work is deemed permissible, workers are entitled to temporary *partial* disability (TPD) payments, but if all work-related activities are prohibited, the worker receives temporary total disability (TTD) payments. After the worker is medically stationary, any persistent disability that reduces capacity to work entitles the worker to *permanent* disability payments that are either partial (PPD)

or total (PTD).

3 Methods

3.1 Data sources

My primary data are administrative records for all accepted disabling claims filed by workers in Oregon for injuries occurring between 2005 and 2018 obtained from DCBS. For each claim, I observe the date the injury occurred, the date the employer was notified, the county where it occurred, the indemnity award, the industry of the employer, the cause of the injury (e.g., transportation accident vs. violent act), the nature of the injury (e.g., traumatic injury vs. infectious disease), and the part of body injured. I also observe information about each claimant, namely their gender, age (binned), occupation, AWW, and the last digit of their SSN, which is the variable used to assign one's SNAP issuance day.[4] Occupations are reported as Standard Occupational Classification (SOC) titles which vary slightly over time, so I standardize all occupations to be consistent with the 2010 version published by the Bureau of Labor Statistics (BLS). Of the claims with non-missing occupation, over 90 percent of occupations could be linked to the most granular (6-digit SOC code) level, while the remaining were linked to 5-digit codes.

Ideally, I would observe whether or not each claimant was a SNAP participant when their injury occurred and if so, whether or not they were the case head for their household. The second-best case would be if I could observe their household size and gross income, which would allow me to directly impute SNAP eligibility. In lieu of this, I instead use the variables I do observe to approximately identify SNAP eligible workers based on data from the American Community Survey (ACS) sourced from IPUMS (Ruggles et al., 2025). The ACS is a nationally representative household survey conducted by the US Census Bureau since 2005, with an annual sample size of more than 3 million households, making it the largest publicly available source of detailed microdata for US households. Among other topics, ACS respondents are asked to report their wage and salary income, the number of weeks worked, and whether anyone in their household received SNAP benefits in the 12 months prior to enumeration. IPUMS also constructs a variable indicating where respondents' household income falls relative to the FPL. I restrict to ACS respondents in the 2005 to 2018 samples who were in the labor force, between the ages of 18 and 64, and who had worked for a wage or

[4] If employers reported wages at a different interval, DCBS automatically converts the reported wage to the weekly equivalent, as the claimant's AWW determines their indemnity payments.

salary in the 12 months prior to enumeration. I calculate each respondent's "weekly wage" by dividing their reported annual wage and salary income by the number of weeks worked. While somewhat imprecise, this is actually a reasonably good approximation of the average weekly wage solicited by WC insurers. I then construct the weekly wage distribution in each year to assign respondents to their respective quintile.

I also rely on several secondary data sources to supplement this work. I use the FNA SNAP Quality Control (QC) database for fiscal years 2005 to 2018 to generate descriptive statistics of SNAP participants in Oregon over my study period (FNA, 2026). The SNAP QC data come from monthly random samples of active SNAP cases generated by each state department to be used in federal program audits. Each case record contains administrative information on the economic and demographic characteristics of all members of the SNAP household, including their age, gender, and employment status, and samples are constructed to be representative at the state level. The number of households sampled in each state-year depends on the average statewide monthly caseload for that fiscal year and whether or not the state was approved to submit a restricted sample. Oregon consistently had caseloads exceeding the maximum threshold, so I observe roughly 1,000 households every year. Since each record in the QC database is rigorously verified for federal use and samples are representative within-state, it can more reliably characterize SNAP households than the ACS or other household surveys.

Finally, I use BLS estimates on the number of injuries by days away from work (DAFW) for each nature of injury and injured body part to construct an injury severity index for claims in my sample. The reported estimates are based on data from the annual Survey of Occupational Injuries and Illnesses (SOII) from 2011 through 2018 (BLS, 2018a,b,c).[5] For each nature and body part classification, BLS reports the total number of injuries in each year and the number of injuries in incremental DAFW bins (1 day, 2 days, 3-5, 6-10, 11-20, 21-30, 31+). The BLS does not publish injury counts by nature $\times$ body part $\times$ DAFW, which would be the ideal metric to determine severity of my observed claims, so I instead use an iterative proportional fitting (IPF) algorithm to generate a fitted distribution that replicates national injury counts by nature $\times$ body part, nature $\times$ DAFW, and body part $\times$ DAFW (Deming and Stephan, 1940). I detail how I construct this index in appendix B.

[5] 2011 is the earliest year these data are available.

3.2 Estimating SNAP eligibility

My data do not allow me to observe individual workers over time, so instead I use the ACS data to estimate the probability of being eligible for SNAP at the gender×age×wage×occupation cell level. I use age in 10-year bins, quintiles of real weekly wage (in 2010$), and SOC groups predominantly at the 5-digit, or broad group, level. A subset of broad codes are aggregated in the public ACS data if there is either no exact match to a Census occupation code or if there were fewer than 10,000 workers nationally. To maintain comparability with ACS observations, I mirror this aggregation in my claims data. I ultimately estimate SNAP eligibility using 384 unique SOC groups; of these, 336 correspond exactly to broad SOC groups (5-digit) and 7 correspond exactly to minor SOC groups (3-digit).[6] The remaining groups are aggregations within minor groups that mirror ACS groups.[7]

For each cell, I nonparametrically estimate the share of workers that are in SNAP eligible households. For simplicity, I assume that households are SNAP eligible if they either reported SNAP receipt in the last year *or* if household income was at or below 185 percent of the FPL, which is the BBCE limit in Oregon. SNAP participation rates among the eligible working poor population in Oregon were consistently over 75 percent between 2005 and 2018, so though I am roughly estimating the probability of being *eligible* for SNAP, this is a reasonably close upper-bound on the probability of SNAP participation.

When cells are comprised of many observations, this is an assumption-free and transparent approach to approximating the probability of SNAP eligibility for workers in that cell, but these estimates are increasingly noisy and unreliable as the number of observations-per-cell falls. To reduce the noise of estimates for sparse cells, I blend the empirical estimate with predictions obtained with a gradient boosting machine learning algorithm based on the same predictors, which I term the "GB estimate". Gradient boosting sequentially builds a set of individual decision trees, where each new tree is trained to correct the residuals of the trees before it. Each successive tree is fit to the gradient of a loss function, and predictions from all trees are then summed to produce the final output, with a learning rate scaling the contribution of each tree. Conversely, a random forest (RF) procedure builds many independent trees in parallel, each on a bootstrapped sample of the data, then averages their predictions. I found that GB performed better than RF at out-of-sample

[6] The 3-digit groups are post-secondary teachers (25-1000), other teachers and instructors (25-3000), other healthcare practitioners and technical occupations (29-9000), supervisors of personal care and service workers (39-1000), funeral service workers (39-4000), fishing and hunting workers (45-3000), and supervisors of transportation and material moving workers (53-1000). These SOC groups are rarely, if ever, reported at more granular levels in the ACS.

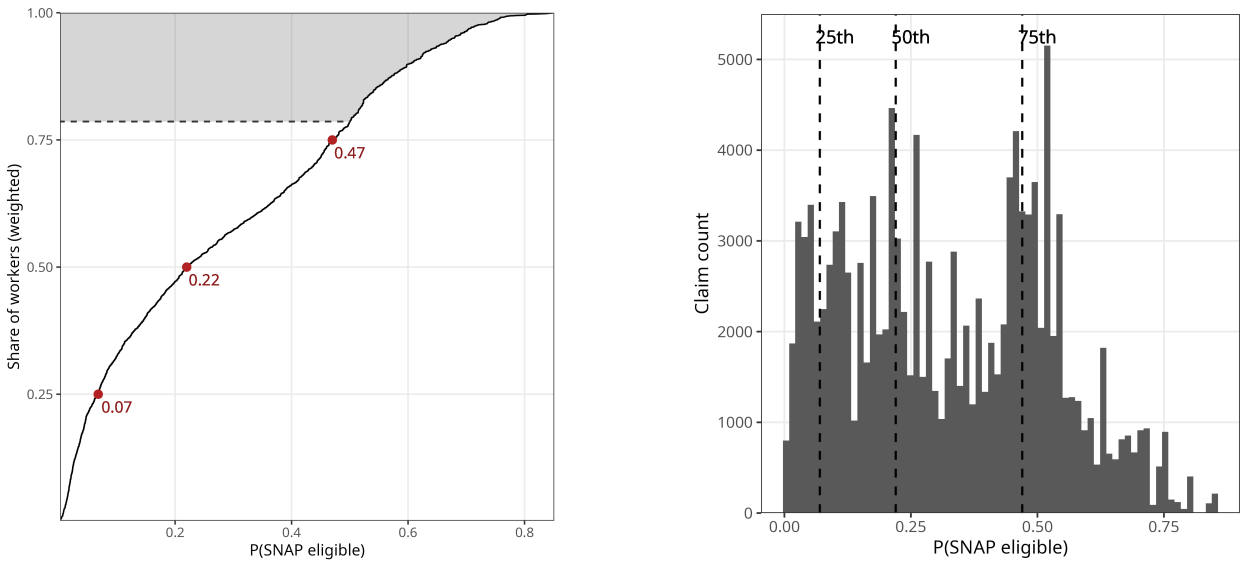
[7] For example, mathematicians (15-2020) and statisticians (15-2040) are aggregated to 15-20XX.

predictions of SNAP eligibility, which made it the preferred choice in my setting.

The weight assigned to the empirical estimate when blending increases smoothly with the number of ACS observations for that cell (N). Specifically, I construct the weights so that the empirical estimate and GB estimate are weighted equally when $N = 25$. The downside of this method is that I only generate eligibility probabilities for cells that exist in the ACS data. Appendix figure A.3 compares the distribution of SNAP eligibility rates across the empirical (panel (a)), the GB (panel (b)), and the blended (panel (c)) specifications. The empirical distribution is relatively smooth, though it is clear that sparse cells generate excessive estimates at 0 and 1. The relative lumpiness of the GB distribution is an algorithmic feature. Like a random forest, the GB algorithm is an ensemble of regression trees so its predicted eligibility is a piecewise-constant function of the covariates and can therefore take only a finite menu of distinct values. The lumpiness in the distribution is more pronounced than a random forest would produce because boosting sums over multiple shallow, shrunken trees that are sequentially fit on a set of coarse predictors, whereas a random forest averages hundreds of deep, independent trees and thus yields a finer grid of output values. The bunching at end-points is resolved in the blended distribution, while the remainder of the distribution closely resembles that of the purely empirical estimates. This is evidence that the blending approach has the intended advantage of generating more plausible estimates for sparse cells.

In panel (a) of figure 1, I show the cumulative distribution function (CDF) of $\hat{p}$ over the sample of ACS workers, weighted at the individual level. Panel (b) shows the distribution for workers in my claims data. The dashed vertical lines in panel (b) indicate the worker-weighted quartiles, which correspond to the labelled points in panel (a). Federal estimates based on ACS data indicate that 18 percent of the Oregon population on average received SNAP over my study period, with approximately 90 percent take-up among eligible working households. This means that my estimates should predict that roughly $18/0.9 = 20$ percent of the worker population was eligible for SNAP. The area of the shaded trapezoid in panel (a) is 21.4 percent, which suggests that I am conservatively identifying SNAP-eligible workers.

I next present additional indirect evidence that my probability measure $\hat{p}$ is positively correlated with SNAP *participation*. Let $b(s)$ denote the assigned SNAP receipt day for individuals with last digit s. Since terminal digits are randomly assigned and the mapping to issuance days is known, one can compute the expected mean value of $b(s)$ on any particular date, which is indicated with a horizontal line at $\overline{b(s)} = 4.6$. If one's terminal SSN digit is wholly exogenous from the probability of injury, which should be true for non-participants, the average value of $b(s)$ among workers injured on any day should always oscillate around



(a) CDF: ACS weighted workers

(b) Histogram: Observed claims

Figure 1: Distribution of P(SNAP eligible) $\hat{p}$

Note: Panel (a) plots the cumulative distribution function (CDF) of $\hat{p}$ over the population of weighted ACS respondents, while panel (b) plots the distribution of $\hat{p}$ over claimants in my sample. The dashed vertical lines in panel (b) denote worker-weighted quartiles corresponding to the labelled points in panel (a). Please refer to the surrounding text in this section for more information about the ACS sample and how $\hat{p}$ was estimated.

$\overline{b(s)}$.

Figure 2 plots the average benefit day among injured workers $\overline{b(s)}$ over the average calendar month for workers above $\hat{p} = 0.5$, which should denote the threshold determining probabilistic eligibility for SNAP. The mean $\overline{b(s)}$ for these workers consistently falls below the expected average, and there is a stark negative deviation from this expected value at the end of the calendar month, precisely when those who receive their benefits on the 1st, 2nd, and 3rd of the month are most likely experiencing food scarcity. I consider this to be descriptive evidence that workplace injuries do respond to the SNAP benefit cycle.

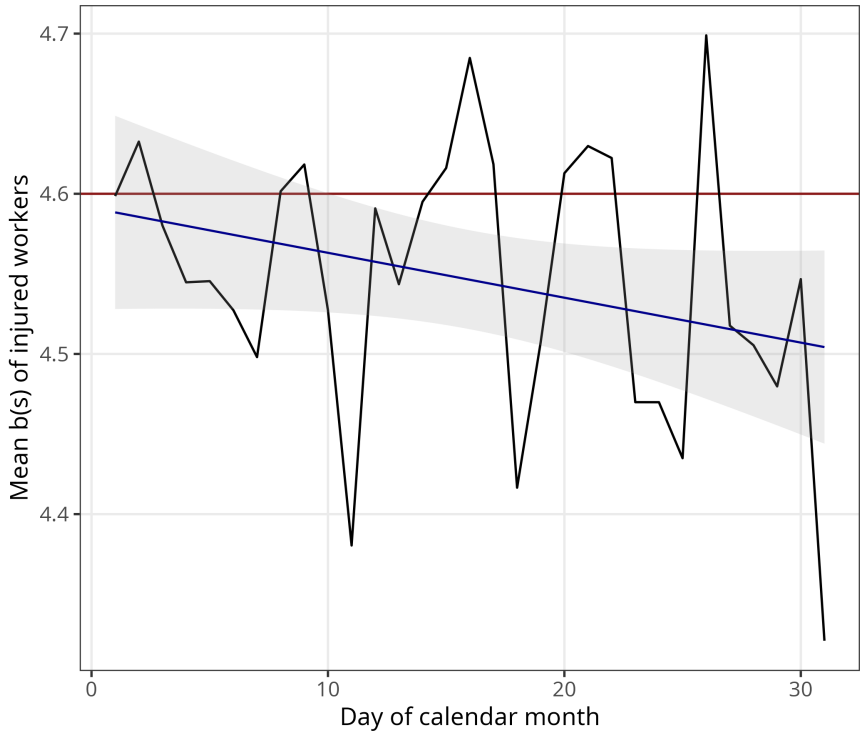


Figure 2: Average issuance day $b(s)$ among injured workers

3.3 Summary statistics

I impose several limitations to arrive at the final set of claims in my analysis sample in order to facilitate my empirical strategy. I discuss the rationale behind each limitation below.

I first restrict to claims filed by Oregon residents between the ages of 18 and 54, as over 90 percent of working female SNAP recipients and roughly 68 percent of working male SNAP recipients in this age range were their respective household head (HH) over this period.[8] This matters for my empirical strategy, which relies on the assumption that I observe the case head in the data. I rely on claimants' observable characteristics - namely gender, age, occupation, and real weekly wage quintile - to determine their probability of SNAP eligibility, so I generally drop claims with missing or zero values.[9]. The magnitudes of reduced-form effects of the SNAP cycle from prior work suggest that the effect size on occupational injuries will be small - recent studies find effects on test scores ranging from 0.01 to 0.17 of a standard deviation (Gassman-Pines

[8]Based on the author's analysis of data from the SNAP QC database, where a household head is considered 'working' if $\texttt{WORKi} = 1$ (available in FY 2015-2018) or if $\texttt{EMPSTAi} > 2$, $\texttt{EMPSTBi} > 1$, and earnings > 0 for FY 2005-2014.

[9]I impute a subset of missing wage values when there were at least 7 other claims in that year for the relevant gender-age bin-occupation cell with non-missing wages, to which I assign the cell-year median. The remaining claims with zero-valued wages are dropped

and Bellows, 2018; Cotti et al., 2018; Bond et al., 2022), and effects on healthcare utilization between 1 and 2 percent (Arteaga et al., 2018; Cotti et al., 2020). In most of these cases, the authors were able to observe SNAP participation status and assigned issuance day for individuals in their respective samples. Given that I do not observe SNAP participation and effect sizes are expected to be small, I limit to claims where injuries were reported within 2 weeks by workers who are most likely eligible for SNAP. I specifically focus on workers in the top quartile of predicted SNAP eligibility ($\hat{p} \geq 0.47$), which leaves me with 36,432 injuries in my final sample.[10]

	Women				Men			
	18–24	25–34	35–44	45–54	18–24	25–34	35–44	45–54
Share traumatic	0.96	0.94	0.91	0.89	0.95	0.93	0.9	0.89
	(0.20)	(0.23)	(0.28)	(0.31)	(0.21)	(0.26)	(0.30)	(0.32)
Severity index	0.5	0.51	0.52	0.53	0.5	0.51	0.53	0.54
	(0.08)	(0.08)	(0.08)	(0.09)	(0.1)	(0.1)	(0.1)	(0.11)
P(SNAP eligible)	0.56	0.61	0.57	0.55	0.52	0.58	0.58	0.56
	(0.05)	(0.1)	(0.08)	(0.07)	(0.05)	(0.09)	(0.1)	(0.08)
Real weekly wage (2010$)	255.86	316.94	329.93	282.54	297.91	365.39	377.51	359.22
	(99)	(109)	(107)	(113)	(113)	(111)	(106)	(113)
Share of claims by award								
Temporary disability	0.92	0.89	0.83	0.79	0.87	0.81	0.75	0.71
Permanent disability	0.05	0.08	0.12	0.15	0.10	0.14	0.18	0.22
Fatality	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
Unknown	0.03	0.04	0.05	0.06	0.03	0.05	0.07	0.07
Share of claims by cause								
Bodily reaction and exertion	0.42	0.52	0.51	0.47	0.32	0.39	0.41	0.42
Contact with objects and equipment	0.22	0.14	0.14	0.13	0.35	0.32	0.25	0.26
Falls, slips, trips	0.21	0.21	0.23	0.31	0.17	0.17	0.2	0.2
Transportation accidents	0.03	0.03	0.02	0.02	0.05	0.03	0.05	0.04
All other causes	0.12	0.11	0.09	0.08	0.11	0.08	0.08	0.08
Share of claims by major SOC group								
Health support	0.12	0.32	0.26	0.12	0.01	0.01	0.03	-
Food preparation, service	0.34	0.17	0.25	0.19	0.26	0.16	0.13	0.12
Building cleaning, grounds maintenance	0.14	0.14	0.21	0.47	0.17	0.06	0.08	0.09
Sales	0.11	0.09	0.07	0.06	0.03	0.05	0.02	0.03
Farm, forestry, fishing	0.01	0.03	0.06	0.02	0.03	0.04	0.13	0.16
Construction	-	-	-	-	0.10	0.16	0.12	0.12
Transportation, material-moving	0.08	0.11	0.05	0.05	0.21	0.31	0.38	0.38
All other	0.20	0.13	0.09	0.09	0.19	0.21	0.12	0.09
Sample size								
N unique SOCs	68	73	42	25	76	77	50	40
N injuries	3,468	5,693	4,799	2,265	3,845	7,749	5,490	3,391

Notes: The baseline sample includes injuries by workers in the top probability quartile of SNAP-eligibility ($\hat{p} \geq 0.47$). Each cell reports the mean across injury years of the within-year statistic, with the standard deviation across years in parentheses. Shares are computed among non-missing observations of the relevant variable.

Table 1: Summary statistics: Analysis sample

[10] Appendix table A.2 provides more detail on how I arrived at my final analysis sample.

Table 1 provides summary statistics for claims in my analysis sample. Between 89 and 96 percent of claims pertain to traumatic injuries that stem from a single incident. The relative share of non-traumatic claims, which include cumulative injuries and systemic illness, generally increases with age, as does injury severity and the share of claims earning permanent disability payments. The average probability of SNAP eligibility among workers in my sample ranges from 0.52 (18-24 men) to 0.61 (25-34 women). For women, average real weekly wages rise with age from $256/week for workers age 18-24 to $330/week for workers age 35-44, then fall to $283/week for workers age 45-54. Real wages similarly rise with age for men from $298/week for workers age 18 to 24 to $378/week for workers between age 35 and 44, but the oldest men in my sample have marginally lower earnings compared to younger men. The most common cause of injury is bodily reaction and exertion, which ranges from 32 to 52 percent of claims across groups. Young men have the highest incidence of injuries caused by contact with objects and equipment, while a mildly larger share of injuries for women stem from falls and other causes.

There is also predictable heterogeneity in the most common occupations by age and gender. Between 31 and 51 percent of women's injuries were for food service workers or healthcare support workers (e.g., nursing and home health aides), with the low end among women aged 45 to 54, nearly half of whose injuries are in building cleaning and grounds maintenance. Injuries for men are concentrated in transportation/material-moving and construction occupations, though over a quarter of men age 18 to 24 were also in food preparation and service occupations.[11] Farm, fishing, and forestry occupations account for approximately 8 percent of men's injuries, and 16 percent of injuries for men age 45 to 54.

3.4 Empirical approach

This section describes my empirical approach to determine how injury risk changes over the SNAP benefit month for workers in the top quartile of probability of SNAP-eligibility. The SNAP cycle literature commonly uses days since last SNAP receipt to proxy for relative food scarcity, though the exact functional form $f(DaysSince)$ is usually setting- and data-dependent. A standard approach is to estimate a continuous effect of days since SNAP issuance with both linear and quadratic terms (e.g., Bond et al., 2022). Another common approach is to estimate differences across weeks of the benefit month using benefit week indicators, which avoids imposing a functional form on the relationship being studied. I follow a similar approach and specify $f(DaysSince)$ as a set of binary indicators that segment the benefit month, and my main spec-

[11] Material-moving is somewhat of an umbrella category, including occupations from warehouse workers to grocery store baggers.

ification compares the first six days of the benefit month (including the day of receipt) to subsequent bins of 5 days. I determine days since SNAP at the claimant level as the number of days between the date of injury and the most recent date of SNAP issuance, which I infer from the claimant's last SSN digit.

Claims in my data are still largely anonymized and individual workers are not uniquely identifiable. To get a daily panel of injury outcomes, I instead aggregate injuries to the date level for groups of workers with key shared characteristics, namely their gender, age bin, occupation, and last SSN digit. The level of analysis is therefore the date (t) $\times$ occupation ($o \in O$) $\times$ demographic group ($d \in D$) $\times$ last SSN digit ($s \in [0, 9]$) level, where D is the set of unique gender $\times$ age combinations in my sample.[12] The basic regression specification is as follows:

$$\begin{aligned} Y_{s,d,o,t} &= \mathbf{X}'\boldsymbol{\beta} + \varepsilon_{s,d,o,t} \\ &\equiv \boldsymbol{\beta}_k' f_k(DaysSince_{st}) + \delta'\mathbf{Z}_d + \gamma_{b(s)} + \gamma_{o,m(t)} + \boldsymbol{\gamma}_{\tau(t)} + \varepsilon_{s,d,o,t} \end{aligned} \tag{1}$$

where $Y_{s,d,o,t}$ is a binary variable equal to 1 if cell sdo experienced an injury on date t ($1(\text{Injuries}_{s,d,o,t} > 0)$) or an injury count. $\mathbf{Z}_d$ is a vector of demographic indicators, namely gender, age, and the interaction. I include occupation $\times$ month-in-year ($\gamma_{o,m(t)}$) fixed effects to capture seasonal differences in injury risk across occupations, as well as a vector of other time fixed effects ($\boldsymbol{\gamma}_{\tau(t)}$) for year, day of month, day of week, and federal holidays. I also include issuance day fixed effects ($\gamma_{b(s)}$), though these are largely nuisance parameters. SSNs are randomly assigned so there is no obvious reason to expect level differences in injury outcomes between issuance day groups, but I include them to follow standard practice in the SNAP cycle literature (e.g., Carr and Packham, 2019; Bond et al., 2022). Standard errors are clustered at the issuance day $\times$ occupation level, so the number of clusters in the pooled sample is $G = 129 \cdot 9 = 1{,}161$.

Regressions are either unweighted or weighted by injury risk at the occupation-level.[13] My preferred specification uses Poisson quasi-maximum likelihood estimation (PQMLE) to measure changes in injury rates and a linear probability model (LPM) to look at probability changes to simplify interpretation. PQMLE coefficients are interpretable as log changes in injury rates, which map closely to percentage changes for $\beta < 0.2$. The dependent variable in this case is the count of injuries for cell s, d, o on date t, and the

[12] $D \in$ {female 18-24, male 18-24, female 25-34, male 25-34, female 35-44, male 35-44, female 45-54, male 45-54}.

[13] The weight assigned to occupation o is the median annual injury rate across years, where the rate is calculated as the number of observed claims (n $= 251,472$) by workers in occupation o in year y divided by estimated Oregon employment in oy. Median values range from 0.06 (miscellaneous personal care and services) to 8.47 (forestry and conservation) per 100 workers, with a mean value of 0.84.

denominator, cell employment in year $y(t)$, is included on the right-hand side with its coefficient fixed to 1.[14] The β_ks from the LPM can be interpreted as the percentage point effect on the probability of *any* injuries on a day in bin k relative to a day in the omitted period simply averaged across all s, d, o cells.

With my regression equation in mind, I can now concretely discuss how measurement error in my dependent (injury rate) and independent (days since SNAP) variables will affect my estimates. The largest source of measurement error in my dependent variable arises from errors in annual employment based on ACS data, but unlike measurement error in independent variables, measurement error in an outcome variable is absorbed into the residual and reduces precision rather than generating bias. Critically, the identifying variation in my design comes from days since benefit issuance, which varies at the st level, whereas the employment denominator varies only at the sdo-year level. Error in the annual employment estimate is therefore common across all days in a year and is orthogonal to the within-cell-year timing of benefit receipt that identifies my coefficients. Measurement error in annual employment would bias my estimates only if it were systematically correlated with the SNAP benefit month within a year, which is not the case by construction.

Two features of my data induce measurement error in treatment status, and both operate to make my estimates conservative. I discuss these sources at length and derive the implications for inference in appendix C. First, I do not observe SNAP participation. The analysis sample consists of workers whose observable characteristics place them in the top quartile of predicted eligibility, so by construction it pools participants with non-participants, whose injury risk should not respond to the benefit cycle. Since one's final SSN digit is orthogonal to the determinants of eligibility and participation in SNAP, this pooling will attenuate the bin coefficients by the participation rate of retained workers, which I estimate to be roughly 0.51.[15] Measurement error will not change the shape of the estimated profile, so the estimates should be interpreted as lower bounds on effects for SNAP participants.

Second, Oregon assigns issuance days using the SNAP household head's SSN, so I will only accurately infer issuance days for injured workers who are the heads of their SNAP household. For other citizen workers in the household, their true issuance day is a random draw from the schedule, and is independent from

[14] Since this outcome relies on annual employment estimates at the cell level, the sample underlying this specification is limited to gender-age-SOC cells that are observed in at least 7 (of 14) years of the ACS. I impute annual employment for missing years using a recency-weighted average of observed employment for top-quartile workers in that cell. I then allocate the worker population proportionally across digit groups $s = \{0, \ldots, 9\}$.

[15] This estimate comes from multiplying the mean eligibility probability among retained workers (≈ 0.57) by the average take-up rate estimated by FNA among eligible working households in Oregon (≈ 0.9 on average over my sample period, and above 0.75 in every year).

the worker's own last digit. Misallocation therefore contaminates every bin identically within a date and cancels in the within-benefit month contrasts that identify the bin effects. The degree of attenuation from this source of measurement error is expected to be heterogeneous across demographic groups due to differences in the probability of headship. Tabulations from the SNAP QC data suggest headship rates between 0.62 and 0.94 for workers across gender-age groups, with markedly lower values for men. Part of the cross-demographic variation in estimated magnitudes shown below therefore reflects differential misallocation rather than differential true effects. Combined, the two sources imply that the reported coefficients understate per-participant effects by a factor of roughly 2.1 to 3.2. However, I note that neither source affects the size of conventional hypothesis tests under the null but both reduce power in proportion to the attenuation factor. This means that my confidence intervals will be wide relative to studies with administrative participation data.

4 Main results

Figure 3 plots the coefficients from estimating equation 1 on the sample of injuries by workers in the top probability quartile of SNAP-eligibility ($\hat{p} \geq 0.47$). I find no effect of the SNAP cycle on the pooled set of workers when regressions are unweighted, but distinct evidence of a late-in-cycle effect when occupations are weighted by injury risk. In the weighted specification, injury rates are lowest in the first six days after benefit issuance and are 5.7 percent above baseline in both of the last two bins of the benefit month ($p = 0.026$ for days 21-25 and $p = 0.028$ for days 26-30).

Figure 4 shows effect heterogeneity across demographic groups. I focus on the entire end-of-benefit month effect relative to the first 15 days to reduce model dimensionality, which also expands the set of dates with identifying variation.[16] The end-of-month rise in injury risk is driven almost entirely by young men (18-24) and, to a lesser extent, older women (45-54). When occupations are weighted by injury risk, I find that the probability of injury is 0.04 percentage points higher in the second half of the benefit month for young men ($p < 0.001$), and PQML estimates imply a 28.5 percent increase in their injury rates ($p < 0.001$). The unweighted estimates are smaller, an increase of about 22.4 percent in injury rates, but remain clearly distinguishable from zero. Effects for older women are smaller and less precise. With occupation weights, their probability of injury rises by 0.05 percentage points at the end of the SNAP month and their injury

[16] Bond et al. (2022) use the same explanatory variable to look at secondary and post-secondary educational outcomes to address data limitations in their setting.

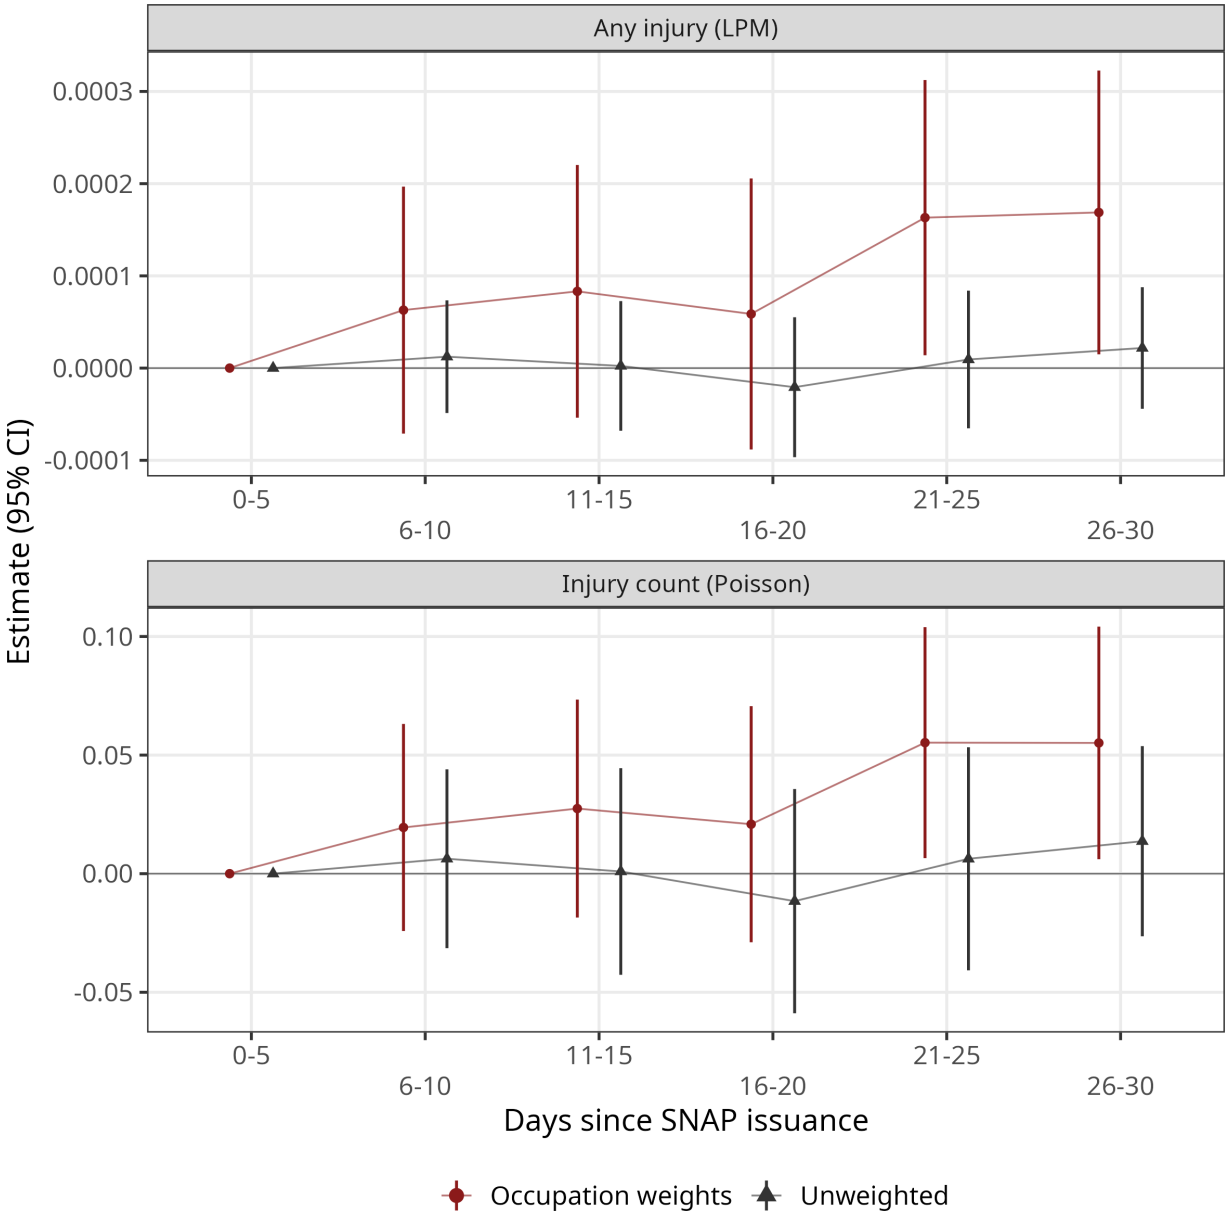


Figure 3: SNAP cycle effects on injury outcomes: Workers with $\hat{p} \geq 0.47$

Notes: Coefficients are obtained from estimating equation 1 on the pooled sample of workers in the top probability quartile of SNAP eligibility (see section 3.3 for a description of how eligibility probability was determined). Estimates in the top panel were obtained from a linear probability model (LPM), while those in the bottom panel were obtained with PQML estimation. Red circles weight occupations by median annual injury rate to upweight riskier occupations; black triangles are unweighted. Coefficients are relative to days 0-5. All regressions include issuance day, year, occupation×month, day of month, day of week, and holiday fixed effects. Standard errors are clustered at the issuance day × occupation level, where $G = 1{,}161$.

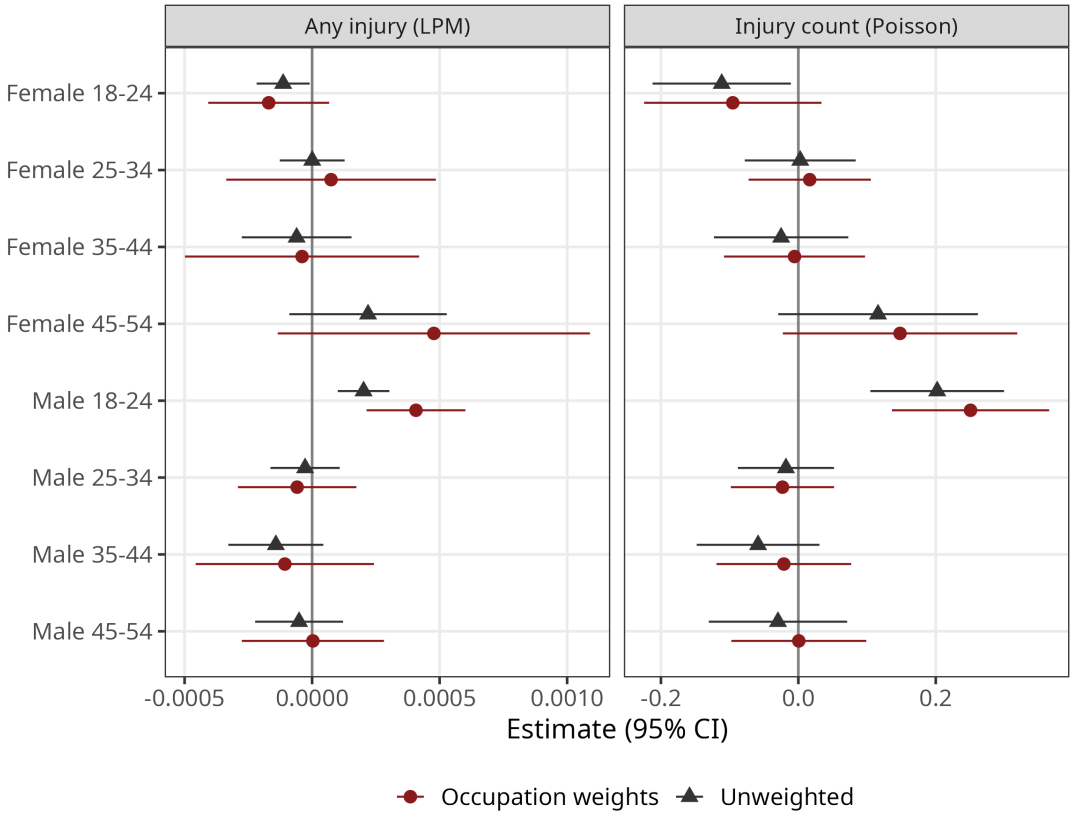


Figure 4: End-of-month effects (Days 16-30): Heterogeneity by age and gender

Notes: Coefficients are obtained from estimating equation 1 on separate samples of workers in the top quartile of probable SNAP-eligibility, $\hat{p} \geq 0.47$ (see section 3.3 for a description of how eligibility probability was determined). Plotted is the coefficient on an indicator for days 16–30 of the benefit month, relative to days 1–15; each regression also includes an indicator for the day of receipt. The left panel shows estimates from a linear probability model (LPM), and the right panel shows PQML estimates, which are interpretable as log changes in injury rates. Red circles are occupation-weighted estimates using the median annual injury rate, while black triangles are unweighted estimates. All regressions include issuance day, year, occupation×month, day of month, day of week, and holiday fixed effects. Standard errors are clustered at the issuance day × occupation level.

rates rise by 16.0 percent, though only the latter is significant at the 10 percent level ($p = 0.13$ and $p = 0.09$ respectively), and neither is significant without weights. For both groups, weighting by occupational injury rates increases the magnitude of the estimates while leaving those for other demographic groups essentially unchanged, which suggests that the end-of-month response is concentrated among workers in occupations with high injury rates.

The youngest women in my sample show a unique and unexpected pattern. Without occupation weights, their injury rates are about 11 percent lower in the second half of the benefit month, and the corresponding LPM estimate is also negative and significant at the 5 percent level. Both estimates lose significance when

occupations are weighted by injury risk (-9.1 percent, $p = 0.15$).

4.1 Robustness

Calendar month effects

Workers in my sample may experience income volatility over the calendar month that is unrelated to SNAP as expenses are paid and paychecks are received. Seligman et al. (2014) find that among the low-income population, nutrition-related health outcomes worsen at the end of the calendar month, which suggests that my estimates could be driven in part by calendar month timing in a way that day-of-month FEs are insufficient to capture. To verify that my main results are indeed driven by the SNAP cycle, figure 5 plots the PQMLE bin profile from figure 3 alongside the corresponding injury profile over the calendar month using identically defined bins. Since much prior work suggests that the first day of the month (Evans and Moore, 2012) and first day of the benefit month (Hastings and Washington, 2010; Cotti et al., 2016) may not be representative baselines, I separate the first 6 days into two bins: day 0 and days 1-5 (day 1 and days 2-6) and all coefficients are estimated relative to days 1-5 (2-6). For the calendar month specification, day-of-benefit month FEs are substituted for day-of-month FEs.

I find that injuries move in an opposite pattern over the calendar month, with injuries actually falling below baseline in the second-to-final bin. Unlike the benefit month case, the shape of the calendar month profile is generally unaffected when weighting by occupational risk, which is suggestive evidence that *food* depletion, rather than resource depletion generally, specifically affects injury rates in high-risk occupations. Appendix figure A.4 confirms that this is also true when splitting the sample by demographic groups.

Adding a control group

My preferred specification focuses only on workers with a high likelihood of SNAP-eligibility, primarily because my probability measure is imprecise so there is not a clear "untreated" group. Moreover, unlike in traditional difference-in-differences settings with true temporal variation, the benefit month is meaningless for non-participants, which means I will only be differencing out noise from non-participants and diluting the treatment effect of participants. However, restricting my sample does constrain my ability to control for date-specific shocks, which in turn reduces precision and can potentially generate bias if correlated with my explanatory variables. To address this concern, I specify a 3 percentage point bandwidth around the 75th

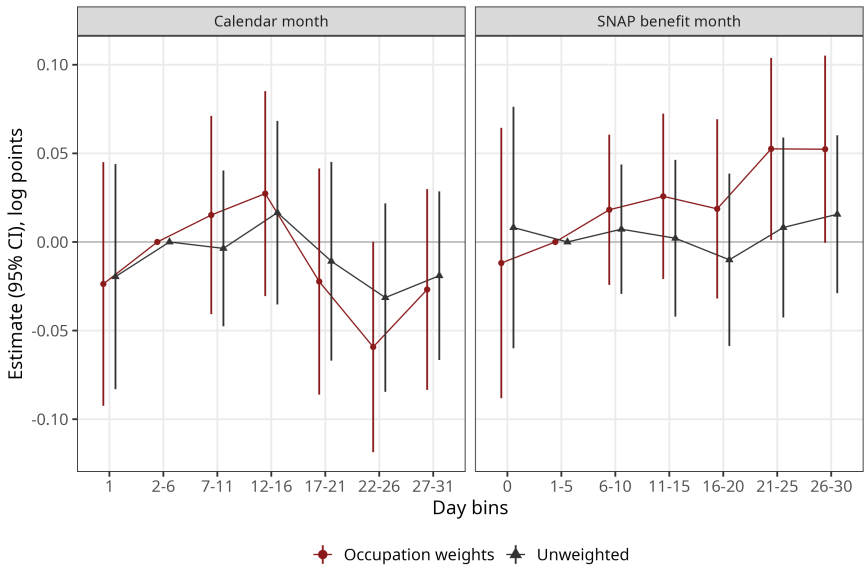


Figure 5: Robustness: Calendar month effects

Notes: All coefficients correspond to PQML estimates and are interpretable as log changes in injury rates. All regressions include issuance day, year, occupation×month, day of week, and holiday fixed effects. In place of day-of-month fixed effects, the calendar month regressions include fixed effects for days since SNAP receipt. Standard errors are clustered at the issuance day × occupation level.

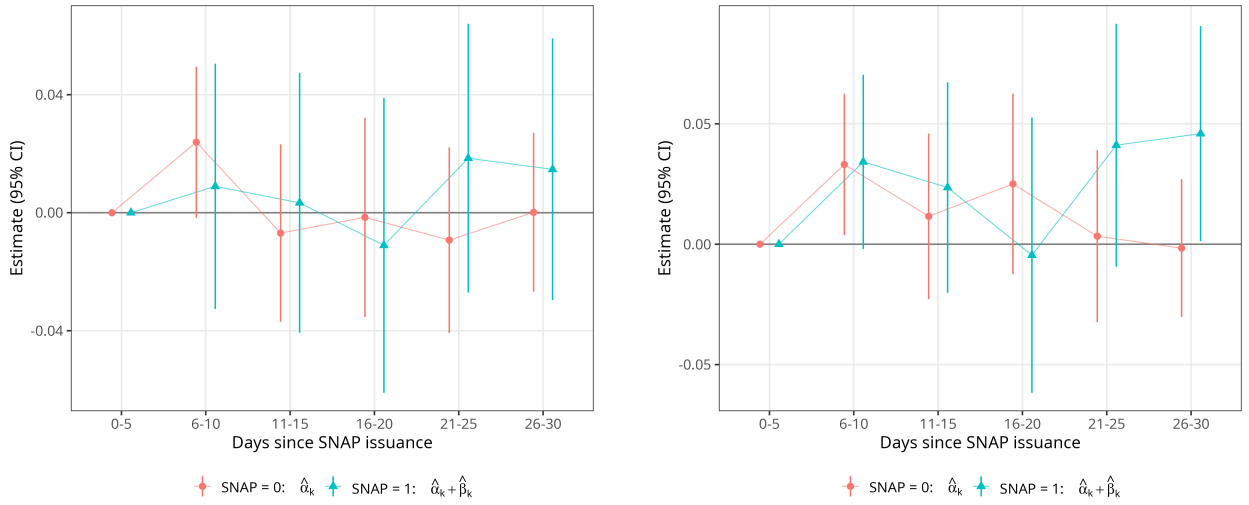
percentile probability value (0.47) and use workers with $\hat{p} < 0.44$ as a pseudo-control group. This restricts my treated sample to workers with $\hat{p} \geq 0.50$ whose observable characteristics predict SNAP-eligibility.

In an effort to minimize confusion from another subscript, I define $SNAP_{d,o} = 1$ for the subset of workers in cell do with $\hat{p} \geq 0.50$ and $SNAP_{d,o} = 0$ for those with $\hat{p} < 0.44$, and estimate the following:[17]

$$Y_{s,d,o,t} = \beta_k f_k(DaysSince_{st}) \times SNAP_{d,o} + \alpha_k f_k(DaysSince_{st}) + \mu SNAP_{d,o} \tag{2}$$
$$+ \gamma_{o,m} + \gamma_{d,o} + \gamma_t + \varepsilon_{s,d,o,t},$$

where $f_k(DaysSince_{st})$ is short-hand notation for the set of bin indicators from equation 1. I keep $\gamma_{o,m}$ but substitute date FEs for the vector of other time controls, and I add a full set of gender-age-occupation interaction terms $\gamma_{d,o}$. I further do not include benefit day fixed effects, as these contributed little value to my main estimates and they would primarily capture noise when estimated over the full population. Finally, I cluster standard errors at the gender-age-occupation level. I chose the 75th percentile of $\hat{p}$ as the threshold to maintain consistency with previous results and because this value is reasonably close to the 0.5 threshold

[17] There are workers on both side of this threshold for nearly all do cells on.



(a) Unweighted

(b) Weighted by occupational injury rates

Figure 6: Robustness: Workers with $\hat{p} < 0.44$ as control group

Notes: Plotted are the level benefit month bin effects from estimating equation 2 for workers in the control group ($\hat{\alpha}_k$) and the combined bin effect for workers in the treated group ($\hat{\alpha}_k + \hat{\beta}_k$), relative to days 0-5. Regressions in panel (a) are unweighted; regressions in panel (b) weight occupations by their median annual injury rate. Coefficients are interpretable as log changes in injury rates, and the vertical axes differ across panels. Regressions include date, occupation×month, and gender×age×occupation fixed effects. Standard errors are clustered at the gender×age×occupation level ($G = 1{,}139$ unweighted and 1,138 weighted).

predicting treatment. However, since I expect there to be contamination on both sides of the threshold, I view these results as secondary to the ones presented in section 3.4.

Figure 6 plots the estimated profiles for both groups, without occupation weights in panel (a) and with them in panel (b). In both panels, injury rates among control workers show no systematic pattern over the benefit month. The control-group estimates for days 21–25 and 26–30 are -0.9 and 0.0 log points unweighted and 0.3 and -0.2 log points weighted. None of the unweighted $\hat{\alpha}_k$ is significant at the 5 percent level (the smallest p-value is 0.07), and only one of the weighted $\hat{\alpha}_k$ (days 6–10) is ($p = 0.03$). The profiles for treated workers mirror my main results. Without weights, injury rates for treated workers rise only slightly at the end of the benefit month, by 1.8 and 1.5 log points in days 21–25 and 26–30 ($p = 0.43$ and $p = 0.52$). With weights, injury rates rise by 4.1 and 4.6 log points in the same bins ($p = 0.11$ and $p = 0.04$), which mirrors the weighted profile in figure 3, and the difference between treated and control workers in days 26–30, $\hat{\beta}_k$, is 4.8 log points ($p = 0.04$). Because the date fixed effects in equation 2 absorb any shock common to all workers on a given day, the absence of a late-month rise among control workers under either weighting scheme suggests that my main results are not driven by unobserved date-level shocks.

5 Mechanisms

In order for an injury to enter the administrative record, the individual must both have been at work and gotten injured, and a worker's labor supply is potentially endogenous to the SNAP benefit cycle. To isolate these channels, I next investigate whether the composition of injury severity changes over the SNAP benefit month. Suppose each worker faces an underlying distribution of injury severity conditional on an event. A reduction in hours worked reduces the probability of *any* injury but all else equal should leave the shape of the distribution unchanged. Likewise, a marginal increase in hours increases exposure to all workplace risk, which should rescale the severity distribution upwards in proportion to its baseline levels. Any change in the observed severity composition of injuries among workers in my sample in response to the SNAP cycle can therefore be attributable to changes in available resources.[18]

For this analysis, I restrict to traumatic injuries which make up about 90 percent of claims. To examine whether resource depletion differentially affects the probability of injuries along the severity distribution, I amend the distribution regression strategy from Foresi and Peracchi (1995) and Chernozhukov et al. (2013) to my setting with count data. The observed severity distribution for traumatic injuries is lumpy and highly-concentrated around the mean, to the extent that half the sample of injuries have index values between 0.48 and 0.57. I define severity thresholds σ_j at local troughs in the severity distribution of observed claims, specifically $\sigma_j = (0.38, 0.47, 0.54, 0.66)$.

I then re-estimate a variant of equation 1 instead summing injuries with severity exceeding σ_j at the s, d, o, t level. The primary explanatory variables are two binary indicators equal to 1 on the day of SNAP receipt (days since SNAP = 0) and on dates more than 20 days since last SNAP receipt (days since SNAP $\geq$ 21). The direction of the day-of-receipt coefficients across the severity distribution has distinct policy relevance; a uniform sign across the severity distribution is indirect evidence of an immediate labor supply response to receipt of non-wage income. Coefficients that differ across thresholds instead indicate a change in the *composition* of injuries, which a pure labor supply response cannot generate. Because the outcomes are nested (every injury above σ_4 is also above σ_1), the difference between adjacent thresholds is informative

[18] One could argue against this logic in an extreme case by noting that the marginal effect of increasing hours worked on injury outcomes is likely non-linear - a fatigued worker working extra overtime may actually face a higher risk of severe injury than a fresh worker who picks up an extra shift at the same resource levels. This is a reasonable objection, but I argue this is unlikely to matter in my setting because many of the workers in my sample are in service industry occupations where shift-work is standard. Shift schedules are often determined by the employer and generally are only flexible on the extensive margin. Moreover, recipient households are regularly re-evaluated to determine eligibility, so marginal workers may have incentive to keep their hours worked within certain bounds in order to meet work requirements (if applicable) and still remain eligible for SNAP.

about where in the distribution the change occurs.

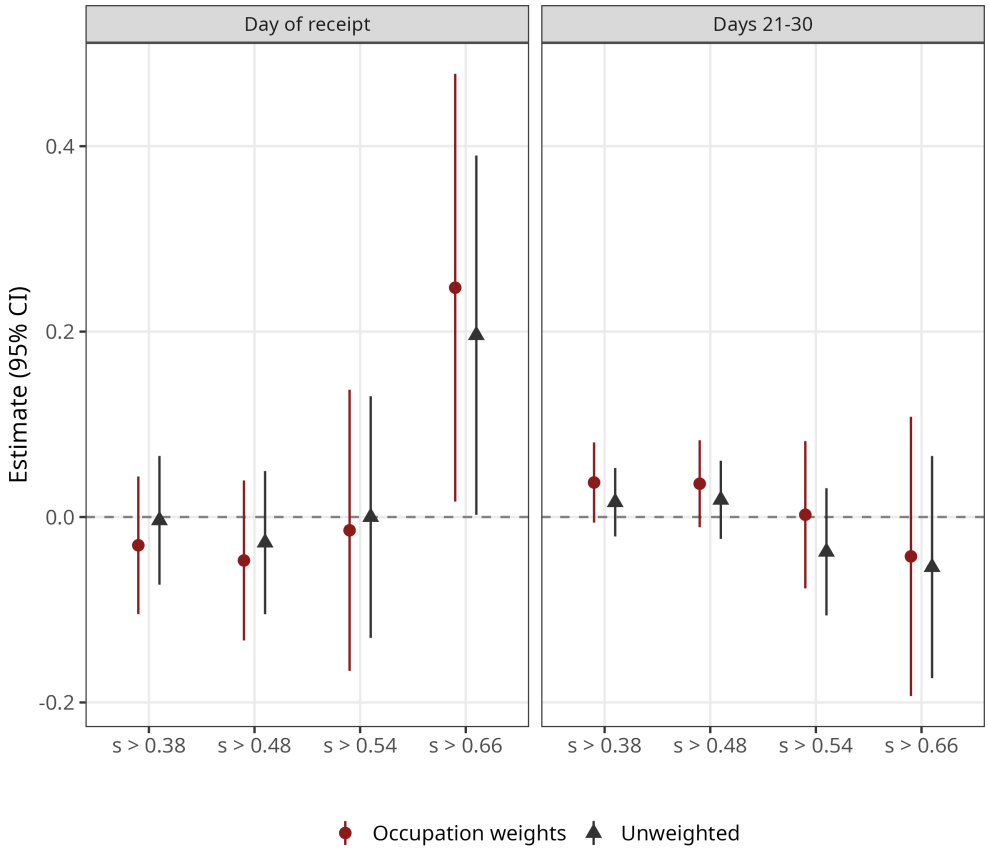


Figure 7: Heterogeneity by severity

Notes: Each pair of points comes from a separate PQML regression of the count of traumatic injuries with severity index above σ_j on indicators for the day of SNAP receipt (left panel) and days 21-30 (right panel), relative to days 1-20, for workers with $\hat{p} \geq 0.47$. Coefficients are interpretable as log changes in injury rates. Red circles weight occupations by median annual injury rate; black triangles are unweighted. All regressions include issuance day, year, occupation×month, day of month, day of week, and holiday fixed effects, with log cell employment as an offset. Standard errors are clustered at the issuance day × occupation level (G ranges from 1,170 to 882, decreasing in σ_j).

Figure 7 plots the day-of-receipt and late-cycle PQML estimates for rate of injury with severity greater than σ_j, with and without occupation weights. Neither set of point estimates is consistent with a proportional shift in risk, though I cannot rule this out conclusively. Late in the benefit month, the weighted rate of injuries above σ_1 and σ_2 rises by about 3.8 percent, while the rate of injuries above σ_3 is unchanged and the rate above σ_4 is, if anything, lower. These late-cycle estimates are imprecise and their confidence intervals include zero, but their pattern places any end-of-month increase among mild to moderate injuries, with index values between 0.38 and 0.54. Over 75 percent of traumatic injuries in this range are sprains or strains, and about a third were caused by overexertion, so part of this effect could reflect changes in work

effort. However, I cannot rule out a proportional change across the distribution as the confidence intervals for the most severe injuries overlap the estimates for mild to moderate injuries.

The day-of-receipt estimates show a different pattern, though I emphasize that this result should be interpreted cautiously as more investigation is required to wholly rule out spuriousness. For the three lower thresholds, the day-of-receipt coefficients are close to zero and slightly negative, so there is no evidence that workers are more exposed to workplace risk on the day benefits arrive. At the top of the distribution, however, the rate of injuries with severity above $\sigma_4 = 0.66$ is about 28.1 percent higher on the day of receipt than over the following 20 days when occupations are weighted, and about 21.7 percent higher when they are not. These injuries are rare (only 9 percent of claims in my sample exceed σ_4), so the estimates are imprecise, but they are several times larger than any late-cycle estimate.

These two results point to different channels. The late-cycle increase among mild to moderate injuries is consistent with a modest change in effort or exposure at the end of the month, when resources are most depleted. The day-of-receipt increase in severe injuries is not. One explanation is that benefits received at midnight are not yet usable for a worker who goes to work before shopping, so the day of receipt is effectively the last day of the previous month's scarcity. If that were the whole story, however, I would also expect severe injuries to be elevated in days 21-30, and the point estimates there are negative. An alternative is that receipt itself changes behavior on that day (e.g., what or when workers eat, distraction, or risk-taking) in ways that raise the risk of serious injury. I cannot distinguish these explanations with the current data, but both imply that the timing of benefit receipt, and not only the depletion of benefits, matters for worker safety.

6 Discussion

This paper asks whether food availability affects the incidence and severity of workplace injuries for low-wage workers. Using Oregon's staggered SNAP issuance schedule and administrative WC claims, I compare injury outcomes for otherwise similar workers assigned to different points in their benefit month. Because I cannot observe participation directly, my estimates are ITT effects for workers in the top quartile of predicted eligibility and likely understate effects for participants.

When predicted-eligible workers are pooled and occupations equally weighted, I find no evidence that injury risk varies over the benefit month. Weighting occupations by their injury risk reveals a late-month

increase of about 5.7 percent, which suggests that the consequences of resource depletion are concentrated in jobs where a lapse is most likely to result in injury. The increase is driven by young men, whose injury rates rise by 28.5 percent in the second half of the benefit month, and to a lesser extent by older women. Unweighted estimates suggest that young women show the opposite pattern, with injury rates that fall over the benefit month, though this decline is not robust to weighting.

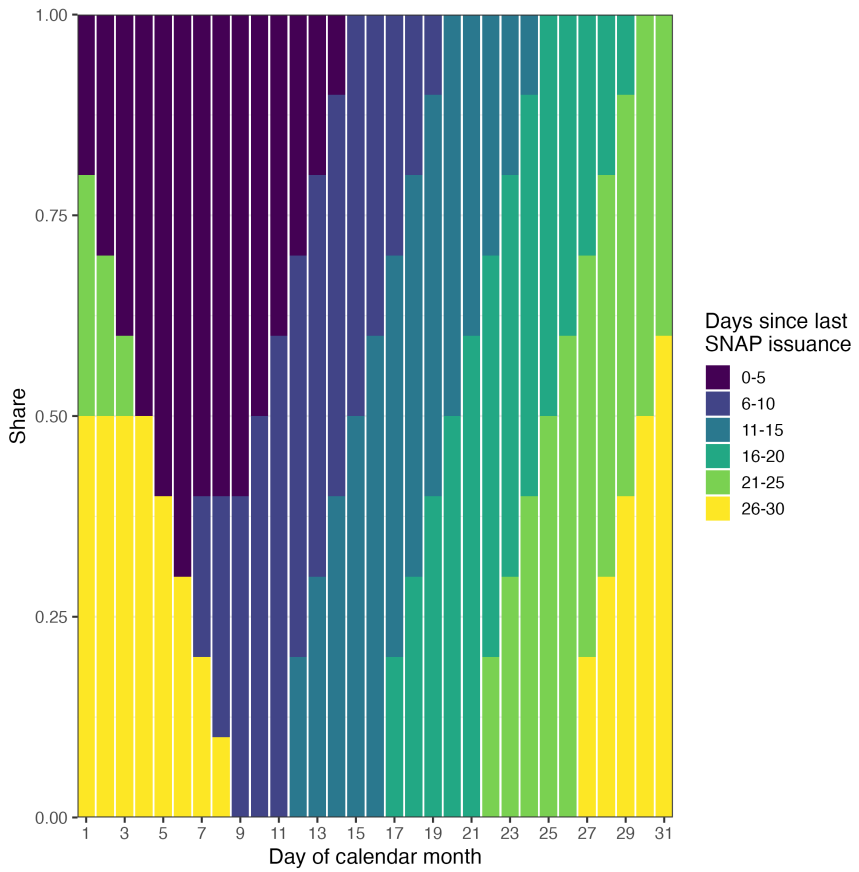
The severity analysis suggests that two channels operate at different points in the benefit month. The point estimates place the end-of-month increase among mild to moderate injuries, mostly sprains and strains from overexertion, which is consistent with workers compensating for depleted resources through additional effort or exposure. The day of receipt tells a different story: the rate of the most severe injuries rises by about 28.1 percent while less severe injuries are unchanged, a change in the composition of injuries that a labor supply response cannot produce. Whether this reflects the last day of the previous month's scarcity or a change in behavior triggered by receipt itself is an important question for future work, since the two imply different policy responses.

I find no consistent injury cycle for workers aged 25 to 44. These workers may be more often in two-adult households or with school-age children, so they may be better able to smooth consumption within the month, though further investigation is needed. The next step to advance this work is to test which policy levers are most effective at mitigating injury risk for hungry workers. The American Recovery and Reinvestment Act was an economic stimulus package enacted in 2009 and in addition to increasing SNAP benefit amounts, it allowed states to more easily waive ABAWD work requirements. Future work will leverage the temporal and geographic variation in exposure to Oregon's ABAWD time limits during and after the Great Recession to link my injury results directly to differential exposure to work requirements across groups. I have also obtained WC claims data from Idaho, a state which changed its SNAP disbursement schedule twice over this period - in 2009, the state transitioned from distributing benefits over the first five days of the month to a single-day disbursement design, and then switched to staggering disbursal over *ten* days in 2016. This offers a complementary empirical design that would allow me to study two distinct shocks to benefit timing on workplace safety outcomes.

The 2025 One Big Beautiful Bill Act (OBBBA) expanded work requirements conditioning access to social safety net programs, including the Supplemental Nutrition Assistance Program (SNAP) and Medicaid. Adults between age 18 and 64 (up from 54) without a documented disability must work at least 80 hours each month in order to remain eligible for SNAP, and failure to meet this requirement for three consecutive

months bars participation for three years. Moreover, the OBBBA revoked exemptions for critical at-risk populations, namely veterans, the unhoused, and individuals between 18 and 24 who have recently aged out of foster care. Recent studies find that work requirements reduce SNAP participation (Cook and East, 2026) and health insurance coverage (Sommers et al., 2019) but do not increase employment, which suggests that financial constraints for part-time, seasonal, and other workers with volatile employment and income will likely be exacerbated.

These findings suggest that the costs of food insecurity reach beyond nutrition into the workplace, and that they fall most heavily on low-wage workers in the most dangerous jobs. This is directly relevant to the OBBBA's expanded work requirements. The workers most likely to lose benefits under an 80-hour monthly threshold are those with volatile hours who remain attached to the labor market. More broadly, if resource depletion raises injury risk, policies that smooth resources within the month, such as more frequent issuance, may improve worker safety at little fiscal cost, and identifying which of these levers is most effective will be the central focus of this work moving forward.



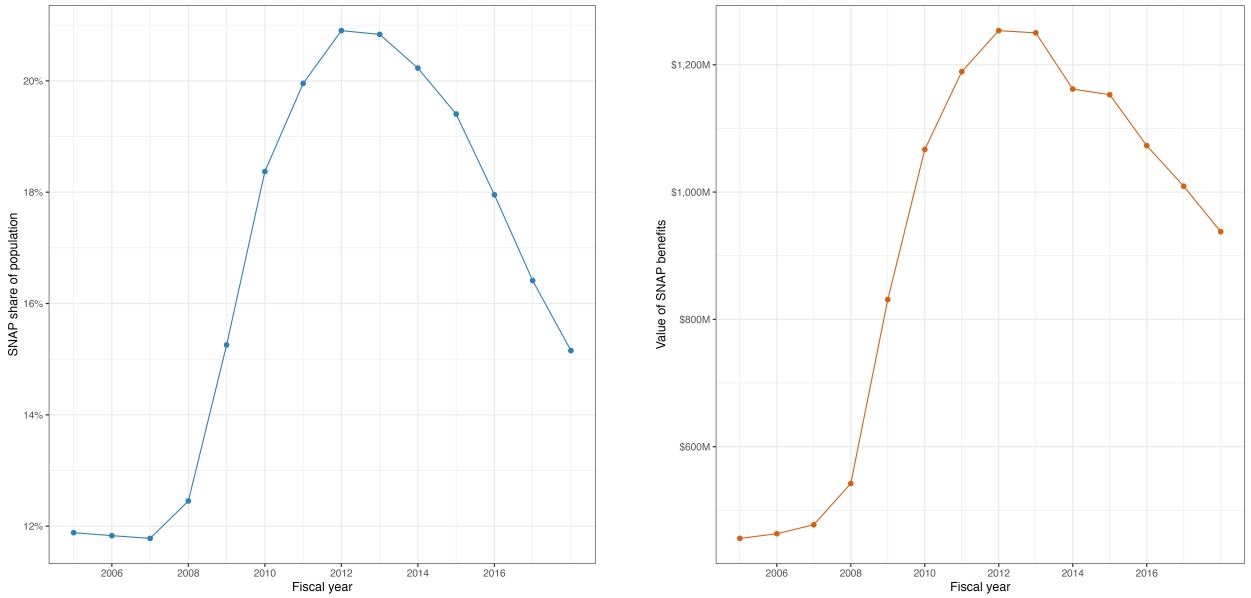
Notes: Oregon staggers distribution of SNAP benefits over the first nine days of the calendar month based on the last digit of the primary account holder's SSN. Last digits that end in 1 to 9 map directly to assigned issuance days, while households whose relevant last digit is 0 also receive benefits on the first of the month.

Figure A.1: Overlap between 31-day calendar month and SNAP benefit month in Oregon

A Additional Tables and Figures

B Constructing the severity index

I measure the severity of each nature $\times$ body part type by how long injuries of that type keep workers away from work, expressed as a rank. Data on injury counts by nature of injury and body part come from the BLS SOII (Table R19), and data on days away from work (DAFW) by nature of injury and by body part injured come from Tables R67 and R68 for 2011 through 2018. The latter tables contain for each injury nature and body part, the number of *national* cases with DAFW across seven duration bins: 1, 2, 3–5, 6–10, 11–20, 21–30, and 31 or more days. A severity index based on the expected DAFW for the injury nature or body part alone would be incomplete; finger injuries generally result in fewer days off because fingers



(a) SNAP participation rate

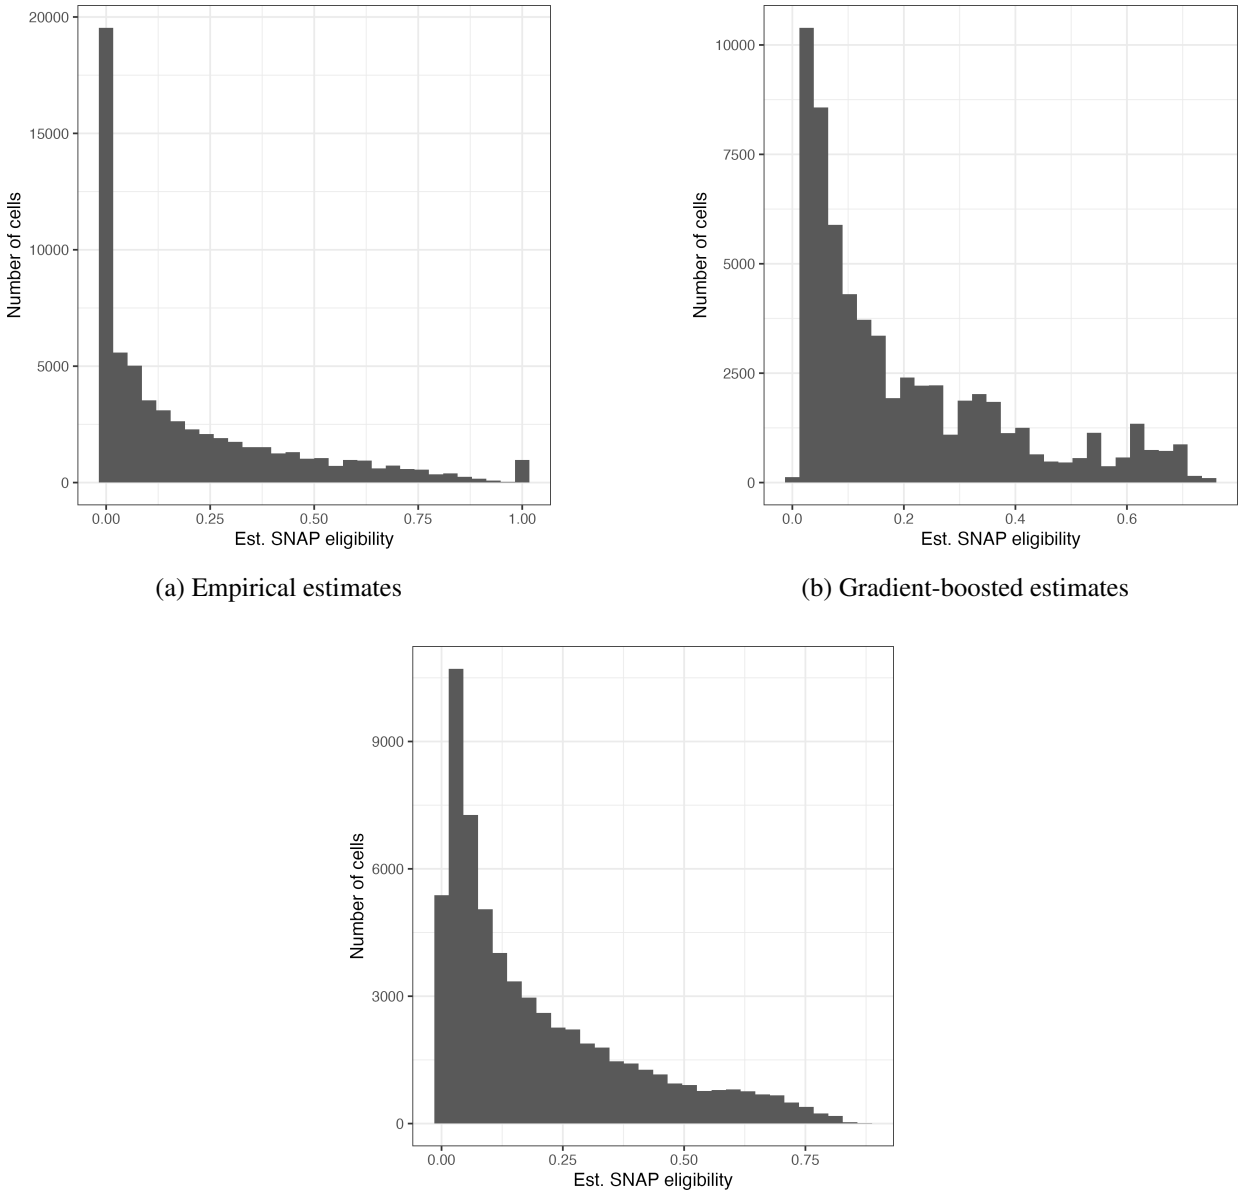
(b) Total value of SNAP benefits

Notes: Data on SNAP participation and benefit amounts come from FNA's SNAP Monthly State Participation and Benefits Summary, while overall population estimates come from FRED.

Figure A.2: SNAP in Oregon: 2005–2018

mostly get cuts and bruises, and fractures alone predict more days off because this category includes major bone fractures. Since a finger fracture is more severe than a finger cut but less severe than a leg fracture, the ideal severity index should account for this nuance. The complication is that BLS does not publish DAFW statistics by nature and body part jointly. I instead estimate the joint distribution of nature $\times$ body part across DAFW bins using an iterative proportional fitting (IPF) procedure (Deming and Stephan, 1940), which uses maximum-likelihood estimation to iteratively generate weights for each two-way combination in order to replicate the observed two-way counts. In practice, this procedure requires the presently untestable assumption that moving an injury from one body part to another changes the odds of a longer work absence by the same factor for every nature. This assumption is more plausible for some natures (e.g., cuts and bruises) than for others; the difference in severity between an amputated finger and an amputated leg is arguably larger than the difference between a bruised finger and bruised leg. However, given that I currently use this index to categorize injuries rather than as an outcome variable, I view this as a second-order concern.

With the fitted distribution for each nature $\times$ body-part combination across duration bins, each type is then ranked relative to the national distribution for all injuries. Let w_{jb} be the share of type j's cases in bin b, and F_b the cumulative share of all national DAFW cases in bins 1 through b. Each bin is assigned



(a) Empirical estimates

(b) Gradient-boosted estimates

(c) Blended estimates

Notes: Plots show the distribution of predicted probability of SNAP eligibility across observed ACS cells, where cells are unique combinations of gender, age bin, SOC group, and real weekly wage quintile. SNAP eligibility is defined at the worker-level and is equal to 1 if the worker's household reported SNAP participation *or* if their household income was at or below the BBCE eligibility threshold of 185 percent of the FPL. The empirical estimates are simple weighted shares of eligibility across all ACS respondents in that cell. The GB estimates are obtained from a gradient-boosting machine learning algorithm based on cell characteristics, and the blended estimates are weighted combinations of the empirical and GB predictions, where the relative weight assigned to the empirical estimates increases with size of the ACS sample.

Figure A.3: Distribution of estimated $P(\text{SNAP eligibility})$ ($\hat{p}$) across ACS cells

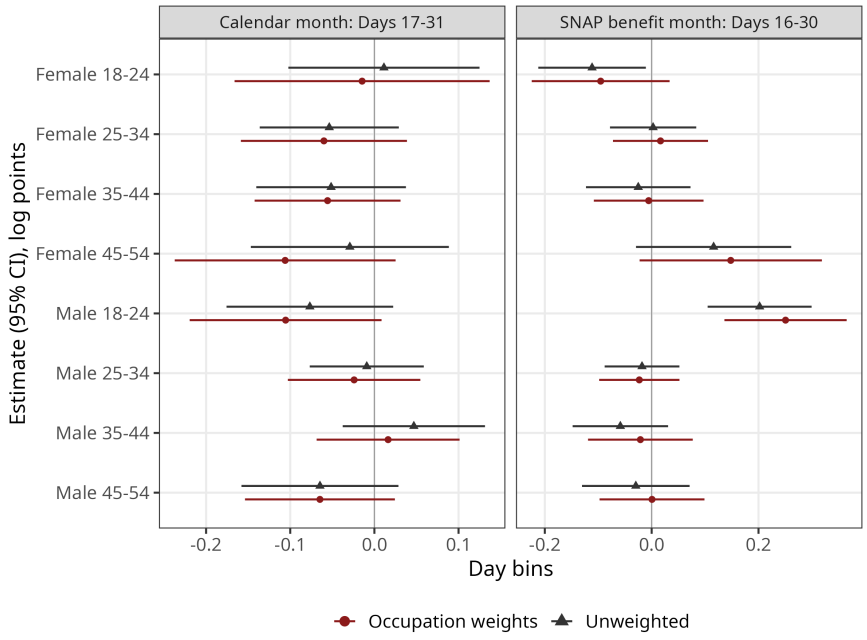


Figure A.4: Robustness: Heterogeneity over calendar month and SNAP benefit month

Note: All coefficients correspond to PQML estimates and are interpretable as log changes in injury rates. All regressions include issuance day, year, occupation×month, day of week, and holiday fixed effects. In place of day-of-month fixed effects, the calendar month regressions include fixed effects for days since SNAP receipt. Standard errors are clustered at the issuance day × occupation level.

	Women			Men		
	N	% of Gender	% of Gender & Age	N	% of Gender	% of Gender & Age
Total workers	881,390	1.00		667,739	1.00	
Head of HH	802,847	0.91		452,499	0.68	
Under 25	153,255	0.17	0.85	69,514	0.10	0.63
25-34	266,832	0.30	0.94	129,335	0.19	0.64
35-44	183,990	0.21	0.92	93,908	0.14	0.62
45-54	123,463	0.14	0.92	88,442	0.13	0.77
55-64	63,481	0.07	0.90	62,923	0.09	0.81
65+	11,826	0.01	0.89	8,377	0.01	0.88
Not head of HH	78,543	0.09		215,240	0.32	
Under 25	26,501	0.03	0.15	41,506	0.06	0.37
25-34	16,591	0.02	0.06	72,729	0.11	0.36
35-44	16,462	0.02	0.08	58,763	0.09	0.38
45-54	10,560	0.01	0.08	26,446	0.04	0.23
55-64	6,988	0.01	0.10	14,665	0.02	0.19
65+	1,441	0.00	0.11	1,131	0.00	0.12

Table A.1: Tabulation of workers in SNAP participating households, Oregon 2005-2018

Note: This table shows total weighted counts of male and female workers participating in SNAP by age and household-head status.

Source: Author's own analysis of SNAP Quality Control data from FNA for fiscal years 2005-2018.

its mid-rank percentile in the national distribution, $\bar{p}_b = (F_{b-1} + F_b)/2$, and define type j's severity as the average percentile of its cases:

$$S_j = \sum_b w_{jb} \bar{p}_b.$$

The bin scores $\bar{p}_b$ are *ridits*, which were named to connote similarity with *probits* and *logits*, and this concept was first introduced by Bross (1958) to construct ordinal ranks for variables "which can neither be optimally described by a nominal classification nor can be subjected to refined measurement systems" (Kantor et al., 1968). S_j is the mean ridit of type j against the national reference distribution. It therefore has a direct probability interpretation. This construction has four advantages over simpler summaries. First, it uses only the ordering of the bins, so it needs no assumed value for the open-ended 31-or-more-days bin. That bin holds 29 percent of national cases, and any mean-days measure would depend on how it was top-coded. Second, it uses each type's full duration distribution rather than its median. Medians cannot separate the most serious injuries, because fractures, meniscus tears, and herniated discs all have medians inside the

Restriction	Dropped	Remaining
All accepted Oregon WC claims, 2005–2018		251,472
Claim restrictions		
Oregon residents aged 18–54	65,173	186,310
Occupation reported and matched to a SOC group	9,416	176,894
Non-zero weekly wage (after imputation)	1,157	175,737
Injury reported within 2 weeks	20,949	154,788
Empirical restrictions		
ACS cell match w/ ≥ 100 weighted workers	4,679	150,109
ACS cell observed in ≥ 7 years	21,399	128,710
Likely SNAP-eligible: $\hat{p} \geq \hat{p}_{75} = 0.47$	92,278	36,432
Analysis sample		**36,432**

Table A.2: Restrictions applied to arrive at analysis sample

top bin. Third, the score is linear in the distribution, so the score of a mix of injury types equals the case-weighted average of their scores. Aggregation to broader categories and claim-weighted summaries are therefore internally consistent. Fourth, natures, body parts, and nature $\times$ body-part combinations are all scored against the same reference distribution, so their scores can be compared directly.

C Sizing attenuation bias

My estimates of the SNAP benefit-cycle effect on occupational injuries are subject to two primary sources of measurement error due to the nature of my data. First, I do not observe SNAP participation. The estimation sample includes all workers whose observable characteristics place them above a threshold of predicted SNAP eligibility, so it necessarily pools participants with non-participants, whose injury risk should not respond to the SNAP benefit cycle. I will refer to this as *dilution*. Even in the case of zero dilution, because households in Oregon are assigned issuance days based on the last digit of the primary applicant's SSN, a second source of measurement error stems from not observing the household in my claims data. For each claim, I assign the issuance day implied by the worker's own SSN, which is correct when the worker is the primary applicant and is an independent draw from the issuance-day distribution when the worker is not. I will refer to this as *misallocation*.

TABLE A.3—SNAP CYCLE EFFECTS ON INJURY OUTCOMES: WORKERS WITH $\hat{p} \geq 0.47$

	Any injury (pp) (1)	Injury rate (per 100,000) (2)	Injury count (3)
	OLS		Poisson
Days 6–10	0.0012	0.0155	0.0063
	(0.0031)	(0.0432)	(0.0192)
Days 11–15	0.0002	0.0046	0.0009
	(0.0036)	(0.0496)	(0.0222)
Days 16–20	−0.0021	−0.0227	−0.0116
	(0.0039)	(0.0525)	(0.0241)
Days 21–25	0.0009	0.0140	0.0063
	(0.0038)	(0.0527)	(0.0240)
Days 26–30	0.0022	0.0301	0.0137
	(0.0034)	(0.0455)	(0.0204)
Fixed effects			
Issuance day	Yes	Yes	Yes
Year	Yes	Yes	Yes
Occupation × Month	Yes	Yes	Yes
Day of month	Yes	Yes	Yes
Day of week	Yes	Yes	Yes
Holiday	Yes	Yes	Yes
Gender × age indicators	Yes	Yes	Yes
Observations	22,241,550	22,241,550	19,341,030
Injuries	36,432	36,432	36,432
Clusters	1,161	1,161	1,161
R^2	0.005	0.002	
Pseudo R^2			0.146

Notes: Coefficients are obtained from estimating equation 1 on the pooled sample of workers with $\hat{p} \geq 0.47$, relative to days 0–5 of the benefit month. Column (1) is an unweighted linear probability model for any injury (in percentage points); column (2) is employment-weighted least squares on the injury rate (per 100,000 worker-days); column (3) is PQML on injury counts with log cell employment as an offset, interpretable as log changes in injury rates. All regressions include issuance day, year, occupation × month, day of month, day of week, and holiday fixed effects. Standard errors clustered by issuance day × occupation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE A.4—HETEROGENEITY BY AGE AND GENDER: WORKERS WITH $\hat{p} \geq 0.47$ (LPM)

	Any injury (pp)							
	18–24		25–34		35–44		45–54	
	Female (1)	Male (2)	Female (3)	Male (4)	Female (5)	Male (6)	Female (7)	Male (8)
Day of receipt	−0.0100	0.0225	−0.0720*	−0.0164	−0.0020	0.0049	−0.0168	0.0047
	(0.0242)	(0.0204)	(0.0422)	(0.0248)	(0.0416)	(0.0336)	(0.0527)	(0.0318)
Days 16–30	−0.0170	0.0407***	0.0074	−0.0059	−0.0039	−0.0107	0.0478	0.0003
	(0.0121)	(0.0099)	(0.0209)	(0.0118)	(0.0233)	(0.0178)	(0.0311)	(0.0142)
Fixed effects								
Issuance day	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation × Month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,374,580	3,885,880	3,527,970	3,783,620	2,147,460	2,454,240	1,227,120	1,840,680
Injuries	3,466	3,845	5,659	7,629	4,799	5,470	2,260	3,304
Clusters	594	684	621	666	378	432	216	324
R^2	0.004	0.004	0.019	0.016	0.013	0.010	0.009	0.009

Notes: Coefficients are obtained from estimating equation 1 separately for each gender × age group, with indicators for the day of SNAP receipt and days 16–30 relative to days 1–15 of the benefit month, and are linear probability model estimates for any injury, in percentage points. Samples include workers with $\hat{p} \geq 0.47$, the top quartile of the worker-weighted population. All regressions include issuance day, year, occupation × month, day of month, day of week, and holiday fixed effects. Standard errors clustered by issuance day × occupation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

To fix ideas, abstract from covariates, weights, and the full set of fixed effects in the main specification and consider the simplified linear model for a particular date t:

$$y_c \;=\; \gamma_{t(c)} \;+\; \sum_{k=1}^{K} \beta_k \, \mathbf{1}\{\kappa_c = k\} \;+\; \varepsilon_c \tag{3}$$

where $c \equiv sdo$ indexes the terminal SSN digit-gender-age-occupation cell, $\gamma_{t(c)}$ is a date effect standing in for the fixed effects in the main specification, and $\kappa_c \in \{0, 1, \ldots, K\}$ is the days-since-issuance bin, with $\kappa = 0$ the omitted reference bin ($\beta_0 = 0$).

Throughout, β_k denotes the causal effect of being k bins from issuance *for a SNAP participant whose issuance day is correctly measured*. All derivations are statements about conditional means and therefore carry over unchanged to the weighted least squares estimates reported in the main text.

Two facts about Oregon's assignment mechanism do most of the work below. SSN last digits are as good as randomly assigned and are independent across individuals, including members of the same household. This assignment mechanism motivates: (1) true issuance days are independent of worker characteristics and of the calendar date of injury, and (2) when a worker is not the primary applicant, the worker's own

TABLE A.5—HETEROGENEITY BY AGE AND GENDER: WORKERS WITH $\hat{p} \geq 0.47$ (PQML)

	Injury count							
	18–24		25–34		35–44		45–54	
	Female (1)	Male (2)	Female (3)	Male (4)	Female (5)	Male (6)	Female (7)	Male (8)
Day of receipt	−0.0628	0.1379	−0.1710*	−0.0602	0.0001	0.0115	−0.0715	0.0244
	(0.1287)	(0.1325)	(0.0948)	(0.0764)	(0.0943)	(0.0859)	(0.1897)	(0.1107)
Days 16–30	−0.0954	0.2506***	0.0166	−0.0231	−0.0056	−0.0211	0.1480*	0.0007
	(0.0658)	(0.0584)	(0.0454)	(0.0383)	(0.0524)	(0.0500)	(0.0871)	(0.0501)
Fixed effects								
Issuance day	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation × Month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,002,360	2,127,430	2,131,350	2,685,800	1,627,200	1,899,790	872,990	1,257,970
Injuries	3,466	3,845	5,659	7,629	4,799	5,470	2,260	3,304
Clusters	594	684	621	666	378	432	216	324
Pseudo R^2	0.076	0.093	0.165	0.147	0.138	0.120	0.112	0.120

Notes: Coefficients are obtained from estimating equation 1 separately for each gender × age group, with indicators for the day of SNAP receipt and days 16–30 relative to days 1–15 of the benefit month, and are PQML estimates on injury counts with log cell employment as an offset, interpretable as log changes in injury rates. Samples include workers with $\hat{p} \geq 0.47$, the top quartile of the worker-weighted population. All regressions include issuance day, year, occupation × month, day of month, day of week, and holiday fixed effects. Standard errors clustered by issuance day × occupation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE A.6—HETEROGENEITY BY SEVERITY

	Injury count							
	Unweighted				Injury-rate weighted			
	$> \sigma_1$ (1)	$> \sigma_2$ (2)	$> \sigma_3$ (3)	$> \sigma_4$ (4)	$> \sigma_1$ (5)	$> \sigma_2$ (6)	$> \sigma_3$ (7)	$> \sigma_4$ (8)
Day of receipt	−0.0036	−0.0276	0.0000	0.1961**	−0.0305	−0.0469	−0.0143	0.2474**
	(0.0354)	(0.0394)	(0.0665)	(0.0989)	(0.0379)	(0.0440)	(0.0774)	(0.1177)
Days 21–30	0.0159	0.0185	−0.0375	−0.0540	0.0372*	0.0359	0.0025	−0.0425
	(0.0188)	(0.0215)	(0.0350)	(0.0612)	(0.0221)	(0.0239)	(0.0406)	(0.0769)
Fixed effects								
Issuance day	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation × Month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Holiday	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	19,664,810	19,158,060	16,142,850	12,429,400	19,664,810	19,158,060	16,142,850	12,429,400
Clusters	1,170	1,161	999	882	1,170	1,161	999	882
Pseudo R^2	0.145	0.141	0.101	0.074	0.140	0.138	0.097	0.073

Notes: Each column is a separate PQML regression of the count of traumatic injuries with severity index above σ_j on indicators for the day of SNAP receipt and days 21–30, relative to days 1–20 of the benefit month, for workers with $\hat{p} \geq 0.47$, with log cell employment as an offset. Coefficients are interpretable as log changes in injury rates. Thresholds are $\sigma_j = (0.38, 0.47, 0.54, 0.66)$. Regressions also include gender × age indicators. All regressions include issuance day, year, occupation × month, day of month, day of week, and holiday fixed effects. Standard errors clustered by issuance day × occupation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	Days away from work (fitted share)							
	1	2	3–5	6–10	11–20	21–30	31+	Index
Sprains × knee	.084	.076	.138	.117	.131	.080	.373	0.577
Sprains, all parts	.106	.098	.179	.128	.116	.066	.306	0.525
Sprains × finger	.168	.130	.217	.141	.123	.056	.165	0.431
Fractures × leg (non-knee)	.043	.037	.077	.082	.094	.081	.585	0.689
Fractures, all parts	.054	.048	.095	.091	.110	.089	.513	0.655
Fractures × finger	.077	.063	.124	.113	.128	.091	.404	0.598
All injuries (raw)	.140	.109	.174	.118	.110	.064	.285	0.500
All injuries (cum.)	.140	.249	.423	.541	.651	.715	1	–
$\bar{p}_b$	.07	.195	.336	.482	.596	.683	.857	=

Table B.1: Subset of fitted DAFW distribution, pooled 2011-2018

Notes: Distributions are estimated using an iterative proportional fitting (IPF) algorithm to generate observed counts of nature × body part, nature × DAFW, and body part × DAFW from BLS Tables R19, R67, and R68 for 2011 through 2018. The severity index for nature j and part p is calculated as $Index_{jp} = \sum_b w_{jpb}\bar{p}_b$, where w_{jpb} is the fitted share of jp type injuries in DAFW bin b and $\bar{p}_b$ is the midpoint of the *cumulative* national share. It therefore has a direct probability interpretation. $S_j = P(X_j > R) + \frac{1}{2}P(X_j = R)$, where X_j is the duration of a random injury of type j and R that of a random national DAFW injury: the probability that an injury of type j keeps a worker away longer than a typical lost-time injury, with ties split evenly. A type whose durations match the nation's scores exactly 0.5. Scores above 0.5 mean durations longer than the national norm.

	Nature only	Eye	Finger	Hand (non-finger)	Wrist	Shoulder	Lumbar back	Knee	Ankle
Fractures	0.655	(0.371)	0.598	0.631	0.673	0.750	0.669	0.725	0.661
Sprains, strains	0.525	(0.262)	0.431	0.461	0.506	0.609	0.503	0.577	0.489
Soreness, pain	0.497	0.242	0.410	0.442	0.488	0.597	0.487	0.563	0.471
Bruises	0.413	0.223	0.353	0.379	0.416	0.513	0.417	0.483	0.401
Cuts, lacerations	0.373	0.218	0.360	0.388	0.427	(0.529)	(0.429)	0.498	0.412
Part only		0.258	0.428	0.431	0.562	0.611	0.501	0.571	0.516

Table B.2: Combined nature × part severity indices

Notes: This table shows estimated severity indices for a subset of nature × body part combinations based on pooled national data from 2011 to 2018 (BLS, 2018a,b,c). Values shown in parentheses are based on fewer than 1,000 injuries (e.g., 340 cases of sprain × eye). The fact that body parts scale each nature index by the same factor is an algorithmic assumption and should not be taken as evidence that this is in fact true.

SSN digit, and thus the implied issuance day, is an independent draw from the schedule, unrelated to the household's true day.

Dilution from unobserved SNAP participation

Bias

Let $S_c \in \{0, 1\}$ indicate SNAP participation for cell c, which is not observed. Since only participants are exposed to the benefit cycle, the outcome equation is then

$$y_c \;=\; \boldsymbol{\gamma}_t \;+\; S_c \sum_k \beta_k \, \mathbf{1}\{\kappa_c = k\} \;+\; \varepsilon_c. \tag{4}$$

For non-participants, the assigned bin is a placebo, as it is constructed from the worker's own SSN digit, but no benefits arrive on any day so it carries no effect.

Define $\pi \equiv \Pr(S_c = 1 \mid c \in \text{sample})$ as the participation share among cells in the analysis sample.

Assumption A1 (Participation is orthogonal to the assigned bin)**.** $\Pr(S_c = 1 \mid \kappa_c = k,\, t) = \pi$ *for all* k *and* t*.*

Assumption A1 is substantiated by random last digit assignment. SNAP participation is determined by income, household structure, and take-up, none of which depend on the last digit of an SSN. Though the sample was selected based on *predicted eligibility*, this does not threaten assumption A1 as the variables used to estimate eligibility are independent of the assigned bin. Rather, it only determines the level of π.

Taking expectations of (4) conditional on the assigned bin and date, and applying Assumption A1,

$$\begin{aligned} E[y_c \mid \kappa_c = k,\, t] &= \boldsymbol{\gamma}_t + \Pr(S_c = 1 \mid \kappa_c = k,\, t)\, \beta_k \\ &= \boldsymbol{\gamma}_t + \pi\, \beta_k. \end{aligned} \tag{5}$$

A saturated regression of y on bin indicators and date effects recovers differences in these conditional means, so the probability limit of the estimated coefficient on bin k is

$$\hat{\beta}_k \;\xrightarrow{p}\; \big(\boldsymbol{\gamma}_t + \pi\beta_k\big) - \big(\boldsymbol{\gamma}_t + \pi\beta_0\big) \;=\; \pi\, \beta_k. \tag{6}$$

Dilution therefore attenuates every bin coefficient by the same proportional factor π. Because the factor is

common across bins, the *shape* of the estimated profile is unbiased; only its scale is compressed. Equation (6) is the standard intent-to-treat result - my estimates are reduced-form effects of *predicted-eligible status* on injury outcomes, and $\beta_k = \hat{\beta}_k/\pi$ recovers the effect for participants.

Implied magnitudes and implications for inference

The participation share decomposes as $\pi = \tau\bar{p}$, where $\bar{p}$ is the mean probability of SNAP eligibility among retained cells and τ is take-up conditional on eligibility. At the 75th-percentile threshold, each worker in cell c satisfies $p_i \geq 0.47$, so $\bar{p}$ is bounded-below at 0.47 and the average across retained cells in my sample is roughly 0.57. For take-up, FNA estimates place Oregon's SNAP participation rate among the eligible working-poor population over the study period near 90 percent, so $\tau \approx 0.9$. These values imply that $\pi \approx 0.51$, so the true effect for participants in bin k is expected to be $1/0.51 = 1.96$ times the estimated coefficient.

Dilution leaves the variance of the estimator essentially unchanged while shrinking its expectation. To see this, decompose the conditional variance of the outcome under (4):

$$\mathrm{Var}(y_c \mid \kappa_c = k,\, t) \;=\; \pi(1-\pi)\,\beta_k^2 \;+\; \sigma_\varepsilon^2, \tag{7}$$

where the first term reflects heterogeneity between participants and non-participants within a bin. In the binary outcome case, β_k is on the order of 10^{-4} while σ_ε^2 is on the order of the outcome variance itself, so $\pi(1-\pi)\beta_k^2$ is smaller than σ_ε^2 by several orders of magnitude and the residual variance is unaffected to first order. The regressors are similarly unchanged, because assigned bins are constructed identically for participants and non-participants. It follows that

$$\mathrm{se}\left(\hat{\beta}_k\right) \;\approx\; \text{unchanged}, \qquad \frac{\hat{\beta}_k}{\mathrm{se}(\hat{\beta}_k)} \;\approx\; \pi \times \frac{\beta_k}{\mathrm{se}(\hat{\beta}_k)}. \tag{8}$$

Under the null $\beta_k = 0$, the diluted effect $\pi\beta_k$ is also zero, so conventional tests will be correctly sized, but the t-statistic against any fixed alternative is deflated by π. Thus the minimum detectable effect for *true participants* is higher by a factor of $1/\pi$. This in part explains the width of my confidence intervals relative to studies with administrative participation data. Finally, note that rescaling the point estimate by $1/\pi$ rescales its confidence interval by the same factor, so while the correction changes the economic magnitude implied

by the point estimate, it does not artificially manufacture statistical precision.

Misallocation of issuance days

Bias

Consider now the subpopulation of participants, and suppose worker i in cell c truly receives benefits on day b_i^*, which is unobserved and dependent on the last digit of their SNAP household head's SSN. Let $m_c \equiv \Pr(\text{worker } i \text{ is not the SNAP household head} \mid c(i) = c)$. By independence of SSN digits across household members,

$$b_i^* = \begin{cases} b_i & \text{with probability } 1 - m_c, \\ \tilde{b} \sim q(\cdot), & \text{with probability } m_c, \end{cases} \tag{9}$$

where b_i is the issuance day implied by i's terminal SSN digit, and $q(\cdot)$ is the population issuance-day distribution. Under Oregon's schedule, $q_1 = 0.2$ for the first of the month, which pools SSN digits 0 and 1, and $q_b = 0.1$ for $b = 2, \ldots, 9$.

Bins are a deterministic function of the calendar date and the issuance day, $\kappa = \kappa(t, b)$, where the mapping varies with t. Define the date-specific bin distribution induced by the SNAP schedule,

$$\phi_j(t) \equiv \sum_b q(b)\, \mathbf{1}\{\kappa(t, b) = j\}, \tag{10}$$

which is the probability that an independent draw from $q(b)$ lands in bin j for a worker injured on date t.

The object in (10) is common to all workers injured on date t, regardless of issuance day, and this property is the foundation for the derivations that follow. From (9), the distribution of the true bin conditional on the observed bin and the date is

$$\Pr\left(\kappa_i^* = j \mid \kappa_i = k,\, t\right) = (1 - m_c)\, \mathbf{1}\{j = k\} + m_c\, \phi_j(t). \tag{11}$$

Since outcomes depend on the *true* bin, taking expectations of the participant analogue of (3) conditional on

the observed bin yields:

$$
\begin{aligned}
E[y_i \mid \kappa_i = k,\, t,\, c] &= \boldsymbol{\gamma}_t + \sum_j \Pr\left(\kappa_i^* = j \mid \kappa_i = k,\, t\right) \beta_j \\
&= \boldsymbol{\gamma}_t + (1 - m_c)\, \beta_k + m_c \sum_j \phi_j(t)\, \beta_j \\
&= \boldsymbol{\gamma}_t + (1 - m_c)\, \beta_k + m_c\, \bar{\beta}(t), \qquad \bar{\beta}(t) \equiv \sum_j \phi_j(t)\, \beta_j, \qquad (12)
\end{aligned}
$$

where $\bar{\beta}(t)$ is the expected bin effect for participating workers injured on date t if their true issuance days were randomly drawn from the population distribution. Note that the contamination term $m_c\, \bar{\beta}(t)$ does not depend on k. Because the redraw in (9) is independent of b_i, each bin is identically contaminated conditional on date t. Differencing (12) against the reference bin *within a date*, the contamination term cancels:

$$
E[y_i \mid \kappa_i = k, t, c] - E[y_i \mid \kappa_i = 0, t, c] \;=\; (1 - m_c)\, \beta_k. \qquad (13)
$$

Thus, as (13) holds on every date, any weighted average across dates equals $(1 - m_c)\beta_k$ as well, and a specification with full date effects would yield estimates proportionately attenuated by a factor of $1 - m_c$. Since my specification only includes year, month, day-of-month, day-of-week, and holiday fixed effects rather than full date effects, the linear combination of these dummy variables spans $\bar{\beta}(t)$ closely but not perfectly. Therefore, $1 - m_c$ should be viewed as a close approximation for the true misallocation bias factor.

Implied magnitudes and implications for inference

Mirroring the main text, let d index the unique combinations of gender and age bin in my sample. I estimate m_d using the sample of working SNAP participants in Oregon observed in the SNAP Quality Control database over my study period to proxy for the expected misallocation rate for gender-age-*occupation* cells (m_c) in my analysis. The QC data include the gender, age, and employment status for all members of the SNAP household, but not workers' occupations, so if head-of-household shares among d workers differed systematically across occupations, the m_d estimate will be inflated or deflated relative to the relevant m_c for cells in my sample. Samples in the QC data are constructed to be representative of SNAP participants at the state-level so by extension, the sampled *workers* should be representative of working SNAP participants in

Oregon. Thus, m_d can be interpreted as an estimate of the average misallocation rate for d workers across all occupations.

Table C.1: Misallocation rates and implied corrections by gender and age

d	Headship rate (QC data)	Misallocation m_d	Attenuation $\lambda_d = 1 - m_d$	Correction $1/\lambda_d$
Women, 18–24	0.85	0.15	0.85	1.18
Women, 25–34	0.94	0.06	0.94	1.06
Women, 35–44	0.92	0.08	0.92	1.09
Women, 45–54	0.92	0.08	0.92	1.09
Men, 18–24	0.63	0.37	0.63	1.59
Men, 25–34	0.64	0.36	0.64	1.56
Men, 35–44	0.62	0.38	0.62	1.61
Men, 45–54	0.77	0.23	0.77	1.30

Note: These estimates are based on the author's own analysis of the SNAP Quality Control database for fiscal years 2005 to 2018. The number of d workers in each household are aggregated using FNA household weights.

Table C.1 reports the derived attenuation factors $\lambda_d = 1 - m_d$ and implied estimate corrections $1/\lambda_d$ for each group d. Two features of table C.1 directly affect the interpretation of the heterogeneity in figure 4 from the main text. First, attenuation differs drastically by gender due to headship rates being higher among female workers across age bins. For example, if true effects were identical for SNAP participants age 35 to 44 regardless of gender, my estimates for men would appear $1 - 0.62/0.92 \approx 33$ percent smaller than the estimates for women in this age group purely due to issuance day misallocation.

Second, attenuation differs across age bins within gender, though less dramatically than the cross-gender comparisons. For men, headship rates are highest among workers age 45 to 54, so in the case of identical true bin effects, the estimates for men age 35 to 44 would appear $1 - 0.62/0.77 \approx 19$ percent smaller than those for men in the oldest group. Conversely, I find that 94 percent of female workers aged 25 to 34 are their SNAP household heads, compared to only 85 percent of workers under age 25, so identical true effects would appear $1 - 0.85/0.94 \approx 10$ percent smaller for the age 18 to 24 workers relative to 25-34 workers.

As with dilution, misallocation will affect the conditional means by amounts of order $\beta \approx 10^{-4}$, which are negligible relative to the residual standard deviation. The standard error of $\hat{\beta}_k$ is therefore largely unaffected by m_c, so the t-statistic scales with the attenuation factor $(1-m_c)$. Measurement error in true issuance days thus costs statistical power proportionally with the attenuation factor and leaves test size unaffected, by the same argument as in the dilution case.

However, unlike in the dilution case, misallocation will contaminate within-cluster correlation in a way that shrinks standard errors clustered at this level. Suppose residual dependence operates at the level of the true issuance day:

$$\mathrm{Cov}(u_i, u_j) = \rho\,\sigma^2\,\mathbf{1}\{b_i^* = b_j^*\}, \qquad i \neq j, \tag{14}$$

while standard errors are clustered on the *assigned* day based on the observed last digit. Workers sharing a true issuance day may be misallocated, which would violate the core assumption that residual dependence exists within but not across clusters. Furthermore, within-cluster covariance would be diluted since workers within a cluster share a true day only probabilistically. For two observations assigned to day b, independence of the redraws across workers gives

$$\begin{aligned}\Pr\left(b_i^* = b_j^* \mid b_i = b_j = b\right) &= \left[(1-m) + m\,q_b\right]^2 + \sum_{b' \neq b} \left(m\,q_{b'}\right)^2 \\ &\approx \left(1 - 0.9\,m\right)^2,\end{aligned} \tag{15}$$

where the approximation uses $q_b \approx 0.1$ and drops the second term, which is at most $0.11\,m^2$ (about three percent of the first term at $m = 0.35$). Expression (15) is the share of the true within-cluster covariance that survives into the estimator. Averaging the misallocation rates in table C.1 across age groups, weighted by the number of injuries in my sample, gives $m \approx 0.09$ for women and $m \approx 0.35$ for men, so approximately 0.85 of the within-cluster covariance survives for women and 0.48 for men. Misallocation therefore attenuates measured within-cluster dependence, though not completely. If headship rates were below 0.32, so $m > 0.68$, this would eradicate 85 percent of the within-cluster covariance, but tabulations based on the QC data suggest this is not the case for any group of workers in my sample.

Two considerations bound the practical importance of this channel. First, the economic content of dependence at the true issuance-day level is the benefit cycle itself, and the specification models that cycle directly through the bin indicators and the issuance-day fixed effects; the residual can carry only the component of the cycle response that the bin partition fails to capture.[19] Second, dilution compounds the bound: with participation share π, only participant pairs can share a true benefit cycle, scaling any residual dependence by a further factor of order π^2. Consistent with both considerations, comparisons of cluster-robust to

[19] One caveat to this is that since the issuance-day fixed effects are defined on *assigned* days, the true-day common component for misallocated workers is not captured, and this component is m-weighted and second order.

heteroskedasticity-robust variance estimates across candidate clustering dimensions (available upon request) show no detectable within-cluster dependence on any dimension with enough clusters for the comparison to be informative.

Combined correction

The two error processes compound multiplicatively. Misallocation operates only within the participant subpopulation, so combining (5) and (12), the conditional mean in cell c is

$$E[y_i \mid \kappa_i = k,\, t,\, c] = \boldsymbol{\gamma}_t + \pi_c \Big[(1 - m_c)\, \beta_k + m_c\, \bar{\beta}(t) \Big], \tag{16}$$

and the within-date contrast identifies

$$\hat{\beta}_k \xrightarrow{p} \pi_c\, (1 - m_c)\, \beta_k \tag{17}$$

in cell c, with the pooled estimate attenuated by the corresponding weighted average. With $\pi \approx 0.51$ and the estimates of misallocation rates m_d in Table C.1, the combined attenuation factor $\pi\lambda_c$ ranges from $0.51 \times 0.94 \approx 0.48$ (women 25–34) to $0.51 \times 0.62 \approx 0.32$ (men 35–44): the estimated coefficients understate per-participant effects by a factor of roughly 2.1 to 3.2. Dilution accounts for the majority of the gap, while misallocation affects the variation in magnitudes across demographic cells.

References

ARTEAGA, I., C. HEFLIN, AND L. HODGES (2018): "SNAP benefits and pregnancy-related emergency room visits," *Population Research and Policy Review*, 37, 1031–1052.

BITLER, M. AND H. HOYNES (2015): "Heterogeneity in the Impact of Economic Cycles and the Great Recession: Effects within and across the Income Distribution," *American Economic Review*, 105, 154–160.

BITLER, M. P., J. B. GELBACH, AND H. W. HOYNES (2005): "Welfare reform and health," *Journal of Human resources*, 40, 309–334.

BLACK, M. M., A. M. QUIGG, J. COOK, P. H. CASEY, D. B. CUTTS, M. CHILTON, A. MEYERS, S. E. DE CUBA, T. HEEREN, S. COLEMAN, ET AL. (2012): "WIC participation and attenuation of stress-related child health risks of household food insecurity and caregiver depressive symptoms," *Archives of pediatrics & adolescent medicine*, 166, 444–451.

BLS (2018a): "Table R19. Number of Nonfatal Occupational Injuries and Illnesses Involving Days Away from Work by Part of Body and Selected Natures of Injury or Illness, Private Industry," Survey of Occupational Injuries and Illnesses, Case and Demographic Characteristics Resource Tables, corrected estimates for reference year 2011.

——— (2018b): "Table R67. Number of Nonfatal Occupational Injuries and Illnesses Involving Days Away from Work by Nature of Injury or Illness and Number of Days Away from Work, and Median Number of Days Away from Work, Private Industry," Survey of Occupational Injuries and Illnesses, Case and Demographic Characteristics Resource Tables, corrected estimates for reference year 2011.

——— (2018c): "Table R68. Number of Nonfatal Occupational Injuries and Illnesses Involving Days Away from Work by Part of Body and Number of Days Away from Work, and Median Number of Days Away from Work, Private Industry," Survey of Occupational Injuries and Illnesses, Case and Demographic Characteristics Resource Tables.

——— (2023): "Coverage in Employer Medical Care Plans among Workers in Different Wage Groups in 2022," The Economics Daily, accessed October 18, 2026.

——— (2026a): "Number and Rate of Nonfatal Work Injuries in Detailed Private Industries, 2024," Injuries, Illnesses, and Fatalities Program, Survey of Occupational Injuries and Illnesses, accessed October 18, 2026.

——— (2026b): "Table 43: Median Weekly Earnings of Full-Time Wage and Salary Workers by Union Affiliation, Occupation, and Industry," Labor Force Statistics from the Current Population Survey, Household Data Annual Averages, annual averages for 2024 and 2025. Last modified February 20, 2026. Accessed October 18, 2026.

BOND, T. N., J. B. CARR, A. PACKHAM, AND J. SMITH (2022): "Hungry for success? SNAP timing, high-stakes exam performance, and college attendance," *American Economic Journal: Economic Policy*, 14, 51–79.

BOONE, J. AND J. C. VAN OURS (2006): "Are recessions good for workplace safety?" *Journal of Health economics*, 25, 1069–1093.

BOONE, J., J. C. VAN OURS, J.-P. WUELLRICH, AND J. ZWEIMÜLLER (2011): "Recessions are bad for workplace safety," *Journal of health economics*, 30, 764–773.

BRAVEMAN, P. A., C. CUBBIN, S. EGERTER, D. R. WILLIAMS, AND E. PAMUK (2010): "Socioeconomic disparities in health in the United States: what the patterns tell us," *American journal of public health*, 100, S186–S196.

BRONCHETTI, E. T. AND M. MCINERNEY (2012): "Revisiting incentive effects in workers' compensation: Do higher benefits really induce more claims?" *ILR Review*, 65, 286–315.

BROSS, I. D. (1958): "How to use ridit analysis," *Biometrics*, 18–38.

CABRAL, M., C. CUI, AND M. DWORSKY (2022): "The demand for insurance and rationale for a mandate: Evidence from workers' compensation insurance," *American Economic Review*, 112, 1621–1668.

CARD, D. AND B. P. MCCALL (1996): "Is workers' compensation covering uninsured medical costs? Evidence from the "monday effect"," *ILR Review*, 49, 690–706.

CARR, J. B. AND A. PACKHAM (2019): "SNAP benefits and crime: Evidence from changing disbursement schedules," *Review of Economics and Statistics*, 101, 310–325.

——— (2021): "SNAP schedules and domestic violence," *Journal of Policy Analysis and Management*, 40, 412–452.

CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, AND B. MELLY (2013): "Inference on counterfactual distributions," *Econometrica*, 81, 2205–2268.

COOK, J. AND C. N. EAST (2026): "The disenrollment and labor market effects of SNAP work requirements on parents," *Journal of Public Economics*, 257, 105631.

COTTI, C., J. GORDANIER, AND O. OZTURK (2016): "Eat (and drink) better tonight: Food stamp benefit timing and drunk driving fatalities," *American Journal of Health Economics*, 2, 511–534.

——— (2018): "When does it count? The timing of food stamp receipt and educational performance," *Economics of Education Review*, 66, 40–50.

COTTI, C. D., J. M. GORDANIER, AND O. D. OZTURK (2020): "Hunger pains? SNAP timing and emergency room visits," *Journal of health economics*, 71, 102313.

DAVIES, M., R. J. PARK, AND A. STANSBURY (2026): "How Do Minimum Wages Affect Workplace Injuries?" *Employment Research Newsletter*, 33, 1.

DEMBE, A. (1999): "Social inequalities in occupational health and health care for work-related injuries and illnesses," *International Journal of Law and Psychiatry*, 22, 567–579.

DEMING, W. E. AND F. F. STEPHAN (1940): "On a least squares adjustment of a sampled frequency table when the expected marginal totals are known," *The Annals of Mathematical Statistics*, 11, 427–444.

DILLENDER, M. (2015): "The effect of health insurance on workers' compensation filing: Evidence from the affordable care act's age-based threshold for dependent coverage," *Journal of Health Economics*, 43, 204–228.

D'ERRICO, A., L. PUNNETT, M. CIFUENTES, J. BOYER, J. TESSLER, R. GORE, P. SCOLLIN, AND C. SLATIN (2007): "Hospital injury rates in relation to socioeconomic status and working conditions," *Occupational and environmental medicine*, 64, 325–333.

EICHER-MILLER, H. A., A. C. MASON, C. M. WEAVER, G. P. MCCABE, AND C. J. BOUSHEY (2009): "Food insecurity is associated with iron deficiency anemia in US adolescents," *The American journal of clinical nutrition*, 90, 1358–1371.

EVANS, W. N. AND T. J. MOORE (2012): "Liquidity, economic activity, and mortality," *Review of Economics and Statistics*, 94, 400–418.

FNA (2026): "Supplemental Nutrition Assistance Program (SNAP) Quality Control Database, Fiscal Years 2005–2018," Data files hosted by Mathematica under contract with USDA FNA, accessed: 2026-10-05.

FNS (2025): "SNAP Program Access Index," Online database, page updated January 24, 2025; accessed via SNAP Quality Control.

FORESI, S. AND F. PERACCHI (1995): "The conditional distribution of excess returns: An empirical analysis," *Journal of the American Statistical Association*, 90, 451–466.

GASSMAN-PINES, A. AND L. BELLOWS (2018): "Food instability and academic achievement: A quasi-experiment using SNAP benefit timing," *American Educational Research Journal*, 55, 897–927.

GENNETIAN, L. A., R. SESHADRI, N. D. HESS, A. N. WINN, AND R. M. GOERGE (2016): "Supplemental Nutrition Assistance Program (SNAP) benefit cycles and student disciplinary infractions," *Social Service Review*, 90, 403–433.

GREGORY, C. A. AND P. DEB (2015): "Does SNAP improve your health?" *Food Policy*, 50, 11–19.

GUARDADO, J. R. AND N. R. ZIEBARTH (2019): "Worker investments in safety, workplace accidents, and compensating wage differentials," *International Economic Review*, 60, 133–155.

HASTINGS, J. AND E. WASHINGTON (2010): "The first of the month effect: consumer behavior and store responses," *American economic Journal: economic policy*, 2, 142–162.

HOYNES, H., D. L. MILLER, AND J. SCHALLER (2012): "Who suffers during recessions?" *Journal of Economic perspectives*, 26, 27–48.

HRADIL, V. (2018): "Does minimum wage affect workplace safety?" *CERGE-EI Working Paper Series*.

KANTOR, S., W. WINKELSTEIN JR, AND M. A. IBRAHIM (1968): "A note of the interpretation of the ridit as a quantile rank," *American Journal of Epidemiology*, 87, 609–615.

KRUEGER, A. B. (1990a): "Incentive effects of workers' compensation insurance," *Journal of Public Economics*, 41, 73–99.

——— (1990b): "Workers' compensation insurance and the duration of workplace injuries," .

KUHN, M. A. (2018): "Who feels the calorie crunch and when? The impact of school meals on cyclical food insecurity," *Journal of Public Economics*, 166, 27–38.

LAFIANDRA, W. AND D. SCHWAB (2023): "A higher minimum wage promotes workplace safety: Evidence from the restaurant industry," in *Proceedings of The 7th International Conference on Applied Research in Management, Economics and Accounting (94-114). https://www. doi. org/10.33422/7th. iarmea*, vol. 125.

NASI (2024): "Workers' Compensation: Benefits, Costs, and Coverage, 2022 Data," Tech. rep., National Academy of Social Insurance, 27th annual report.

OHLSON, M. (2020): "Effects of socioeconomic status and race on access to healthcare in the United States," *Perspectives*, 12, 2.

RUGGLES, S., S. FLOOD, M. SOBEK, D. BACKMAN, G. COOPER, J. A. RIVERA DREW, S. RICHARDS, R. RODGERS, J. SCHROEDER, AND K. C. W. WILLIAMS (2025): "IPUMS USA: Version 16.0 [dataset]," .

SELIGMAN, H. K., A. B. BINDMAN, E. VITTINGHOFF, A. M. KANAYA, AND M. B. KUSHEL (2007): "Food insecurity is associated with diabetes mellitus: results from the National Health Examination and Nutrition Examination Survey (NHANES) 1999–2002," *Journal of general internal medicine*, 22, 1018–1023.

SELIGMAN, H. K., A. F. BOLGER, D. GUZMAN, A. LÓPEZ, AND K. BIBBINS-DOMINGO (2014): "Exhaustion of food budgets at month's end and hospital admissions for hypoglycemia," *Health affairs*, 33, 116–123.

SMITH, R. S. (1990): "Mostly on Monday: is workers' compensation covering off-the-job injuries?" in *Benefits, costs, and cycles in workers' compensation*, Springer, 115–127.

SOMMERS, B. D., A. L. GOLDMAN, R. J. BLENDON, E. J. ORAV, AND A. M. EPSTEIN (2019): "Medicaid work requirements—results from the first year in Arkansas," *New England Journal of Medicine*, 381, 1073–1082.

STEEGE, A. L., S. L. BARON, S. M. MARSH, C. C. MENÉNDEZ, AND J. R. MYERS (2014): "Examining occupational health and safety disparities using national data: A cause for continuing concern," *American journal of industrial medicine*, 57, 527–538.

STEENLAND, K., W. HALPERIN, S. HU, AND J. T. WALKER (2003): "Deaths due to injuries among employed adults: the effects of socioeconomic class," *Epidemiology*, 14, 74–79.