Accounting for the Survivors: Traffic Injuries, Match Destruction, and the Cost of Nonfatal Accidents*

Kuan-Ming Chen† Chen-An Lien‡

October 7, 2026

Abstract

Most crash victims survive, yet road-safety appraisal prices a nonfatal injury as a fixed fraction of a fatality, scaled by its medical severity, because no data follow the injured beyond their treatment. We measure directly what a nonfatal traffic injury, the most common acute health shock to working-age adults, costs the worker who survives it. Linking the universe of police-reported crashes in Taiwan to 24 years of health, employment, and mortality registries, we estimate dynamic effects with a staggered event-study design that compares earlier- to later-injured workers. Health and economic trajectories diverge: medical utilization spikes in the accident year and largely subsides, while earnings and full-time employment losses widen to roughly 9 percent by year five. The loss runs through exit from full-time work and separation from the pre-accident employer, persists among workers whose health has recovered, and depends more on the job the injury interrupts than on the injury itself. Over a working life, the measured medical and earnings cost is about USD 13,000 per hospitalized injury. Rules that fix an injury's value at the time of injury match this average but miss its timing and its incidence, so fatality-centered evaluations understate the returns to investments that reduce crash severity.

*We thank Yu-Chang Chen, Ming-Jen Lin, Lin-Tung Tsai, and seminar participants in the NBER Economics of Transportation in the 21st Century Conference, NTHU, NTPU, Academia Sinica, UTokyo, Hitotsubashi, and AASLE for helpful comments. The data is kindly provided by the Health and Welfare Data Science Center and the National Police Agency, Ministry of the Interior in Taiwan. Kuan-Ming Chen acknowledges the support from the National Science and Technology Council grant NSTC 111-2628-H-002-019 and the Yushan Fellow Program by the Ministry of Education (MOE) NTU-113V1037-2.

†National Taiwan University. Email: kuanmingchen@ntu.edu.tw.

‡University of Chicago. Email: ahn890209@gmail.com.

Keywords: traffic injuries; health shocks; value of a statistical life; road-safety appraisal. *JEL codes:* I12, I18, J63, R41, H54.

1 Introduction

Most victims of traffic crashes survive. Road crashes kill about 1.19 million people each year and injure 20 to 50 million more (WHO, 2023), so a nonfatal traffic injury is one of the most common acute health shocks that working-age adults face. Governments spend close to 1 percent of GDP building and maintaining inland transport infrastructure (OECD International Transport Forum, 2023), and they justify that spending mainly by the fatalities it avoids, priced with the value of a statistical life (VSL). Nonfatal injuries enter the same calculation as fractions of a fatality. U.S. transportation guidance values an injury at a fixed share of the VSL determined by its medical severity score, from 0.3 percent for the mildest class to 59 percent for the most critical (U.S. Department of Transportation, 2016); European appraisal handbooks use two tiers, serious and slight (CE Delft et al., 2019).

None of these methods directly measure what an injury costs the person who survives it. Human-capital accounting adds a hospital bill per injury class to lost production, computed from a wage rate and a scheduled number of work-loss days for that diagnosis (Blincoe et al., 2015; Wijnen and Stipdonk, 2016); severity weighting scales a fatality value by a medical score. In both cases the labor-market loss is read off the medical severity of the injury. This is largely a data constraint. Discharge files, surveys, and compensation claims record the treatment and stop there, so analysts cannot see what happens to the injured person's earnings and employment in the years that follow. The U.S. cost inventory itself notes that no recent source exists for the probability of permanent disability by diagnosis or for medical costs beyond eighteen months after injury (Blincoe et al., 2015).[1]

This paper measures what a hospitalized traffic injury costs a surviving worker, in health-care utilization and in earnings, over the years after the accident. A newly linked administrative panel lets us observe these outcomes rather than impute them. We follow the injured in records that contain their medical claims, their earnings, their employers, and their survival, so the labor-market loss is estimated directly. Following the same people on both margins also answers two questions the imputation methods cannot ask. First, how do the medical

[1] "A major limitation of the costs presented is that some cost components are unavoidably quite old. In particular, no recent source exists for the percentage of lifetime medical costs that is incurred more than 18 months post-injury, probabilities of permanent disability by detailed diagnosis and whether hospital admitted, or the ratio of household work days lost to wage work days lost" (Blincoe et al., 2015, p. 24).

and the economic consequences unfold over time, and do they move together? Second, do workers whose health recovers also recover economically, or does the economic loss persist after the medical one has passed?

We link the universe of police-reported traffic accidents in Taiwan (2003–2023) to 24 years (2000–2023) of National Health Insurance (NHI) enrollment, medical claims, employment records, and mortality registries. We estimate the dynamic effects of a first hospitalized injury with a staggered event-study design that compares earlier-injured workers to observably similar workers injured in later years (Callaway and Sant'Anna, 2021). Defining the event from police records rather than claims means the sample is not selected on care-seeking or treatment intensity, and the length of the panel lets us follow trajectories long after statutory protections expire. The design accounts for life-cycle differences and avoids the biases of conventional two-way fixed effects estimators, which we show understate the long-run losses.

Our first finding is that the health consequences of a hospitalized injury are concentrated in the accident year, while the labor-market consequences build up over time. NHI-covered medical spending rises by roughly 360 percent relative to the pre-accident control mean in the accident year, falls sharply in the following year, and remains only modestly above baseline through year five, with the remaining spending shifting from acute trauma care to follow-up care. Earnings follow a different path. The drop is modest in the accident year and then widens steadily. Five years after the accident, annual earnings and full-time employment are both roughly 9 percent below their counterfactual levels, and the extensive margin accounts for 80 to 95 percent of the earnings loss. The cost of a nonfatal injury is therefore back-loaded, and it consists of workers who are no longer in full-time work, not workers who earn less while working. One possible explanation of this divergence is that the injury is simply long-lasting and that, after the first year, there is little more that medical care can do for it. We test that explanation below and find that most of the loss falls on workers whose health has recovered.

Our second finding is that the labor-market loss runs through separation from the pre-accident employer. Explaining why the loss keeps growing after health has recovered requires knowing what happens to the job, so we follow each worker's employer. Separations from the pre-accident employer are concentrated in the months after the crash and decline smoothly

afterward. Five years later, injured workers are 2 percentage points less likely to be with their pre-accident employer, against a 3 percentage point decline in full-time employment at any employer. Among workers attached to an employer before the accident, the 36 percent who leave that employer within two years account for 91 percent of the year-five earnings decline, while workers who stay earn 1.3 percent less than before the accident.

Lasting work incapacity accounts for only a bounded share of this loss. Among workers whose medical utilization has returned to the range of not-yet-injured controls by the second post-accident year, roughly three-quarters of the year-five earnings and employment losses remain. Among all injured individuals, 6.1 percent are newly certified for catastrophic illness in the accident year or the year after, and their year-five earnings loss is about four times the average. Multiplying the 6.1 percent certified share by that group's own year-five loss attributes about one quarter of the aggregate year-five earnings loss to new certification. Even if every newly certified worker lost all pre-accident earnings, the share would be at most 64 percent. Cumulative mortality rises by 1.6 percentage points by year five, against 1.34 percent in the comparison group, so survivors are not healthy on every margin, but most of the economic loss is borne by workers whose measured health has recovered and who are not working full time.

Our third finding is that the size of the loss depends on the job the injury interrupts more than on the injury itself. More severe injuries generate larger losses, but across groups with similar medical responses, labor-market trajectories diverge sharply. After standardizing every subgroup to a common injury-severity distribution, the year-five earnings loss in the top quartile of pre-accident earnings is more than eight times the loss in the bottom quartile, and the full-time employment losses are 24 and 5 percentage points. Men and car drivers lose more than women and other road users at the same severity.

Taken together, these results supply what current practice imputes: a cost per nonfatal injury measured on the injured population itself. Extended to the end of a working life and discounted at 3 percent, the medical spending and lost earnings we observe come to about TWD 393,000 (USD 13,000) per hospitalized working-age injury. To see what the imputation misses, we apply the two rulebooks used in the United States and Europe to the same Taiwanese inputs, so that the comparison is between methods rather than between

countries. Human-capital rules add a medical cost to a scheduled production loss; VSL-based rules take a fixed fraction of the value of a statistical life. The human-capital rules, which schedule a fixed number of years of lost production per injury class, get the average about right: they assign 1.3 to 1.6 years of production to an injury of this severity, and we observe 1.4 years of the injured population's own earnings. What no rule that fixes the value at the time of injury can see is where the loss falls and when it arrives: on workers who recover medically, years after the medical bills that severity weights are built on. The rules that appraisal actually uses multiply a VSL by a severity fraction meant to cover pain, suffering, lost quality of life, and the risk of later death. Against those values the tangible components we measure are a few percent, and a quarter to a third once the excess mortality we estimate is valued at the Taiwan VSL, so the official values are built almost entirely on the part of the cost that no record measures.

The decomposition also says who bears the cost. In the framework of the marginal value of public funds, the government's saving from a prevented injury, the NHI spending plus the taxes and contributions on the lost earnings, is large enough that a measure preventing a hospitalized injury for less than roughly TWD 160,000 pays for itself in fiscal terms; the larger private loss accrues years after the accident. In Taiwan, hospitalized nonfatal injuries outnumber traffic fatalities by roughly eight to one, so the survivor margin is large enough to change the benefit side of road-safety appraisal, most of all for interventions that make injuries milder rather than prevent crashes.

Our findings speak first to the literature on the economic consequences of acute health shocks. That injuries and hospitalizations cause persistent earnings losses is established: Dobkin et al. (2018) show it for hospital admissions among insured adults, and register-based studies document it after road and commuting injuries in Denmark, Austria, and Chile (Dano, 2005; Halla and Zweimüller, 2013; García-Gómez et al., 2013; Parro and Pohl, 2021; Barr et al., 2025). Bíró et al. (2025) show that mild accidents, with full-recovery diagnoses and short absences, reduce wages by about 1 percent through forgone wage growth among workers who remain employed, with no employment response.

We complement and extend this literature by observing both sides of the process for the same individuals. Each existing study observes one side, and that leaves the policy practice,

reading the economic loss off a medical severity score, untested. Labor registers document earnings losses but cannot trace medical recovery, so they cannot say how well the severity of an injury predicts its economic cost, the relationship that valuation methods assume when they infer the labor-market loss from a medical score. Defining the event from the accident rather than from treatment is the second thing our data allow. Nearly all existing studies identify the injury from a hospital admission or an insurance record rather than from the accident itself.[2] Claims-based designs observe treatment but condition the sample on care-seeking; in our data, diagnosis-based transport-injury codes recover only about 40 percent of police-reported accidents, and conditioning on them understates the effects in our data (Appendix N). We observe health, employer identity, earnings, and survival within the same individuals at population scale, and we show that at hospitalization severity the loss runs through exit from full-time work among workers who recover medically, with new catastrophic-illness certification accounting for a bounded and minority share. Set beside the mild-injury evidence, the pattern is that as severity rises, the binding margin shifts from wage growth among stayers to detachment itself, and the losses become an order of magnitude larger and widen instead of fading.

The detachment we document also connects acute health shocks to the literature on job displacement (Jacobson et al., 1993) and on the value of specific employment relationships (Jäger and Heining, 2022; Johnsen et al., 2024). A hospitalization acts as an idiosyncratic event that separates workers from their employers, and the matches that break are the ones with the least accumulated specific capital, at the smallest firms and the shortest tenures. Unlike mass layoffs, however, the loss here shows up mainly as non-employment, not as re-employment at lower pay. Finally, we contribute to the transportation and urban economics literature that measures the social cost of road crashes through cost inventories, human-capital calculations, and VSL or willingness-to-pay valuations (Elvik, 2000; Finkelstein et al., 2006; Blincoe et al., 2015; Wijnen and Stipdonk, 2016; Bougna et al., 2022). We provide the longitudinal parameters needed to price the survivor burden, and we show why severity-weight valuation misprices it.

[2]Dano (2005), García-Gómez et al. (2013), Dobkin et al. (2018), and Parro and Pohl (2021) define the event from hospital records, and Halla and Zweimüller (2013) and Bíró et al. (2025) from social- or health-insurance records. Barr et al. (2025) start from police crash reports of minor crashes.

The remainder of the paper proceeds as follows. Section 2 describes Taiwan's road-safety environment, health insurance, and employment protections. Section 3 describes the linked administrative data and sample construction. Section 4 presents the empirical framework. Section 5 reports the average dynamics, the margins on which the loss occurs, the incapacity bounds, heterogeneity, mortality, and robustness. Section 6 reviews how nonfatal injuries are priced now, sets our estimate beside those rules applied to Taiwanese inputs, and reads the cost through public finance and road-safety appraisal. Section 7 concludes.

2 Background

2.1 Research Context

2.1.1 Reporting and Classification of Traffic Accidents

When a traffic accident occurs in Taiwan, the parties involved report it to the police. Reporting is highly incentivized for accidents with injuries because every vehicle owner must carry Compulsory Automobile Liability Insurance (CALI), and the claim regulations under the CALI Act require the insured to report an accident to the police immediately and treat the police accident report as the basis for settling claims (Financial Supervisory Commission, Taiwan, 2022b, Arts. 15 and 17). Officers from the nearest station document the incident, including the identities of those involved, road and weather conditions, and any apparent injuries. The resulting record documents the health shock independently of the individual's subsequent care-seeking or hospital admission, which lets us define accident exposure from police records rather than from medical claims. In 2023 the police recorded 403,085 injury accidents, 539,535 injured persons, and 3,023 deaths, about 2,300 injured persons per 100,000 residents and 178 injured persons per road death (National Police Agency, 2024); the corresponding figures are roughly 430 and 130 in Germany and 290 and 140 in Japan (Statistisches Bundesamt (Destatis), 2024; National Police Agency of Japan, 2024).[3] Part of

[3]Germany, 2023: 291,890 accidents involving personal injury, 366,557 injured, 2,839 killed, population about 84.7 million. Japan, 2023: 307,930 accidents, 365,595 injured, 2,678 killed, population about 124.4 million. Taiwan's population in 2023 was about 23.4 million. A national count of property-damage-only (A3) crashes is unavailable because local governments retain the A3 records and most do not publish aggregate figures. Injured persons per road death is the more informative ratio for reporting completeness, because deaths are recorded nearly completely everywhere while minor injuries are the first to go unreported.

the gap reflects Taiwan's motorcycle-heavy traffic, but the number of recorded injuries per fatality indicates that the police record reaches further down the severity distribution than in peer countries.

2.1.2 Health Coverage and Compensation Mechanisms

Once the police have recorded an accident, two mechanisms determine who pays for the injuries: the NHI system and the CALI program whose reporting rule underlies the police record.

Taiwan's NHI is a single-payer, compulsory system that covers the entire population, so an injured person is treated immediately, before liability is settled, and NHI pays about three-quarters of the medical spending that traffic accidents generate (MOHW, 2022). CALI covers most of the rest: it pays the injured party fixed benefits regardless of fault, medical expenses up to TWD 200,000 per person per accident and lump sums for disability, on a fifteen-grade scale, and for death (Financial Supervisory Commission, Taiwan, 2022a, 2026, Arts. 7 and 27).[4] Neither program replaces lost earnings. Claims for lost earnings run through civil liability and voluntary third-party insurance.

Beyond immediate medical care, the most severely injured individuals may qualify for ongoing income support and disability services. In the Ministry of Health and Welfare's annual disability register, traffic accidents were the recorded cause of approximately 3% of all certified disabilities at the end of 2023 (38,194 of 1,214,668), the subset of injured individuals whose impairments are severe enough to meet formal disability criteria (Ministry of Health and Welfare, Department of Statistics, 2026).

For those who suffer complete loss of working capacity, CALI provides up to TWD 2 million in disability compensation, covering roughly half of their estimated economic losses. Benefits are scaled across 15 disability grades, with the least severe injuries receiving approximately 30% of their projected losses.[5]

[4]The medical cap was unchanged over our sample period. The death and first-grade disability benefits rose in steps from TWD 1.4 million in 2003 to TWD 2 million from 2012 onward, and to TWD 3 million from July 2026, after our sample ends.

[5]For comparison, the U.S. Social Security Disability Insurance (SSDI) program provides ongoing monthly payments to workers whose disabilities limit their ability to work. In January 2026, the estimated average monthly benefit for disabled workers was USD 1,630, or approximately USD 19,600 annually (Social Security

2.1.3 Employment Protection after Work-Related Injuries

Beyond medical and disability compensation, Taiwan's labor regulations provide a buffer against immediate job loss, which helps explain why the labor-market detachment we observe is gradual. For traffic injuries that qualify as occupational injuries, including commuting accidents that occur along a standard route, the Labor Standards Act requires employment protections (Ministry of Labor, Taiwan, 2024). Employers are generally prohibited from terminating a worker during their medical treatment period and must continue to pay their pre-injury wage. If the worker has not recovered after two years, employers can opt for a lump-sum settlement to release this wage-compensation obligation. For a large share of injured workers, the employment relationship is therefore legally preserved in the months after the crash. Whether the protection merely postpones dismissals is testable, since the enrollment records identify each worker's employer month by month: the monthly hazard of leaving the pre-accident employer declines smoothly through the twenty-fourth month with no jump when the obligation expires (Section 5.2 and Appendix K).

Once the acute treatment phase ends, the Labor Occupational Accident Insurance and Protection Act requires employers to assist workers in returning to their original roles or offer suitable alternative positions with necessary workplace accommodations. The closest U.S. analogues are the Family and Medical Leave Act (FMLA), which grants up to 12 weeks of unpaid, job-protected leave, and the Americans with Disabilities Act (ADA), which requires reasonable workplace accommodations (U.S. Department of Labor, 2025).

These protections are partial and temporary. Wage continuation ends when the statutory period expires or a lump sum is paid, and compensation is keyed to impairment grades, so the persistent earnings losses documented below fall largely on injured workers and their families.

Administration, 2026). Unlike CALI compensation, SSDI benefits are based on prior earnings histories and paid monthly. See also Deshpande and Lockwood (2022) for an analysis of the insurance value of U.S. disability programs.

3 Data

Our empirical analysis links two administrative data sources. We use police records to define traffic-injury events and NHI records to measure medical utilization and labor-market outcomes. We then construct the person-level analysis sample around each individual's first recorded accident.

3.1 Accident Data from the National Police Agency

We construct the core accident dataset using administrative records from the National Police Agency (NPA), which document all officially reported traffic accidents from 2003 to 2023. Our data include the universe of accidents involving a death or an injury and exclude property-damage-only cases.[6] Because reporting is closely tied to compulsory insurance claims, as discussed in Section 2.1, NPA records provide near-universal coverage of medically relevant accidents. Police records identify the occurrence, timing, and circumstances of accidents.

The NPA data are organized at two levels. The case-level file contains a unique case identifier, the date of the accident, and the accident case type. The individual-level file contains one record for each person involved in an accident and provides the case identifier, an encrypted individual identifier that links to the NHI records, gender, vehicle type, road-user role (driver, passenger, or pedestrian), injury status, and the anatomical location of injury. Linking the two files by case identifier gives a person-level record of accident involvement covering over 5.45 million accident cases and more than 7.3 million individuals.

3.2 National Health Insurance Data

To measure post-accident labor-market and health trajectories, we use the National Health Insurance Research Database (NHIRD), the administrative record of the NHI program described in Section 2. Because NHI is compulsory and universal, more than 99 percent of Taiwan's residents are enrolled (National Health Insurance Administration, Ministry of Health

[6]The police classify accidents as A1 if someone dies at the scene or within 24 hours, A2 if someone is injured but no one dies within 24 hours, and A3 if there is property damage only. Our data cover A1 and A2 accidents.

and Welfare, 2023), and because copayments are low, utilization is high: working-age individuals in our sample average more than ten outpatient visits per year before any accident (Table 1). The NHIRD provides enrollment and claims data for the entire population from 2000 to 2023, which lets us measure employment status, earnings, and medical utilization for everyone in the police records.

The NHIRD has two main files, the enrollment file and the claims file. The enrollment file spans 2000 to 2023 and covers more than 23 million individuals a year. It records sex, date of birth, residential township, insured salary, employment status, and firm identifier. Our employment measure covers public-sector workers and other employer-based insurance categories as well as private-sector employees, but excludes the self-employed and military personnel. We deflate all salary and earnings measures using Taiwan's consumer price index (CPI) and express them in 2021 New Taiwan dollars (TWD). Employment status is updated monthly, so we observe job separations, employment spells, and moves across firms.

The claims file records every reimbursed inpatient, outpatient, and emergency service, with ICD-coded diagnoses, procedures, prescriptions, dates of service, and itemized expenditures. From it we measure the intensity and timing of treatment.

The claims files could in principle also identify traffic injuries via ICD codes, and most register-based studies define injury events from hospital or claims records (Dano, 2005; García-Gómez et al., 2013; Bíró et al., 2025). We instead define accident exposure from the police records. In our linked data, diagnosis-based transport-injury codes recover only about 40 percent of police-reported injury accidents, that is, accidents in which anyone was injured or killed, so a claims-based definition would condition the sample on care-seeking and on how the treating physician coded the visit.[7] Appendix N compares estimates on the claims-identifiable subsample with the police-defined sample.

Together the two files follow every individual month by month on both the labor-market and the health side, before and after the accident.

[7] Officially, classes A1 and A2 (Section 3.1). We identify motor vehicle traffic accidents in claims using ICD-10-CM external cause-of-injury codes (V-codes), restricting the definition to motor vehicle traffic crashes.

3.3 Sample Construction

Our main analysis sample consists of working-age individuals, aged 25 to 65, who were hospitalized within two days of their first police-reported accident, about 432,000 people. We reach it in three steps. We start from every individual ever recorded as injured in a police-reported accident, keep only each individual's first accident so that later accidents do not confound the event, and exclude the few who died within 30 days of it, using the NHI death registry, because the paper is about survivors. This yields approximately 5.2 million injured individuals; Appendix P summarizes their accidents. We then restrict to those injured at working age, about 2.8 million, and finally to those hospitalized within two days of the accident. Police records include many minor injuries; the hospitalization restriction selects injuries severe enough to interrupt work and makes the health shock comparable across individuals. Outcomes are aggregated to calendar years, and event time counts calendar years relative to the accident year, so the accident year itself contains both pre- and post-accident months. Results for the two broader injured samples are reported alongside for comparison.

For description only, we also construct a matched comparison group of people who had not yet had a recorded traffic accident. The estimation in Section 4 does not use this group; it compares earlier-injured with later-injured workers. The matched group serves to show how the eventually injured differ from everyone else before any accident. For each of the 5.2 million injured individuals, we select one counterpart from the same calendar year, matched on age (in years), gender, and residential district (at the zip code level, an area with about 60,000 residents on average), and we require that the assigned year precede any later accident of the matched person, so that the comparison group is observed before any accident. For instance, if a 32-year-old male residing in Da'an District of Taipei City was injured in 2010, we select a non-injured individual from the same age-gender-district group in that year. For the hospitalized main sample, the comparison group is the subset of these matched individuals who were themselves hospitalized in the matched year, about 194,000 people. It is smaller than the main sample because the injured are selected on hospitalization and their matches are not; only the matches who happened to be hospitalized that year are kept, so

TABLE 1: Summary Statistics of Injured Individuals and Matched Comparisons

Period:	A year before accidents					
Sample:	Injured	Non-Injured	Injured	Non-Injured	Injured	Non-Injured
Age restriction (25–65)			✓	✓	✓	✓
Hospitalization					✓	✓
Panel A: Demographics						
Mean Age	36.5	36.4	42.9	42.9	46.1	46.2
Male Ratio	52%	52%	49%	49%	48%	51%
Panel B: Employment Status						
Mean Earning (TWD/year)	144,695	150,233	229,110	248,344	185,724	193,843
Mean Earning of Employee	369,313	421,985	431,029	485,563	423,592	469,042
Work Rate	27.7%	26.8%	44.1%	43.9%	35.8%	34.7%
Panel C: Health Outcomes						
Hospitalized Days	0.61	0.8	0.75	0.9	1.1	7.7
Outpatient Visits	12.1	10.1	12.3	10.5	13.7	18.7
Applied NHI Points	16,427	16,040	18,354	17,250	22,433	70,845
Number of Individuals	5,151,050	5,151,033	2,759,225	2,733,940	432,449	193,593

Note: This table reports descriptive statistics for accident-injured individuals and their matched counterparts in the year prior to the accident (or pseudo-accident year for the matched set). *Injured (All)* includes all injured individuals; *Injured (25–65)* restricts to those injured at working age; and *Injured (25–65, Hosp.)* further limits to those hospitalized within two days of the accident. Each treated group is paired with a matched set of non-injured individuals from the same calendar year, age, gender, and residential district. For the hospitalized subsample, the comparison group is the subset of matched individuals who were themselves hospitalized in the matched year; it is smaller than the injured group because the injured are selected on hospitalization and their matches are not. The slight shortfall in the other two comparison columns reflects cases with no same-year match, for instance very old individuals in sparsely populated districts. These statistics describe baseline differences between the eventually injured and the never injured; the event-study design in Section 4 does not use the matched group.

that both groups experienced a hospitalization.

3.4 Summary Statistics and a Naive Comparison

Table 1 reports summary statistics for the main outcomes in the year before the accident, or the matched pseudo-accident year, for the injured samples and their matched never-injured comparison groups.

Since we focus on the labor-market costs of traffic injuries, we gradually restrict attention to injured working-age individuals. In the summary statistics, *Injured (All)* includes all individuals injured in traffic accidents, *Injured (25–65)* restricts to those injured between ages 25 and 65, and *Injured (25–65, Hosp.)* further limits the sample to those hospitalized

within two days after the accident. Each group is set against its matched comparison group (*Non-Injured*) from Section 3.3. For the hospitalized subsample, the comparison group is the subset of matched individuals who were hospitalized in the matched year, which is why it is smaller than the injured group.

We find little difference in pre-accident employment rates between future injured individuals and never-injured individuals, but a substantial gap in earnings. Among working-age individuals, the employment rate is 44.1% for future injured individuals and 43.9% for their never-injured counterparts.[8] Average annual earnings, by contrast, are lower for future injured individuals (TWD 229,110) than for the never-injured group (TWD 248,344). The gap remains among the employed, so it is a difference in pay rather than in employment, consistent with sorting into occupations with more road exposure, such as transportation and manual work.

Health care use also differs before the accident. The injured group has higher outpatient visits (12.3 vs. 10.5), more NHI points used (18,354 vs. 17,250), and similar or higher hospitalization days depending on the subsample. In the hospitalized subsample the comparison group uses far more care even in the year before: 7.7 hospital days against 1.1, 18.7 outpatient visits against 13.7, and more than three times the NHI points. The reason is what put each group in the hospital. The injured were admitted for a traffic injury, which arrives without warning and to people whose prior health is ordinary. Their matches were admitted for whatever else sends people to the hospital, largely chronic and progressive conditions such as cancer, heart disease, and kidney failure, which are preceded by heavy care and followed by more. Being hospitalized in a given year is therefore a much stronger signal of poor health for the comparison group than for the injured, and this column should not be read as a like-for-like health comparison.

The eventually injured are therefore not directly comparable to the never injured, even after matching on basic demographics. Whether a person is ever injured is related to labor-market and health characteristics before the accident, which is why the design in the next

[8]These employment rates are below Taiwan's overall employment rate for two reasons. Our employment measure is based on employer-linked NHI enrollment and covers public-sector workers and other employer-based insurance categories as well as private-sector employees, but excludes the self-employed and military personnel. The sample also includes individuals at the margins of the working-age range.

section does not compare the two.

4 Method

Our goal is to identify the dynamic causal effect of a first hospitalized traffic injury on subsequent labor-market and health outcomes. Motivated by the pre-accident differences documented in Section 3.4, we do not compare the injured with the never-injured. Instead we compare individuals injured in a given year with individuals who will be injured in a later year but have not yet been, cohort by cohort: each accident cohort is followed for several years before and after its accident, and at each point in time its outcomes are set against those of the not-yet-injured.

The identifying assumption is that, absent the accident, the earlier- and later-injured would have followed parallel outcome paths, conditional on the included covariates. This is weaker than assuming that being injured is random. The probability of ever being injured may well depend on occupation, exposure, or risk preferences, and the summary statistics show that it does. What the design requires is that, among people who are eventually injured, the timing of the first injury is not driven by differential trends in their potential outcomes. If the year of the first accident were randomly assigned among the eventually injured, this condition would hold automatically; random timing is sufficient for it but not necessary. Under this quasi-random timing assumption, the later-injured provide the counterfactual for the earlier-injured.

We implement the comparisons with the estimator of Callaway and Sant'Anna (2021). For each accident cohort and each year relative to the accident, we estimate the average treatment effect on the treated with the doubly robust estimator that uses the not-yet-injured as controls, conditioning on birth year to account for life-cycle differences across cohorts, and we then average the cohort-specific estimates at each year since the accident to obtain the dynamic profile. The main analysis uses a balanced window from five years before to five years after the accident, so that every event-time estimate averages over the same cohorts and sample composition does not change across event times; because the age restriction is imposed before the panel is balanced, the estimation sample contains only

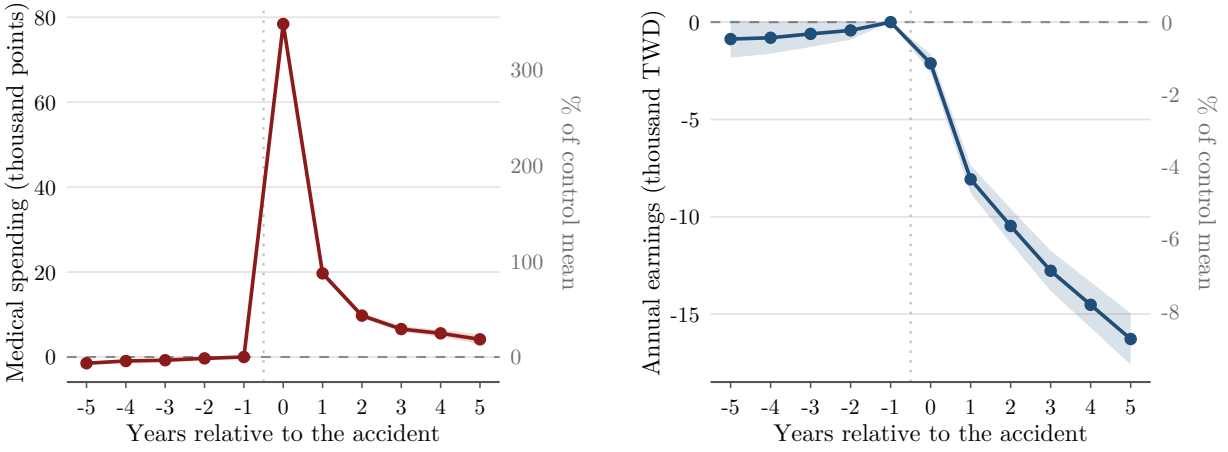
individuals observed between ages 25 and 65 in every sample year. Standard errors are clustered at the individual level. Appendix A states the estimand, the estimator, and the aggregation, and Appendix G reports estimates over longer windows.

Three pieces of evidence speak to the timing assumption. The first is the pre-accident coefficients of the event studies themselves, reported with the results: if earlier- and later-injured workers were on different trajectories before the accident, the estimates for the years before it would not be flat. The second is a direct comparison of pre-accident outcomes. For each accident cohort and its later-treated comparison cohort, measured in the year before the earlier cohort's accident, the standardized differences in earnings, employment, and medical utilization are small, within the conventional threshold of 0.10 for every cohort pair (Appendix A). The never-injured, by contrast, differ systematically from the eventually injured before any accident occurs, in earnings and in medical utilization (Section 3.4 and Appendix A), which is why the design does not use them. The third is a placebo. We reassign accident years at random among the eventually injured, keeping the number of workers per cohort fixed, and re-estimate the event studies on the placebo cohorts. The placebo estimates show none of the dynamics in the actual estimates, so the cohort and calendar-time structure of the design does not generate the results on its own (Appendix B).

5 Results

As the paper aims to measure what a traffic injury costs a surviving worker, and where that cost comes from, this section supplies the building blocks and Section 6 assembles them into a cost per injury and sets it beside existing valuation rules. Three inputs are needed. The first is the time path of the two cost components, medical spending and earnings (Section 5.1). The second is the margin on which the labor-market loss occurs, because a loss that runs through workers leaving full-time work is priced, and remedied, differently from a loss that runs through lower pay among those who keep working (Section 5.2), and whether that loss is a health loss in disguise (Section 5.3). The third is how the loss varies across injured workers, since an average cost per injury hides differences that matter for targeting (Section 5.4). The section closes with mortality (Section 5.5), which enters the accounting separately, and with

the robustness checks (Section 5.6).



(A) NHI-covered medical spending

(B) Earnings

FIGURE 1: Dynamic Effects of Traffic Injuries on Medical Spending and Annual Earnings

Notes: Each panel reports cohort-time estimates using later-treated controls with a balanced event window of $[-5, 5]$ years. Shaded bands show 95% confidence intervals. Estimates are normalized to $t = -1$ as the reference period. The left vertical axis shows the raw estimated values, while the right vertical axis shows the corresponding percentage change relative to the baseline. Earnings are expressed in thousands of 2021 TWD. Medical spending is measured in thousands of NHI points, a standardized reimbursement unit under Taiwan's NHI (≈TWD 0.9 per point).

5.1 Average Dynamic Effects

Figure 1 reports the dynamic effects of a hospitalized injury on the two primary outcomes, NHI-covered medical spending and annual earnings, from the later-treated design with birth-year controls. Earnings are in 2021 TWD. Estimates for the remaining outcomes (inpatient days, outpatient visits, full-time employment, job separation, and attachment to the pre-accident employer) are reported in Appendix E, and estimates that identify accidents from diagnosis codes in the claims rather than from police records are reported in Appendix N.

The medical response is concentrated in the accident year. NHI-covered spending rises by about 80,000 points in the accident year, roughly 360 percent of the control mean, falls sharply in the following year, and remains modestly above baseline through year five. Under universal, low-copayment insurance, this spending measures care received rather than ability to pay. Its composition also changes: acute trauma care accounts for most of the accident-year increase, and follow-up care for long-term complications takes over from the first post-

accident year (Appendix F). Inpatient days spike in the accident year and return close to baseline; outpatient visits stay elevated through year five (Appendix E). Measured by care received, the health shock is largely resolved within a year or two.

The labor-market loss follows a different course: it starts small and grows. The pre-period coefficients are small and statistically indistinguishable from zero, consistent with the checks in Section 4. The earnings effect is modest in the accident year and then widens in every subsequent year; by year five, annual earnings are about 9 percent below the control mean. Full-time employment follows the same path and is also roughly 9 percent lower by year five. The small accident-year effect partly reflects timing: outcomes are measured by calendar year, so a worker injured in December contributes eleven pre-accident months to the accident-year total. Timing cannot explain the years that follow, however: the loss keeps growing long after the accident year, because more and more injured workers are out of full-time work, as the next subsection shows.

For scale, we apply the same design to first all-cause hospital admissions, the event studied by Dobkin et al. (2018). The raw year-five earnings loss after a general admission is about 1.8 times the loss after a traffic injury, but the general-admission event studies show clear pre-trends in earnings and utilization, consistent with anticipated admissions and health that was already deteriorating; the police-defined, unanticipated traffic injury does not (Appendix D).

5.2 Margins of Response: Employment and the Employer Relationship

The full-time employment rate falls in the same way as earnings (Figure OA6a), and both keep falling after medical spending has returned toward baseline. The next question is on which margin the loss occurs and what happens to the injured worker's job. We first separate the earnings loss into an extensive margin, workers who leave full-time employment, and an intensive margin, lower earnings among those who keep working. We then follow the relationship with the pre-accident employer: whether it survives, when it ends, and how earnings evolve for those who leave and those who stay. We close by asking which employment relationships are most likely to break.

5.2.1 Extensive and Intensive Margins

The earnings loss could come from workers leaving full-time work, from lower earnings among those who remain, or from both. To separate the two, we write average earnings as the full-time employment rate times average earnings among the employed. The extensive-margin component is the change in the employment rate multiplied by baseline earnings among the employed; the intensive-margin component is the change in earnings among those still employed, scaled by the baseline employment rate. The intensive component has a selection problem: if the workers who leave are those with the lowest earnings, the workers who remain look better paid than they would otherwise, and a naive comparison understates the intensive-margin loss. We address this with the bounds of Lee (2009). The idea is to make the two groups comparable by dropping, from the comparison group, the share of employed workers that the injury pushed out of employment; trimming the highest earners gives one bound and trimming the lowest gives the other, so the intensive-margin effect is bracketed under the assumption that the injury only moves workers out of employment (Appendix B).

The results show that the extensive margin dominates. After the first post-accident year, the fall in the full-time employment indicator accounts for roughly 80 to 95 percent of the annual earnings loss. The Lee bounds show that the intensive-margin loss can be considerably larger than the naive comparison suggests, but employment exit still explains most of the decline.

5.2.2 The Pre-Accident Employer

The employment rate alone can mask two things. Some workers leave and later return, even though on net more and more are out of full-time work over time, and some of those who remain employed may have changed employer. Because the enrollment records identify each worker's employer, we can look at both. Job separation, an indicator for leaving one's employer, spikes in the accident year and returns close to baseline afterward (Figure OA6b): the disruption is concentrated at the time of the injury rather than spread over a sequence of later layoffs. For some workers, however, the disruption is permanent. Attachment to the pre-accident employer remains persistently lower, and five years after the accident injured

workers are still about 2 percentage points, or 6 to 7 percent of the control mean, less likely to be with that employer (Figure OA6c).

One reason separations do not all happen at once is the wage-continuation protection that covers occupational injuries for up to two years. If protection were merely postponing dismissals, separations should bunch when it expires. They do not: the monthly hazard of leaving the pre-accident employer declines smoothly through month 24 with no jump (Appendix K).

Separations account for most of the earnings decline. Among the employer-attached workers, 36 percent first leave the pre-accident employer within two years of the accident and 64 percent remain. Figure 2 plots the raw mean earnings of the two groups relative to their own pre-accident level; unlike the earlier figures, it is a descriptive trend, not an event-study estimate against a comparison group. By year five, mean earnings are 31 percent below the pre-accident level among separators but only 1.3 percent lower among stayers, and with the two groups' population shares held fixed, separators account for 91 percent of the aggregate year-five decline. Because separation within two years is itself a post-treatment outcome, this is an accounting result rather than a causal subgroup effect, but it locates the loss: it sits almost entirely with workers whose relationship with the pre-accident employer ended.

5.2.3 Which Employment Relationships Break

Which relationships end is informative about why they end. Figure 3 splits the year-five effects by pre-accident tenure and firm size among workers employed before the accident. Full-time employment losses are largest for the shortest-tenured workers and at the smallest firms, the relationships with the least accumulated specific capital and the employers least able to absorb a long absence; in proportional terms, earnings losses are similar across tenure groups. In this respect a hospitalization resembles the displacement events studied in the job-loss literature (Jacobson et al., 1993; Jäger and Heining, 2022; Johnsen et al., 2024), with the difference that the loss here shows up mainly as non-employment rather than as re-employment at lower pay. Section 5.4 takes up heterogeneity across the full set of worker and injury characteristics.

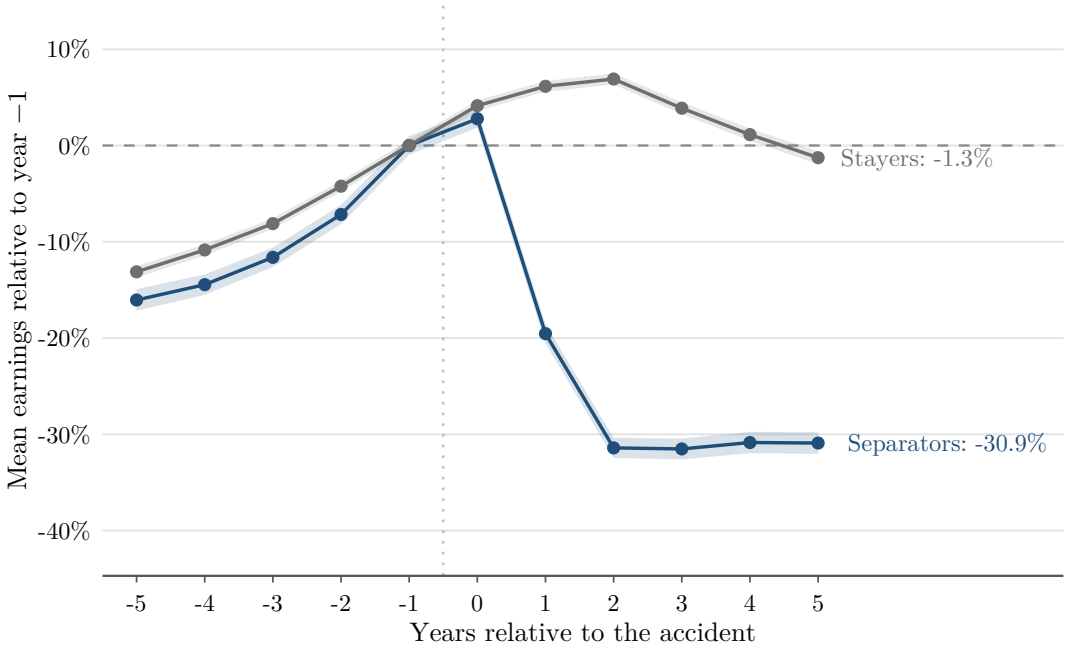


FIGURE 2: Raw Earnings Trajectories by Separation from the Pre-Accident Employer

Notes: Unlike the event-study figures, this figure plots raw means: each series is the mean annual earnings of the group relative to its own mean in event year −1, with no comparison group. Separators first leave the pre-accident employer during event months 0–23; stayers remain attached through that window. Shaded regions are 95% confidence intervals for the means. Separation status is a post-treatment classification, so the paths are descriptive rather than causal subgroup effects.

5.3 Does Lasting Incapacity Explain the Loss?

Section 5.1 showed medical spending settling within a year or two while the labor-market loss kept growing. One possible explanation is that the injury is in fact long-lasting: medical care is concentrated at the start because there is little more that treatment can do, and the workers who leave full-time employment do so because they can no longer work. Under this explanation the labor-market loss is a health loss in disguise. We assess it in two ways, by asking whether the loss persists among workers whose measured health has recovered, and by quantifying how much of the loss is associated with new catastrophic-illness certification.

If lasting incapacity were the binding constraint, the loss should be concentrated among workers whose health fails to recover. We do not find this. Following the corroboration strategy described in Appendix L, we identify injured workers whose total medical spending by the second post-accident year has returned to within the range typical of not-yet-injured

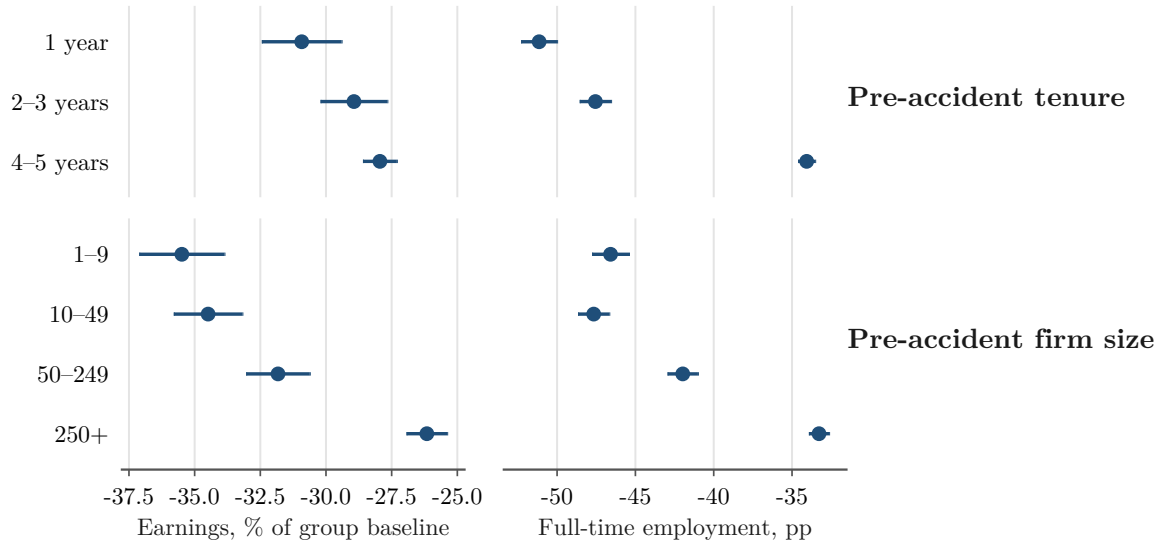


FIGURE 3: Year-Five Labor-Market Effects by Pre-Accident Firm Attachment

Notes: The figure reports event-time-5 effects on earnings (left column, in percent of each subgroup's mean earnings at $t = -1$) and full-time employment (right column, percentage points) by pre-accident firm tenure (top) and pre-accident firm size (bottom), among workers employed in the year before the accident. Horizontal whiskers show 95% confidence intervals.

controls (at or below the control 75th percentile in the same calendar year); about two-thirds of classifiable injured workers meet this criterion. Even within this low-utilization group, labor-market losses continue to widen: the year-five earnings loss is about TWD 13,000 and the full-time employment loss is 2.4 percentage points, retaining 79 and 77 percent of the corresponding main-sample effects. We caution that this split is not a causal decomposition, because post-accident utilization is itself an outcome of the accident. The direction of the resulting selection is informative, however. Workers whose utilization returns to the control range are, if anything, healthier and better placed in the labor market than the average injured worker, so their counterfactual earnings are likely above those of the not-yet-injured controls and the estimated loss in this group is biased toward zero. The retained shares are therefore more likely understated than overstated.

An accounting exercise quantifies the share of the loss associated with new catastrophic-illness certification; Appendix M gives the details and the subgroup event studies. Among the injured, 6.1 percent are newly certified in the accident year or the year after. Their estimated year-five earnings loss is TWD 67,978, compared with TWD 16,280 in the main sample. Ap-

plying the certified share to the subgroup estimate gives TWD 4,144 per injured individual, or 25.5 percent of the aggregate year-five loss. Under the more extreme assumption that every newly certified worker loses all pre-accident earnings, the corresponding bound is 63.9 percent. Either way, most of the economic loss is borne outside this formally recognized severe-illness group, and the divergence between the medical and labor-market trajectories cannot be explained by catastrophic-illness certification alone.

5.4 Heterogeneity in Post-Accident Consequences

The average effects hide large differences across injured workers. This subsection asks whether the economic loss scales with the physical severity of the injury or depends on the worker's position before the accident.

Severity accounts for part of the variation. As Figure OA17 shows, more severe injuries produce larger medical responses and larger labor-market losses, and milder injuries produce smaller effects on both. Injury location tells the same story (Figure OA21): injuries to the head or neck, or to several body regions at once, have the largest and most persistent effects.

Severity does not explain the whole distribution, however. Along several dimensions, groups with similar medical responses have very different labor-market outcomes. Among individuals with positive pre-accident earnings, Figure OA20 shows that higher-income groups lose more earnings and more full-time employment, with the largest effects in the top income quartile, while the health responses of the income quartiles are similar.

The same contrast appears across road users: Figure OA22 and Figure OA23 show that car drivers have especially large earnings and employment losses even though the differences in medical utilization are small. Figure OA24 shows larger labor-market losses among men than among women, again with similar health care utilization and medical spending.

A natural concern is that these subgroup contrasts partly reflect differences in injury-severity composition rather than differences in labor-market exposure. We address this by standardizing every subgroup to the common main-sample severity distribution, using hospitalization-day severity quartiles; the severity composition of each subgroup and the standardized estimates for every dimension are reported in Appendix O (Figure OA19). The labor-market gradients survive standardization. After standardization, the year-five

earnings loss is about TWD 211,000 in the highest positive pre-accident earnings quartile against about TWD 25,000 in the lowest, and the full-time employment losses are 24 and 5 percentage points, while the standardized medical spending responses of the two quartiles are of similar size. Men lose more than women on both labor-market outcomes at the same severity mix, again with similar medical responses, and the same holds across road users and accident types. The heterogeneity in economic losses is therefore not a disguised severity gradient.

Two explanations are consistent with the mechanism documented in Section 5.2, though we cannot test them directly. The first is the job the injury interrupts. Men and car drivers are more often in jobs that depend on driving or on physical capacity, so an injury that a desk worker could work through ends their match, and a high earner has more to lose from any match that ends. The second is the employment relationship. The matches that break after an injury are those with the least accumulated specific capital, at small firms and short tenures, and if men, car drivers, and manual workers are more concentrated in such relationships, the same injury would end more of their matches. We offer these as conjectures; the evidence establishes that the loss varies with the worker's labor-market position at a given severity, not why.

If the labor-market loss were driven by medical severity alone, the groups with the larger earnings losses would also have the larger medical responses, and they do not. What the injury interrupts, the worker's earnings profile, job tasks, and employment relationship, determines how much it costs. A single cost per injury, applied to every survivor, therefore misses most of the variation. The complete set of heterogeneity figures is in Appendix O.

5.5 Mortality and Cumulative Health Hazards

Survival at the time of the accident does not necessarily end the health consequences of injury. Damaged health may increase later mortality, even among individuals who survive the immediate crash. Unlike every estimate so far, the mortality estimates compare the injured with the never-treated population rather than with the later-treated: because death is an absorbing outcome, later-treated controls mechanically exit the risk set. This switch requires justification, since the never-treated differ from the eventually injured in levels, and

the relevant comparison is the counterfactual survival path.

As shown in Figure 4, the mortality gap emerges gradually rather than mechanically in the accident year. The estimate is close to zero at $t = 0$, rises to approximately 0.005 in the first post-accident year, and continues to increase thereafter. By year five, the cumulative increase reaches roughly 0.016, compared with a five-year cumulative mortality of 0.0134 in the comparison group; the excess is therefore about 120 percent of the comparison-group mean. At ten years, the gap reaches approximately 0.03, compared with comparison-group cumulative mortality of 0.0205, an excess of about 146 percent (Appendix G). This delayed and widening trajectory is consistent with continued health deterioration among survivors over the observed horizon. Because the excess mortality keeps rising years after the accident, it is a cost of the injury in its own right, and a large one: valued at the lower bound of the hedonic-wage VSL for Taiwan of Hammitt et al. (2022), TWD 188 million at 2021 prices, the ten-year excess of roughly 3 percentage points is worth about TWD 6 million per injured worker, more than ten times the tangible cost of Section 6.2. The tangible cost does not include this component; Section 6.2 monetizes it only where the comparison with official values requires like totals, with the caveat that a VSL prices marginal changes in risk.



FIGURE 4: Cumulative Mortality after a Traffic Injury, Compared with the Never-Injured

Notes: Unlike the other event-study figures, the comparison group here is the never-injured population rather than the later-injured, because death is an absorbing outcome. The outcome is an indicator that turns on once an individual dies and remains equal to one thereafter, so the coefficients are changes in the cumulative probability of having died by each event time, not contemporaneous mortality rates. The ten-year path is reported in Appendix G.

5.6 Robustness of the Dynamic Patterns

The dynamic patterns are robust to the main specification choices. Extending the balanced window to seven and ten years after the accident shows the earnings and employment losses continuing to widen rather than reversing (Appendix G). Defining the sample by hospitalization within 24 hours of the accident, rather than within 48 hours, ties the hospitalization more tightly to the crash and changes the estimates little (Appendix I). Restricting the comparison group to workers injured within three years of the treated cohort, so that treated and comparison cohorts are closer in calendar time, also leaves the patterns intact (Appendix H). Dropping the birth-year controls changes little (Appendix J). The placebo that reassigns accident years at random among the eventually injured finds no effects on any outcome (Appendix B).

The comparison with a conventional estimator also matters. A pooled two-way fixed effects event study on the same data suggests a partial recovery of earnings, whereas the cohort-by-cohort later-treated design reproduces the deepening losses (Appendix C); the

recovery in the pooled estimator is an artifact of comparing early cohorts with already-treated later ones.

6 Valuing the Cost of a Nonfatal Injury

The preceding sections establish that the persistent cost of a hospitalized traffic injury is primarily a labor-market cost, borne largely by workers who recover medically. This section asks what that finding implies for the price a road-safety appraisal should put on a nonfatal injury. Section 6.1 reviews how agencies price nonfatal injuries now. Section 6.2 puts our estimate on the same lifetime footing and sets it beside those rules applied to Taiwanese inputs. Section 6.3 sorts the cost into what the government and what the injured worker bear and draws the implications for appraisal. Note that the exercise is not a complete welfare valuation. The data do not record pain, suffering, informal care, or out-of-pocket spending, and we put a money value on excess mortality only when we compare our total with official values that include it.

6.1 Existing Valuation Approaches

Two families dominate international road-safety appraisal, and both commonly index the value to the medical severity of the injury (Ashenfelter, 2006). Willingness-to-pay approaches start from a VSL and assign each injury a fraction of it. The U.S. Department of Transportation sets the fractions by Abbreviated Injury Scale level at 0.3, 4.7, 10.5, 26.6, and 59.3 percent of a VSL of USD 9.6 million in 2015 dollars (U.S. Department of Transportation, 2016).[9] Human-capital approaches instead sum medical costs and production losses. The U.S. cost inventory computes both per surviving victim by severity level, using hospital discharge samples for medical costs and assumed work-loss durations for production (Blincoe et al., 2015).[10] Under either approach, the valuation of an injury does not depend on what

[9]European practice is the same with two tiers. The European Commission's handbook on external costs values a serious injury at 13 percent and a slight injury at 1 percent of the VSL, a convention that dates from the HEATCO project; the official values in the United Kingdom, Sweden, Norway, Germany, France, and New Zealand place serious injuries between 10 and 17 percent of the VSL and slight injuries between roughly 0.5 and 3 percent (CE Delft et al., 2019).

[10]The European handbook uses standard values of about EUR 44,000 in lost production and EUR 17,000 in medical costs per serious injury, at 2015 prices, and values production at the capacity to work rather than

happens to the injured person's employment after treatment ends.

Taiwan's official schedule, in 2018 prices, values a fatality at TWD 9.73 million and a nonfatal injury at TWD 0.73 million, and MOTC applies these unit values directly in project appraisal (Chou et al., 2019; Ministry of Transportation and Communications, 2024).[11]

The broader Taiwanese literature produces widely varying estimates. Hedonic-wage VSL estimates span roughly TWD 5–50 million (Liu et al., 1997; Liu and Hammitt, 2004),[12] Earlier compensation-based estimates put the cost per fatality at roughly TWD 2.7–3.0 million (Chou et al., 2019); for nonfatal injuries, WTP surveys imply about TWD 0.6 million per injured individual while accounting estimates from court-verdict data are closer to TWD 1.0 million. The court figure deserves little weight: it averages the small and highly selected set of cases that were litigated, and it says nothing about the typical injury. As in the United States and Europe, none of these approaches observes the post-accident dynamics of the injured (see also Tsai et al., 2015). The variation matters for policy: when shadow prices differ this much across methods, the evaluation of road-safety programs becomes sensitive to the choice of valuation metric, and compensation-based benchmarks risk understating the realized costs borne by survivors, particularly the post-acute labor-market losses documented above.

6.2 Our Estimate Against Existing Valuation Rules

Our estimate adds up the two components the linked records measure, excess NHI spending and lost earnings, both taken from the event-study estimates. Over the accident year and the five years after it, a hospitalized working-age injury costs about TWD 176,000, most of it medical. Through year ten the earnings loss overtakes the medical component and is still growing. Because official values are lifetime constructs, we extend the observed path to the end of a working life, holding the year-ten earnings loss constant to age 65 and discounting

at what the injured person actually earns afterward (CE Delft et al., 2019; Wijnen and Stipdonk, 2016).

[11] IOT combines the two approaches: the fatality value adds the present value of lifetime income lost, by the human-capital method, to a willingness-to-pay value for the nonmarket loss, while the injury value is drawn directly from willingness-to-pay evidence. The injury value happens to equal about 7.5 percent of the fatality value, but it is not derived as a fraction of a VSL, so it is not the Taiwanese counterpart of the U.S. or European severity fractions.

[12] Using an approximate exchange rate of TWD 30 per USD, this range corresponds to about USD 0.17–1.67 million.

at 3 percent. The result is TWD 393,000 per injury, about USD 13,000, one third medical spending and two thirds lost earnings. Appendix C gives the construction, the sensitivity to the discount rate, and the values by severity quartile.

This figure is a cost measured on the injured population itself; current practice imputes it from the medical severity of the injury. To see what the imputation misses, we apply the rules used in the United States and Europe to the same Taiwanese inputs, so that the comparison is between methods applied to the same country and the same injury rather than between countries. Every input the rules need, the earnings base, the medical spending, the excess mortality, and the VSL, is held the same across bars, and each is documented in Appendix C. For institutional context, the IOT schedule discussed above assigns TWD 0.73 million to a nonfatal injury and MOTC applies that value in project appraisal (Chou et al., 2019; Ministry of Transportation and Communications, 2024). We treat it as a separate Taiwanese benchmark rather than classify it as a VSL-based rule, because IOT derives the injury value from direct willingness-to-pay evidence rather than as a fixed share of a fatality value.

The human-capital rules, used in the U.S. cost inventory and as the tangible part of the European handbook, are a schedule. Each injury class is assigned a number of years of lost production, priced at an average wage or at output per head, plus a medical cost taken from discharge data (Blincoe et al., 2015; CE Delft et al., 2019). Their strength is that they are cheap to apply and transparent about what they count. Their assumptions are that the duration of the loss is a property of the injury class, that production is valued at what a person could earn rather than at what the injured person earns afterward, and that nothing that happens after treatment changes the value. Figure 5 shows what these assumptions buy and what they miss. Priced at the injured population's own pre-accident earnings and combined with our observed medical spending, the European and U.S. schedules land close to our twenty-year figure: they assign 1.3 to 1.6 years of production, and we observe 1.4 years of the injured population's earnings. On duration, the rules and the data agree. Priced at employees' earnings instead, the same schedules roughly double, so the earnings base moves the value by more than the choice of rule. And a schedule is blind to what the severity quartiles in the figure show, a cost that varies threefold across injuries in the same

administrative class, and to what Section 5.3 showed, a loss that falls on workers who have recovered medically and arrives years after the bills a severity class is built on.

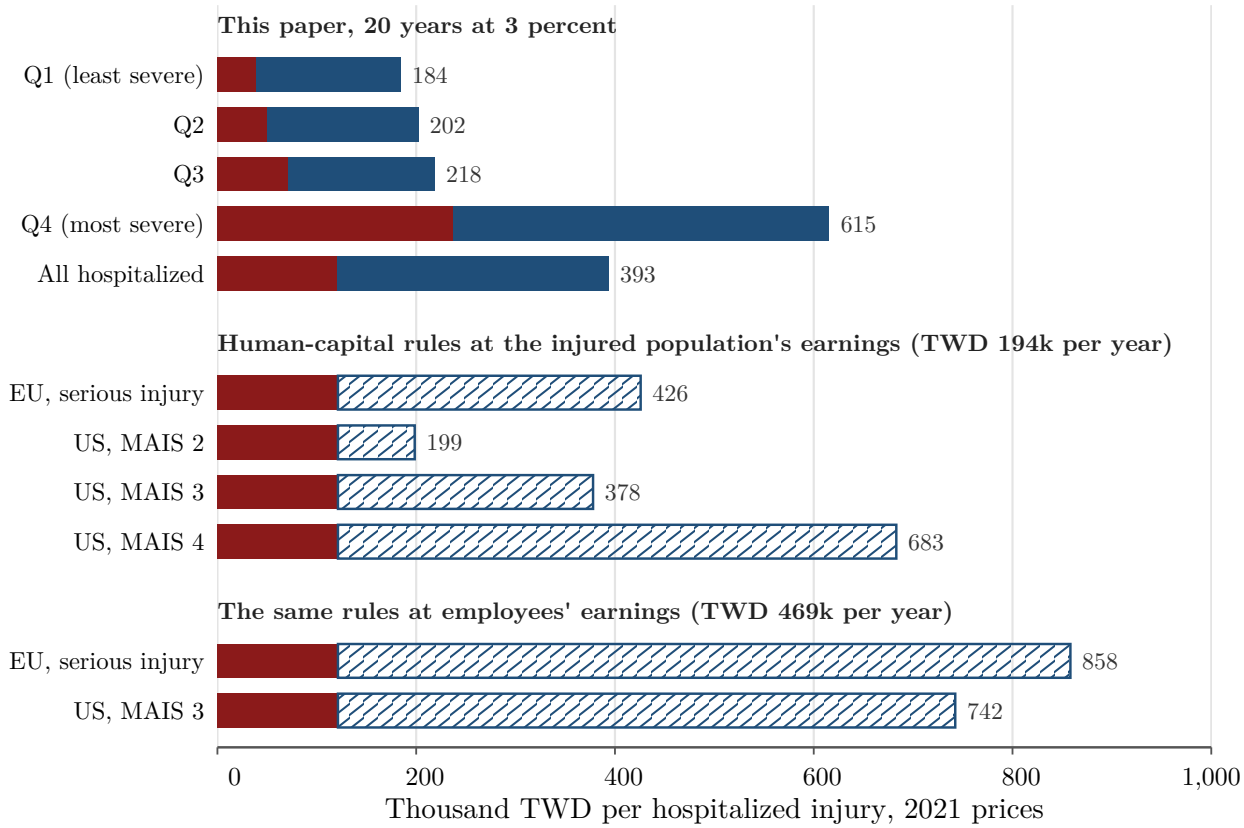


FIGURE 5: Tangible Cost of a Hospitalized Injury: This Paper and Human-Capital Rules on Taiwanese Inputs

Notes: Thousand 2021 TWD per hospitalized working-age injury; medical spending in dark red, lost earnings or production in blue. Solid bars are this paper's estimates over event years 0 to 19 discounted at 3 percent, with the last estimated annual earnings loss held constant through year 19. Hatched bars are the production losses that the European and U.S. human-capital rules impute, expressed as years of one person's annual production (1.57 years for a serious injury in Europe; 0.40, 1.32, and 2.90 years for MAIS 2, 3, and 4 in the United States) and priced at the pre-accident mean earnings of the matched never-injured group, TWD 193,843 including non-workers and TWD 469,042 among employees, plus this paper's observed medical spending. Sources, conversions, and every plotted value are in Appendix C.

The VSL-based rules, which road-safety appraisal in the United States and Europe actually uses, value an injury as a fixed fraction of the VSL (U.S. Department of Transportation, 2016; CE Delft et al., 2019). Their strength is coverage: the fraction is meant to include pain, suffering, lost quality of life, and the risk of later death, none of which an administrative record measures. Their assumptions are that a severity fraction estimated elsewhere

transfers across countries and injuries, that a VSL estimated for small changes in risk can be scaled to an injury, and that the whole value is known at the time of injury. Figure 6 applies the U.S. and European fractions to the only VSL estimated for Taiwan on a national sample of workers, the lower bound of Hammitt et al. (2022) at 2021 prices, and places our estimate beside them with the excess mortality of Section 5.5 valued at the same VSL. The tangible components we measure are about 2 percent of the official values; with mortality included, our total is a quarter to a third of them. The gap is the intangible fraction, and it rests on two numbers that are hard to pin down: the VSL itself, which for Taiwan spans an order of magnitude across studies (Section 6.1), and a severity fraction that no study has estimated for Taiwan. Whether it is the right size for Taiwan is a question about willingness to pay that these data cannot answer, and the mortality term is indicative only, because a VSL prices marginal risks.

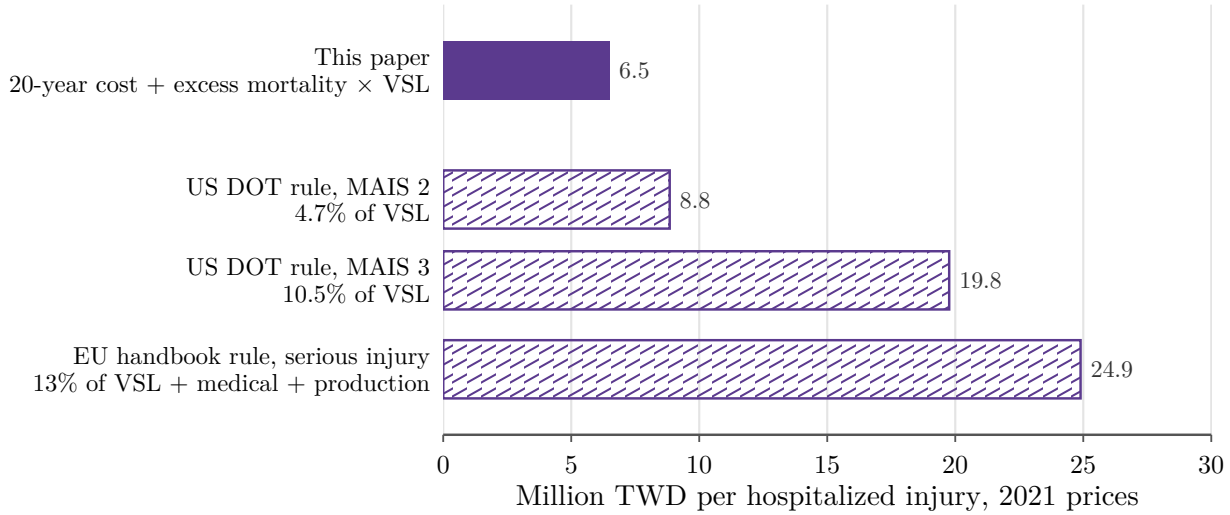


FIGURE 6: Official Values Built on the Taiwan VSL and This Paper's Cost with Mortality

Notes: Million 2021 TWD per hospitalized working-age injury. Hatched bars apply an external rule; the solid bar is this paper's estimate. All bars use the lower bound of the Taiwan VSL of Hammitt et al. (2022), TWD 188 million at 2021 prices with no adjustment for income growth since 1997. The U.S. rows multiply the VSL by the Department of Transportation fractions of 4.7 and 10.5 percent (U.S. Department of Transportation, 2016). The European row adds 13 percent of the VSL to the medical and production components of the European row of Figure 5 (CE Delft et al., 2019). This paper's bar adds the twenty-year tangible cost to the excess mortality of 3.2 percentage points by year 10 valued at the VSL, an indicative term because a VSL prices marginal changes in risk. Construction and sources are in Appendix C.

The comparison therefore does not show that the official values are too high. It shows

that the part of the cost that can be measured is the part those values pay least attention to, that the human-capital rules get its size roughly right on average, and that both families of rules, by fixing the value at the time of injury, cannot see its timing, its incidence, or its variation within a severity class.

6.3 Who Bears the Cost, and What It Implies for Appraisal

The components map onto the marginal value of public funds (MVPF), the framework now standard for evaluating public spending, which divides the beneficiaries' willingness to pay for a policy by its net cost to the government (Hendren and Sprung-Keyser, 2020; Finkelstein and Hendren, 2020). Preventing a hospitalized injury saves the government the NHI spending and the social-insurance contributions on the lost earnings.

How much is illustrated by 2021 statutory rates. Employees and employers together pay 90 percent of the 11.5 percent labor- and employment-insurance premium and 90 percent of the 5.17 percent NHI premium, a combined rate of $0.9 \times (11.5 + 5.17) = 15.003$ percent (Bureau of Labor Insurance, 2020; National Health Insurance Administration, 2021); we set the income-tax component to zero rather than impute household tax liabilities. Applying 15 percent to the TWD 272,000 present value of lost earnings and adding TWD 121,000 in excess NHI spending gives about TWD 160,000 per injury avoided. The calculation is illustrative. It holds the 2021 rates fixed over the sample period and ignores insured-salary ceilings, dependent coverage, the change in NHI enrollment category after job loss, and household income-tax liabilities.

Any measure that prevents a hospitalized injury for less than TWD 160,000 pays for itself in fiscal terms, and its MVPF is unbounded in the sense of Hendren and Sprung-Keyser (2020). Measures that cost more are judged on the private side: the after-tax earnings loss, which our estimate supplies as a floor, and the mortality risk and intangible losses above it. Because most of the private loss accrues years after the accident and to workers who have recovered medically, the willingness to pay that matters here is not the one a severity-weighted valuation assigns at the time of injury. We do not compute an MVPF for any program, since that requires the program's cost per injury avoided; the contribution is the per-injury inputs and their timing.

For infrastructure appraisal the implication is that a fatality-centered evaluation misses two margins: the nonfatal injuries avoided and the reduction in severity among crashes that still occur. A benefit calculation that includes them needs a cost per nonfatal injury and a value for shifting an injury to a lower severity class, and our estimates supply both, the twenty-year cost for the first and the quartile spread in Figure 5 for the second. In Taiwan, hospitalized nonfatal injuries outnumber traffic fatalities by roughly eight to one, so even a per-injury cost far below the VSL materially alters the benefit side, and it does so most for interventions whose main effect is to make injuries milder rather than to prevent crashes: speed management, intersection design, protected facilities for vulnerable road users, and post-crash response. The heterogeneity documented above adds a distributional margin. Injuries with similar health impacts carry very different labor-market costs, so the return to prevention depends on who is exposed, and reducing employment detachment among vulnerable workers carries a social-insurance value that the average earnings figure understates.

7 Conclusion

Acute health shocks destroy employment relationships, not just health. This paper documents that process for the most common acute health shock to working-age adults, linking 24 years of police-reported traffic accidents to population health, employment, and mortality records and estimating the dynamic effects of a hospitalized injury with a staggered event-study design that compares earlier-injured workers to later-injured ones.

Medical care is concentrated in the accident year and largely settles within a year or two. The labor-market loss starts small and grows: by year five, annual earnings and full-time employment are about 9 percent below their counterfactual levels, and the loss keeps widening through year ten. It runs through the extensive margin, through workers who leave full-time work and, most often, the employer they had before the accident. It persists among workers whose medical utilization has returned to the range of uninjured controls, while workers newly certified for catastrophic illness account for 25.5 percent of the aggregate year-five earnings loss under the accounting attribution. Its size depends on the job the injury interrupts more than on the injury itself: at the same severity, high-income workers, men,

and car drivers lose far more. Survivors are not healthy on every margin, since cumulative mortality keeps rising years after the accident.

For valuation, the estimates put a measured cost on the survivor: about TWD 393,000 (USD 13,000) in medical spending and lost earnings over a working life, discounted at 3 percent. Applied to Taiwanese inputs, the human-capital rules used in cost inventories schedule a production loss of about the same duration, so on average they are not wrong. What they and the VSL-based rules used in appraisal cannot see is that the loss arrives years after the medical bills, falls on workers who have recovered, and varies threefold within a severity class. The official values are an order of magnitude larger because of the intangible fraction they add, so the part of the cost that can be measured is the part those values attend to least. The measure excludes pain and suffering, informal care, and private costs outside the administrative data, and it is not a lower bound on a welfare loss.

For public finance and transportation policy, the split of the cost matters as much as its size. Preventing a hospitalized injury saves the government its medical spending and the taxes and contributions on the lost earnings, enough that inexpensive measures pay for themselves in fiscal terms, while the larger private loss accrues years later. Because hospitalized nonfatal injuries outnumber traffic fatalities by roughly eight to one in our setting, assessments that count only fatalities miss a large and persistent source of social cost, most of all for investments that make injuries milder or keep injured workers attached to their employers.

A Estimation Details

This appendix states the estimand and estimator behind the event-study results, following Callaway and Sant'Anna (2021).

Setup. Let G_i denote the calendar year of individual i's first hospitalized traffic injury, and let $Y_{i,t}(g)$ denote the outcome in year t if the first injury occurs in year g, with $Y_{i,t}(0)$ the outcome absent injury through year t. Write $G_{i,g} = \mathbb{1}\{G_i = g\}$ and $D_{i,t} = \mathbb{1}\{G_i \leq t\}$ for having been injured by year t. The group-time average treatment effect is

$$\mathrm{ATT}(g,t) = \mathbb{E}\left[Y_t(g) - Y_t(0) \,|\, G = g\right].$$

Identification. For cohort g at a post-treatment time $t \geq g$, the comparison group consists of individuals not yet injured by t ($G_i > t$), who will experience their first injury in a later year. Let X denote the covariates, in the main specification birth year. Under conditional parallel trends with respect to the not-yet-injured, $\mathbb{E}[Y_t(0) - Y_{g-1}(0) \mid X, G_g = 1] = \mathbb{E}[Y_t(0) - Y_{g-1}(0) \mid X, D_t = 0, G_g = 0]$, together with no anticipation and overlap, $\mathrm{ATT}(g,t)$ is identified for $t \geq g$ with $g-1$ as the base period.

Estimator. Let $p_{g,t}(X) = \Pr(G_g = 1 \mid X,\, G_g + (1-D_t)(1-G_g) = 1)$ be the generalized propensity score of belonging to cohort g rather than to the not-yet-injured comparison group at t, and let $m_{g,t}(X) = \mathbb{E}[Y_t - Y_{g-1} \mid X, D_t = 0, G_g = 0]$ be the outcome-change regression among the comparison group. The doubly robust estimand, Equation (2.2) of Callaway and Sant'Anna (2021), is

$$\mathrm{ATT}(g,t) = \mathbb{E}\left[\left(\frac{G_g}{\mathbb{E}[G_g]} - \frac{\dfrac{p_{g,t}(X)(1-D_t)(1-G_g)}{1-p_{g,t}(X)}}{\mathbb{E}\left[\dfrac{p_{g,t}(X)(1-D_t)(1-G_g)}{1-p_{g,t}(X)}\right]}\right)\left(Y_t - Y_{g-1} - m_{g,t}(X)\right)\right],$$

which is consistent if either the propensity score or the outcome regression is correctly specified. We estimate $p_{g,t}$ by logit and $m_{g,t}$ by least squares, both in a centered birth-year

polynomial. The implementation uses the universal base period $g-1$. Thus, for a pre-treatment event time $t < g-1$, the placebo coefficient compares $Y_t - Y_{g-1}$ between cohort g and cohorts injured after g; it is not an adjacent-period contrast. Event time -1 is the normalized reference period.

Aggregation. The dynamic profile at event time $e = t - g$ is

$$\theta(e) = \sum_g \mathbb{1}\{g + e \leq T\} \Pr(G = g \mid g + e \leq T) \, \mathrm{ATT}(g, g + e),$$

where T is the last year of the panel. Setting `balanced_event_time=5` restricts the dynamic aggregation to treated cohorts observable through event time $e = 5$, so estimates from $e = 0$ through $e = 5$ use the same treated-cohort composition. It does not further restrict the comparison group to cohorts injured at least five years after cohort g. For each group-time contrast, controls must remain untreated through both the target and base periods, so with zero anticipation they satisfy $G_i > \max\{t, g-1\}$; in particular, post-treatment controls satisfy $G_i > t$. Inference uses the individual-level influence function of the aggregated estimator, which accounts for dependence across a person's annual observations.

B Extensive and Intensive Margin Decomposition

This appendix decomposes the labor-market earnings effects into extensive and intensive margins. We define employment as holding a full-time job, corresponding to the full-time employment indicator equal to one. The baseline period is the year before the accident, normalized to event time $t = -1$. For the main hospitalized working-age sample, baseline earnings among employed workers are TWD 423,592 and the baseline full-time employment rate is 0.359.

The decomposition is based on the accounting identity

$$\Delta E[Y_i] = E[Y_i \mid emp_i = 1]_0 \cdot \Delta \Pr(emp_i = 1) + \Pr(emp_i = 1)_0 \cdot \Delta E[Y_i \mid emp_i = 1].$$

The first term is the extensive-margin component: the earnings loss mechanically implied by changes in full-time employment. The second term is the intensive-margin component: the earnings change among those who remain employed, scaled by the baseline employment rate.

Simple Accounting Decomposition

Table A1 reports the simple accounting decomposition using the main event-study estimates with birth-year controls. The total earnings effect is the event-study estimate for annual earnings including zeros. The employment effect is the event-study estimate for full-time employment, and the conditional-earnings effect is estimated among employed workers.

TABLE A1: Simple Extensive and Intensive Margin Decomposition

Event time	Total ATT	Extensive	Intensive	Residual
0	-2,119	-3,267	-494	1,643
1	-8,075	-7,774	-1,567	1,266
2	-10,472	-8,866	-750	-856
3	-12,774	-10,430	-776	-1,569
4	-14,524	-11,369	-1,156	-1,999
5	-16,280	-13,066	-831	-2,384

Notes: The extensive component equals baseline employed earnings multiplied by the full-time employment effect. The intensive component equals the baseline full-time employment rate multiplied by the conditional-earnings effect among employed workers. The residual is the difference between the total earnings effect and the sum of the two components. Values are in TWD.

The main pattern is clear: the extensive margin accounts for most of the earnings loss. By five years after the accident, the total annual earnings effect is about TWD 16,300, while the extensive component alone is about TWD 13,100. The intensive component is negative but much smaller, below TWD 1,000 in absolute value at year five. The residual reflects that the decomposition is applied to aggregated event-time estimates rather than to an exact individual-level accounting identity within each cohort-time cell.

Lee Bounds for Selection into Employment

The simple conditional-earnings comparison may be biased because post-accident employment is itself affected by treatment. If the most severely injured workers are more likely to leave employment, then earnings among the remaining employed workers are positively selected. To address this concern, we implement Lee bounds for the intensive-margin effect. Within each cohort-time cell, we compare treated and later-treated controls after trimming the control group's conditional earnings distribution to equalize employment rates. We then aggregate the bounds across event times.

Table A2 reports the implied bounds for the intensive contribution and the corresponding accounting total. These results use the non-balanced event-time Lee bounds.

TABLE A2: Lee-Bounds Decomposition

Event time	Extensive	Intensive lower	Intensive upper	Total bounds
0	-3,267	-5,600	-2,708	[-8,867, -5,976]
1	-7,774	-8,283	2,728	[-16,058, -5,047]
2	-8,866	-8,927	1,989	[-17,793, -6,877]
3	-10,430	-9,681	-352	[-20,111, -10,782]
4	-11,369	-10,094	-5,137	[-21,464, -16,506]
5	-13,066	-10,813	-7,952	[-23,879, -21,018]

Notes: Intensive bounds equal the Lee-bound conditional-earnings effects multiplied by the baseline full-time employment rate. Total bounds add the extensive component to the intensive bounds. Values are in TWD.

The Lee bounds suggest that the naive employed-only comparison likely understates intensive-margin losses. Once selection into employment is allowed for, the intensive component can be substantially more negative. Nevertheless, the extensive margin remains quantitatively central: employment exit alone explains a large share of the observed earnings decline, especially in later post-accident years.

Pre-Accident Employed Subsample

As a complementary exercise, we restrict the sample to individuals who were employed in the year before the accident and re-estimate the event-study effects for subsequent earnings and full-time employment. This exercise holds the baseline workforce fixed and is closer to the approach commonly used in the job-displacement literature.

TABLE A3: Pre-Accident Employed Subsample

Event time	Earnings ATT	Employment ATT	Exit contribution
0	-27,231	-0.140	-59,374
1	-60,593	-0.220	-93,106
2	-83,379	-0.273	-115,521
3	-104,852	-0.320	-135,618
4	-124,057	-0.361	-152,957
5	-142,174	-0.404	-170,921

Notes: The sample is restricted to individuals employed in the year before the accident. The exit contribution equals baseline employed earnings multiplied by the full-time employment effect. Earnings and exit contributions are in TWD.

Among the pre-accident employed, the post-accident decline in full-time employment is

much larger. By year five, full-time employment falls by about 40 percentage points, implying a mechanical exit-related earnings loss of roughly TWD 171,000. This exit contribution is larger than the observed earnings loss, indicating that some workers retain partial earnings outside full-time employment or that lower-intensity work and compositional changes partly offset the mechanical full-time exit calculation. The exercise therefore reinforces the central interpretation: persistent earnings losses are closely tied to labor-market detachment after injury.

C Construction of the Valuation Benchmarks

This appendix documents every input behind Figure 5 and Figure 6 in Section 6.2. All values are in 2021 Taiwan dollars. Table A4 lists the constants and their sources. The calculation is reproduced by a single script in the replication package, which also writes a table of every plotted value.

This paper's twenty-year cost. The event-study estimates give an annual earnings loss and an annual NHI medical cost for each event year, with birth-year controls, through year 10 for the full hospitalized working-age sample and through year 5 for the severity quartiles of Appendix O. Medical spending converts NHI points at TWD 0.9 per point. To put the cost on the lifetime footing that the external rules use, we hold the last estimated annual earnings loss constant through event year 19, the year in which a worker of the mean age at injury (46.1, Table 1) reaches 65, assume no further medical spending after the last estimated year, and discount at 3 percent, the rate used in the U.S. cost inventory (Blincoe et al., 2015). For the full sample this gives TWD 121,000 of medical spending and TWD 272,000 of lost earnings, TWD 393,000 in total, the twenty-year cost of Section 6.2. For reference, the observed sums are TWD 176,000 through year 5 (TWD 112,000 medical and TWD 64,000 earnings) and TWD 291,000 through year 10 (TWD 128,000 and TWD 163,000), and the year-ten annual earnings loss held constant in the projection is TWD 23,000. Undiscounted, the twenty-year total is TWD 502,000; at 5 and 7 percent it is TWD 341,000 and TWD 301,000. The quartile rows apply the same rule from year 5; because the quartile estimates stop earlier, their projections carry more of the extrapolation. The measure includes formal earnings and public medical spending only. It excludes nonmarket losses, informal care, out-of-pocket spending, and pain and suffering, and it is not a lower bound on a welfare loss.

Human-capital rules as years of production. Both external human-capital rules fix a production loss per injury class in the currency and price level of a particular year. We convert each to a unit that can be transferred across countries without a price adjustment: the number of years of one person's annual output that the rule assigns to the injury, obtained by dividing the published production loss by output per head in the same country and year.

The European standard value is the SafetyCube estimate of EUR 43,627 of lost production per serious (hospitalized) injury at 2015 prices (Wijnen et al., 2017, Table 6-2), which the European Commission handbook adopts as its default for member states (CE Delft et al., 2019); the handbook values production at the injured person's capacity to work rather than at realized earnings. European Union output per head in 2015 was USD 30,809 (International Monetary Fund, 2026), or EUR 27,766 at the 2015 annual average rate of USD 1.1096 per euro (Board of Governors of the Federal Reserve System, 2022), so the rule assigns 1.571 years. The U.S. cost inventory reports market-productivity losses per surviving victim by maximum AIS level, discounted at 3 percent in 2010 dollars: USD 19,359 for MAIS 2, 64,338 for MAIS 3, and 140,816 for MAIS 4 (Blincoe et al., 2015, Table 2-3). U.S. output per head in 2010 was USD 48,586 (International Monetary Fund, 2026), so the rule assigns 0.398, 1.324, and 2.898 years. Our hospitalized sample spans MAIS levels; MAIS 3 (serious, not life-threatening, typically requiring admission) is the closest single level, and MAIS 2 and MAIS 4 bracket it. The years are then priced at the earnings the injured population would have had: the mean earnings in the year before the accident of the matched never-injured group in the hospitalized working-age sample, TWD 193,843 including non-workers (the same base as our earnings estimates, which are also averages over workers and non-workers) and TWD 469,042 among employees (Table 1). Each rule row adds this paper's observed medical spending, TWD 121,000, rather than the rules' own medical values, so that the rows differ only in the production component.

The VSL. The only VSL estimated for Taiwan on a national sample of workers is the hedonic-wage estimate of Hammitt et al. (2022), whose Table 4 reports, for 1997, a range of USD 4.806 to 9.350 million in 1991 dollars depending on the fatality-risk measure. We use the lower bound. Following the authors, Taiwan dollars convert at the 1991 rate of TWD 25.75 per U.S. dollar, giving TWD 123.8 million at 1991 prices; the Taiwan consumer price index rose by a factor of 1.521 between 1991 and 2021 (chained from the annual inflation series in International Monetary Fund, 2026), giving TWD 188.2 million at 2021 prices. We make no adjustment for real income growth between 1997 and 2021. With the income elasticity of the VSL of 0.65 to 0.91 that the same paper estimates from its 1982 to 1997 panel, such an

adjustment would raise the VSL substantially, and with it every bar in Figure 6, including the mortality term in ours; the ordering of the bars is unaffected.

The VSL-based rules. The U.S. Department of Transportation values a nonfatal injury at a fixed fraction of the VSL by maximum AIS level: 0.3, 4.7, 10.5, 26.6, and 59.3 percent for MAIS 1 to 5 (U.S. Department of Transportation, 2016). The fraction is the whole value; no medical or production component is added. Applied to the Taiwan VSL, MAIS 2 gives TWD 8.8 million and MAIS 3 TWD 19.8 million. The European handbook values the human cost of a serious injury at 13 percent of the VSL, a convention inherited from the HEATCO project, and adds the medical, production, and other tangible components to it (CE Delft et al., 2019; Wijnen et al., 2017). We add the medical and production components of the European row of Figure 5 (TWD 426,000) and omit property damage and administrative costs, which our data do not measure and which the U.S. rule does not include either; the row is TWD 24.9 million.

This paper with mortality. Our bar in Figure 6 is the twenty-year tangible cost plus the cumulative excess mortality by event year 10 from Section 5.5, 3.23 percentage points in the never-treated comparison, multiplied by the VSL: TWD 0.39 million plus TWD 6.08 million, or TWD 6.5 million. A VSL prices small changes in mortality risk, and a 3.23 percentage point excess over ten years in an already-injured population is not small, so the mortality term is indicative rather than a welfare measure. We include it because the severity fraction in the official rules is meant to cover the risk of later death along with pain, suffering, and lost quality of life, and a comparison that omitted it would set unlike totals against each other.

TABLE A4: Inputs to the Valuation Benchmarks

Input	Value	Source
This paper		
Discount rate	3%	Blincoe et al. (2015)
Projection horizon	event years 0–19	mean age at injury 46.1
NHI points to TWD	0.9	Section 6
Pre-accident mean earnings, matched never-injured, all	TWD 193,843	Table 1
Pre-accident mean earnings, matched never-injured, employees	TWD 469,042	Table 1
Cumulative excess mortality by year 10	0.0323	Section 5.5
European human-capital rule		
Production loss per serious injury, 2015 prices	EUR 43,627	Wijnen et al. (2017, Table 6-2); CE Delft et al. (2019)
EU output per head, 2015	USD 30,809	International Monetary Fund (2026), NGDPDPC
USD per EUR, 2015 annual average	1.1096	Board of Governors of the Federal Reserve System (2022)
Implied years of production	1.571	
Human-cost share of VSL, serious injury	13%	CE Delft et al. (2019)
U.S. human-capital rule		
Market-productivity loss per surviving victim, 3%, 2010 dollars, MAIS 2/3/4	USD 19,359; 64,338; 140,816	Blincoe et al. (2015, Table 2-3)
U.S. output per head, 2010	USD 48,586	International Monetary Fund (2026), NGDPDPC
Implied years of production, MAIS 2/3/4	0.398 / 1.324 / 2.898	
VSL fractions, MAIS 2/3	4.7% / 10.5%	U.S. Department of Transportation (2016)
Taiwan VSL		
VSL, 1997, lower bound, 1991 dollars	USD 4.806 million	Hammitt et al. (2022, Table 4)
TWD per USD, 1991	25.75	Hammitt et al. (2022)
Taiwan CPI, 2021 relative to 1991	1.521	International Monetary Fund (2026), PCPIPCH
VSL at 2021 prices	TWD 188.2 million	

Notes: Output per head is GDP per capita in current U.S. dollars (IMF series NGDPDPC). The CPI factor chains the annual average consumer-price inflation rates for Taiwan (IMF series PCPIPCH) from 1992 through 2021. No adjustment is made for real income growth in the VSL.

References

Ashenfelter, O. (2006). Measuring the value of a statistical life: Problems and prospects. *The Economic Journal 116*(510), C10–C23.

Barr, A., J. Eggleston, S. Mello, and A. Smith (2025, October). The big economic consequences of small financial shocks: Evidence from minor car crashes. Working paper, Texas A&M University, U.S. Census Bureau, Dartmouth College, and West Point.

Bíró, A., M. Bisztray, J. Galindo da Fonseca, and T. Molnár (2025). Accident-induced absence from work and wage growth. *Journal of Labor Economics*. Forthcoming.

Blincoe, L. J., T. R. Miller, E. Zaloshnja, and B. A. Lawrence (2015). The economic and societal impact of motor vehicle crashes, 2010. Technical Report DOT HS 812 013, National Highway Traffic Safety Administration.

Board of Governors of the Federal Reserve System (2022). Foreign exchange rates: G.5A annual. `https://www.federalreserve.gov/releases/g5a/`. Annual averages of daily noon buying rates; accessed September 2026.

Bougna, T., G. Hundal, and P. Taniform (2022). Quantitative analysis of the social costs of road traffic crashes literature. *Accident Analysis and Prevention 165*, 106282.

Bureau of Labor Insurance (2020). Labor insurance premium rate adjusted to 11.5 percent from january 1, 2021 [自 110 年 1 月 1 日起，勞工保險普通事故保險費調高為 11.5%]. Ministry of Labor, Taiwan, `https://www.bli.gov.tw/0104915.html`. The rate includes the 1 percent employment-insurance premium; accessed 2026-09-18.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

CE Delft, INFRAS, and Fraunhofer ISI (2019). Handbook on the external costs of transport: Version 2019. Technical report, European Commission, Directorate-General for Mobility and Transport (DG MOVE), Luxembourg. Version 2019 update of the 2014 Handbook on the External Costs of Transport.

Chou, J.-C., Y.-C. Chiu, C.-W. Kuo, M.-C. Wang, T.-Y. Chen, C.-W. Hsieh, K.-K. Chang, T.-H. Yeh, and K.-P. Chen (2019). 道路交通事故成本推估之研究. Technical report, Institute of Transportation, MOTC.

Dano, A. M. (2005). Road injuries and long-run effects on income and employment. *Health Economics 14*(9), 955–970.

Deshpande, M. and L. M. Lockwood (2022). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica 90*(4), 1781–1810.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2018). The economic consequences of hospital admissions. *American Economic Review 108*(2), 308–352.

Elvik, R. (2000). How much do road accidents cost the national economy? *Accident; Analysis and Prevention 32*(6), 849–851.

Financial Supervisory Commission, Taiwan (2022a). Compulsory automobile liability insurance act (強制汽車責任保險法), articles 7 and 27. `https://law.moj.gov.tw/ENG/LawClass/LawAll.aspx?pcode=G0390060`. As amended June 15, 2022. Accessed September 5, 2026.

Financial Supervisory Commission, Taiwan (2022b). Regulations for underwriting and claim adjustment procedures under compulsory automobile liability insurance, articles 15 and 17. `https://law.moj.gov.tw/ENG/LawClass/LawAll.aspx?pcode=G0390062`. As amended November 25, 2022. Accessed September 5, 2026.

Financial Supervisory Commission, Taiwan (2026). Payment standards for compulsory automobile liability insurance (強制汽車責任保險給付標準), articles 2, 3, and 6. `https://law.moj.gov.tw/LawClass/LawAll.aspx?pcode=G0390067`. As amended May 29, 2026. Accessed September 5, 2026.

Finkelstein, A. and N. Hendren (2020). Welfare analysis meets causal inference. *Journal of Economic Perspectives 34*(4), 146–167.

Finkelstein, E. A., P. S. Corso, and T. R. Miller (2006, 04). Lifetime productivity losses due to injuries. In *The Incidence and Economic Burden of Injuries in the United States*. Oxford University Press.

García-Gómez, P., H. van Kippersluis, O. O'Donnell, and E. van Doorslaer (2013). Long-term and spillover effects of health shocks on employment and income. *Journal of Human Resources 48*(4), 873–909.

Halla, M. and M. Zweimüller (2013). The effect of health on earnings: Quasi-experimental evidence from commuting accidents. *Labour Economics 24*, 23–38.

Hammitt, J. K., J.-T. Liu, and J.-L. Liu (2022). Is survival a luxury good? income elasticity of the value per statistical life. *Journal of Risk and Uncertainty 65*(3), 239–260.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *Quarterly Journal of Economics 135*(3), 1209–1318.

International Monetary Fund (2026). World economic outlook database. `https://www.imf.org/en/Publications/WEO/weo-database`. Series NGDPDPC (GDP per capita, current U.S. dollars), PPPPC, and PCPIPCH (inflation, average consumer prices); accessed September 2026.

Jacobson, L. S., R. J. LaLonde, and D. G. Sullivan (1993). Earnings losses of displaced workers. *American Economic Review 83*(4), 685–709.

Jäger, S. and J. Heining (2022). How substitutable are workers? evidence from worker deaths. Working Paper 30629, National Bureau of Economic Research.

Johnsen, J. V., H. Ku, and K. G. Salvanes (2024). Competition and career advancement. *Review of Economic Studies 91*(5), 2954–2980.

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies 76*(3), 1071–1102.

Liu, J.-T. and J. K. Hammitt (2004). Effects of disease type and latency on the value of mortality risk. *Journal of Risk and Uncertainty 28*(1), 73–95.

Liu, J.-T., J. K. Hammitt, and J.-L. Liu (1997). Estimated hedonic wage function and value of life in a developing country. *Economics Letters 57*(3), 353–358.

Ministry of Health and Welfare, Department of Statistics (2026). Table 1.1.7: The disabled population by cause and classification (身心障礙者人數按障礙成因及類別分). https://dep.mohw.gov.tw/dos/cp-5224-62359-113.html. Data updated March 10, 2026. Accessed September 5, 2026.

Ministry of Labor, Taiwan (2024). Labor standards act, articles 13 and 59. https://laws.mol.gov.tw/FLAW/FLAWDAT0201.aspx?id=FL014930. Accessed May 31, 2026.

Ministry of Transportation and Communications (2024, July). 智慧運輸系統發展建設計畫（114–117 年）[intelligent transportation systems development and construction plan, 2025–2028]. Technical report, Ministry of Transportation and Communications, Taiwan. Approved plan; accessed 2026-09-18.

MOHW (2022). 111 年度全民健康保險財務評估報告 [national health insurance annual financial evaluation report, 2022]. Technical report, National Health Insurance Administration, Ministry of Health and Welfare, Taipei, Taiwan.

National Health Insurance Administration (2021). Premium calculation formulas for categories 1 to 3 [第 1 類到第 3 類保險費計算公式及釋例]. Ministry of Health and Welfare, Taiwan, https://www.nhi.gov.tw/ch/cp-3277-6c895-2588-1.html. Documents the 5.17 percent premium rate and contribution shares; accessed 2026-09-18.

National Health Insurance Administration, Ministry of Health and Welfare (2023). 2022–2023 national health insurance in taiwan annual report. Technical report, National Health Insurance Administration. Accessed September 5, 2026.

National Police Agency (2024). Police statistical yearbook, 2023 [112 年警政統計年報]. Ministry of the Interior, Taipei, https://www.npa.gov.tw/static/ebook/Y112/mobile/index.html. Road traffic accident statistics; accessed 2026-09-18.

National Police Agency of Japan (2024). Traffic accidents in 2023 [令和 5 年中の交通事故の

発生状況]. https://www.npa.go.jp/publications/statistics/koutsuu/toukeihyo_e.html. Accessed 2026-09-13.

OECD International Transport Forum (2023). Global transport infrastructure investment data update. OECD.

Parro, F. and R. V. Pohl (2021). The effect of accidents on labor market outcomes: Evidence from chile. *Health Economics 30*, 1787–1803.

Social Security Administration (2026). 2026 cost-of-living adjustment (cola) fact sheet. https://www.ssa.gov/news/en/cola/factsheets/2026.html. Accessed May 31, 2026.

Statistisches Bundesamt (Destatis) (2024). Road traffic accidents: Long term series. https://www.destatis.de/EN/Themes/Society-Environment/Traffic-Accidents/Tables/liste-traffic-accidents.html. Accessed 2026-09-13.

Tsai, J.-F., C.-P. Chu, and S.-R. Hu (2015). Road pricing for congestion and accident externalities for mixed traffic of motorcycles and automobiles. *Transportation Research Part A: Policy and Practice 71*, 153–166.

U.S. Department of Labor (2025). Fact sheet #28: The family and medical leave act. https://www.dol.gov/agencies/whd/fact-sheets/28-fmla. Revised March 2025; accessed May 31, 2026.

U.S. Department of Transportation (2016, August). Guidance on treatment of the economic value of a statistical life (VSL) in U.S. Department of Transportation analyses: 2016 adjustment. Technical report, Office of the Secretary of Transportation. Memorandum dated August 8, 2016. Injury values expressed as fractions of the VSL by Abbreviated Injury Scale (AIS) level: 0.003, 0.047, 0.105, 0.266, 0.593 for AIS 1–5.

WHO (2023). Global status report on road safety 2023. Technical report, World Health Organization, Geneva.

Wijnen, W. and H. Stipdonk (2016). Social costs of road crashes: An international analysis. *Accident Analysis and Prevention 94*, 97–106.

Wijnen, W., W. Weijermars, W. Vanden Berghe, A. Schoeters, R. Bauer, L. Carnis, R. Elvik, A. Theofilatos, A. Filtness, S. Reed, C. Perez, and H. Martensen (2017). Crash cost estimates for European countries, deliverable 3.2 of the H2020 project SafetyCube. Technical report, European Commission, Horizon 2020 project SafetyCube, Brussels.

Online Appendix

Accounting for the Survivors:
Traffic Injuries, Match Destruction, and the Cost of Nonfatal Accidents

Kuan-Ming Chen Chen-An Lien

Not for publication

A Balance Check

Table OA1: T-test of Labor Outcomes among Injured Individuals before Accident Occurrence

Dependent Variable	Earning	Full-time Job	Lost-Job
Constant	111,704.4	0.2030	0.058
	(96.97)	(0.0001)	(0.0004)
Injured	-7,382.2	-0.0011	-0.012
	(130.2)	(0.0002)	(0.0003)
S.E.: Clustered	by: ID	by: ID	by: ID
Observations	145,722,410	145,722,410	145,722,410
R2	0.18921	0.19705	0.10775

Table OA2: T-test of Health Outcomes among Injured Individuals before Accident Occurrence

Dependent Variable	Inpatient	Outpatient	NHI Points
Constant	0.6803	10.00	11,994.9
	(0.0027)	(0.0041)	(22.04)
Injured	-0.2001	0.9213	-1,048.0
	(0.0031)	(0.0061)	(29.23)
S.E.: Clustered	by: ID	by: ID	by: ID
Observations	145,722,410	145,722,410	145,722,410
R2	0.00373	0.10022	0.01795

Note: Here, we present the results of a t-test between the never-accident group and injuries in accidents. "Injured" is a dummy variable that equals 1 if the observation belongs to an individual injured in a traffic accident and 0 otherwise. Similarly, we use data from the year prior to the accident for this analysis. The table shows that even before the accident, these injured individuals already had worse outcomes than that of the never-accident group. In other words, using the never-accident group as a control group is not an appropriate choice. Therefore, we select the later-treated group as the control group.

Later-Treated Cohort Balance and Timing Predictability

This part of the appendix reports the cohort-balance diagnostic summarized in Online Appendix 4, together with a timing-predictability regression. Figure OA1 compares each accident cohort g with its later-treated comparison cohort $g + 5$, measuring both groups in calendar year $g - 1$, across 16 cohort pairs (294,418 observations, controlling for a centered birth-year polynomial). Because the administrative sample is large enough for economically negligible differences to be statistically significant, the informative diagnostic is the standardized difference; across cohorts and covariates the standardized differences lie well within

the conventional ± 0.10 threshold. Table OA3 reports the timing-predictability regressions: among 433,360 eventually injured individuals, a person-level regression of first-accident year on sex and a quadratic in birth year has an R^2 of 0.1171; adding 2002 levels of earnings, full-time work, inpatient and outpatient utilization, and medical spending raises it to 0.1191, and adding their 2000–2002 linear trends raises it to 0.1197. Although the added regressors are jointly significant in this large sample, a one-standard-deviation change in any individual level or trend shifts predicted accident timing by at most 0.24 year.

Table OA3: Predicting the Year of First Traffic Injury

	(1) Demographics	(2) + Levels	(3) + Trends
Standardized pre-accident coefficients			
Earnings level	—	0.208	0.235
Full-time-work level	—	−0.067	−0.038
Inpatient level	—	−0.081	−0.117
Outpatient level	—	−0.117	−0.089
Medical-spending level	—	−0.099	−0.147
Earnings trend	—	—	−0.120
Full-time-work trend	—	—	−0.017
Inpatient trend	—	—	0.048
Outpatient trend	—	—	−0.054
Medical-spending trend	—	—	0.077
Demographic controls	Yes	Yes	Yes
Pre-accident levels	No	Yes	Yes
Pre-accident trends	No	No	Yes
Joint F: levels	—	192.9	—
Joint F: trends	—	—	61.5
R^2	0.1171	0.1191	0.1197
Incremental R^2 relative to (1)	—	0.0020	0.0026
Individuals	433,360	433,360	433,360

Notes: The dependent variable is the calendar year of an individual's first recorded traffic injury. All models include sex and a quadratic in centered birth year. Levels are measured in 2002, and trends are individual linear changes over 2000–2002. Displayed coefficients rescale each pre-accident predictor to one standard deviation, so their units are predicted accident years. The joint tests use the homoskedastic specification; both have $p < 0.001$. Point estimates and R^2 do not depend on the variance estimator.

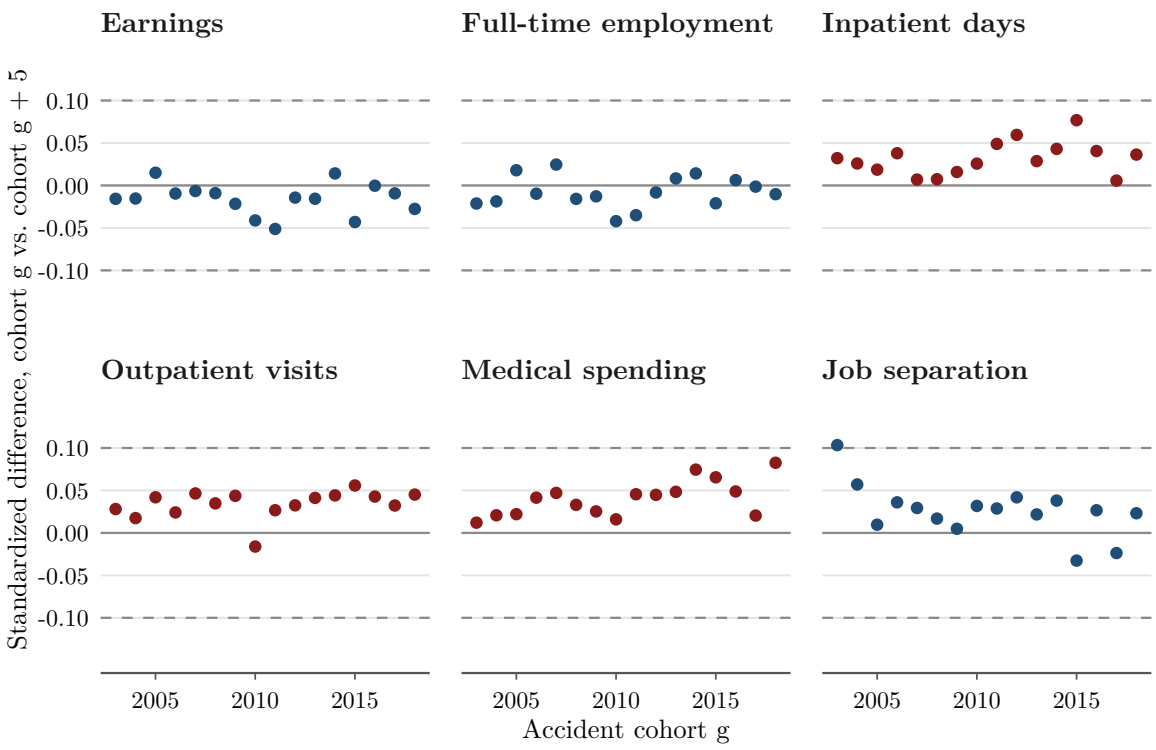


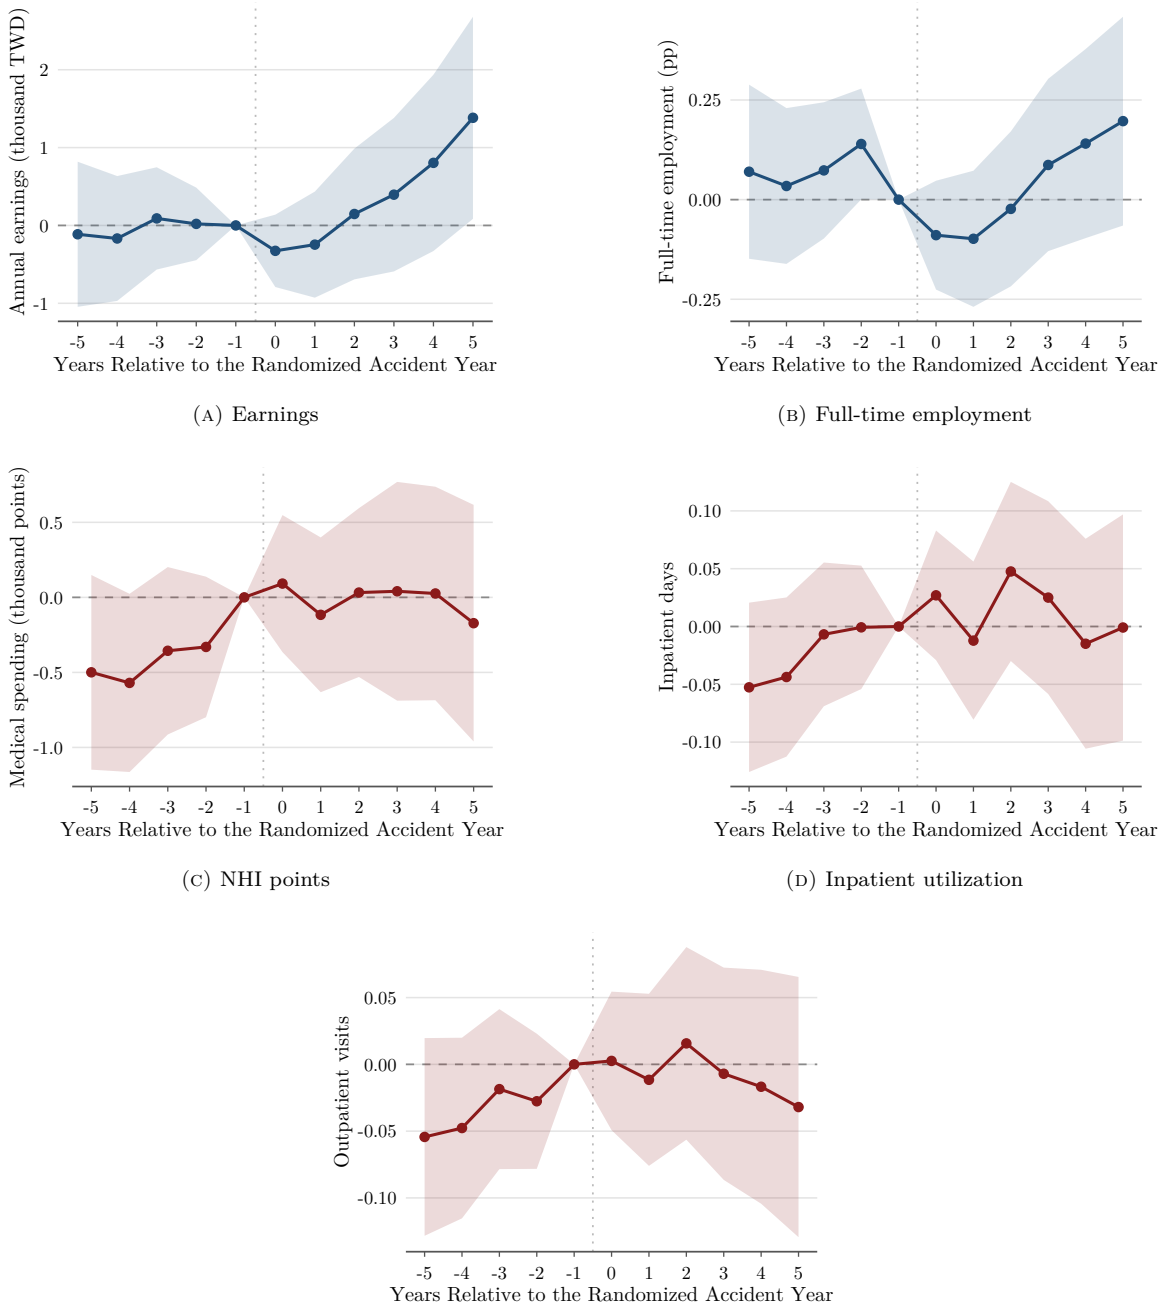
FIGURE OA1: Cohort-by-Cohort Standardized Differences: Treated vs. Later-Treated

Notes: Each point reports the standardized difference between accident cohort g and its later-treated comparison cohort $g + 5$, with both groups measured in calendar year $g - 1$. Dashed lines mark ± 0.10, a conventional threshold for small imbalance. The regression controls for a centered birth-year polynomial.

B Placebo: Randomized Accident Years

This appendix reports a placebo exercise that tests whether the main dynamic estimates could be generated by the cohort or calendar-time structure of the design rather than by the accident itself. Within the same main sample, we randomly reassign accident years across individuals, preserving the overall distribution of accident timing, and re-run the identical later-treated specification with birth-year controls.

Figure OA2 reports the results. The randomized accident years reproduce none of the true post-accident dynamics. At the placebo event year, the earnings effect is TWD -326 and the full-time employment effect is -0.09 percentage points, compared with TWD $-2{,}119$ and -0.77 percentage points at $t = 0$ in the main sample. Unlike the main estimates, the placebo effects also do not widen in subsequent years. On the health side, NHI-covered medical spending moves by only about $+92$ points at the placebo event year, compared with the spike of roughly $+80{,}000$ points following a true accident.



(A) Earnings

(B) Full-time employment

(C) NHI points

(D) Inpatient utilization

(E) Outpatient visits

FIGURE OA2: Placebo Event Studies with Randomized Accident Years

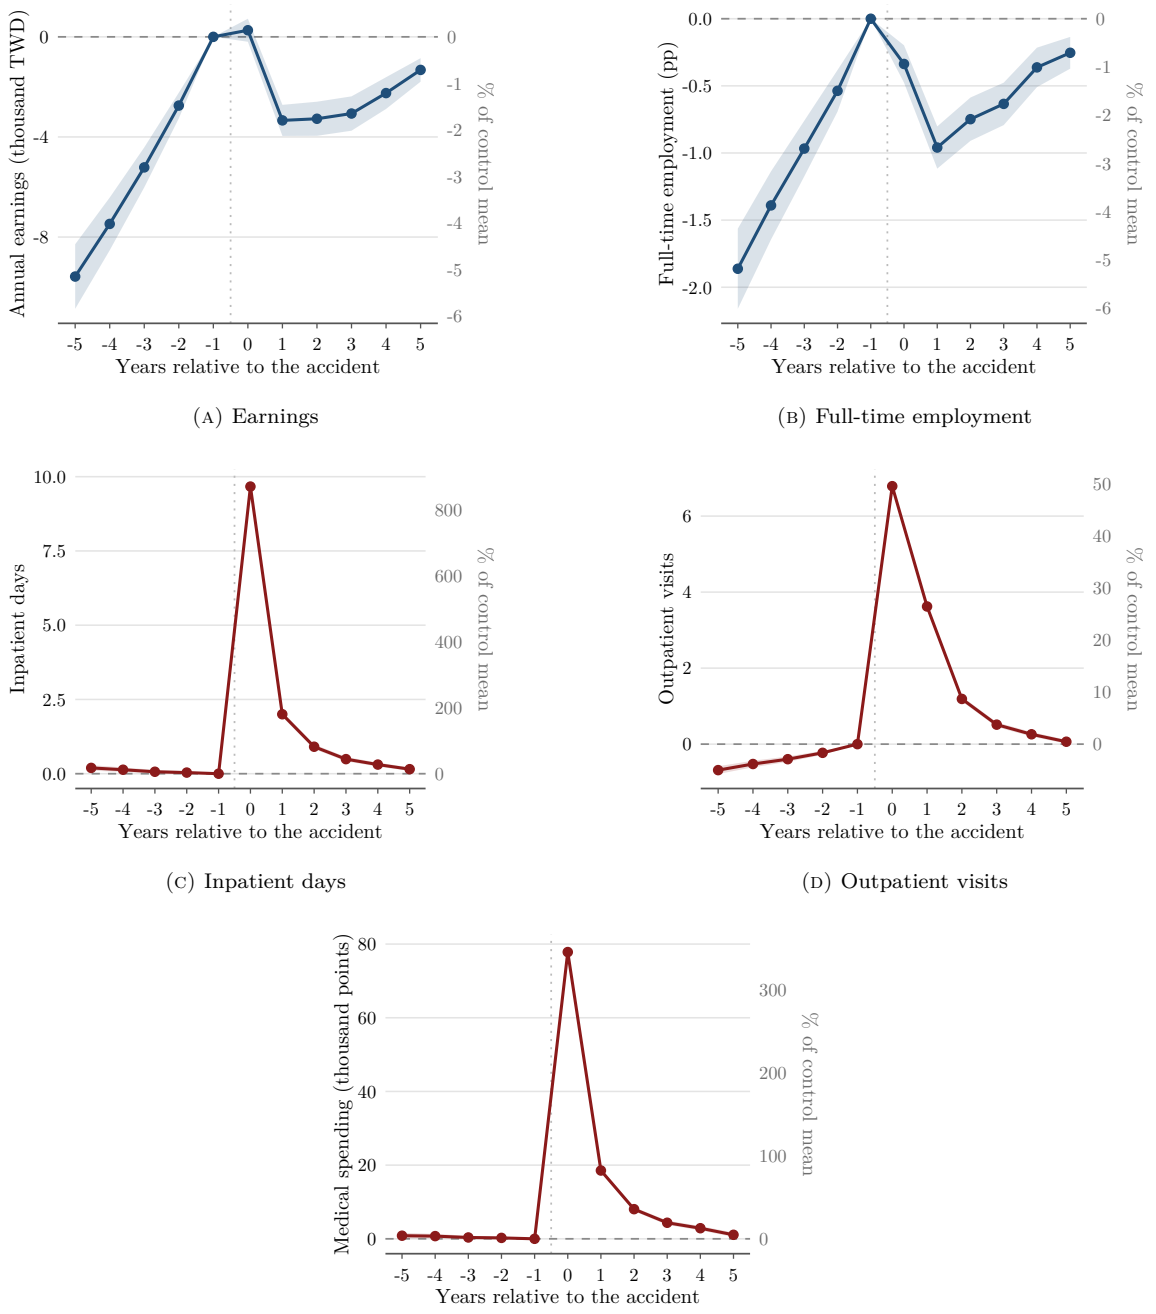
Notes: Each panel reports event-study estimates from the main specification after randomly reassigning accident years within the main sample, preserving the overall distribution of accident timing. Estimates are normalized to the year before the placebo accident year. Shaded bands show 95% confidence intervals.

C Conventional TWFE Event-Study Estimates

This appendix reports conventional TWFE event-study estimates for the main outcomes. The specification uses the same event-time normalization as the main dynamic analysis and includes birth-year controls to account for life-cycle differences across accident cohorts. These controls do not eliminate the pre-trend patterns in the TWFE estimates, suggesting that the discrepancy is not driven solely by life-cycle composition.

We also estimate cohort-specific TWFE specifications using later-treated individuals as the comparison group, paralleling the comparison structure in the main CS-DiD design. These cohort-specific estimates are broadly similar to the main dynamic results: earnings and full-time employment losses generally deepen after the accident rather than recover quickly. The divergence arises when cohorts are pooled in a conventional TWFE event-study, where already-treated observations can implicitly serve as controls for other treated cohorts and treatment effects are mixed across event times.

Figure OA3 reports the pooled TWFE estimates. The labor-market path differs sharply from both the CS-DiD estimates and the cohort-specific later-treated TWFE estimates: pooled TWFE suggests that earnings and full-time employment losses begin to recover after the second post-accident year, whereas the other two approaches show losses that continue to deepen through the fifth post-accident year. This comparison illustrates why the main analysis relies on the staggered design rather than the pooled TWFE event-study.



(A) Earnings

(B) Full-time employment

(C) Inpatient days

(D) Outpatient visits

(E) NHI points

FIGURE OA3: Conventional TWFE Event-Study Estimates

Notes: Each panel reports conventional two-way fixed effects event-study estimates for the main outcomes. Estimates are normalized to the year before the accident. Shaded bands show 95% confidence intervals. Earnings are expressed in 2021 TWD. NHI points are standardized reimbursement units under Taiwan's NHI.

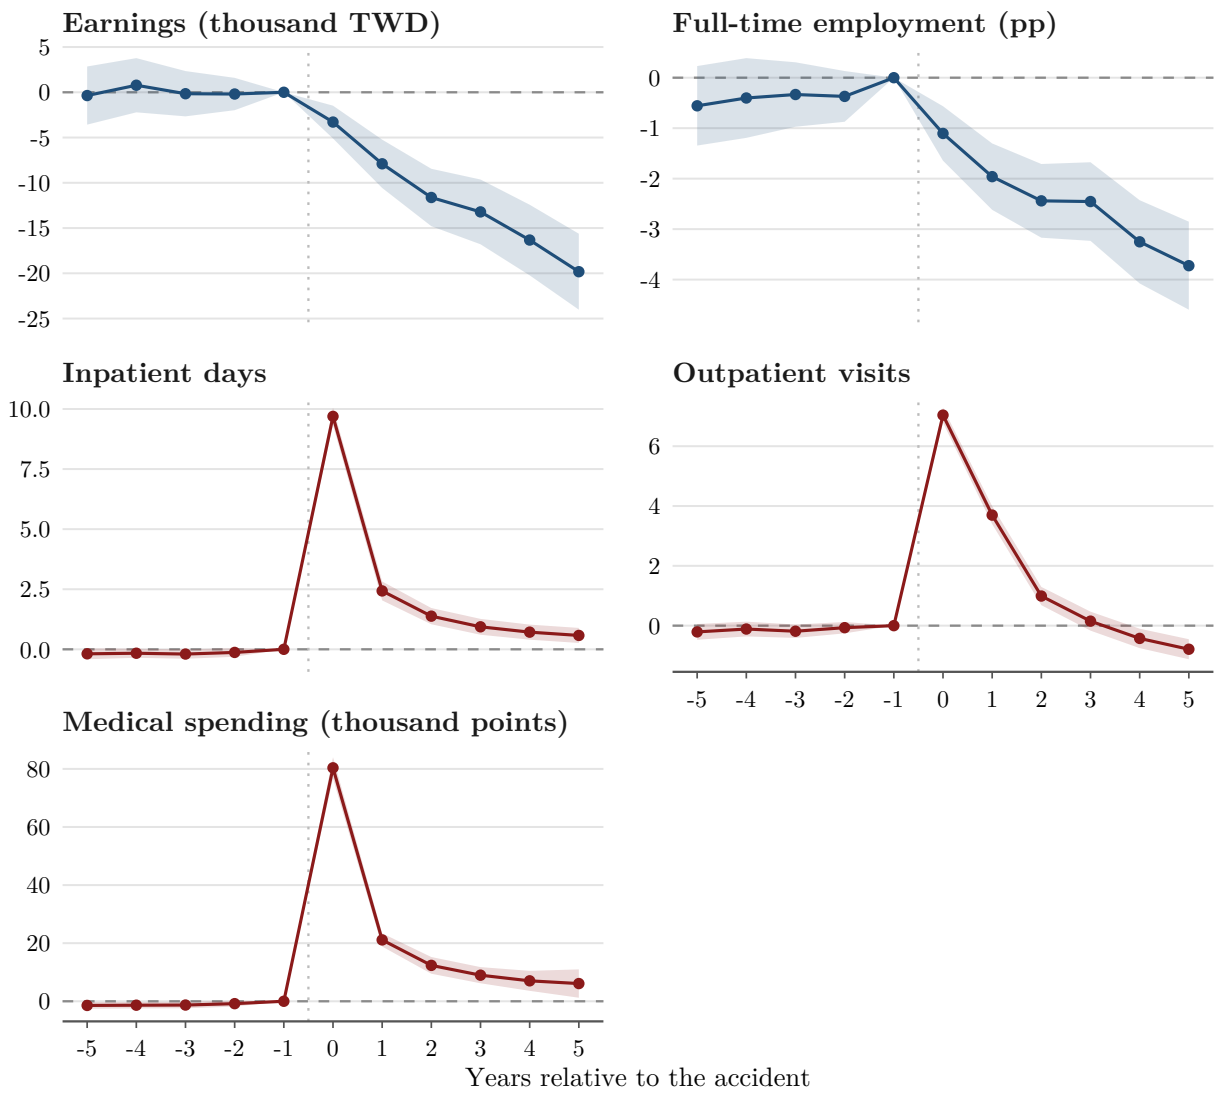


FIGURE OA4: Cohort-Specific Later-Treated TWFE Event-Study Estimates

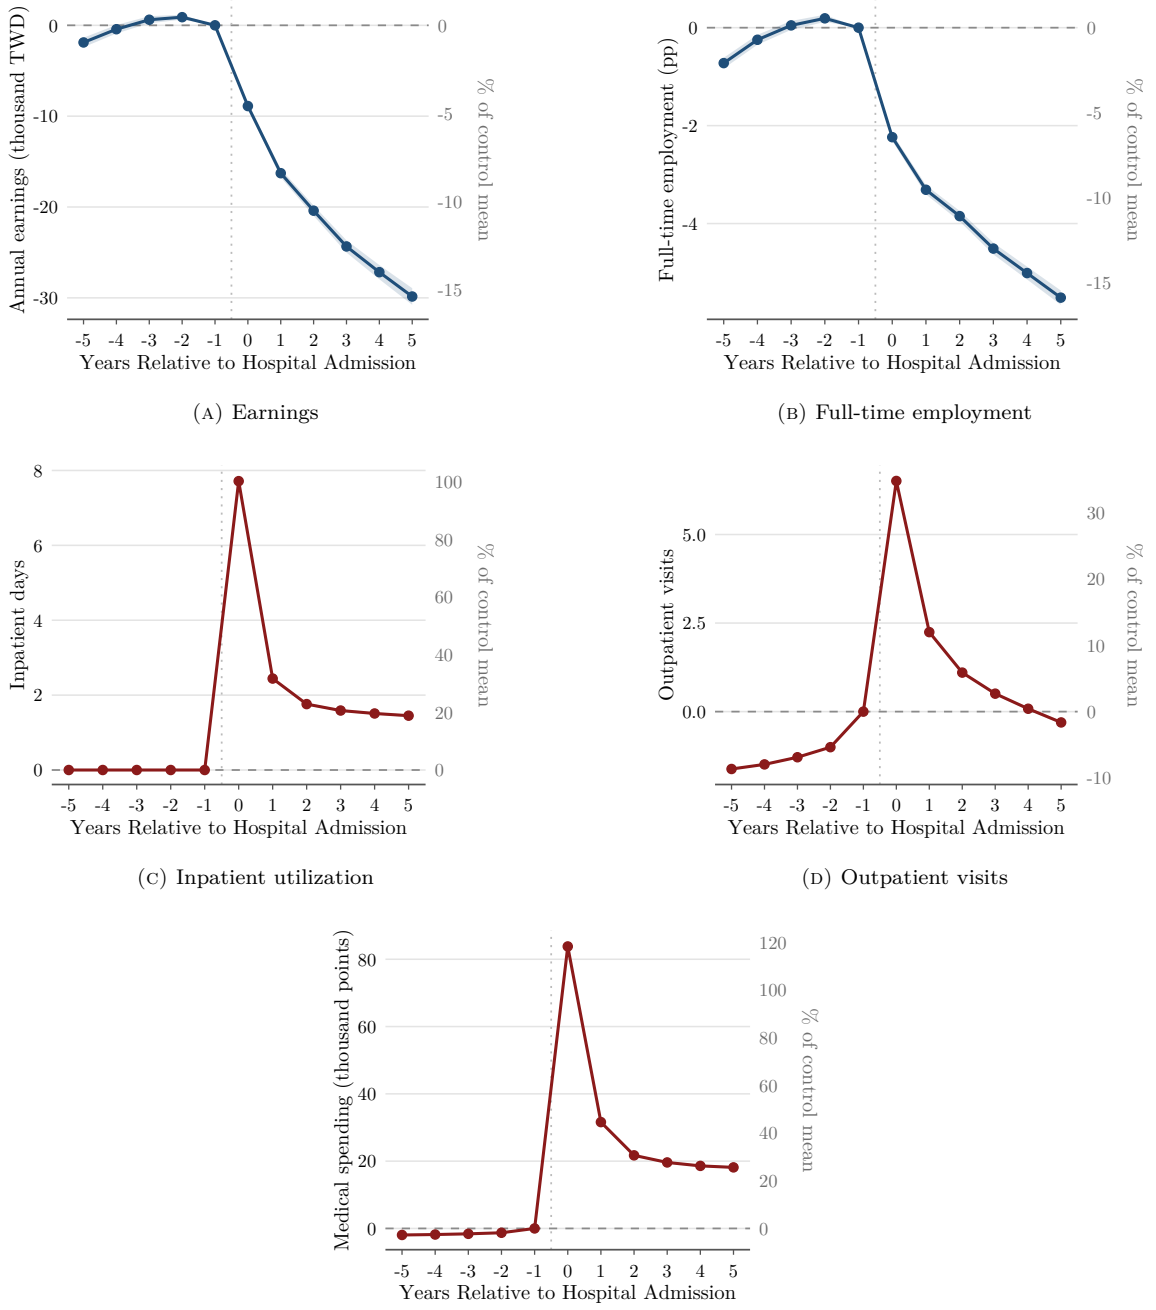
Notes: This figure reports a cohort-specific TWFE event-study for individuals first injured in 2011, using later-treated individuals as the comparison group. Estimates are normalized to the year before the accident. Shaded bands show 95% confidence intervals.

D Benchmark: General Hospital Admissions

This appendix benchmarks the traffic-injury estimates against general (all-cause) hospital admissions, placing our shock in the same framework as the hospital-admissions literature (Dobkin et al., 2018). We apply the identical later-treated event-study specification to individuals experiencing a first hospital admission of any cause, with birth-year controls.

Labor-market outcomes. Raw year-five labor-market losses are larger after a general admission than after a traffic injury: the year-five earnings effect is TWD −29,852 for general admissions versus TWD −16,280 for traffic injuries (a ratio of 1.83), and the full-time employment effect is −5.51 percentage points versus −3.08 (a ratio of 1.79). These raw comparisons, however, should not be read as causal differences between shock types, because the general-admission event studies display clear pre-trends: earnings move from TWD −1,887 at $t = -5$ to +896 at $t = -2$, and full-time employment is significantly different from zero in several pre-periods. Outpatient utilization and medical spending also drift upward before admission (the flat pre-period for inpatient care is mechanical, reflecting the first-admission definition). All-cause admissions therefore include chronic deterioration and anticipated hospitalizations. This is exactly the selection-on-timing concern that a police-defined, plausibly unanticipated traffic injury avoids. We consequently retain this exercise as a qualitative appendix benchmark and do not use the raw ratios as causal comparisons in the main text.

Health outcomes. At $t = 0$, a general admission increases inpatient utilization by about 7.7 days, outpatient visits by about 6.5, and NHI-covered spending by about 83,804 points, an admission-year spike similar in magnitude to the roughly 80,000-point spike after traffic injuries. Unlike traffic injuries, however, general-admission medical spending remains elevated by about 18,145 points at year five, consistent with the underlying chronic conditions that many all-cause admissions reflect.



(A) Earnings

(B) Full-time employment

(C) Inpatient utilization

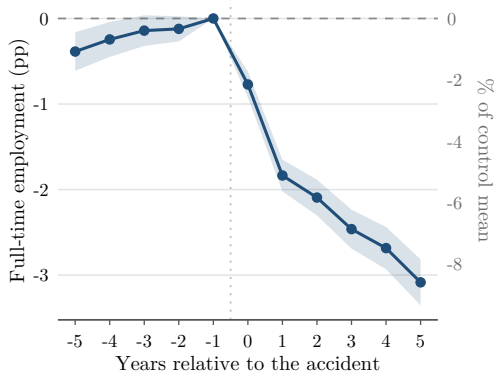
(D) Outpatient visits

(E) NHI points

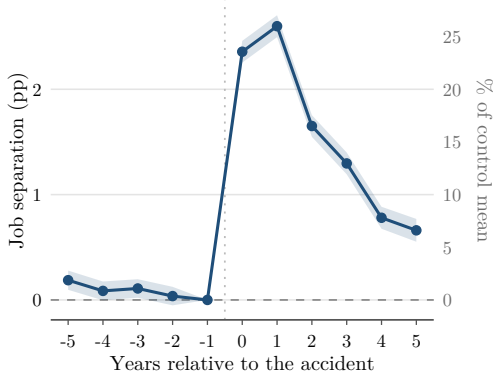
FIGURE OA5: Event Studies for General (All-Cause) Hospital Admissions

Notes: Each panel reports event-study estimates applying the main later-treated specification with birth-year controls to first all-cause hospital admissions. Estimates are normalized to the year before admission. Shaded bands show 95% confidence intervals. The general-admission sample displays non-negligible pre-trends in labor-market and outpatient outcomes; see the text for interpretation.

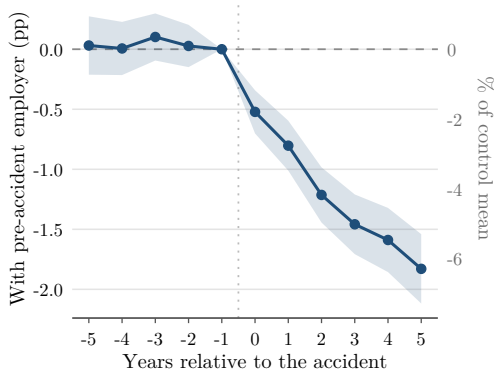
E Additional Employment and Utilization Outcomes



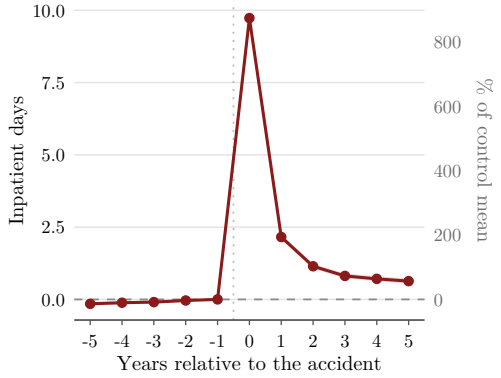
(A) Full-time employment (birth-year controls)



(B) Losing Job (birth-year controls)



(C) Staying in the Same Firm (birth-year controls)



(D) Inpatient days (birth-year controls)



(E) Outpatient visits (birth-year controls)

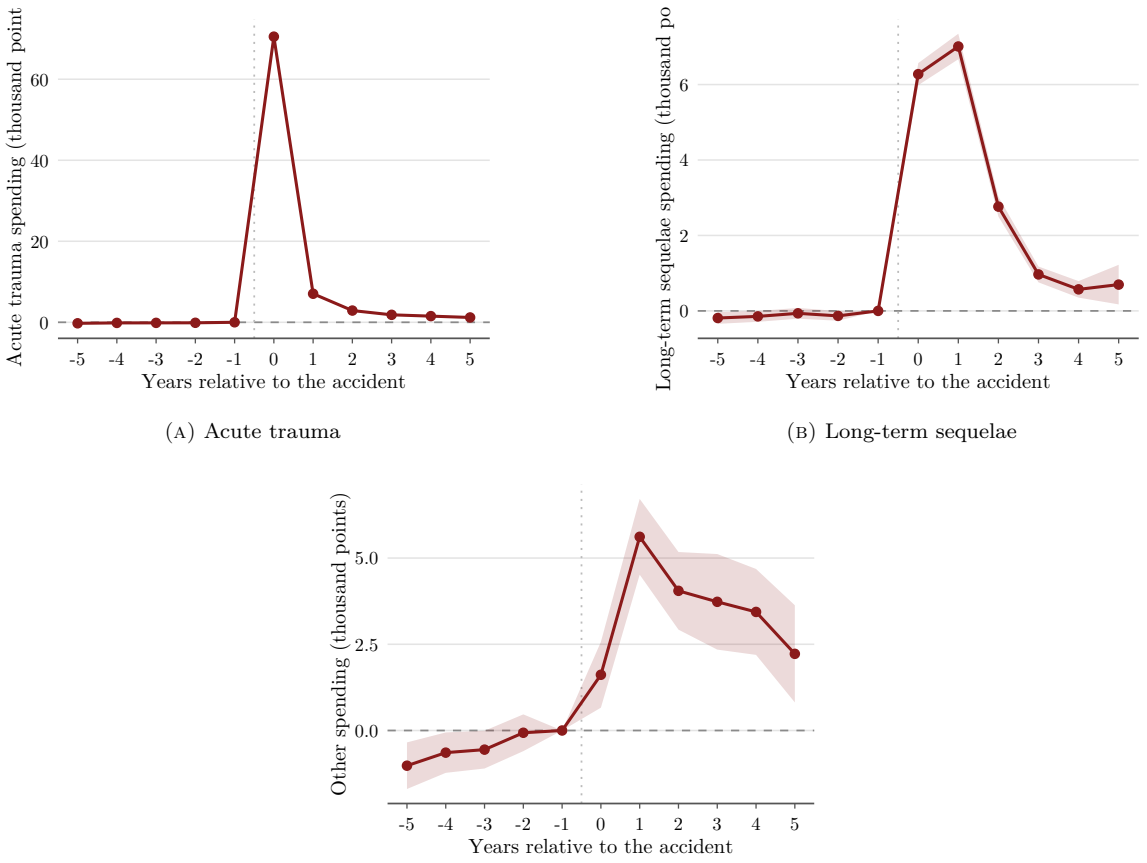
FIGURE OA6: Dynamic Effects of Traffic Injuries On Labor and Health Outcomes

Notes: This figure reports additional outcomes from the main specification.

F Medical Expense Decomposition

This appendix reports event-study estimates that decompose excess NHI points by medical spending category. The main text summarizes the composition of excess spending, showing that the accident-year spike is driven primarily by acute trauma care, while subsequent spending includes a larger role for long-term sequelae. The residual category combines remaining medical spending categories.

Figure OA7 plots the event-study estimates for acute trauma, long-term sequelae, and the residual other-spending category. The patterns reinforce the interpretation in the main text: trauma-related spending spikes sharply in the accident year and declines afterward, while long-term sequelae remain elevated beyond the initial injury episode.



(A) Acute trauma

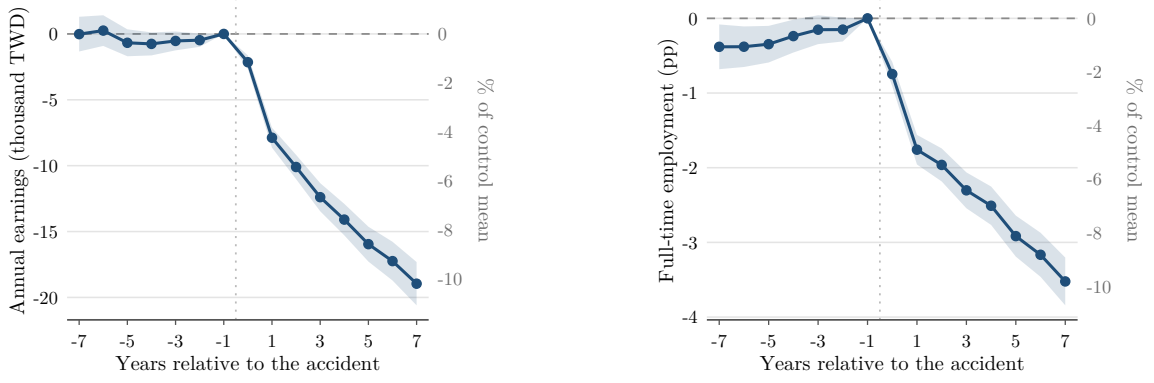
(B) Long-term sequelae

(C) Other medical spending

FIGURE OA7: Dynamic Effects by Medical Spending Category

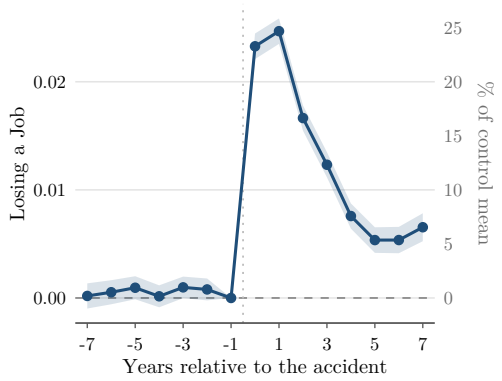
Notes: Each panel reports event-study estimates using later-treated controls with birth-year controls. Estimates are normalized to $t = -1$. Shaded bands show 95% confidence intervals. Medical spending is measured in NHI points.

G Extended Window: 7 & 10 Years

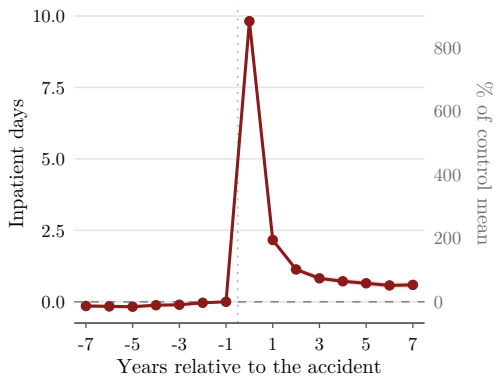


(A) Earnings (birth-year controls)

(B) Full-time employment (birth-year controls)



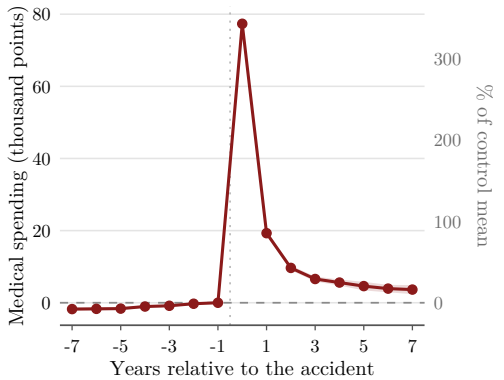
(C) Losing Job (birth-year controls)



(D) Inpatient days (birth-year controls)



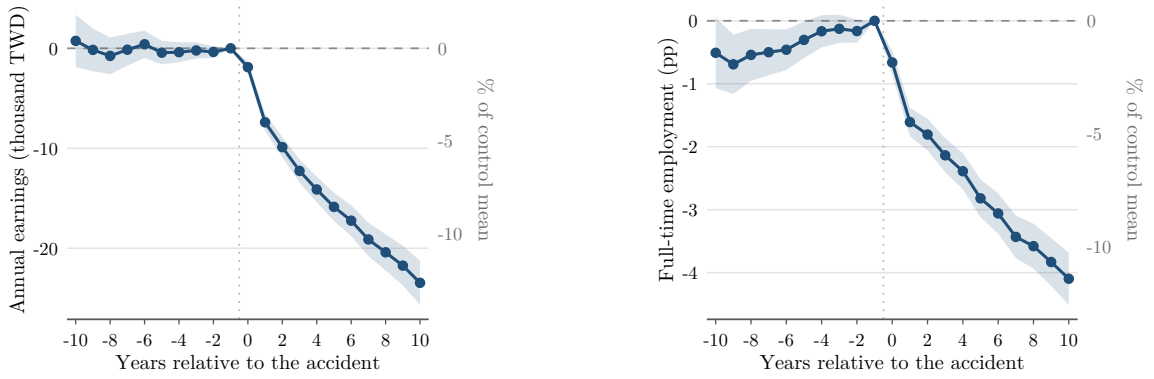
(E) Outpatient visits (birth-year controls)



(F) NHI points (birth-year controls)

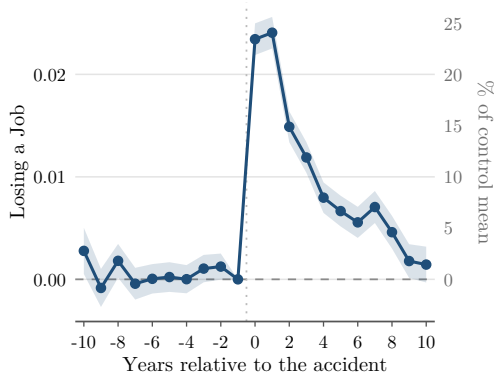
FIGURE OA8: Dynamic Effects over a 7-Year Window

Notes: This figure replicates the main specification using a balanced event window of $[-7, 7]$ years.

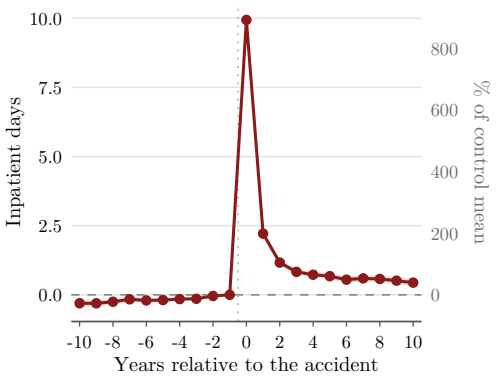


(A) Earnings (birth-year controls)

(B) Full-time employment (birth-year controls)



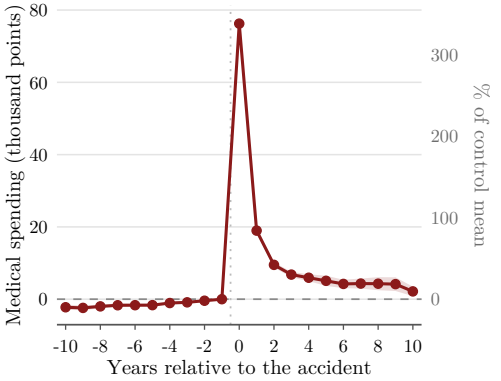
(C) Losing Job (birth-year controls)



(D) Inpatient days (birth-year controls)



(E) Outpatient visits (birth-year controls)



(F) NHI points (birth-year controls)

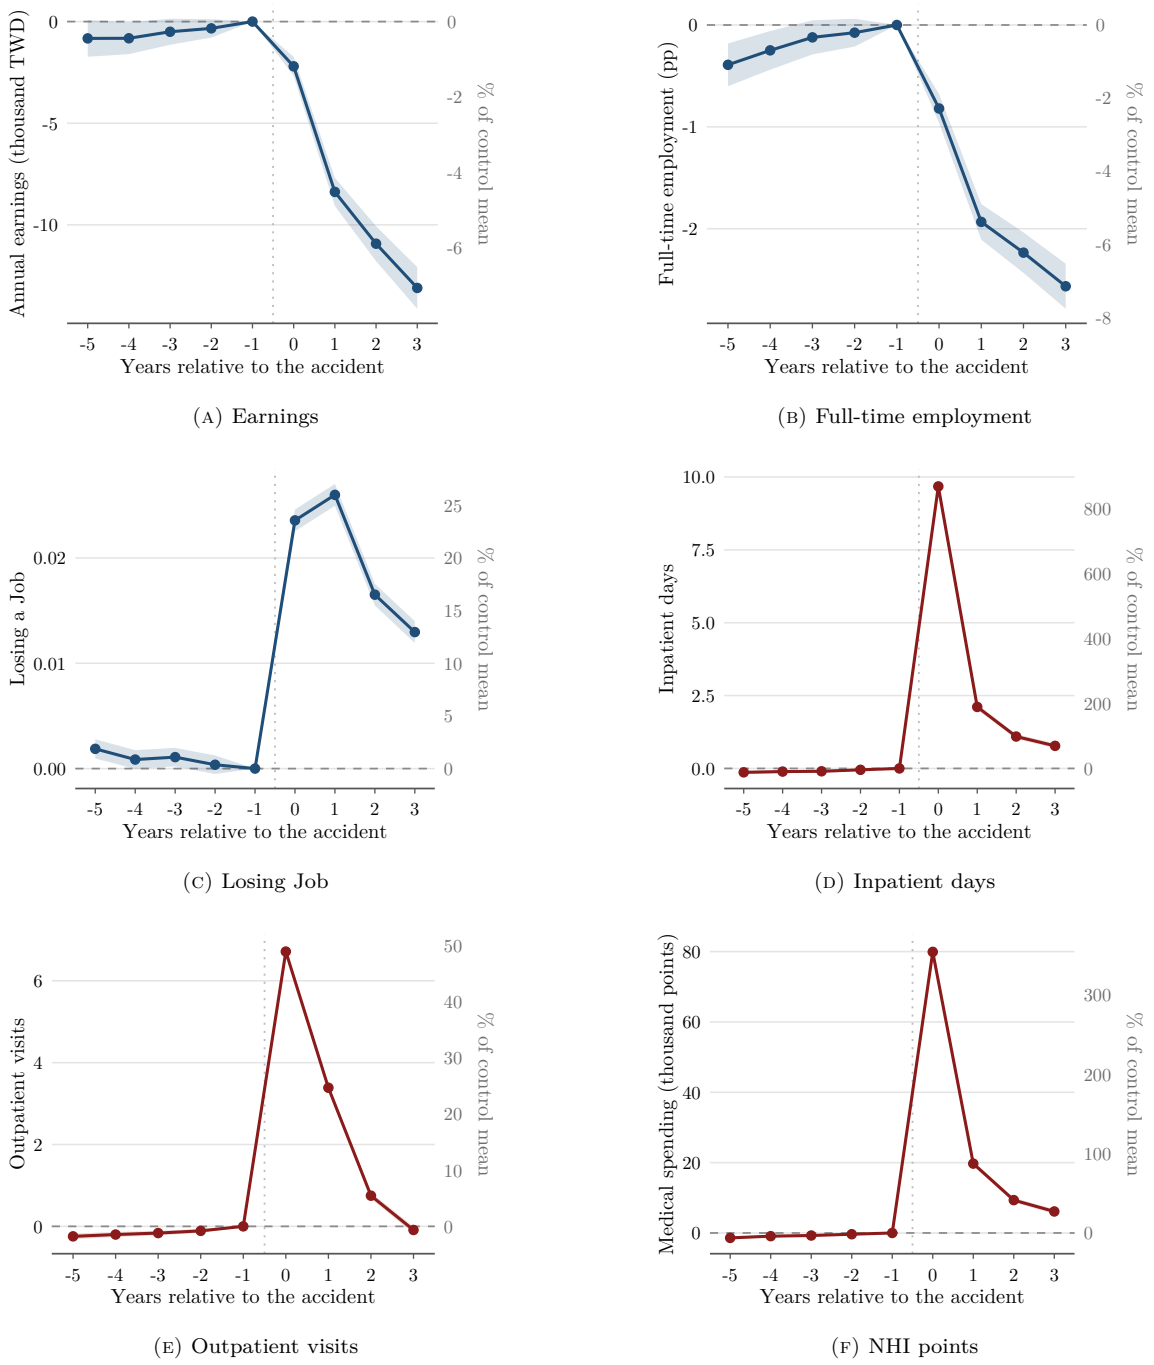
FIGURE OA9: Dynamic Effects over a 10-Year Window

Notes: This figure replicates the main specification using a balanced event window of $[-10, 10]$ years.

H Robustness to Narrow Treatment-Cohort Windows

This appendix examines whether the main dynamic results are sensitive to using a narrower set of comparison cohorts. The baseline specification uses later-treated individuals as controls. As a robustness check, we restrict the comparison group to cohorts treated within a three-year window, making treated and comparison cohorts closer in calendar time. This restriction reduces concerns that the estimates are driven by broad secular differences between early- and late-treated cohorts.

Figure OA10 reports the results for the main outcomes with birth-year controls. The qualitative patterns are similar to the baseline estimates: earnings and full-time employment decline persistently after the accident, while inpatient days, outpatient visits, and NHI points rise sharply around the accident and remain elevated afterward.



(A) Earnings

(B) Full-time employment

(C) Losing Job

(D) Inpatient days

(E) Outpatient visits

(F) NHI points

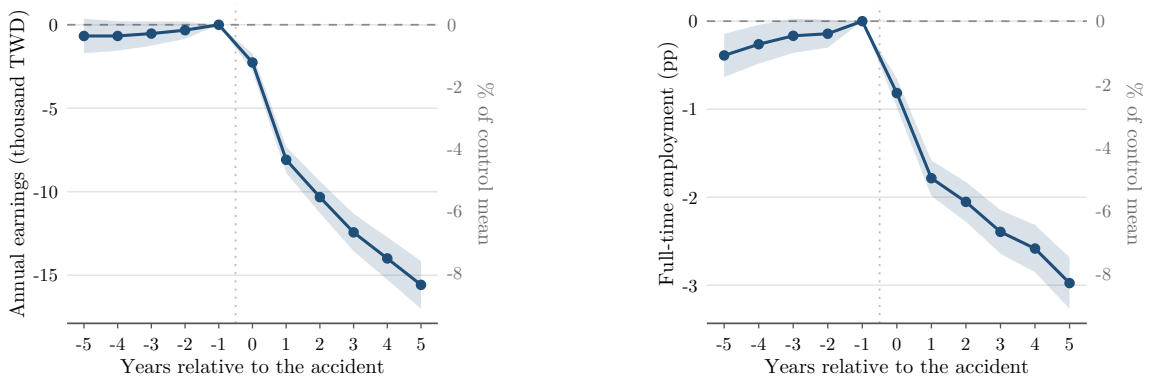
FIGURE OA10: Robustness to Three-Year Treatment-Cohort Window

Notes: Each panel reports event-study estimates using later-treated controls restricted to cohorts treated within a three-year window. Estimates control for birth-year fixed effects and are normalized to $t = -1$. Shaded bands show 95% confidence intervals.

I Robustness to Hospitalization-Window Definition

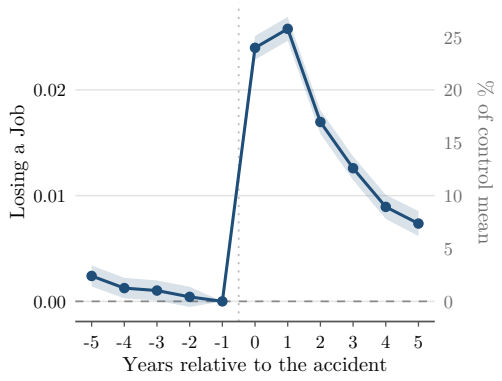
This appendix examines whether the main results are sensitive to the hospitalization-window definition used to construct the main sample. The baseline sample includes injured individuals hospitalized within 48 hours after the traffic accident. As a robustness check, we restrict the sample to accidents followed by hospitalization within 24 hours. This alternative definition focuses on cases with a tighter temporal link between the police-reported accident and the hospitalization episode.

Figure OA11 reports the results for the main outcomes with birth-year controls. The qualitative patterns are similar to the baseline estimates: earnings and full-time employment decline after the accident, job separation rises sharply around the accident, and medical utilization and NHI points increase substantially.



(A) Earnings

(B) Full-time employment



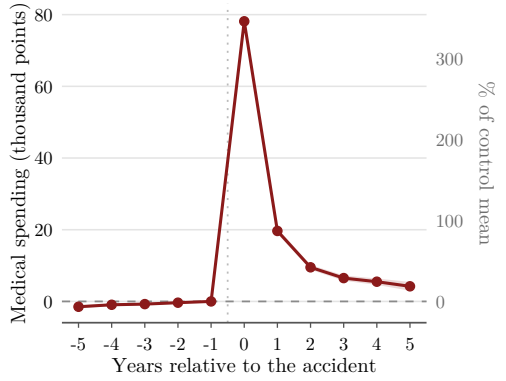
(C) Losing job



(D) Inpatient days



(E) Outpatient visits

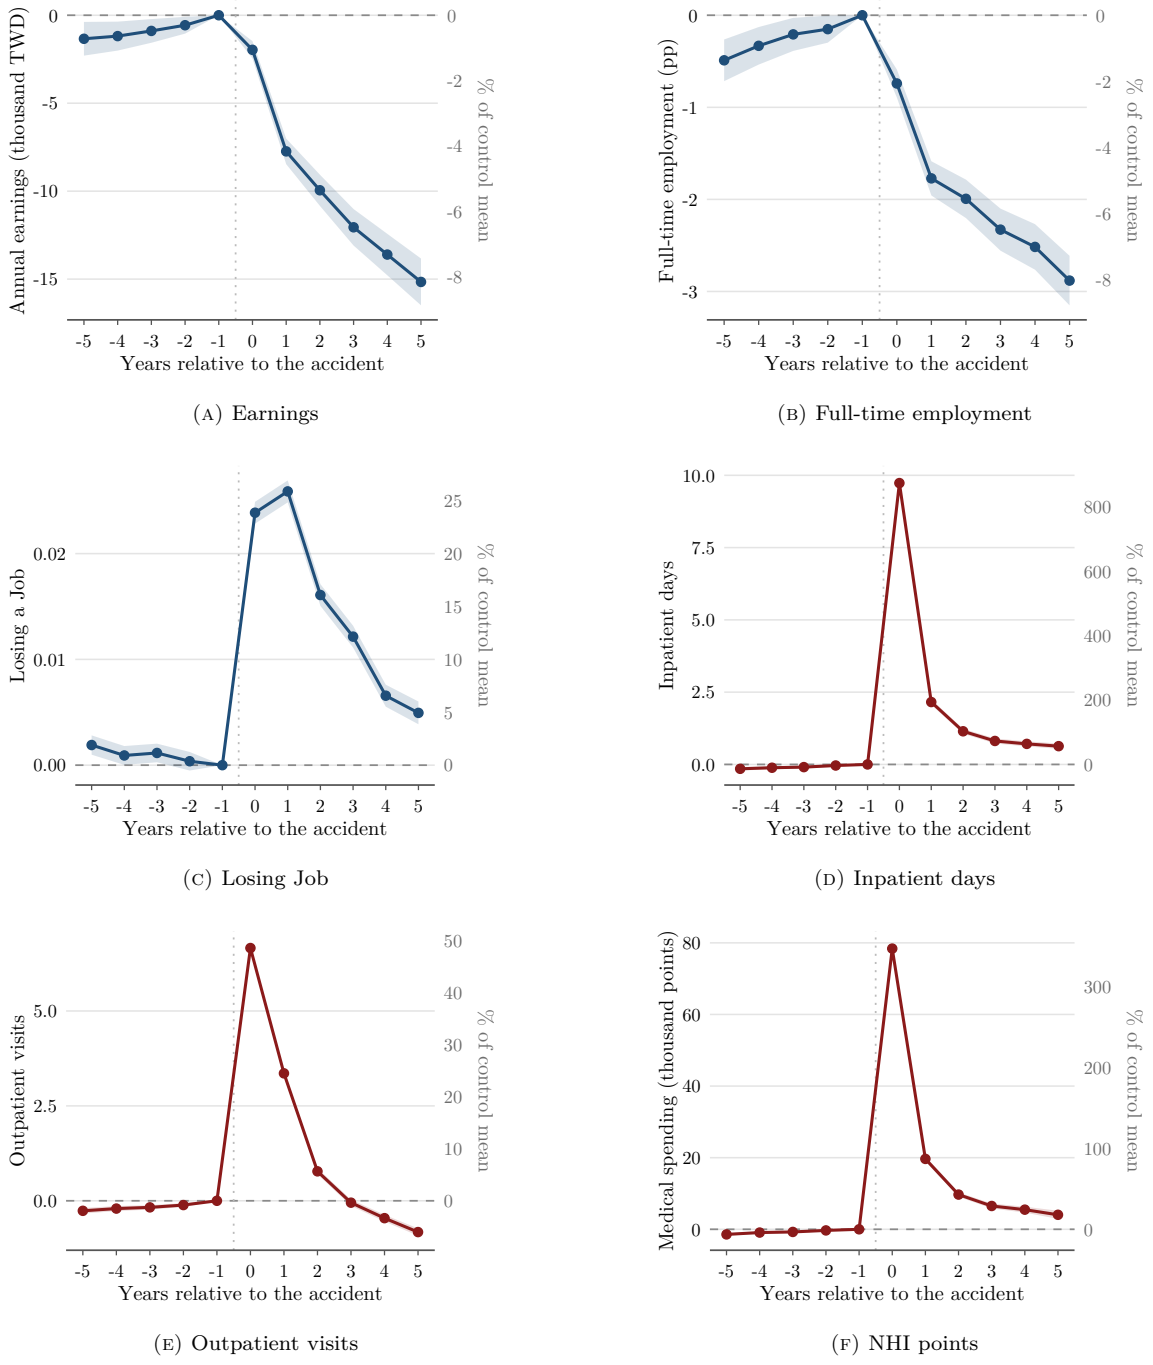


(F) NHI points

FIGURE OA11: Robustness to Hospitalization within 24 Hours

Notes: This figure replicates the main specification after restricting the sample to individuals hospitalized within 24 hours after the traffic accident. Estimates control for birth-year fixed effects and are normalized to $t = -1$. Shaded bands show 95% confidence intervals.

J Main Results without Controlling Birth Year



(A) Earnings

(B) Full-time employment

(C) Losing Job

(D) Inpatient days

(E) Outpatient visits

(F) NHI points

FIGURE OA12: Dynamic Effects of Traffic Injuries without Covariates

Notes: This figure replicates the main specification but excludes all covariates.

K Separation Timing and the Expiry of Employment Protection

This appendix reports the timing of separations from the pre-accident employer discussed in Online Appendix 5.2. Under the Labor Standards Act, an employer must continue paying a worker injured in an occupational accident, including many commuting accidents, for up to two years, after which the obligation can be discharged with a lump sum. If this protection merely postponed dismissals, separations should bunch at month 24. Among the 60,465 workers attached to an identifiable employer in the month before the accident, Figure OA13 plots the monthly hazard of first leaving that employer over months 12 to 36. The hazard is 1.23 percent in month 24, compared with 1.28 percent for months 18 to 23 pooled and 1.16 percent for months 25 to 30 pooled: the series declines smoothly through month 24 with no discrete jump. Table OA4 reports the separator–stayer accounting decomposition behind Figure 2.

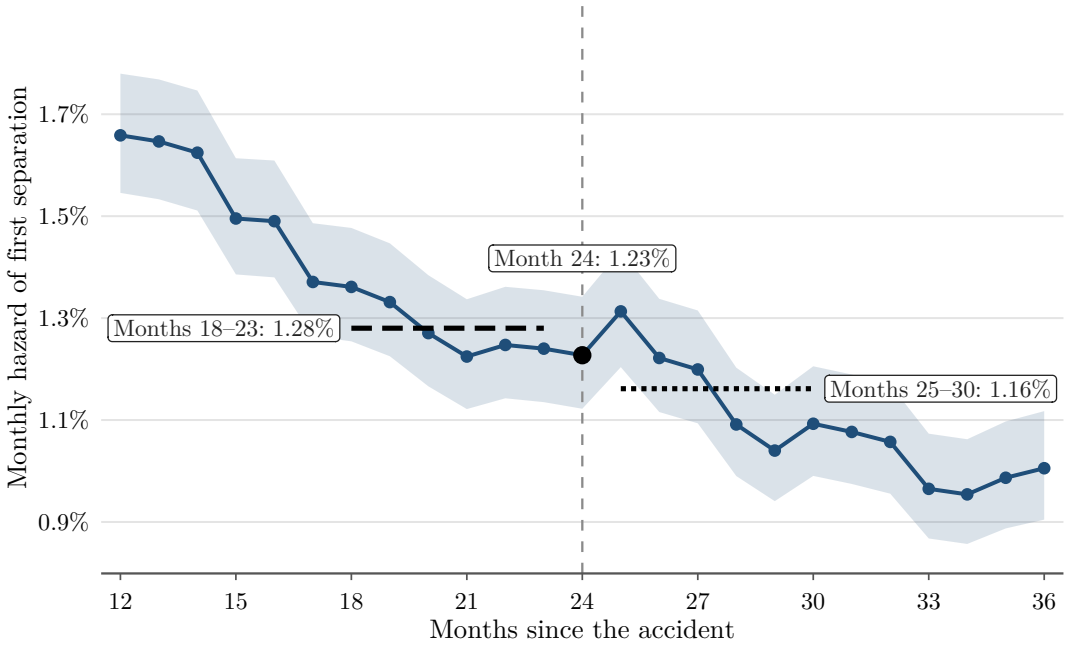


FIGURE OA13: Monthly Hazard of First Separation from the Pre-Accident Employer

Notes: The sample contains 60,465 workers attached to an identifiable employer in event month −1. The figure reports the monthly hazard of first leaving that employer, conditional on remaining attached at the start of the month. The shaded region is the 95% confidence interval. The vertical dashed line marks month 24; the horizontal segments report the pooled hazards for months 18–23 and 25–30.

TABLE OA4: Descriptive Year-Five Earnings Decomposition by Separation Status

Group	Individuals	Share	$t = -1$ earnings	$t = 5$ earnings	Change	Loss share
Separator	21,839	36.1%	336,790	232,697	−30.9%	90.6%
Stayer	38,626	63.9%	483,493	477,382	−1.3%	9.4%

Notes: Separators first leave the pre-accident employer during event months 0–23; stayers remain attached through that window. Earnings are in 2021 TWD. The contribution column holds group shares fixed and decomposes the aggregate descriptive year-five earnings change of TWD −41,501. Because separation status is defined after the accident, the table is an accounting decomposition rather than a causal subgroup comparison.

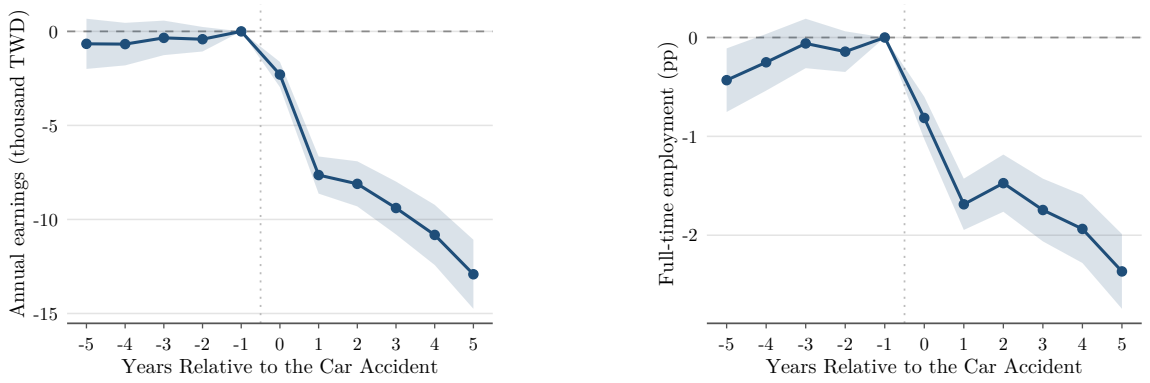
L Low-Utilization (Recovered) Subsample

This appendix reports the low-utilization ("recovered") subsample analysis referenced in Online Appendix 5.2. The purpose is corroboration: if labor-market losses persist among workers whose medical utilization has returned to the range typical of uninjured individuals, then medical recovery and economic recovery are not the same object.

Definition. For each accident cohort, we compute the 75th percentile of total medical spending (`T_DOT`) at event time $t = 2$ among later-treated controls in the same calendar year. An injured worker enters the low-utilization subsample if their own $t = 2$ medical spending is at or below this control-based threshold. Of 388,118 classifiable injured workers, 260,453 (67.1 percent) meet the criterion.

Results. Figure OA14 reports the event studies. Labor-market losses in the low-utilization subsample follow the same widening trajectory as in the main sample: the earnings effect grows from TWD −2,295 at $t = 0$ to TWD −12,915 at $t = 5$, and the full-time employment effect from −0.81 to −2.37 percentage points. Relative to the main-sample year-five effects (TWD −16,280 and −3.08 percentage points), the subsample retains 79 percent of the earnings loss and 77 percent of the employment loss.

Interpretation. Because post-accident medical utilization is a post-treatment outcome, conditioning on it can induce selection, and these estimates should not be read as causal effects for "recovered workers." We therefore present the exercise as corroboration of the decoupling documented in the main text, not as identification. The direction of selection is, if anything, conservative for our claim: workers whose utilization normalizes quickly are plausibly positively selected on both health and labor-market prospects, yet more than three-quarters of the average loss remains within this group. Taiwan's institutional setting also rules out one mechanical alternative: because NHI coverage is universal, low-copayment, and not tied to employment, job loss does not restrict access to care, so low post-accident utilization is unlikely to reflect lost insurance rather than recovery.



(A) Earnings

(B) Full-time employment

FIGURE OA14: Event Studies for the Low-Utilization-at-Year-Two Subsample

Notes: Each panel reports event-study estimates from the main later-treated specification with birth-year controls, restricted to injured workers whose total medical spending at $t = 2$ is at or below the 75th percentile of same-calendar-year later-treated controls. Estimates are normalized to $t = -1$. Shaded bands show 95% confidence intervals.

M New Catastrophic-Illness Certification Subsample

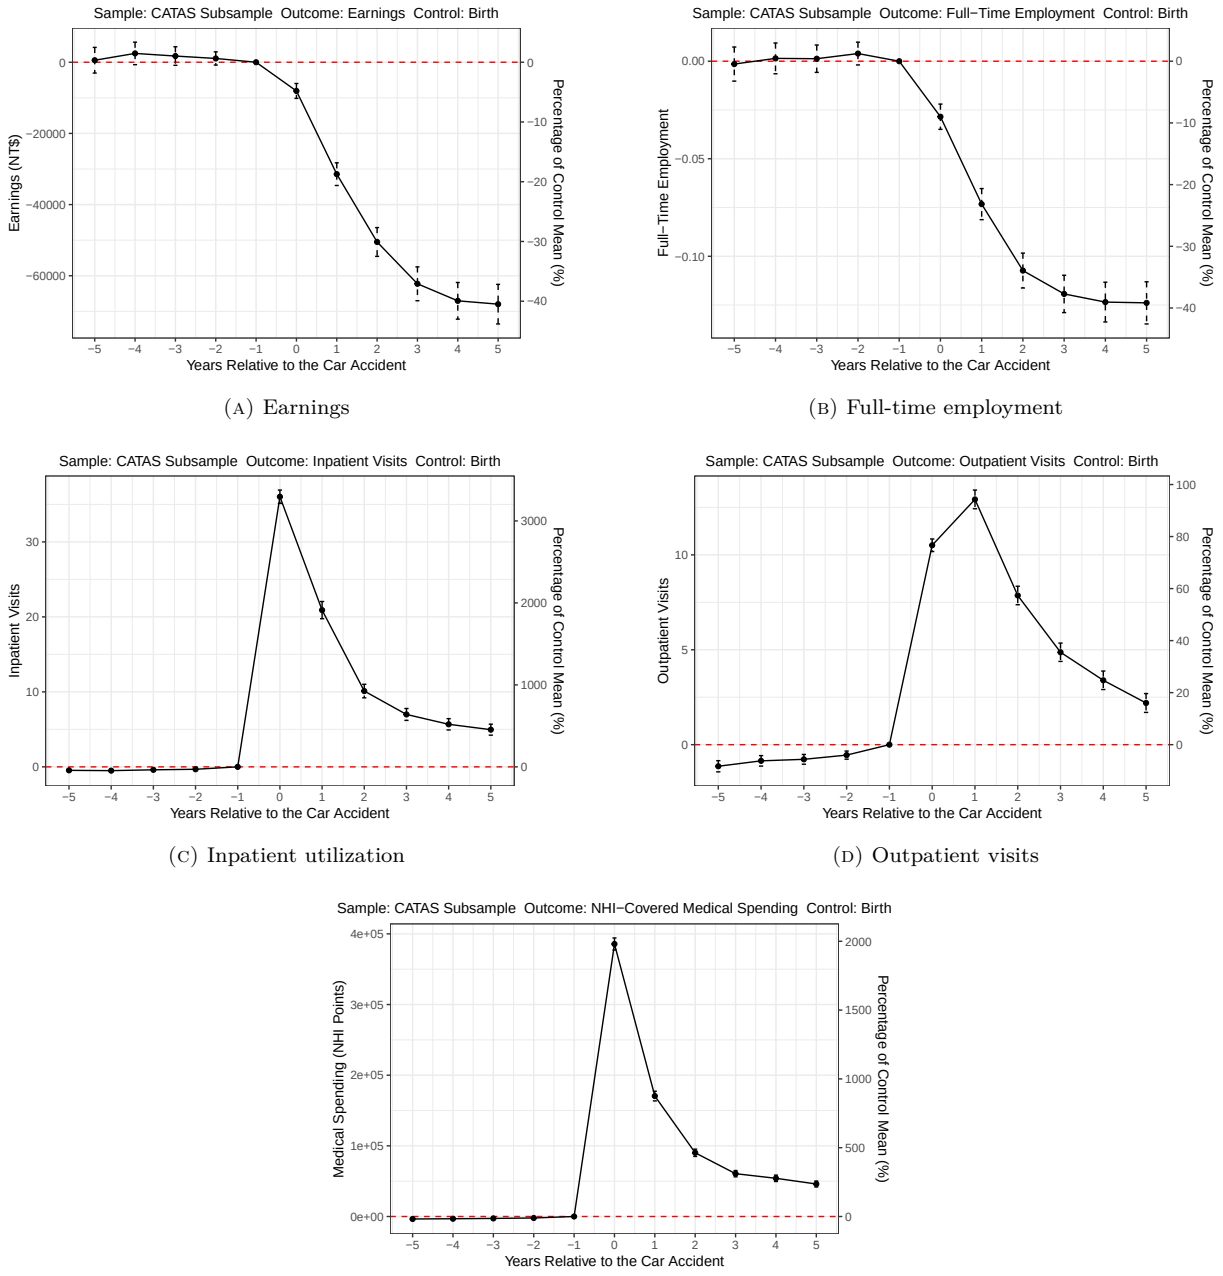
This appendix reports the event studies underlying the catastrophic-illness accounting exercise in Online Appendix 5.3. We define the subgroup as injured individuals whose first National Health Insurance catastrophic-illness certification occurred in the accident year or the following calendar year. We apply the main later-treated specification with birth-year controls and the same balanced five-year event window used for the main estimate.

Accounting. Among all injured individuals contributing to the balanced year-five estimate, 6.1 percent entered the newly certified subgroup. At year five, their estimated annual earnings loss is TWD 67,978 and their full-time employment probability is 12.4 percentage points lower. The corresponding main-sample earnings loss is TWD 16,280. Weighting the subgroup loss by its share gives an aggregate contribution of approximately TWD 4,144 per injured individual, or

$$\frac{0.060963 \times 67{,}978}{16{,}280} = 0.2545,$$

so the subgroup accounts for 25.5 percent of the aggregate year-five earnings loss under this accounting attribution. Under the more extreme assumption that every newly certified worker loses all nonnegative pre-accident earnings, the implied bound is 63.9 percent.

Interpretation. Figure OA15 shows that the subgroup experiences large and persistent reductions in earnings and full-time employment, together with substantial medical utilization. Catastrophic-illness certification is determined after the accident, so these event studies describe a post-treatment subgroup and should not be interpreted as causal effects of certification or as a causal decomposition of the main effect. The NHI certification also records formally recognized catastrophic illness rather than possession of a disability certificate. The accounting exercise therefore quantifies how much of the aggregate loss is associated with this severe-illness marker; it does not measure the total contribution of disability or lasting impairment.



(A) Earnings

(B) Full-time employment

(C) Inpatient utilization

(D) Outpatient visits

(E) NHI points

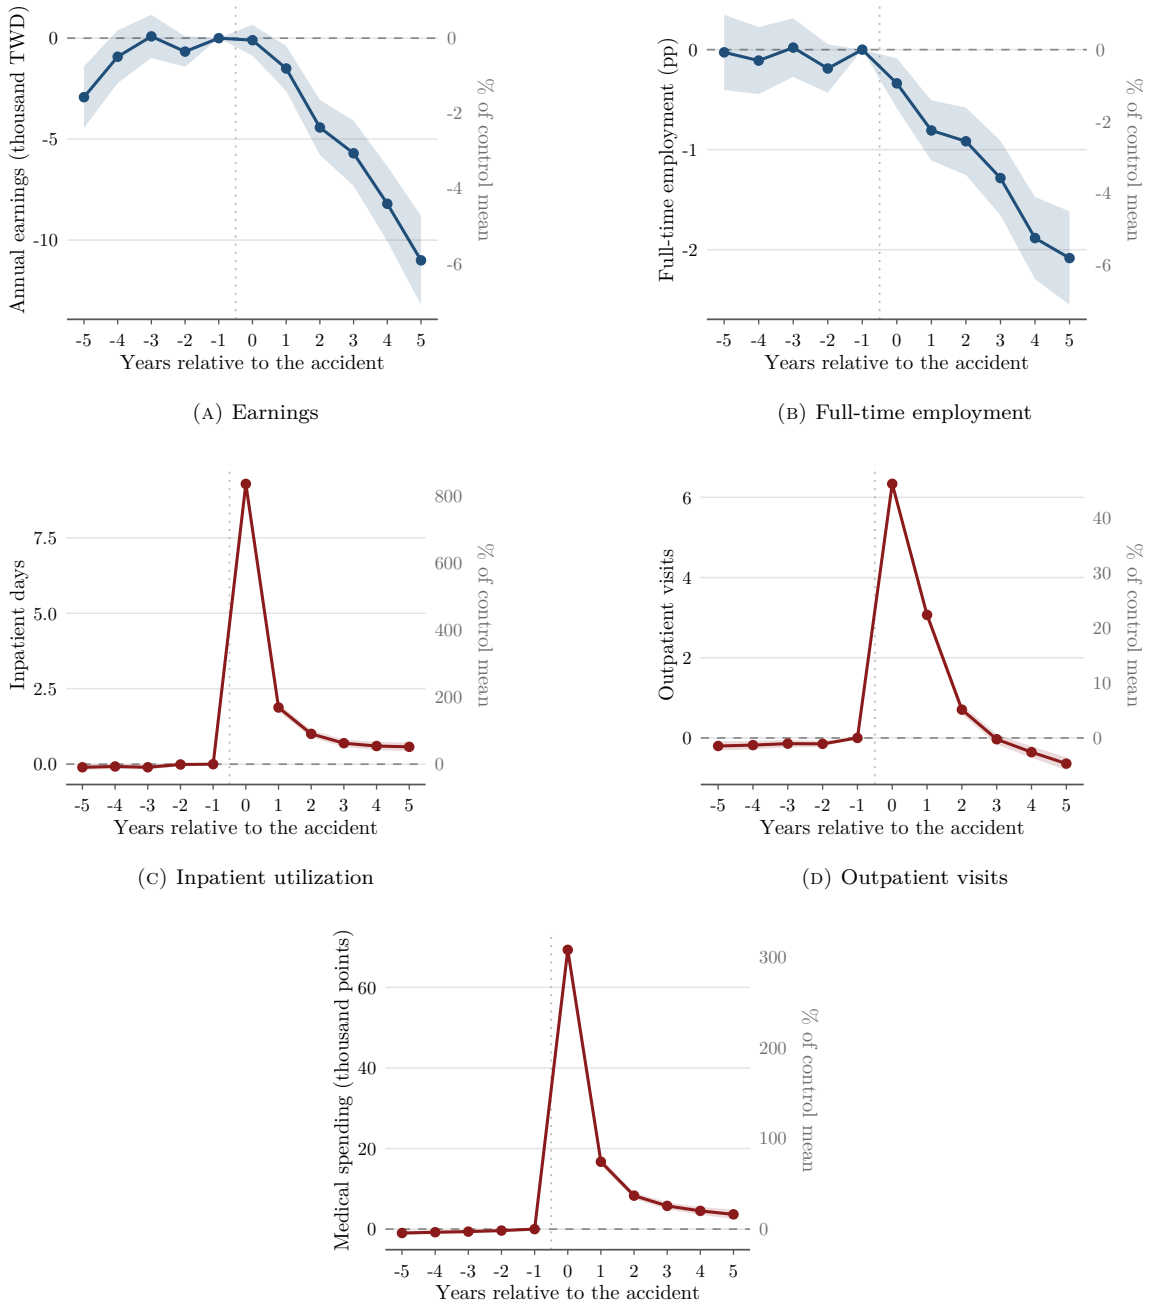
FIGURE OA15: Event Studies for the Newly Certified Catastrophic-Illness Subsample

Notes: The sample consists of injured individuals whose first NHI catastrophic-illness certification occurred in the accident year or the following calendar year. Each panel applies the main later-treated specification with birth-year controls and a balanced event window of $[-5, 5]$ years. Estimates are normalized to $t = -1$; whiskers show 95% confidence intervals. Because certification is observed after the accident, these post-treatment subgroup estimates are descriptive and are used only for the accounting exercise.

N Claims-Identifiable Traffic-Injury Subsample

This appendix examines how the common practice of identifying traffic injuries from medical claims, rather than police records, affects the estimated consequences. In our linked data, diagnosis-based transport-injury codes recover only about 40 percent of police-reported accidents. We re-estimate the main event studies on the subsample of injured workers who are identifiable from NHI claims and compare them with the police-defined main sample.

Conditioning on identification in medical claims does not select a subgroup with larger estimated responses in this hospitalized sample. The accident-year medical-spending effect is approximately 69,292 NHI points in the claims-identifiable subsample, compared with 78,405 points in the police-defined main sample, a difference of about 12 percent. At year five, the earnings effect is TWD −15,100 rather than TWD −16,280, and the full-time employment effect is −2.88 rather than −3.08 percentage points. Thus, claims identification attenuates the measured spending and labor-market responses by roughly 7–12 percent here. This comparison does not imply that claims-based samples are always less severe; it shows that conditioning on a transport-injury diagnosis changes sample composition and, in our setting, modestly understates the effects obtained from the police-defined event universe.



(A) Earnings

(B) Full-time employment

(C) Inpatient utilization

(D) Outpatient visits

(E) NHI points

FIGURE OA16: Event Studies for the Claims-Identifiable Traffic-Injury Subsample

Notes: Each panel reports event-study estimates from the main later-treated specification with birth-year controls, restricted to traffic injuries identifiable from diagnosis-based transport-injury codes in NHI claims. Estimates are normalized to $t = -1$. Shaded bands show 95% confidence intervals.

O Heterogeneity Analysis

This appendix reports additional heterogeneity analyses that complement the main results in Online Appendix 5.

O.1 Heterogeneity by Injury Severity Level

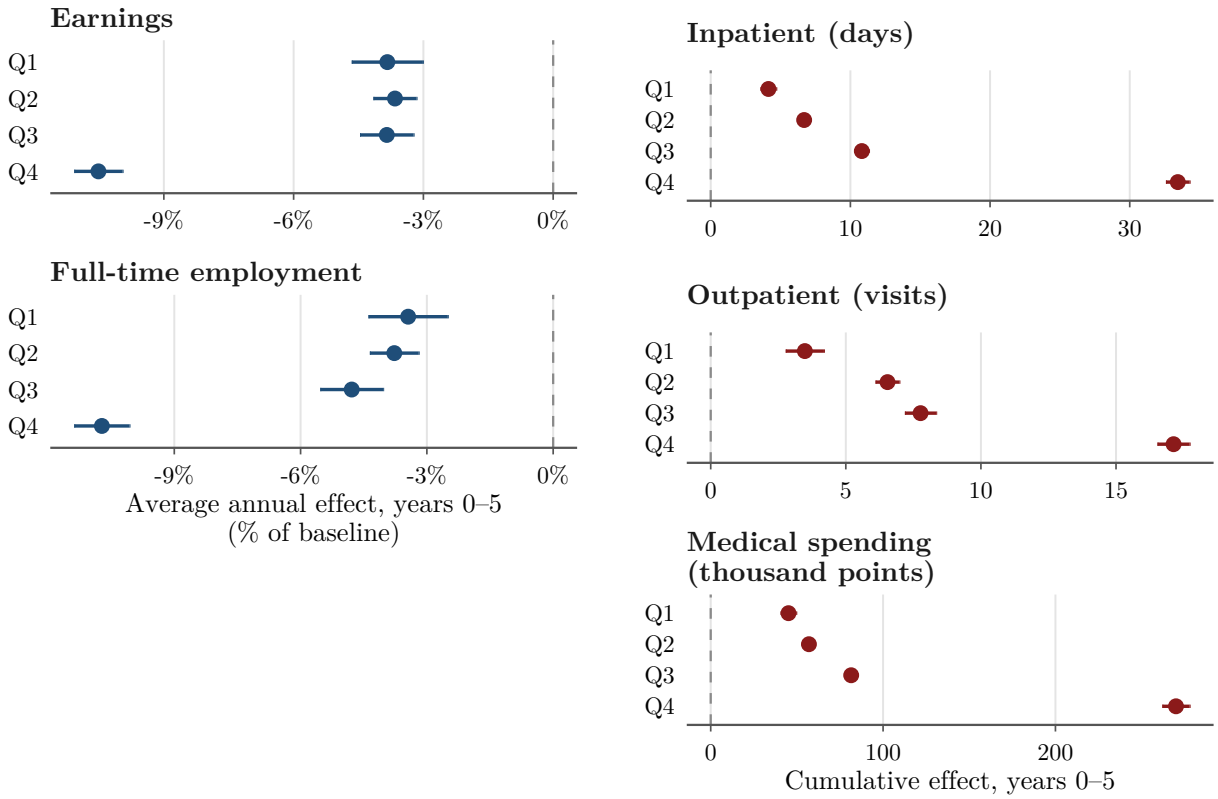


FIGURE OA17: Heterogeneity by Injury Severity Level

Aggregate DiD estimates by injury severity quartile. The left panel reports labor outcomes as average annual percentage changes over $t \in [0, 5]$. The right panel reports cumulative health care utilization over $t \in [0, 5]$. All estimates are reported with 95% confidence intervals.

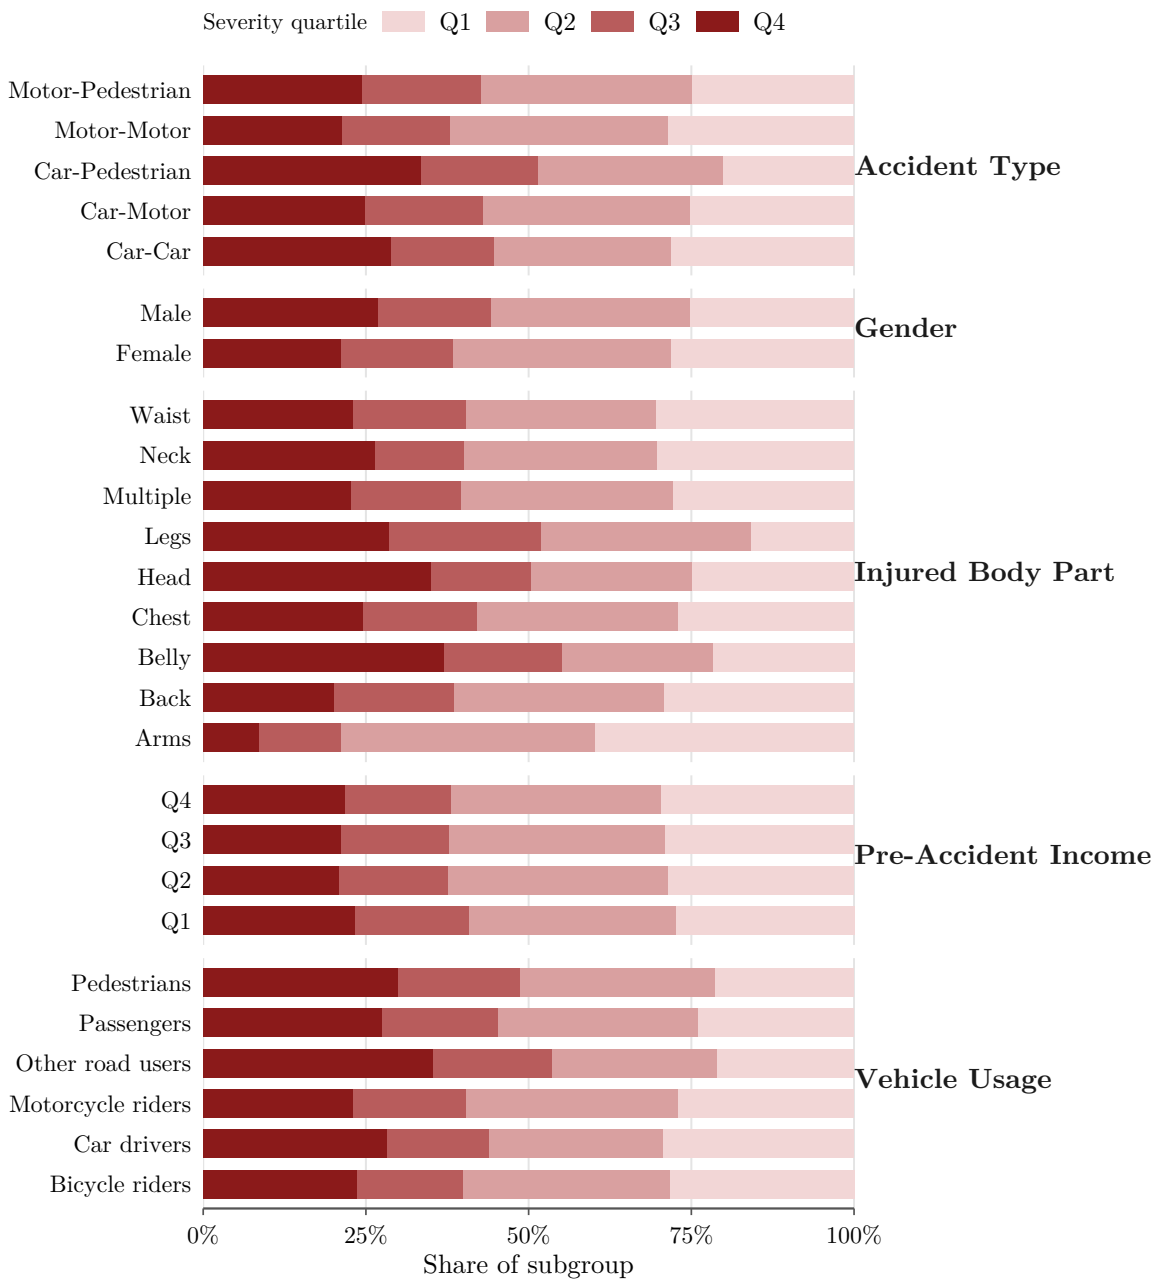


FIGURE OA18: Severity Composition by Subgroup

Notes: The figure reports the distribution of hospitalization-day severity quartiles within each subgroup used in the heterogeneity analysis. Differences in severity composition motivate the severity-standardized estimates reported in Figure OA19.

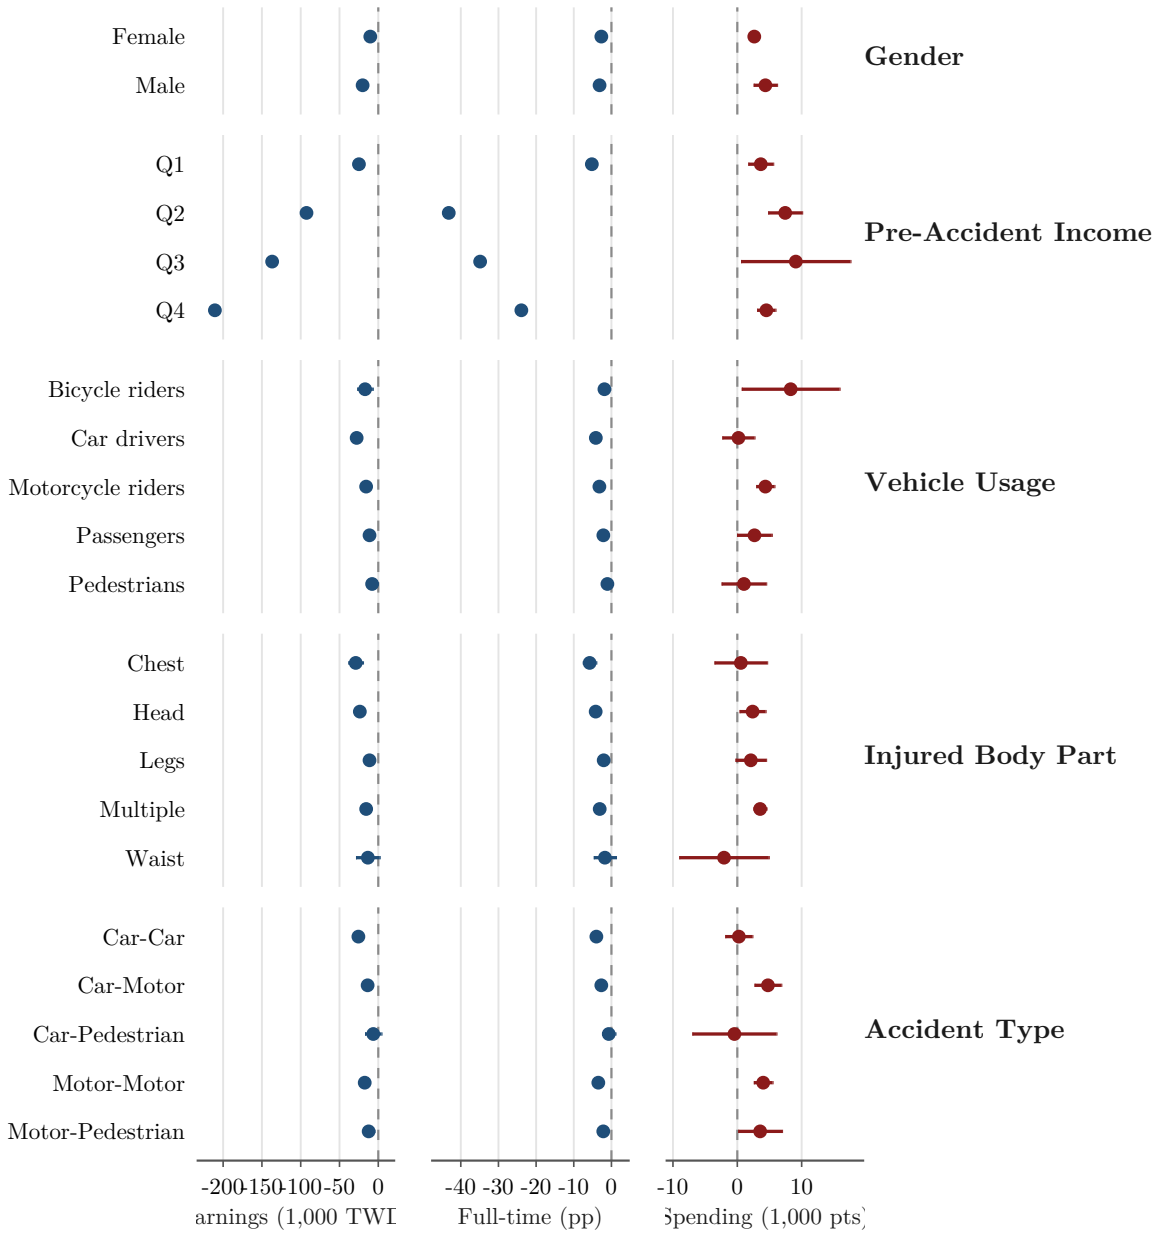


FIGURE OA19: Severity-Standardized Subgroup Effects at Event Time 5, All Dimensions

Notes: The figure reports year-five effects on earnings (left column, thousand TWD), full-time employment (middle column, percentage points), and NHI-covered medical spending (right column, thousand points) by subgroup, after standardizing each subgroup to the common main-sample distribution of hospitalization-day severity quartiles. Horizontal whiskers show 95% confidence intervals.

O.2 Heterogeneity by Pre-Accident Income Level

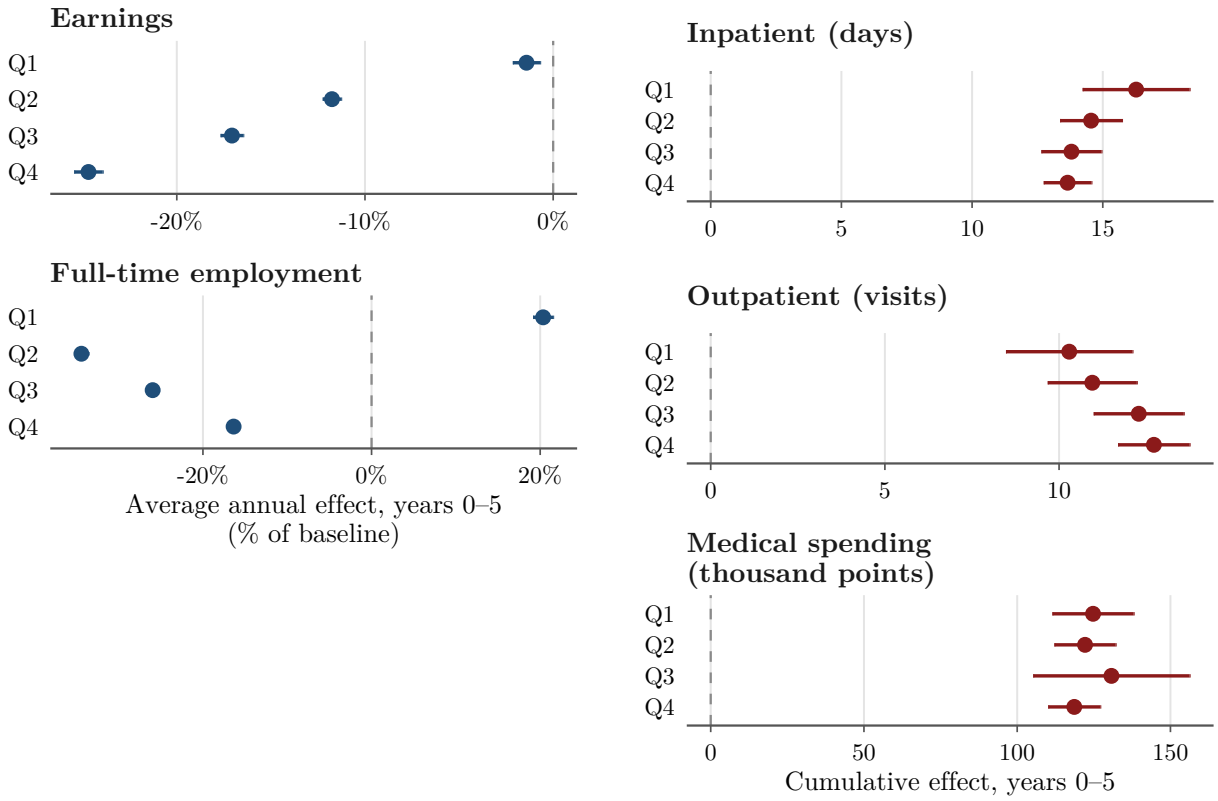


FIGURE OA20: Heterogeneity by Pre-Accident Income Level

Aggregate DiD estimates by pre-accident income level. The left panel reports labor outcomes as average annual percentage changes over $t \in [0, 5]$. The right panel reports cumulative health care utilization over $t \in [0, 5]$. All estimates are reported with 95% confidence intervals.

O.3 Heterogeneity by Injured Body Part

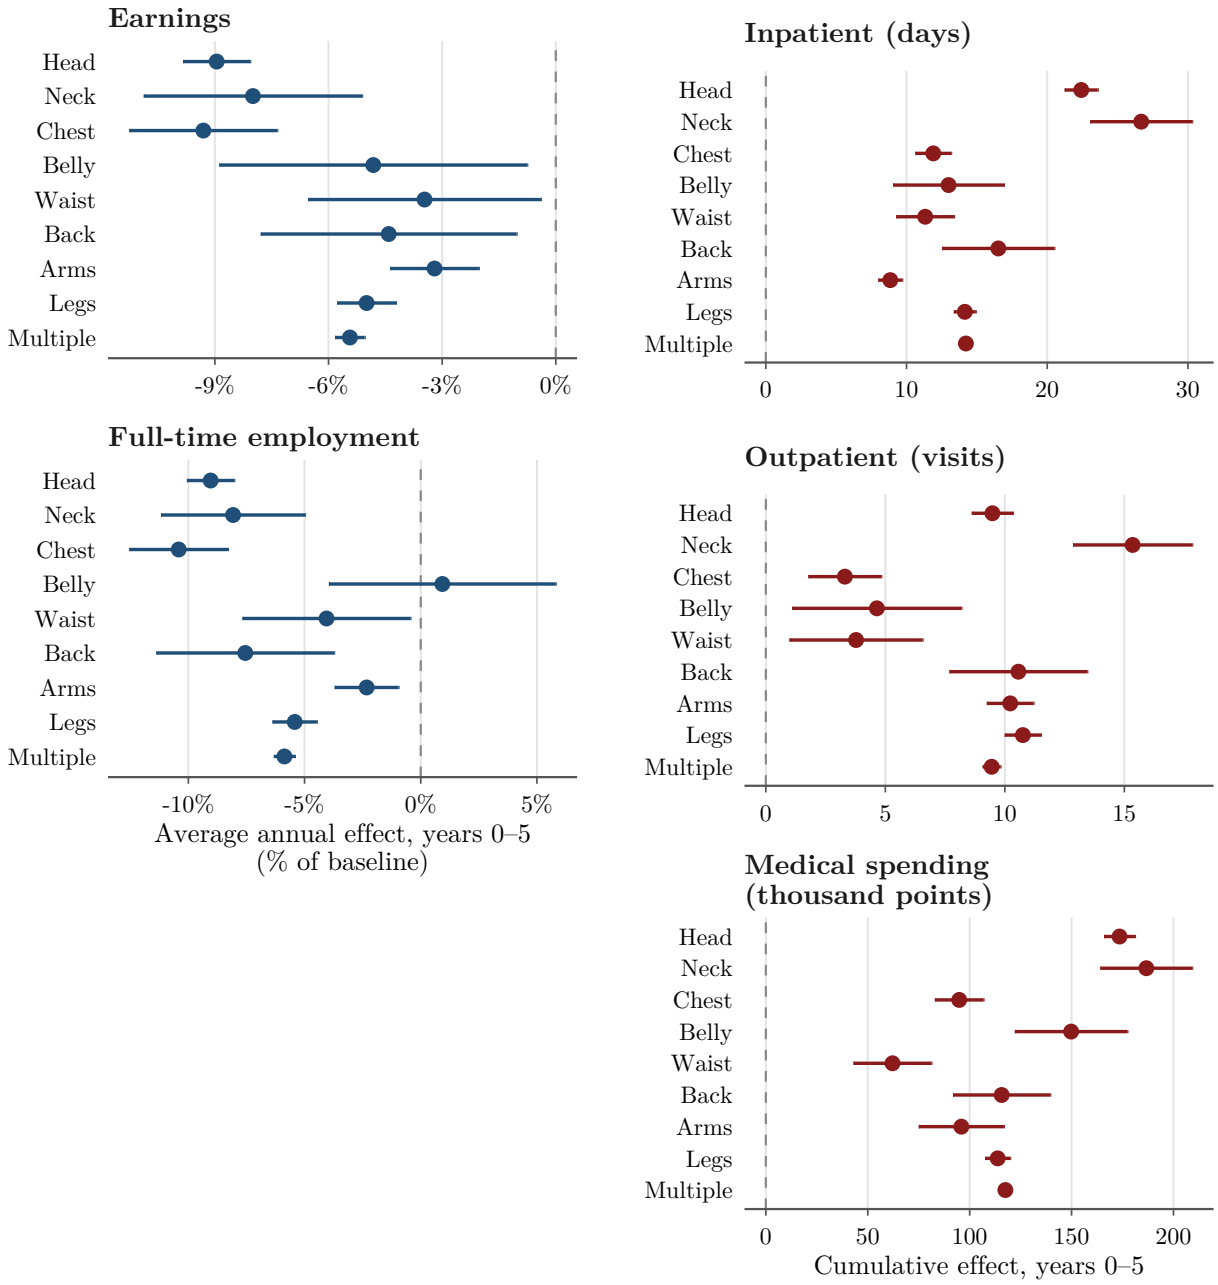


FIGURE OA21: Heterogeneity by Injured Body Part

Aggregate DiD estimates by injured body part. The left panel reports labor outcomes as average annual percentage changes over $t \in [0, 5]$. The right panel reports cumulative health care utilization over $t \in [0, 5]$. All estimates are reported with 95% confidence intervals.

O.4 Heterogeneity by Vehicle Usage

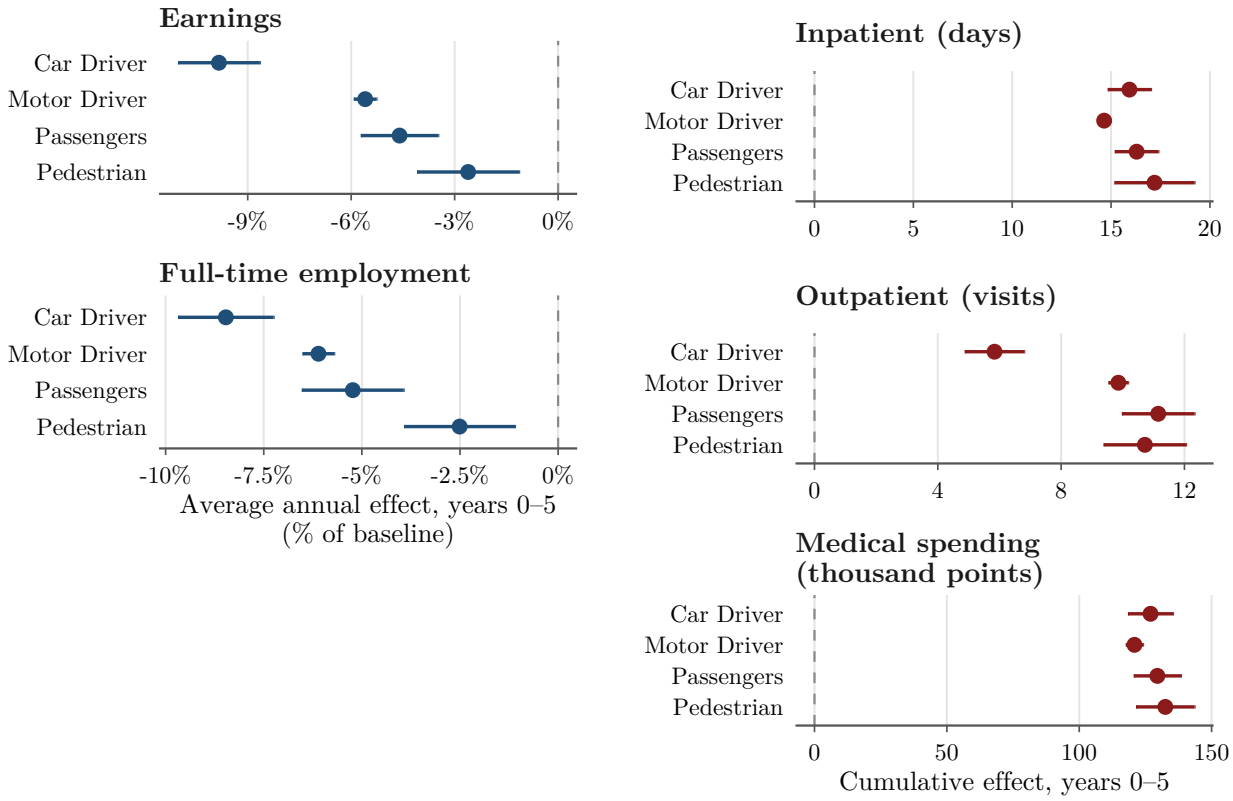


FIGURE OA22: Heterogeneity by Vehicle Usage

Aggregate DiD estimates by vehicle usage. The left panel reports labor outcomes as average annual percentage changes over $t \in [0, 5]$. The right panel reports cumulative health care utilization over $t \in [0, 5]$. All estimates are reported with 95% confidence intervals.

O.5 Heterogeneity by Accident Case Type

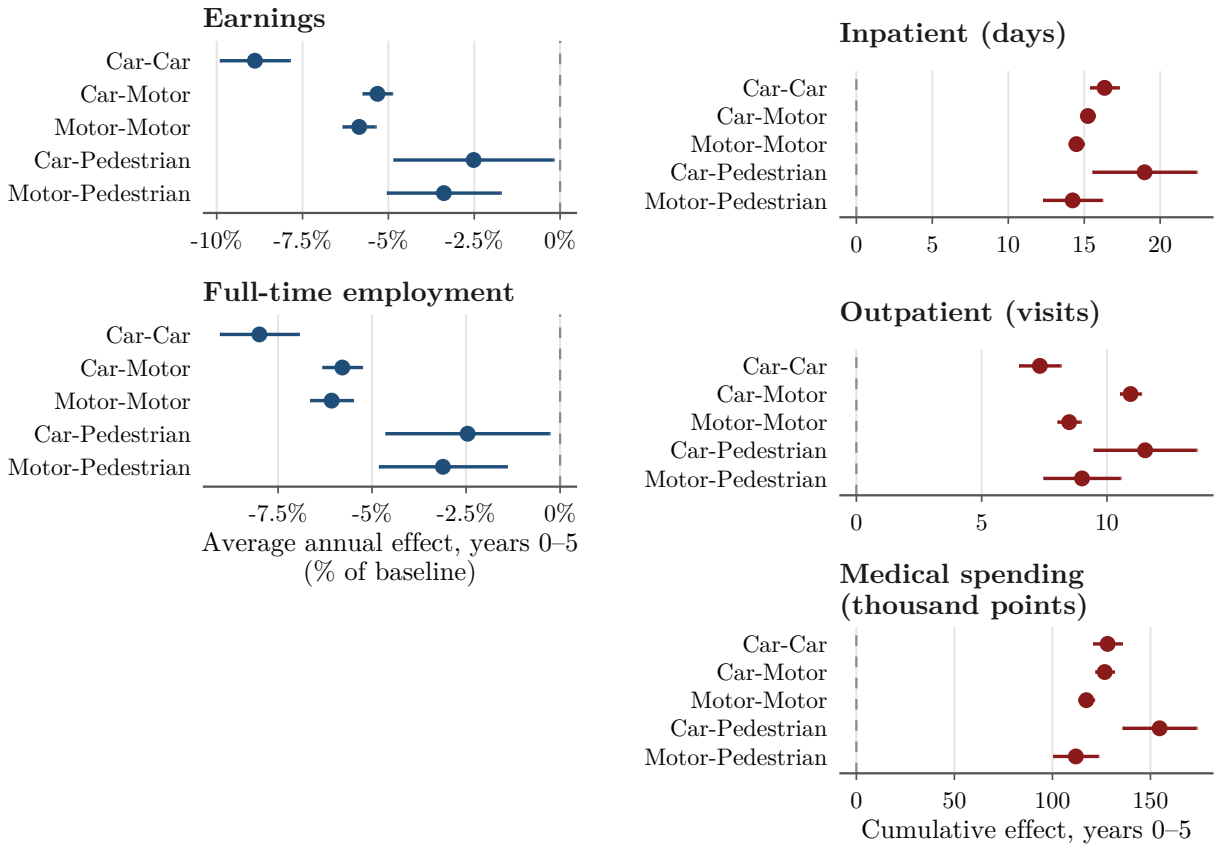


FIGURE OA23: Heterogeneity by Accident Case Type

Aggregate DiD estimates by accident case type. The left panel reports labor outcomes as average annual percentage changes over $t \in [0, 5]$. The right panel reports cumulative health care utilization over $t \in [0, 5]$. All estimates are reported with 95% confidence intervals.

O.6 Heterogeneity by Gender

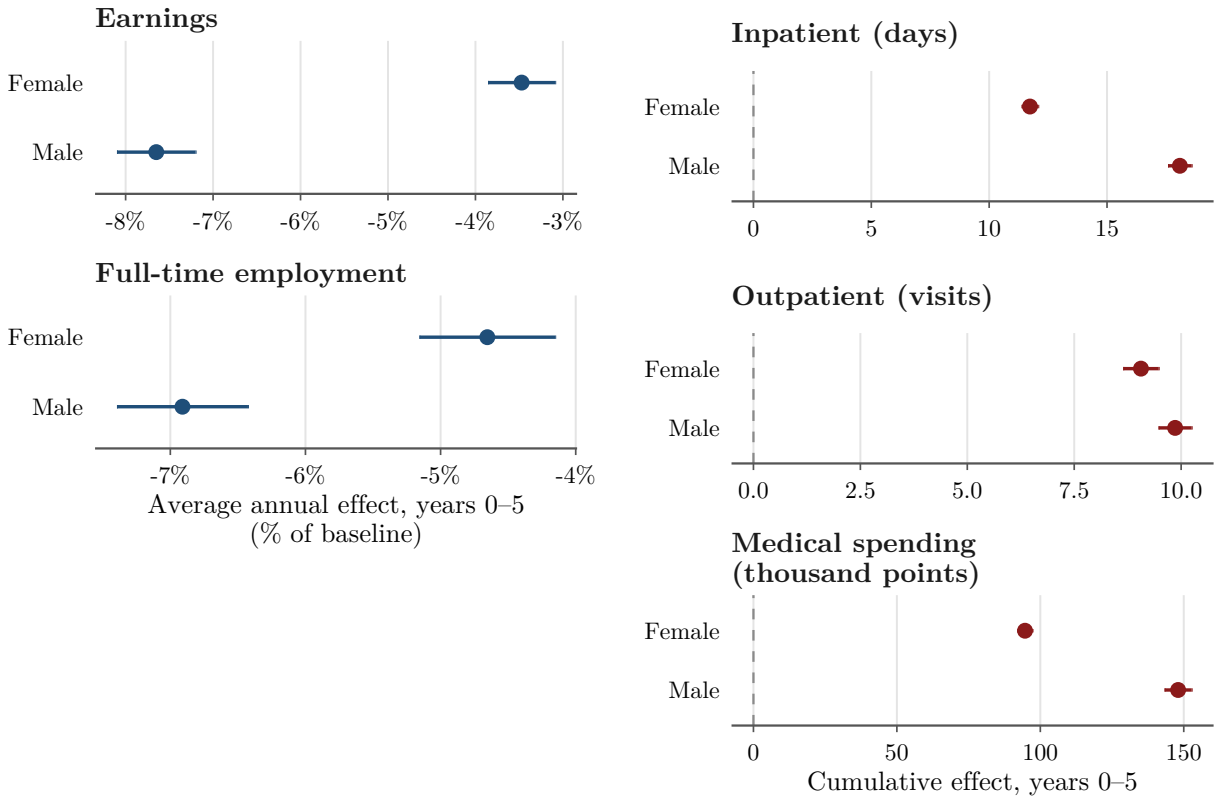


FIGURE OA24: Heterogeneity by Gender

Aggregate DiD estimates by gender. The left panel reports labor outcomes as average annual percentage changes over $t \in [0, 5]$. The right panel reports cumulative health care utilization over $t \in [0, 5]$. All estimates are reported with 95% confidence intervals.

P Accident-Level Summary Statistics

Table OA5: Summary Statistics of Traffic Accidents in 2003–2023

Sample: First-time Accidents in 2003–2023	Accident Type A1	A2
Panel A: Casualty Status		
People Involved	28,252	7,192,923
Mean	2.75	1.63
Number of Injuries	17,136	6,226,141
Mean	1.66	1.41
Number of Death	11,116	2,085
Mean	1.08	0.0005
People Hospitalized	6,915	966,782
Mean	0.67	0.21

Table OA5 (continued)

Sample: First-time Accidents in 2003–2023	Accident Type A1	A2
Panel B: Accident Cause		
Car and Car	31%	7%
Car and Motor	35%	50%
Motor and Motor	23%	38%
Car and Pedestrian	3%	2%
Motor and Pedestrian	8%	3%

Table OA5 (continued)

Sample: First-time Accidents in 2003–2023	Accident Type A1	A2
Panel C: Vehicle Usage		
Car Driver	31.8%	31.1%
Motor Driver	26.8%	54.1%
Bike Driver	1.8%	2.8%
Other Driver	1.8%	0.6%
Passenger	32.7%	8.4%
Pedestrian	4.8%	2.8%
Number of Accidents	10,263	4,412,102

Note: We present summary statistics for "first-time" accidents (at the incident level) occurring between 2003 and 2023. "First-time" accidents refer to the initial accidents retained after our data construction; incidents in which all individuals involved had already experienced two or more prior accidents are excluded. "A1" refers to traffic accidents in which at least one person died at the scene or within 24 hours. "A2" refers to traffic accidents in which at least one person was injured or died more than 24 hours after the crash. Panels A–C report casualty status, accident causes, and road-user roles, respectively.