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

Kuan-Ming Chen† Chen-An Lien‡

September 4, 2026

Abstract

Road-safety policy is evaluated largely through fatalities, yet most crash victims survive. We study how a nonfatal traffic injury, the most common acute health shock to working-age adults, translates into long-run economic losses. Linking the universe of police-reported crashes in Taiwan to 24 years of population health, employment, and mortality registries, we estimate dynamic effects with a staggered event-study design comparing earlier- to later-injured individuals. 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 losses operate almost entirely through the extensive margin and reflect the destruction of pre-accident worker–firm matches. Under an extreme accounting assumption, new catastrophic-illness certification could account for at most 63.9 percent of the year-five earnings loss. Workers with similar medical trajectories experience sharply different economic losses depending on pre-accident labor-market position. Combining the ten-year estimates with a transparent earnings extrapolation yields a twenty-year administrative-cost benchmark of approximately USD 16,700 per hospitalized working-age injury. Evaluations centered on fatalities therefore 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 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.

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 the most common acute health shock 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).

Human-capital accounting, the other common method, is built the same way. Analysts take a medical bill per injury class from discharge records, add lost production computed as a wage rate times an assumed number of work-loss days for that diagnosis, and, for the share of cases assumed permanently disabled, multiply out the remaining working years (Blincoe et al., 2015; Wijnen and Stipdonk, 2016). The inputs that determine the long-run cost are the least reliable: the U.S. inventory 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). Neither method follows the injured person, because discharge records, surveys, and compensation files do not record earnings and employment in the years that follow. Both therefore rest on an untested assumption: that the lasting economic cost of an injury is proportional to the lasting damage to health.

This paper asks where the economic cost of a hospitalized traffic injury comes from in the years after the accident, which directly tests the assumption that cost is proportional to lasting health damage. One possibility is continuing medical need and reduced capacity to work. The other is the labor market: workers whose health recovers but who leave full-time employment and do not return. The distinction matters for both policy and valuation. If the cost reflects lasting incapacity, the relevant responses are rehabilitation and disability insurance, and severity-weighted valuation is approximately right. If the cost instead falls on workers who recover medically but do not return to full-time work, the relevant responses are return-to-work support and employment protection, and severity weights mismeasure the cost by construction, since they assign the smallest weights to the injuries that recover best. Separating the two requires observing medical recovery, employer identity, earnings, and survival for the same individuals over many years.

We link the universe of police-reported traffic accidents in Taiwan (2003–2023) to 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 24-year 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 backloaded, and it consists of workers who are no longer in full-time work, not workers who earn less while working.

Our second finding characterizes this detachment and bounds how much of it lasting incapacity can explain. Separations from the pre-accident employer are concentrated in the months after the crash and decline smoothly afterward. They do not bunch when statutory wage-continuation protection expires at two years, so the timing of detachment follows the injury and not the protection calendar. Five years later, injured workers are still 6 percent less likely to work for their pre-accident 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. Detachment also outlasts recovery. 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. Lasting incapacity accounts for a bounded share of the aggregate loss. 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. If we attribute their entire

loss to incapacity, certified incapacity accounts for roughly one quarter of the aggregate year-five earnings loss; under the extreme assumption that every certified worker loses all pre-accident earnings, the share is at most 64 percent. **[TODO: Verify the one-quarter figure: 6.1 percent times the certified-subsample year-five ATT of TWD 67,978, divided by the main-sample ATT of TWD 16,280, gives 25.5 percent; confirm that the certified share and the subsample ATT refer to the same balanced sample.]** Cumulative mortality rises by about 1.6 percentage points by year five, against 1.34 percent in the comparison group, so survivors are not healthy on every margin. Most of the economic loss, however, 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 TWD 211,000, against TWD 25,000 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. The same health shock produces very different economic losses depending on the labor-market position it disrupts.

These results change how a nonfatal injury should be priced. Combining the ten-year estimates with a transparent extrapolation yields a twenty-year administrative-cost benchmark of about TWD 502,000 (USD 16,700) per hospitalized working-age injury, or USD 13,100 at a 3 percent discount rate. The benchmark excludes pain, suffering, informal care, and other private burdens, and it is anchored in observed trajectories. Its composition matters more than its level. The dominant component is a labor-market loss borne largely by workers who recover medically, so valuations that scale fatality values by medical severity weights understate exactly the injuries that leave health intact and working lives disrupted. 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.

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. This literature establishes that losses persist, but each study observes one side of the process. Labor registers document earnings losses but cannot trace medical recovery. 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. 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 certified incapacity 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.

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 detachment evidence and incapacity accounting, heterogeneity, and mortality. Section 6 translates the estimates into an administrative-cost benchmark, and Section 7 concludes.

2 Background

2.1 Valuation Framework

Road-safety policies rely on cost-benefit analysis to prioritize investments in infrastructure and enforcement, making estimates of the economic costs of traffic accidents central to policy evaluation. Existing methods disagree widely. VSL and WTP approaches scale fatality values to nonfatal injuries using severity weights, quality-adjusted life years, or stated-preference responses, while human-capital and accounting approaches monetize medical expenditures and imputed productivity losses from compensation schedules and short-run records (Ashenfelter, 2006). Within Taiwan, hedonic-wage VSL estimates span roughly NT$5–50 million (Liu et al., 1997; Liu and Hammitt, 2004),[1] official per-fatality accident costs based on compensation data are roughly NT$2.7–3.0 million (Chou et al., 2019), and for nonfatal injuries, WTP surveys imply about NT$0.6 million per injured individual while accounting estimates from court-verdict data are closer to NT$1.0 million. 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 we document below.

2.2 Research Context

2.2.1 Reporting and Classification of Traffic Accidents

When a traffic accident occurs in Taiwan, the involved parties typically report to the police as insurance companies generally require a police report to initiate the claims process. When a report is made, officers from the nearest police station are dispatched to the scene to document the incident, including the identities of those involved, road and weather conditions, and any apparent injuries.

The police classify traffic accidents into three categories: A1, A2, and A3, which are based on the severity of personal injury and property damage. A3 accidents involve only property damage and no physical injuries or fatalities. In contrast, A2 accidents involve at least one injured person but no deaths at the scene or within 24 hours. An accident is classified as A1 if any involved individual dies within 24 hours of the event.

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

This reporting system creates a standardized administrative trail for accidents involving personal injuries or deaths, particularly in A1 and A2 cases. Importantly, this police-generated trail provides a record of the health shock that is independent of the individual's subsequent medical-seeking behavior or hospital admission status. Once documented by the police, injured individuals typically enter the healthcare system, where the NHI program plays a central role in facilitating medical access and recording treatment episodes.

2.2.2 Health Coverage and Compensation Mechanisms

Once an accident has been reported and recorded by the police, injured individuals typically enter the medical system, where the question shifts from documenting the event to determining who ultimately bears the cost of treatment. Two institutional mechanisms are central to this allocation in Taiwan: the NHI system and the Compulsory Automobile Liability Insurance (CALI) program.

Taiwan's NHI is a single-payer, compulsory health insurance system that covers the entire population, allowing injured individuals to receive medical care immediately after a traffic accident even before legal liability is determined. According to MOHW (2022), total medical expenditures associated with traffic accidents amounted to approximately NT$8.94 billion in 2022, of which the NHI directly covered about NT$6.65 billion (roughly 74%). Most residual spending is covered by the CALI program, Taiwan's mandatory liability insurance for vehicle owners, which provides standardized coverage for medical expenses and disability or death benefits up to statutory limits. Together, NHI and CALI form the financial infrastructure supporting injured individuals: NHI functions as the primary payer for medical care, while CALI serves as a secondary compensation layer that distributes remaining costs between compulsory and private insurance.

Beyond immediate medical care, the most severely injured individuals may qualify for ongoing income support and disability services. As of 2023, traffic-related injuries accounted for approximately 3% of all certified disabilities in Taiwan, reflecting the subset of accident injured individuals whose impairments are severe enough to meet formal disability criteria.

For those who suffer complete loss of working capacity, CALI provides up to NT$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.[2]

[2]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

2.2.3 Employment Protection after Work-Related Injuries

Beyond medical and disability compensation, Taiwan's labor regulations provide an initial buffer against immediate job loss, which helps explain the gradual nature of the labor market detachment we observe. For traffic injuries that qualify as occupational injuries, including commuting accidents that occur along a standard route, the Labor Standards Act mandates strict 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. Consequently, for a substantial subset of injured workers, the formal employment relationship is legally preserved in the immediate aftermath of the crash.

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. To provide a benchmark for international readers, these provisions share conceptual similarities with protections in the United States, where the Family and Medical Leave Act (FMLA) grants up to 12 weeks of unpaid, job-protected leave, and the Americans with Disabilities Act (ADA) requires reasonable workplace accommodations (U.S. Department of Labor, 2025).

This institutional structure acts as a temporary shock absorber. However, as statutory grace periods expire or practical accommodation challenges prove insurmountable, the formal employment relationship often severs. Critically, institutional compensation schedules are based on physical impairment grades, largely ignoring the specific worker-firm match destruction that drives much of the economic penalty. These compensation mechanisms and legal protections provide partial, short-term income protection but ultimately leave substantial residual losses, particularly the widening earnings gap, borne by injured individuals and their families.

3 Data

Our empirical analysis links two administrative data sources. We use police records to define traffic-injury events and National Health Insurance records to measure medical utilization

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.

and labor-market outcomes. We then construct the person-level analysis sample around each individual's first recorded accident. This section describes each step in turn.

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 A1 and A2 accidents during this period, while minor A3 cases are excluded. Because reporting is closely tied to insurance claims and legal documentation, as discussed in Section 2.2, NPA records provide near-universal coverage of medically relevant accidents. As noted earlier, we define accident exposure from these NPA records rather than from NHI claims because diagnosis-based transport-injury codes recover only about 40% of police-reported A1 and A2 accidents in our linked data. 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 and contextual information, including the date and location of the accident, the presumed cause recorded by police officers, and other accident characteristics. The individual-level file contains one record for each person involved in an accident and provides the case identifier, encrypted individual identifier, gender, vehicle type, road-user role (e.g., driver, passenger, pedestrian), injury status, and anatomical location of injury. Linking these files by case identifier allows us to construct a person-level panel of accident involvement. The resulting dataset includes over 5.45 million unique traffic accident cases and more than 7.3 million individuals involved.

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 NHIRD provides comprehensive enrollment and claims data for the entire Taiwanese population from 2000 to 2023, enabling the measurement of employment status, earnings, and medical utilization.

The NHIRD contains two major files: The enrollment file and the claim file. The enrollment files span the years 2000 to 2023 and cover more than 23 million individuals annually. These records include basic demographic characteristics such as sex, date of birth, and residential township. In addition, they provide labor-relevant information including insured salary, employment status, and firm identifier. We deflate all salary and earnings measures using Taiwan's GDP deflator and express them in 2021 New Taiwan Dollars. Employment

status is updated monthly, enabling us to observe job separations, employment spells, and transitions across firms.

In parallel, the NHIRD includes claims files that record all reimbursed inpatient, outpatient, and emergency services. These data contain ICD-coded diagnoses, treatment procedures, prescribed medications, dates of service, and itemized expenditures. The claims records allow us to construct rich measures of health care utilization, including the intensity and timing of medical treatment episodes.

These files from NHIRD allow us to track individuals on a monthly basis, covering both labor market and health trajectories. This longitudinal scope is central to our analysis, as it enables the construction of pre- and post-accident outcome measures with precise timing and granularity.

3.3 Sample Construction

We construct our main sample to include all individuals ever involved in police-reported accidents. To concentrate on survivors with potentially recoverable outcomes, we exclude individuals who died within 30 days of the accident, using mortality information from the NHI death registry. We also retain only each individual's first recorded accident to avoid confounding effects from repeated exposures. This yields a cohort of approximately 5.2 million unique accident-injured individuals. See Appendix A for a summary of accident-level characteristics among these first-time accident cases.

To understand how these ever-accident people compared with other individuals, we supplement the main accident-injured sample by constructing a matched comparison group who had not yet experienced a recorded traffic accident. For each injured person, we select one non-injured counterpart from the same calendar year, matched on age (in years), gender, and residential district (at the zip code level). 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. To ensure the reference set represents individuals prior to any accident involvement, we restrict eligibility to those whose assigned year precedes any later accident they may experience. This procedure yields approximately 3.8 million matched individuals.

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 (NT$/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 matched pool additionally requires a hospitalization record in that year. These statistics highlight baseline differences between injured and never-injured individuals prior to the accident; they are not intended as the basis for our main identification strategy. One may notice that the slight shortfall in the number of non-injured set reflects practical constraints in matching. For instance, it is sometimes impossible to locate a same-year match for very old individuals in sparsely populated districts. Our final dataset comprises roughly 7.7 million individuals: 3.9 million with accident-related injuries and 3.8 million matched non-injured comparisons.

4 Method

4.1 Estimation Target and Naïve Comparison

Next, we present the summary statistics of the sample we constructed in Table 1. For the sake of comparison, we also present the summary for the matched never-accident sample. The table presents summary statistics for key outcomes in the year prior to the accident (or pseudo-accident year for never-injured individuals), across matched samples.

Since we are focusing on the costs of injured individuals from traffic accidents on the labor market, we gradually restrict our focus 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 matched to a corresponding control group (*Non-Injured*), constructed from individuals with the same age, gender, and residential location in the accident year. For the hospitalized subsample, the matched group is drawn from individuals aged 25–65 who were hospitalized in the corresponding year, ensuring comparability in the nature of the health shock.

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.[3] In contrast, average annual earnings are markedly lower for future injured individuals (NT$229,110) than for the never-injured group (NT$248,344). This earnings gap persists even when conditioning on being employed, suggesting that differences arise primarily on the intensive margin. This pattern may reflect sorting into different occupations, with those in higher-risk professions, such as transportation or manual labor, being more exposed to traffic accidents.

In addition to these labor-market differences, health utilization patterns also reveal divergence between groups prior to 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 control group exhibits higher hospitalization days even before the accident year. This reflects the fact that

[3]These employment rates are somewhat lower than Taiwan's overall employment rate, reflecting two key features of our data. First, our sample includes younger individuals who have not yet entered the labor force. Second, our employment measure is based on NHI employment records, which exclude self-employed workers, certain agricultural workers, and employees of irrigation associations, among other categories not covered by standard employer contributions.

hospitalization is itself a serious health event; hence, for individuals hospitalized due to an exogenous shock like a traffic accident, there is no corresponding signal of hospitalization in the pre-accident year.

Taken together, these findings suggest that individuals who are eventually injured in traffic accidents are not directly comparable to those who are never injured, even after controlling for basic demographics. Although traffic accidents are often viewed as exogenous shocks, these comparisons show that the margin of ever experiencing an accident is systematically related to pre-accident labor-market and health characteristics.

4.2 Identification Strategy

Our goal is to identify the dynamic causal effect of a first traffic injury on subsequent labor-market and health outcomes. Motivated by the pre-accident differences documented above, our identification strategy uses a later-treated comparison group, comparing individuals injured in a given year to individuals injured in later years. Treatment is defined as an individual's first traffic accident injury in calendar year g, which determines the treatment cohort. Event time t is measured relative to g, with $t < 0$ denoting pre-accident years and $t > 0$ post-accident years. For cohort g at event time t, the comparison group consists of individuals who have not yet been injured but will experience their first accident in a later year.

The key identifying assumption is that, absent the accident, outcomes of treated and later-treated individuals would have followed parallel trajectories. Conditional on the included covariates, this assumption is plausible if the *timing* of an individual's first accident is as good as random. The identifying assumption is weaker than random assignment of accident ever-occurrence, but still requires that earlier- and later-injured individuals would have followed parallel outcome paths absent injury. While the probability of ever being injured may correlate with risk preferences, occupation, or exposure, our design assumes that, conditional on covariates and eventual accident involvement, the timing of first injury is not driven by differential potential outcome trends. Under this quasi-random timing assumption, later-treated individuals provide a credible counterfactual. The next subsection describes how we implement these comparisons using the Callaway and Sant'Anna (2021) framework.

We stress-test the quasi-random timing assumption directly. First, Figure 1 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 for small imbalance. Second, the randomized accident-year placebo in Appendix J shows that the cohort and calendar-time structure of the design alone reproduces none of the estimated dynamics. Third, 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 the R^2 to 0.1191; adding their 2000–2002 linear trends raises it to 0.1197. Thus, predetermined outcome levels and trends add only 0.0026 to explained variation beyond demographics. Although they are jointly statistically 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. This small incremental predictive content supports the conditional, rather than unconditional, quasi-random timing assumption used in the design.

By contrast, comparisons against the never-injured are systematically imbalanced even before the accident (Appendix B), which is why our design does not rest on them.

4.3 Estimation Framework

We estimate group-time average treatment effects, $\mathrm{ATT}(g,t)$, using the later-treated comparison group described above and the doubly robust estimator in Equation (2.2) of Callaway and Sant'Anna (2021). Our main specification controls for birth year to account for life-cycle differences across accident cohorts. We aggregate the cohort-specific estimates by event time to recover the overall dynamic treatment profile.

Given potential treatment effect heterogeneity, our main analysis focuses on a balanced event-time window from $t = -5$ to $t = +5$. This choice ensures that comparisons at each relative time are based on the same set of cohorts, thereby avoiding changes in sample composition that could arise if the window were allowed to vary across event times. This restriction enhances comparability across periods and improves the credibility of pre-trend assessments. Estimates using the full available event window are reported in the Appendix E.

Standard errors are based on individual-level influence functions, thereby accounting for within-person dependence across annual observations. We apply this inference procedure uniformly to the dynamic estimates reported in the main tables and figures. Because the age restriction is imposed before enforcing panel balance, the estimation sample includes

TABLE 2: 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.

only individuals observed between ages 25 and 65 in every sample year; sample composition therefore remains fixed across event times.

5 Results

We organize the empirical results around three questions. First, how long do the consequences of a nonfatal traffic injury last, and how much do they cost? We trace the average dynamic trajectories of public health expenditure, medical utilization, earnings, and employment. Second, what converts an acute health shock into a persistent economic loss? We examine employment margins, attachment to the pre-accident employer, and workers whose measured health recovers, assembling the evidence for match destruction as the operative mechanism. Third, who bears the largest dynamic losses? We explore heterogeneity

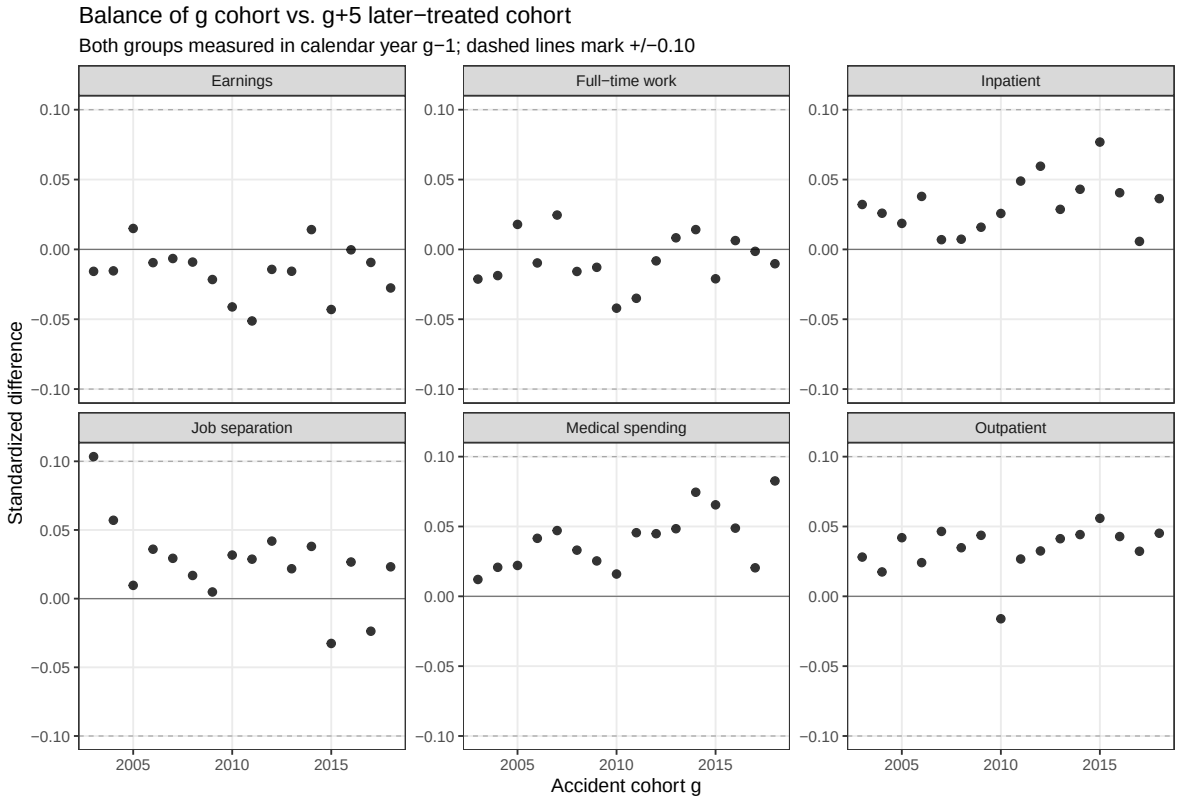
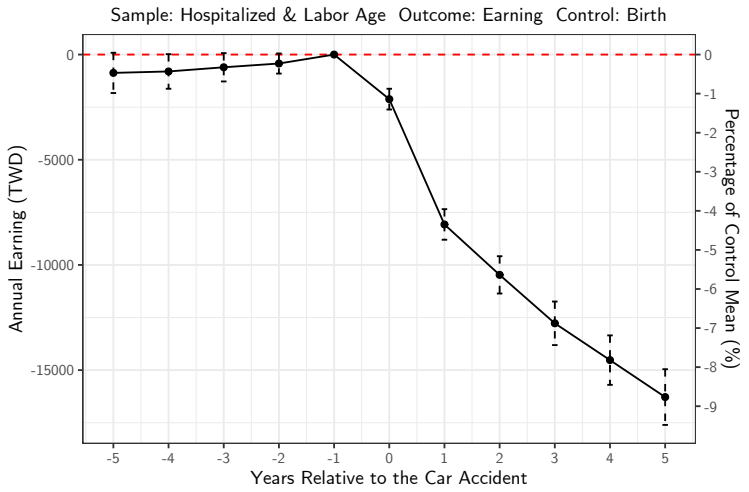


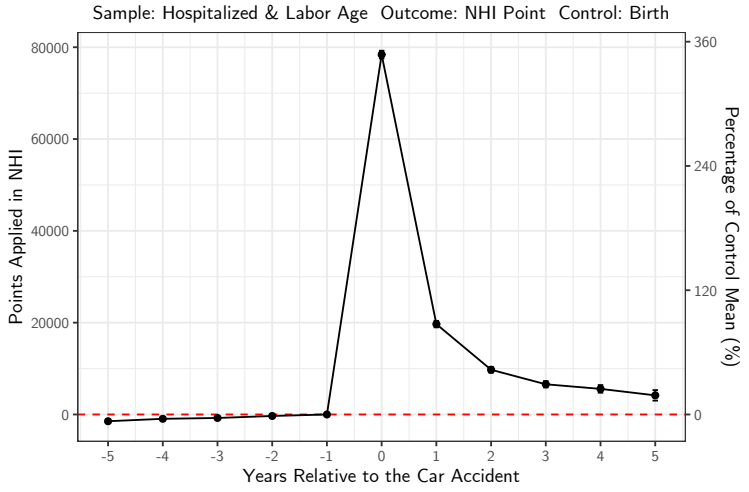
FIGURE 1: Cohort-by-Cohort Standardized Differences: Treated vs. Later-Treated

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.

across injured individuals, distinguishing costs that scale with physical severity from those that depend on workers' pre-accident labor-market positions. We then document mortality effects, report robustness checks, and translate the estimates into an administrative-cost benchmark.



(A) Earnings (birth-year controls)



(B) NHI points (birth-year controls)

FIGURE 2: Dynamic Effects of Traffic Injuries On Annual Earning and Health Expenditure

Notes: Each panel reports cohort-time estimates using later-treated controls with a balanced event window of $[-5, 5]$ years. Whiskers 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 2021 TWD using the GDP deflator. Medical spending is measured in NHI points, a standardized reimbursement unit under Taiwan's NHI (≈NT$0.9 per point).

5.1 Average Dynamic Effects: Labor-Market Detachment and Medical Utilization

To understand the trajectory of a nonfatal injury, we jointly observe the medical and economic outcomes of survivors. Figure 2 presents the dynamic effects on our two primary outcomes: earnings and public health expenditures.

Traffic injuries generate persistent labor-market losses. The estimates show little evidence of differential pre-trends over the five-year pre-period. Notably, the earnings effect is relatively modest in the accident year. Accidents occur throughout the calendar year, so accident-year earnings mechanically mix pre- and post-accident months, diluting the measured $t = 0$ effect. Rather than recovering as the initial shock passes, the decline in earnings widens steadily from the first post-accident year onward, consistent with return-to-work frictions delaying rather than eliminating employment losses. By year five, annual earnings are approximately 9 percent below the pre-accident control mean. Full-time employment follows a similar trajectory. It declines modestly immediately after the accident, continues to fall throughout the five-year window, and is also roughly 9 percent lower by year five.

These labor-market dynamics contrast sharply with the health outcomes, which have a different time profile. The medical response begins with a large acute-care intervention that scales down significantly but does not fully disappear. In the main text, we capture this trajectory primarily through public health expenditure. Under Taiwan's universal health insurance and low-copayment system, this government spending serves as a measure of actual healthcare utilization and the severity of the underlying health shock. Total public health expenditure increases by approximately 80,000 points at $t = 0$, a roughly 360 percent increase relative to the pre-accident control mean. Spending falls substantially after the initial spike but remains above baseline throughout the five-year horizon.

The specific healthcare utilization metrics underlying this spending are reported in Appendix Figure 10d and Figure 10e. Inpatient days increase by roughly 10 days on average in the accident year, reflecting the initial hospitalization episode, but decline sharply after $t = 0$. Outpatient visits rise by approximately 42 percent in the first post-accident year and remain elevated through $t = 5$, indicating continued demand for follow-up care. The diagnosis-based decomposition in Appendix D further explains this persistence: acute trauma care accounts for about 90 percent of the accident-year increase, but by the first post-accident year, the trauma share falls to 36 percent, while long-term sequelae management rises from 8 percent to 36 percent. Post-accident medical costs therefore shift in composition toward follow-up care rather than reflecting a prolonged acute crisis.

5.2 The Mechanism: Destruction of Worker–Firm Matches

The divergence between medical recovery and widening economic losses points to a mechanism operating through the labor market rather than through lasting incapacity. This subsection assembles the evidence. The losses are concentrated on the extensive margin;

they run through severed attachment to the pre-accident employer; and they persist among workers whose measured health has recovered. A hospitalization, in this reading, functions as an idiosyncratic displacement event (Jacobson et al., 1993), and its cost depends on the specific employment relationship it destroys (Jäger and Heining, 2022; Johnsen et al., 2024).

5.2.1 Employment Margins

The earnings decomposition asks whether earnings losses operate through reduced employment or lower earnings conditional on work. We evaluate the extensive margin directly in Appendix C and a decomposition in Appendix O, where the outcome variable is an indicator for holding full-time employment. After the first post-accident year, the reduction in this full-time employment indicator accounts for roughly 80 to 95 percent of the annual earnings loss, while the conditional-earnings component is much smaller. The persistent decline in earnings operates primarily through the extensive margin: workers are dropping out of full-time employment entirely rather than merely working fewer hours or taking pay cuts upon returning. Lee bounds reported in Appendix O address selection into post-accident employment and confirm that the extensive margin remains quantitatively central.

5.2.2 Firm Attachment

We further unpack this extensive-margin response by examining job separation and firm attachment in Appendix C. Job separation spikes sharply in the accident year but returns close to baseline afterward, suggesting a concentrated disruption rather than a continuing sequence of layoffs. However, this initial friction alters the employment relationship for a subset of workers. Firm attachment, measuring whether injured workers remain connected to their pre-accident employer, remains persistently lower. Five years after the accident, injured individuals remain roughly 2 percentage points less likely to be with their pre-accident employer, corresponding to a 6 to 7 percent reduction relative to the control mean. The temporary absence therefore disrupts pre-accident employment relationships in ways that persist well after the immediate physical recovery.

Which matches are destroyed is informative about the mechanism. Figure 3 splits year-five labor-market effects by pre-accident tenure and pre-accident firm size among workers employed before the accident. Full-time employment losses are largest for workers with the shortest tenure and for workers at the smallest firms. These are the matches with the least accumulated specific capital, at the employers least able to accommodate an extended absence. Earnings losses are larger in absolute terms for longer-tenured workers, but this

gradient reflects their higher baseline earnings. Relative to each tenure group's own mean earnings at $t = -1$, the year-five losses are 30.9 percent for workers with one year of tenure, 28.9 percent for those with two to three years, and 28.0 percent for those with four to five years. Longer-tenured workers therefore lose more TWD, but slightly less in proportional terms.

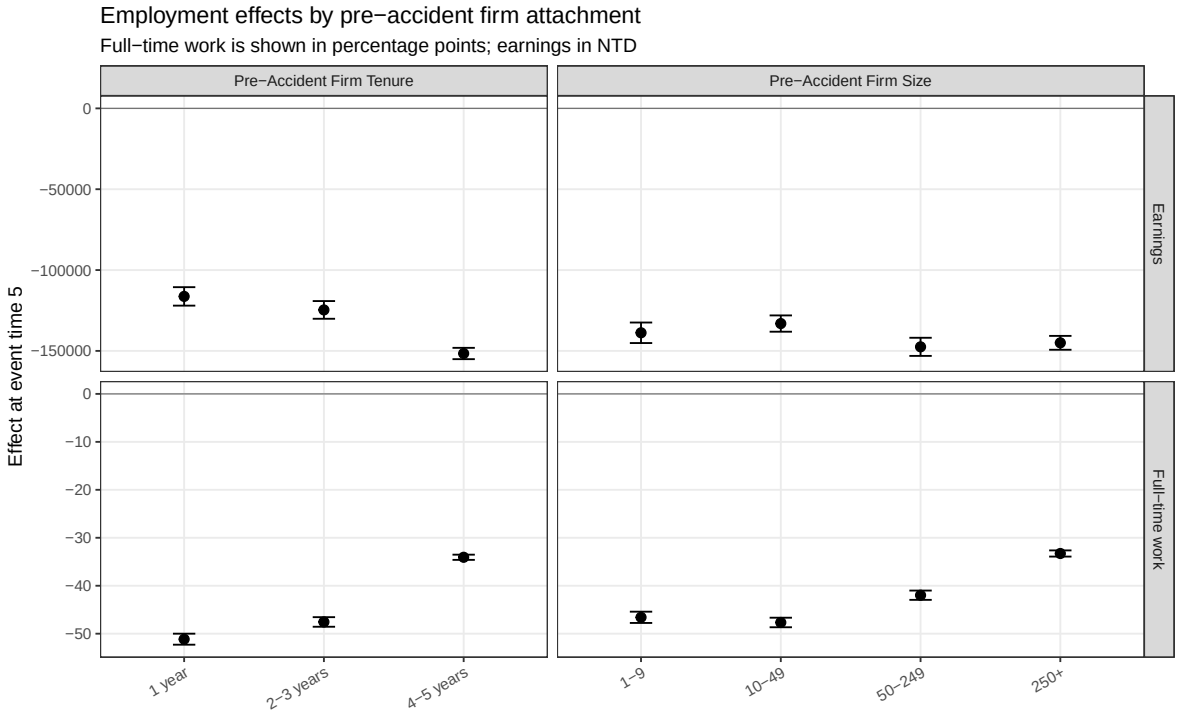


FIGURE 3: Year-Five Labor-Market Effects by Pre-Accident Firm Attachment

Notes: The figure reports event-time-5 effects on earnings (TWD, top row) and full-time employment (percentage points, bottom row) by pre-accident firm size (left column) and pre-accident firm tenure (right column), among workers employed in the year before the accident. Whiskers show 95% confidence intervals.

5.2.3 Medical Recovery Does Not Deliver Economic Recovery

If lasting incapacity were the operative constraint, losses should be concentrated among workers whose health fails to recover. We find the opposite. 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 controls (at or below the control 75th percentile in the same calendar year); 67.1 percent 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 TWD 12,915 and the full-time

employment loss is 2.37 percentage points, retaining 79 and 77 percent of the corresponding main-sample effects. Because post-accident utilization is itself a post-treatment outcome, we present this split as corroboration rather than as a causal decomposition; the selection runs against us, however, since strictly-recovered workers are, if anything, positively selected on both health and labor-market prospects.

Among all injured individuals contributing to the year-five estimate, 6.1 percent first received catastrophic-illness certification in the accident year or the following year. Their mean nonnegative annual earnings in event year -1 were TWD 170,628. To construct a deliberately conservative accounting bound, we attribute the entire baseline earnings of every newly certified worker to lasting incapacity, equivalent to TWD 10,402 per injured individual. Relative to the observed year-five earnings loss of TWD 16,280 per injured individual, new catastrophic-illness certification could therefore account for at most 63.9 percent of the aggregate loss. This is an upper-bound accounting exercise rather than a causal decomposition: certification is measured after the accident, and it captures formally recognized catastrophic illness rather than every possible form of lasting impairment.

5.2.4 Separation Dynamics

We next describe the timing of match dissolution among 60,465 workers attached to an identifiable employer one month before the accident. Figure 4 plots the monthly hazard of first leaving that employer over months 12–36. The hazard is 1.23 percent in month 24, compared with 1.28 percent when months 18–23 are pooled and 1.16 percent when months 25–30 are pooled. The series changes smoothly through month 24 and displays no discrete two-year discontinuity.

A descriptive stayer–separator decomposition shows that severed matches account for most of the subsequent earnings decline. Among these employer-attached workers, 36.1 percent first leave the pre-accident employer during event months 0–23, while 63.9 percent remain through that window. Figure 5 shows that by year five, mean earnings are 30.9 percent below the $t = -1$ level among separators but only 1.3 percent lower among stayers. Holding the two groups' population shares fixed, separators account for 90.6 percent of the aggregate descriptive year-five earnings decline (Table 3). Because separation within two years is itself a post-treatment outcome, these are accounting results rather than causal subgroup effects.

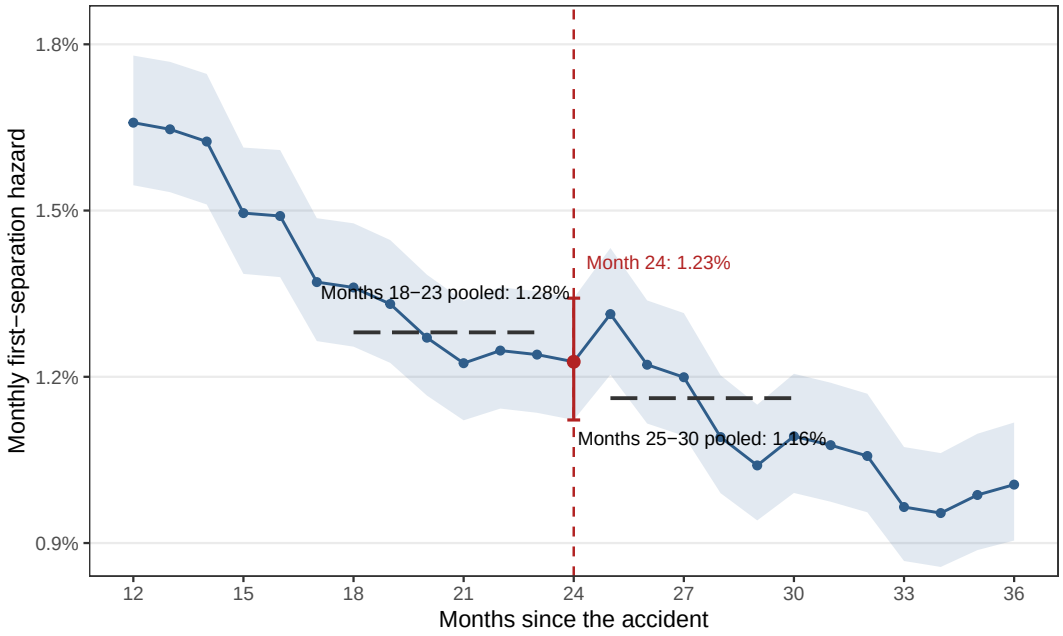


FIGURE 4: Monthly Hazard of First Separation from the Pre-Accident Employer

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; horizontal dashed segments report the pooled hazards for months 18–23 and 25–30.

TABLE 3: 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.

5.3 Heterogeneity in Post-Accident Consequences

The average effects reveal persistent labor-market detachment and health burdens, but these costs are not distributed equally. We next test whether the economic consequences of traffic injuries simply scale with physical severity, or if they depend on the injured worker's baseline characteristics.

Physical severity is an undeniable determinant of cost. As shown in Figure 6, more severe

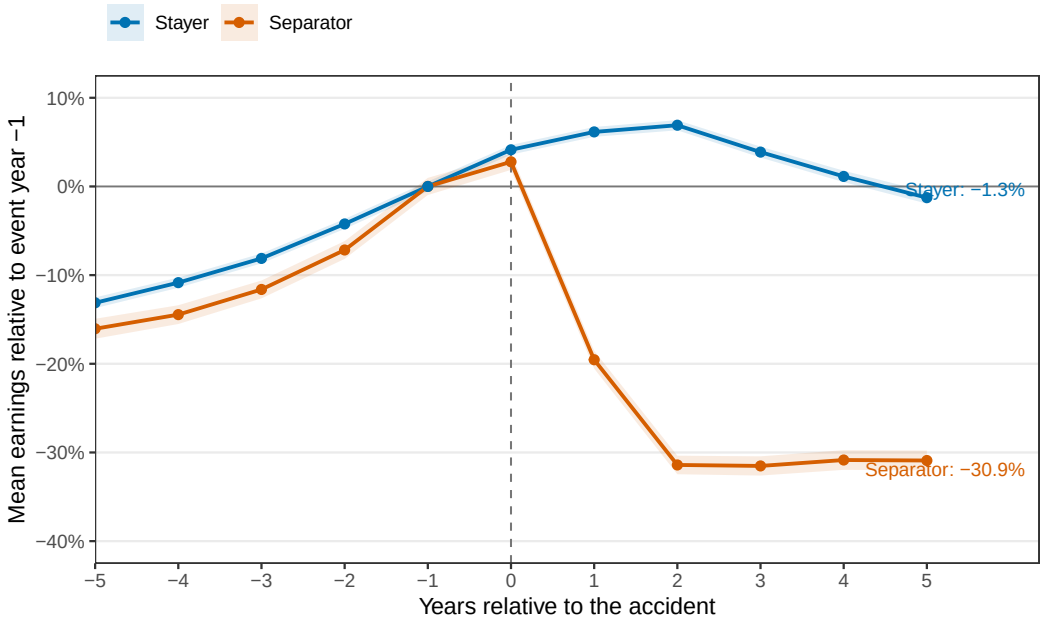


FIGURE 5: Earnings Trajectories by Separation from the Pre-Accident Employer

Notes: Separators first leave the pre-accident employer during event months 0–23; stayers remain attached through that window. Each series reports mean annual earnings relative to its own event-year-−1 mean; shaded regions are 95% confidence intervals. Separation status is a post-treatment classification, so the paths are descriptive rather than causal subgroup effects.

injuries generate substantially larger healthcare utilization and deeper labor-market losses, while milder injuries generate comparatively modest effects. Injury location provides a complementary view of functional impairment (Figure 26). Injuries to critical or multiple body regions, such as the head, neck, or multiple sites, generate the largest and most persistent effects. The nature of the physical impairment clearly constrains sustained work capacity.

However, observed medical severity does not fully explain the distribution of labor-market losses. Across several dimensions, groups with broadly similar medical responses experience markedly different labor-market outcomes. Among individuals with positive pre-accident earnings, Figure 7 shows that higher-income groups experience larger earnings and full-time employment losses, with the most severe effects concentrated in the top income quartile. Health outcomes across these income quartiles, by contrast, are broadly similar.

A similar contrast appears across road users: Figure 27 and Figure 28 show that car drivers experience especially large earnings and employment declines even though differences in medical utilization are more limited. Finally, Figure 29 shows substantially larger labor-market losses among men than among women, while health care utilization and medical spending are broadly similar.

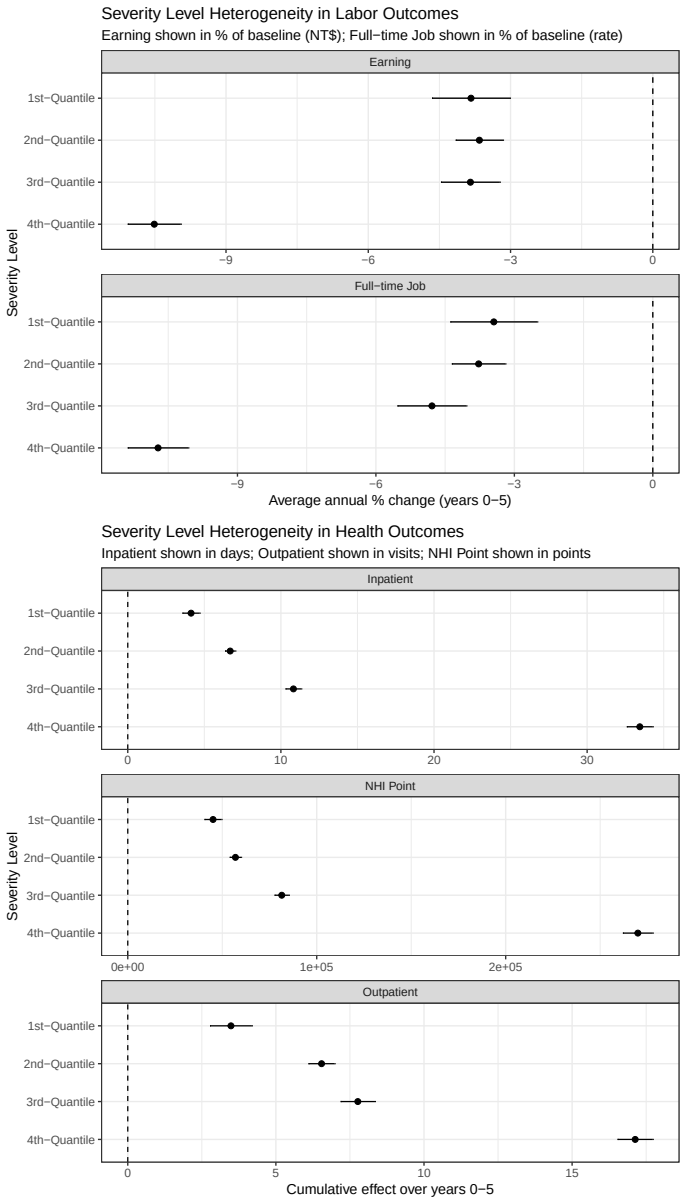


FIGURE 6: Heterogeneity by Injury Severity Level

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

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 directly by standardizing every subgroup to the common main-sample severity distribution, using hospitalization-day severity quartiles (Figure 8; the underlying severity composition of each subgroup is reported in Appendix N). The labor-market gradients survive standardiza-

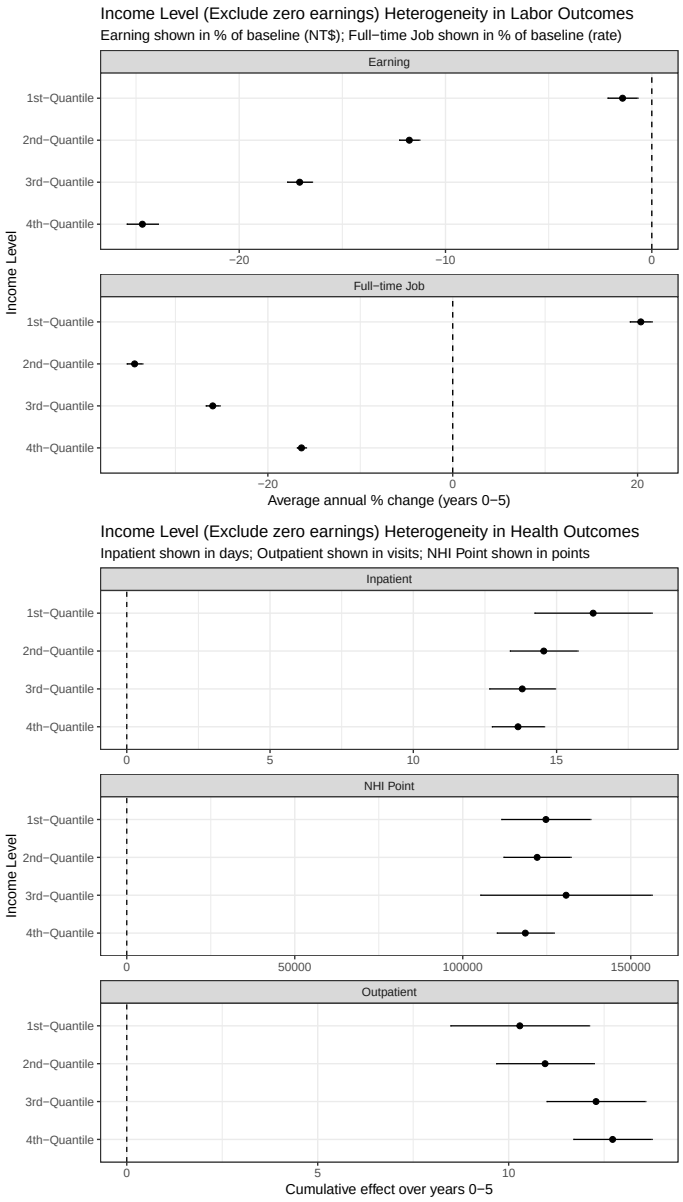


FIGURE 7: Heterogeneity by Pre-Accident Income Level

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

tion: groups differ sharply in year-five earnings and full-time employment effects even when compared at the same severity mix, while severity-standardized medical spending responses are broadly similar across subgroups. For example, after standardization the year-five earnings effect is TWD −210,623 in the highest positive pre-accident earnings quartile, compared with TWD −24,921 in the lowest; the corresponding full-time employment effects are −23.9

and −5.2 percentage points. The heterogeneity in economic losses is therefore not a disguised severity gradient.

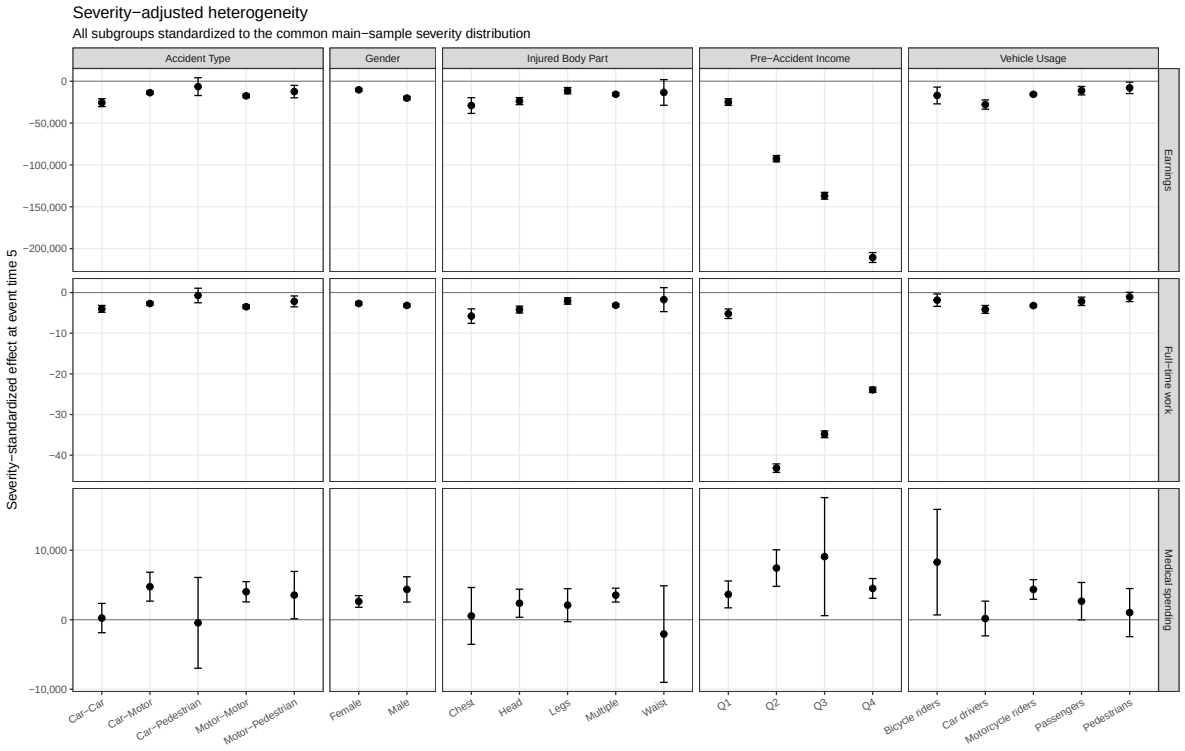


FIGURE 8: Severity-Standardized Subgroup Effects at Event Time 5

Notes: The figure reports year-five effects on earnings (top row), full-time employment (middle row), and medical spending (bottom row) by subgroup, after standardizing each subgroup to the common main-sample distribution of hospitalization-day severity quartiles. Whiskers show 95% confidence intervals.

These patterns help explain why average earnings losses deepen after the accident. If labor-market effects were driven only by medical severity, groups with larger earnings losses would also display systematically larger medical responses. Instead, similar health shocks translate into different economic trajectories. Traffic injuries interact heavily with workers' pre-accident positions, including their earnings profiles, job tasks, and employment relationships. The injury shock therefore does more than impair health: it can dislodge workers from labor-market trajectories that differ substantially across groups. This challenges policy evaluations that apply uniform costs to survivor outcomes. The complete set of heterogeneity figures is reported in Appendix N.

5.4 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. Because death is an absorbing outcome, we estimate the mortality effects using the never-treated population as the comparison group rather than the later-treated design used above. This switch requires justification, since we argue elsewhere that the never-treated differ from the eventually-injured in levels: for an absorbing outcome, later-treated controls mechanically exit the risk set, and the relevant comparison is the counterfactual survival path. **[TODO: FILL: (i) show pre-period balance in death rates between injured and never-treated comparison groups; (ii) re-estimate mortality with later-treated controls over the feasible balanced horizon (deaths within the window are observable) and report agreement with the never-treated estimates.]**

As shown in Figure 9, 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. This delayed and widening trajectory is consistent with continued health deterioration among survivors over the observed horizon.

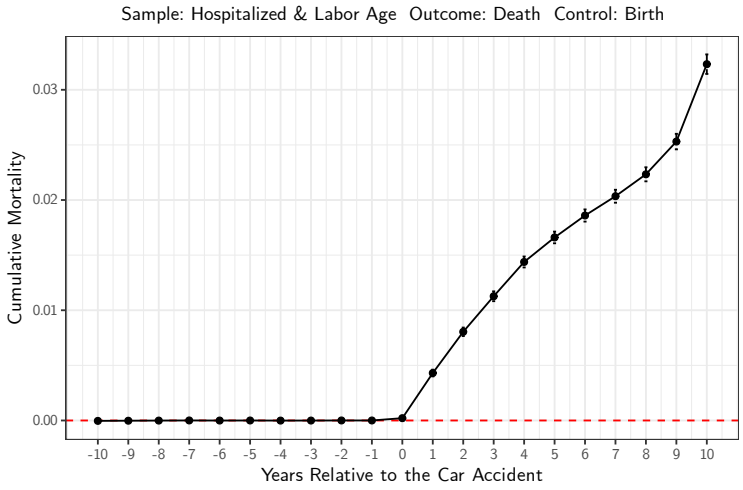


FIGURE 9: Dynamic Effects of Traffic Injuries on Mortality (Birth-Year Controls)

Notes: The outcome, *Death*, is constructed as an absorbing indicator that turns on once an individual dies and remains equal to one thereafter. Accordingly, the event-study coefficients should be interpreted as changes in the cumulative probability of having died by each event time, rather than as contemporaneous mortality rates or hazards.

5.5 Robustness of the Dynamic Patterns

These dynamic patterns are robust to a variety of alternative specifications. First, the qualitative trajectories hold when using longer event windows (Appendix E), a stricter hospitalization-window definition (Appendix G), and narrower treatment-cohort comparisons (Appendix F). Second, the results are stable across specifications that omit covariate adjustments (Appendix H), indicating that the patterns are not driven by the birth-cohort controls.

Third, a placebo exercise confirms that the estimated dynamics are generated by the accident itself rather than by the cohort or calendar-time structure of the design. Randomly reassigning accident years across individuals, keeping the overall distribution of accident timing fixed, and re-running the identical specification produces estimates close to zero on every outcome: the placebo accident-year earnings effect is TWD −326 (versus TWD −2,119 in the main sample), the full-time employment effect is −0.09 percentage points (versus −0.77), and medical spending moves by about +92 NHI points (versus the true spike of roughly +80,000 points). None of the widening post-accident dynamics is reproduced. Appendix J reports the full placebo event studies.

The comparison with conventional estimators also illustrates the importance of our empirical design. A conventional pooled two-way fixed effects (TWFE) event study suggests

a recovery pattern that differs from the cohort-specific estimates (Appendix I). By contrast, our cohort-specific later-treated design broadly reproduces the deepening losses. This demonstrates that the persistent labor-market detachment is a robust feature of the data, whereas static TWFE models mask these underlying dynamic trajectories.

Finally, Appendix K benchmarks the traffic-injury estimates against general (all-cause) hospital admissions estimated in the same framework, in the spirit of Dobkin et al. (2018). Raw year-five labor-market losses after a general admission are larger than after a traffic injury, but the general-admission event studies display clear pre-trends, consistent with anticipated admissions and deteriorating health before hospitalization. This contrast underscores the identification value of a police-defined and plausibly unanticipated injury event.

6 Valuing the Dynamic Costs of Nonfatal Injuries

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 those dynamics imply for the shadow price of a nonfatal injury. The goal is not to produce a complete welfare valuation of nonfatal traffic injuries. A complete estimate would need to include pain and suffering, quality-of-life losses, informal caregiving, out-of-pocket spending, and psychological distress. Instead, our estimates provide an administrative-cost benchmark that captures the components recorded directly in earnings and public medical-spending data. Mortality is analyzed separately and is not monetized here.

6.1 An Administrative-Cost Benchmark

We begin with the components that can be measured directly in the linked administrative data. Earnings are expressed in 2021 TWD using the GDP deflator, and all USD conversions in this section use a rounded exchange rate of TWD 30 per USD. For the main sample of working-age individuals hospitalized after traffic accidents, aggregating excess medical spending and earnings losses from the accident year through event year five yields a per-injury cost of approximately TWD 176,000 (USD 5,900). This total combines about TWD 112,000 (USD 3,700) in excess NHI-covered medical spending and TWD 64,000 (USD 2,100) in lost earnings. Medical reimbursement units are converted at TWD 0.9 per unit. Both components are calculated from the dynamic treatment effects rather than from cross-sectional differences or retrospective reports.

The five-year estimate captures the burden over the main post-accident window, but the event-study patterns indicate that costs persist beyond it. To construct a transparent longer-

horizon scenario, we use the year-specific estimates from the balanced ten-year specification reported in Appendix E. From the accident year through event year ten, these estimates imply cumulative earnings losses of TWD 163,365 and excess medical spending of 141,707 reimbursement units, equivalent to TWD 127,536. For event years 11–19, we hold the annual earnings loss constant at its event-year-ten estimate of TWD 23,458. We conservatively set additional medical spending after event year ten to zero rather than extrapolating the smaller but still positive long-run estimates. This construction produces projected twenty-year earnings losses of TWD 374,490 and an undiscounted administrative-cost total of TWD 502,026, or approximately USD 16,700. It is an illustrative projection, not a realized twenty-year effect or a full welfare estimate.

TABLE 4: Discount-Rate Sensitivity of the Illustrative Twenty-Year Projection

Annual discount rate	Present value (TWD)	Present value (USD)
0%	502,000	16,700
1%	460,700	15,400
2%	424,700	14,200
3%	393,200	13,100
4%	365,600	12,200
5%	341,300	11,400
6%	319,900	10,700
7%	300,800	10,000

Notes: For event years 0–10, the calculation uses the year-specific earnings and medical-spending estimates from the balanced ten-year specification. Medical reimbursement units are converted at TWD 0.9 per unit. For event years 11–19, the annual earnings loss is held constant at the event-year-ten estimate and additional medical spending is set to zero. Each annual component is discounted according to its event year. USD values use TWD 30 per USD. Values are rounded to the nearest TWD 100 and USD 100.

This projection should be interpreted as a conservative administrative-cost benchmark. It includes formal earnings losses and public medical spending but deliberately excludes nonmarket losses and private costs not captured in administrative records. A conventional accounting estimate may appear precise, but it is often built from imputed productivity losses, short-run medical expenses, or statutory compensation schedules. Our calculation is deliberately narrower and anchored in observed post-accident trajectories. It should be compared with other valuation methods as a partial administrative measure, not interpreted as a lower bound that necessarily captures every omitted component.

6.2 Comparison with Existing Valuation Approaches

The administrative-cost benchmark complements, rather than replaces, existing valuation approaches. As discussed in Section 2, Taiwan-based estimates for nonfatal traffic injuries vary substantially across methods. Survey-based willingness-to-pay estimates imply costs of roughly USD 20,000 per injured individual, while accounting-based estimates derived from court-verdict data are closer to USD 33,000. Our undiscounted twenty-year projection of about USD 16,700 falls below this range, consistent with its deliberately narrower scope: the benchmark excludes mortality, pain and suffering, informal care, quality-of-life losses, and other private burdens.

This comparison has two implications. First, the administrative components are economically meaningful even without monetizing intangible losses. Second, one-time measures may miss persistent labor-market scarring and post-acute medical needs that are difficult to recover from surveys, hospital bills, or legal awards. Because the dominant component is a labor-market loss among workers who largely recover medically, scaling fatality values by medical severity weights mismeasures this cost by construction.

6.3 Implications for Infrastructure Investment

These estimates matter for infrastructure investment because conventional safety evaluations are overwhelmingly centered on fatalities. A simplified benefit calculation values an intervention by multiplying the expected number of fatalities avoided by a VSL or an official fatality-cost benchmark. That approach misses two margins that are central in our setting: the number of nonfatal injuries avoided and the reduction in injury severity among crashes that still occur.

A more complete benefit calculation would include at least three components:

$$\text{Benefits} = \Delta\text{Fatalities} \times VSL + \Delta\text{Nonfatal Injuries} \times C_{\text{injury}} + \Delta\text{Severity} \times C_{\text{severity}},$$

where C_{injury} represents the long-run cost of a nonfatal injury and C_{severity} captures the value of shifting injuries from more severe to less severe categories. Our estimates provide an empirical benchmark for these latter two terms.

In Taiwan, nonfatal inpatient injuries outnumber traffic fatalities by a ratio of roughly eight to one. The multiplication therefore uses hospitalized nonfatal injuries as its injury base, matching the sample underlying the administrative-cost benchmark rather than the broader all-injured population. As a result, even a per-injury cost far below the value of

a statistical life can materially alter the benefit side of infrastructure assessments. This point is especially relevant for interventions whose primary benefit is not eliminating crashes entirely but reducing injury severity. Speed management, safer intersection design, protected facilities for vulnerable road users, and post-crash response improvements generate benefits by converting severe injuries into milder injuries or by reducing the dynamic consequences conditional on survival. Fatality-centered evaluations may undervalue such projects because they assign little or no weight to persistent earnings losses, employment detachment, and post-acute medical spending among survivors.

The labor-market heterogeneity documented above further refines these policy implications. Groups with similar health impacts can experience substantially different earnings and employment losses, implying that economic costs depend heavily on who is exposed to injury risk and how injuries interact with pre-accident jobs. Policy evaluation should account for both efficiency and distributional margins. Preventing injuries among workers with steep expected earnings trajectories may generate large measured economic gains, while reducing employment detachment among more vulnerable workers carries profound social-insurance value.

Taken together, the results suggest that nonfatal injuries are too frequent, too persistent, and too heterogeneous to be treated as a simple fixed fraction of fatalities. Incorporating dynamic realized costs into infrastructure assessment would raise the estimated returns to projects that prevent serious injuries, mitigate severity, or protect groups whose injuries translate into large labor-market disruptions.

7 Conclusion

Acute health shocks destroy employment relationships, not just health. This paper documents that mechanism for the most common acute health shock to working-age adults, linking 24 years of police-reported traffic accidents to population-level health, employment, and mortality records. Using a staggered event-study design with later-treated controls, we trace the economic and health consequences of hospitalization after traffic crashes among working-age survivors, providing longitudinal evidence on how a transitory medical event becomes a persistent economic loss.

The results show that nonfatal injuries are not merely short-run medical events. Annual earnings and full-time employment decline by approximately 9 percent by year five. This persistent scarring operates primarily through the extensive margin, driven by detachment from full-time work and severed attachment to pre-accident employers. It persists even among

workers whose measured medical utilization returns to the range of uninjured controls. While the medical burden is heavily concentrated in the accident year, healthcare utilization, public health spending, and cumulative mortality remain elevated well after the acute trauma passes. Furthermore, these costs are unevenly distributed. Severe and critical-body-region injuries generate predictably larger losses, but identical physical shocks also yield substantially different economic trajectories depending on a worker's initial labor-market position. High-income workers, men, and car drivers experience especially large labor-market declines relative to their medical responses.

Combining the ten-year estimates with a transparent earnings extrapolation yields a twenty-year administrative-cost benchmark of approximately TWD 502,000 (USD 16,700) per hospitalized working-age injury. This estimate is deliberately narrower than a full welfare valuation: it excludes pain and suffering, informal care, quality-of-life losses, and private costs outside the administrative data. It is a projected administrative-cost benchmark anchored in observed dynamic post-accident trajectories, rather than a complete social-cost estimate or a lower bound that necessarily captures every omitted component.

For transportation and urban policy, the central implication is that survivor costs must be treated as a core component of road-safety evaluation rather than a residual category. Because nonfatal inpatient injuries substantially outnumber traffic fatalities (by a ratio of roughly eight to one in our setting), assessments that focus exclusively on fatality avoidance may miss a large and persistent source of social cost. This is especially true for infrastructure investments that reduce injury severity or help injured workers remain attached to employment. Improving road safety yields benefits not only through lives saved, but also by preventing the persistent economic losses borne by the much larger population of injured survivors.

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.

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.

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.

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.

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.

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

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.

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.

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.

A Accident-Level Summary Statistics

Table 5: 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 5 (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 5 (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.

B Balance Check

Table 6: T-test of Labor Outcomes among Injured Individuals before Accident Occurrence

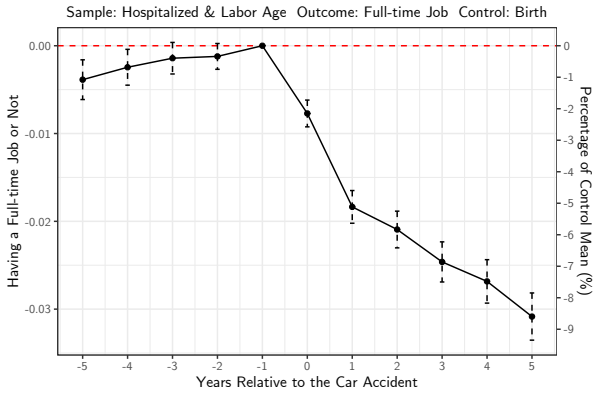
Dependent Variable	Earning	Full-time Job	Lost-Job
Constant	111,704.4	0.2030	0.058
	(96.97)	(0.0001)	(0.0004)
Injured	-7,382.2	-0.0011	-0.012
	(130.2)	(0.0002)	(0.0003)
S.E.: Clustered	by: ID	by: ID	by: ID
Observations	145,722,410	145,722,410	145,722,410
R2	0.18921	0.19705	0.10775

Table 7: 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.

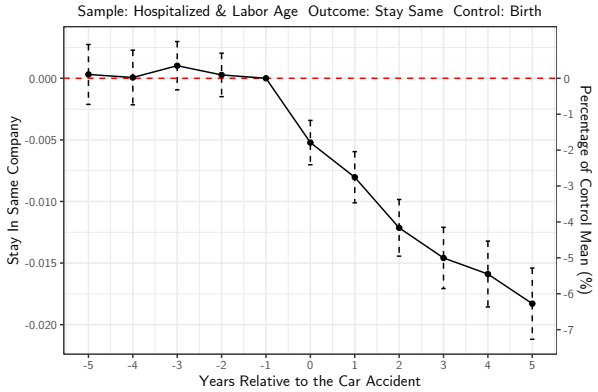
C More Main Results



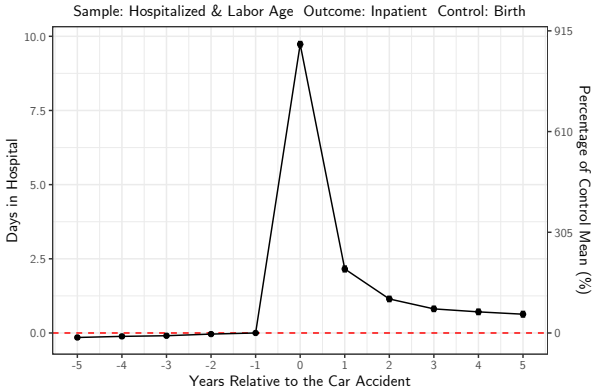
(A) Full-time employment (birth-year controls)



(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)

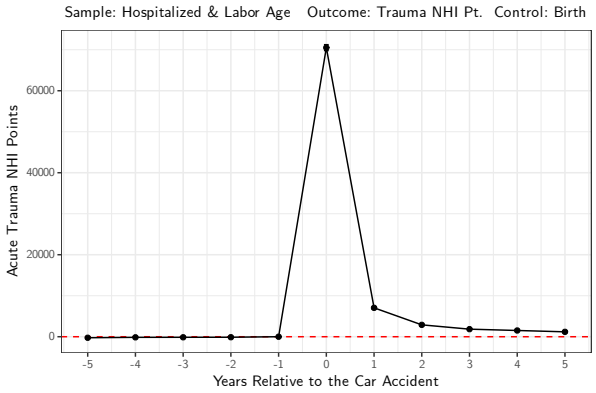
FIGURE 10: Dynamic Effects of Traffic Injuries On Labor and Health Outcomes

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

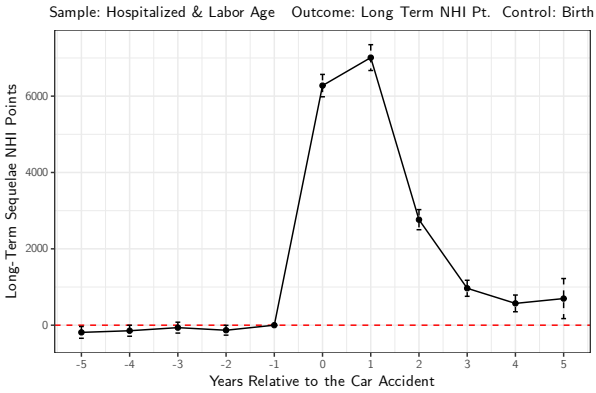
D Medical Expense Decomposition

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.

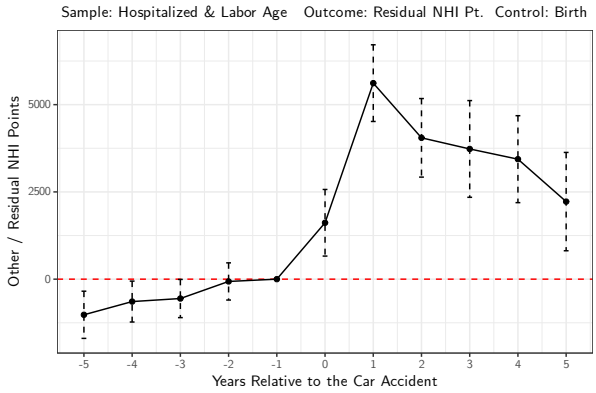
Figure 11 plots the event-study estimates for acute trauma, long-term sequelae, and the residual other-spending category. The patterns reinforce the interpretation in the main text: trauma-related spending spikes sharply in the accident year and declines afterward, while long-term sequelae remain elevated beyond the initial injury episode.



(A) Acute trauma



(B) Long-term sequelae

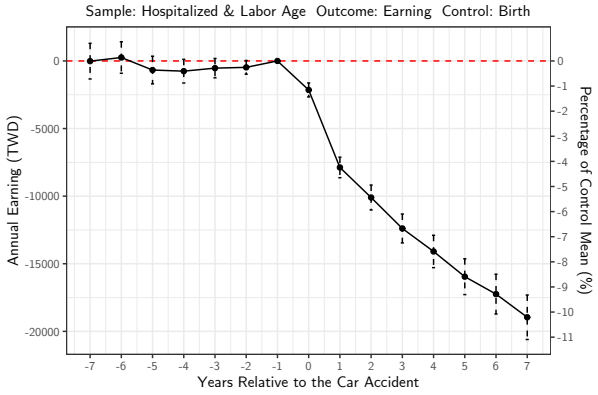


(C) Other medical spending

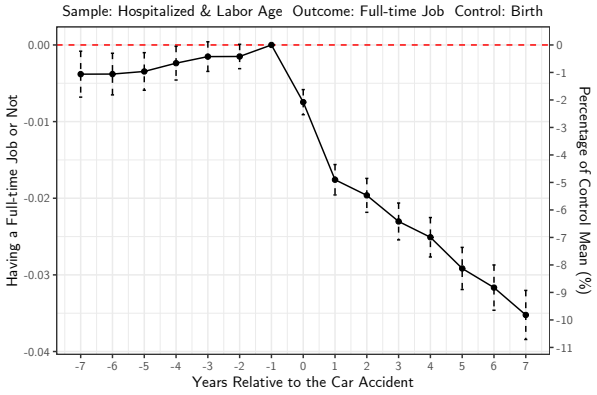
FIGURE 11: Dynamic Effects by Medical Spending Category

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

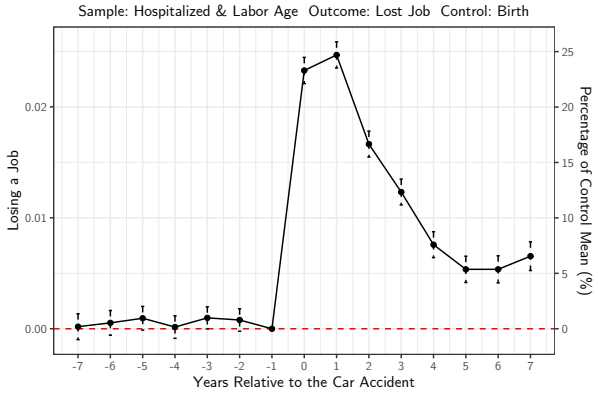
E Extended Window: 7 & 10 Years



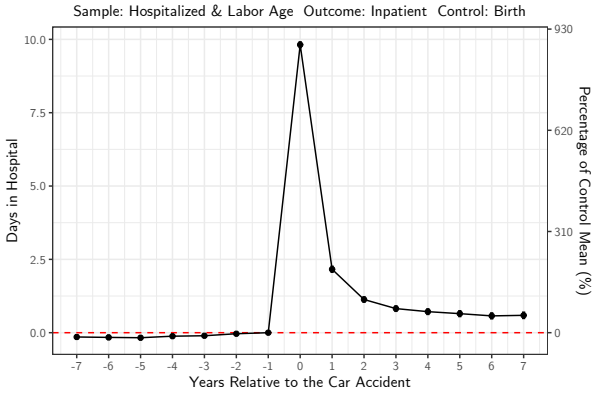
(A) Earnings (birth-year controls)



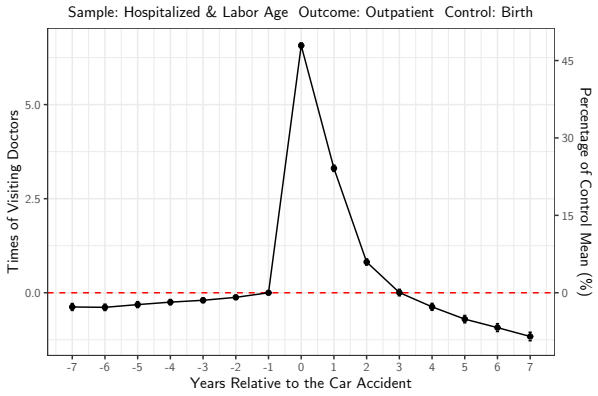
(B) Full-time employment (birth-year controls)



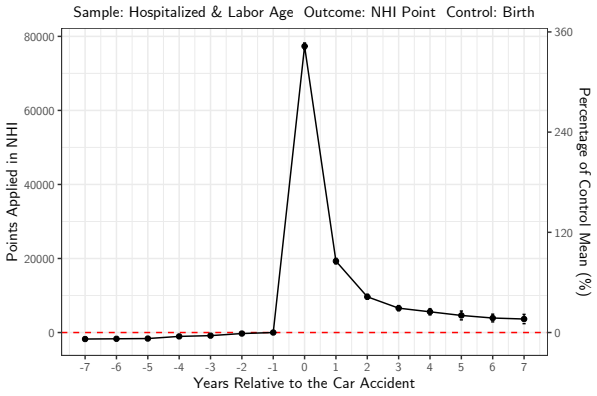
(C) Losing Job (birth-year controls)



(D) Inpatient days (birth-year controls)



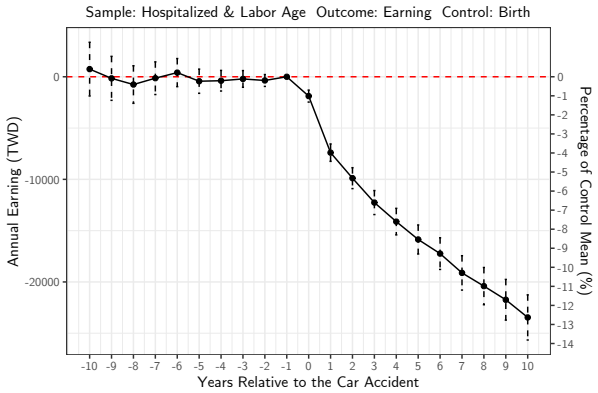
(E) Outpatient visits (birth-year controls)



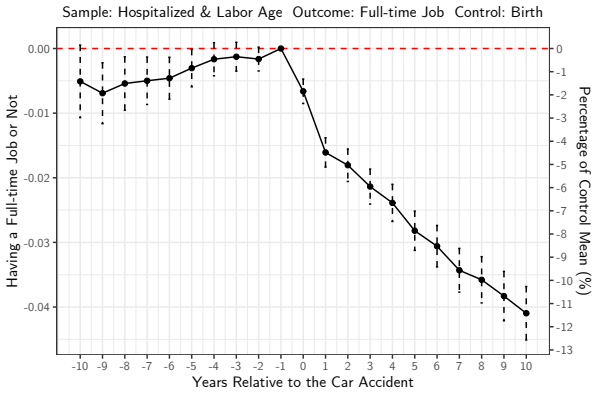
(F) NHI points (birth-year controls)

FIGURE 12: Dynamic Effects over a 7-Year Window

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



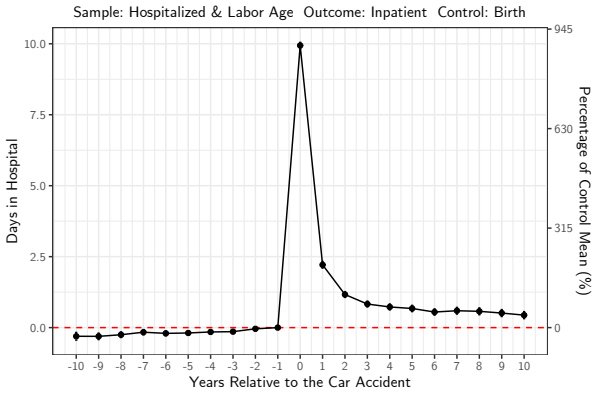
(A) Earnings (birth-year controls)



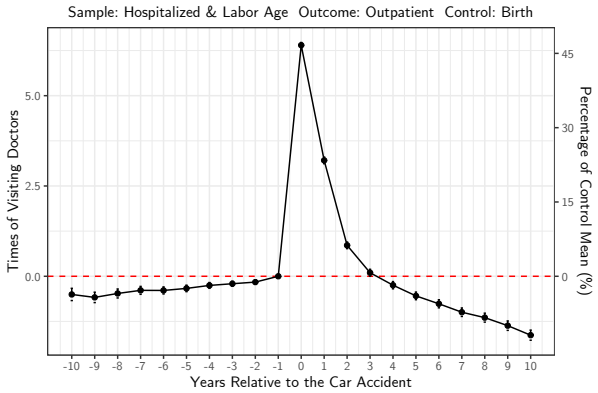
(B) Full-time employment (birth-year controls)



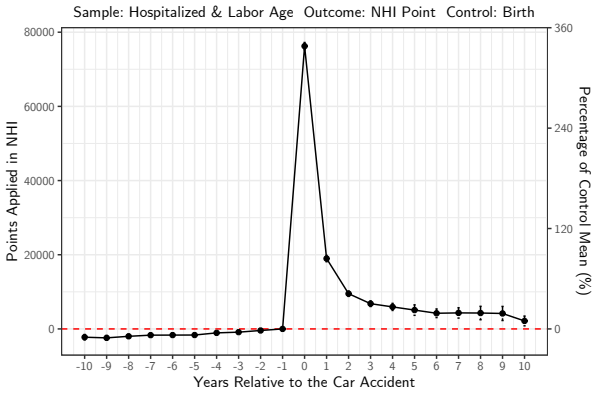
(C) Losing Job (birth-year controls)



(D) Inpatient days (birth-year controls)



(E) Outpatient visits (birth-year controls)



(F) NHI points (birth-year controls)

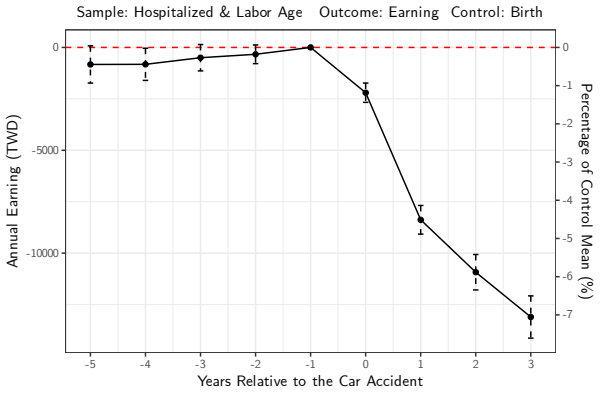
FIGURE 13: Dynamic Effects over a 10-Year Window

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

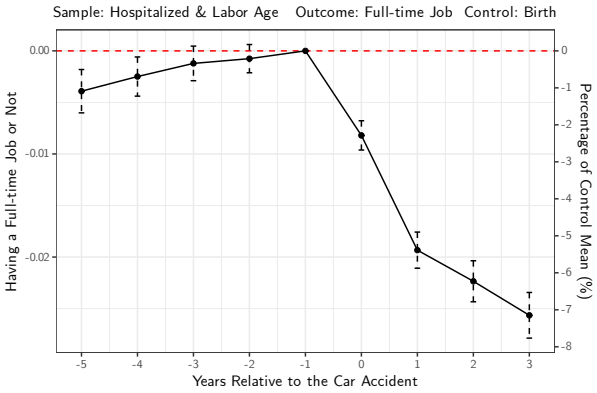
F Robustness to Narrow Treatment-Cohort Windows

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.

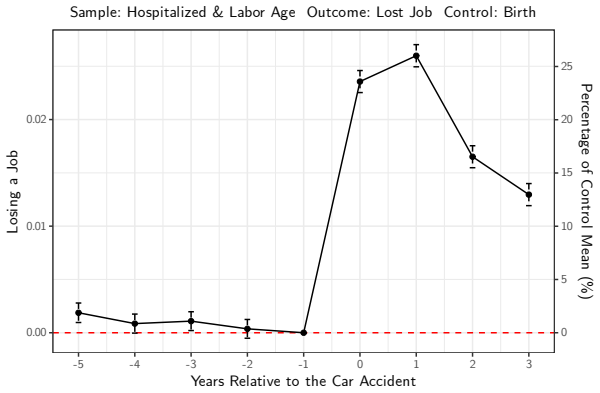
Figure 14 reports the results for the main outcomes with birth-year controls. The qualitative patterns are similar to the baseline estimates: earnings and full-time employment decline persistently after the accident, while inpatient days, outpatient visits, and NHI points rise sharply around the accident and remain elevated afterward.



(A) Earnings



(B) Full-time employment



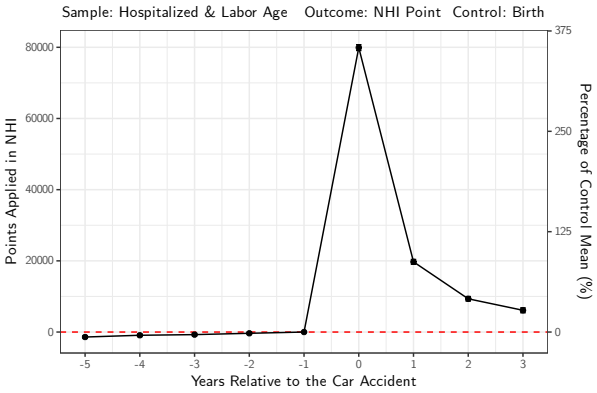
(C) Losing Job



(D) Inpatient days



(E) Outpatient visits



(F) NHI points

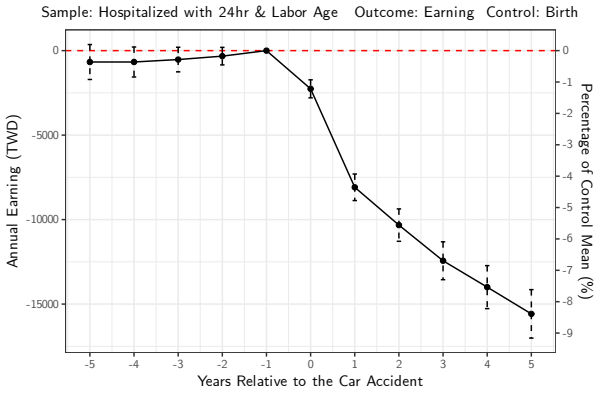
FIGURE 14: Robustness to Three-Year Treatment-Cohort Window

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 $t = -1$. Whiskers show 95% confidence intervals.

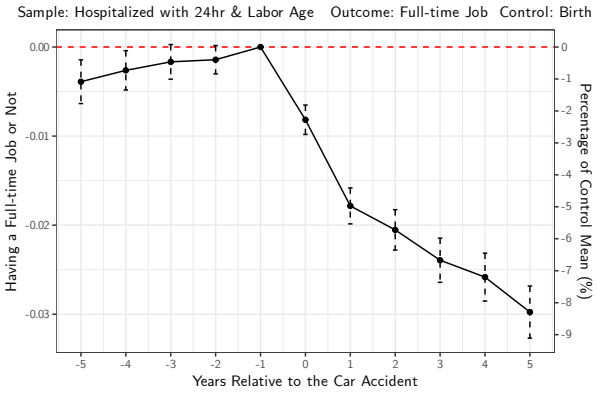
G Robustness to Hospitalization-Window Definition

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 15 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



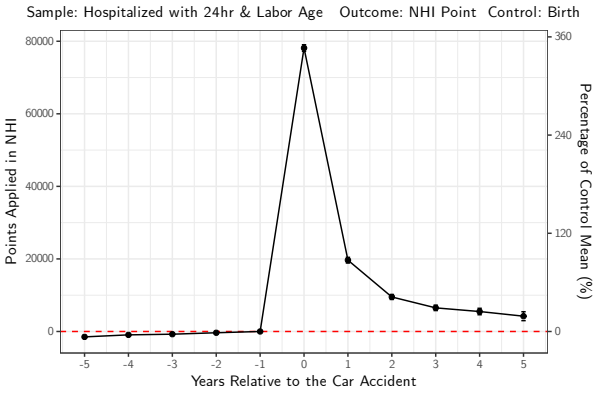
(C) Losing job



(D) Inpatient days



(E) Outpatient visits

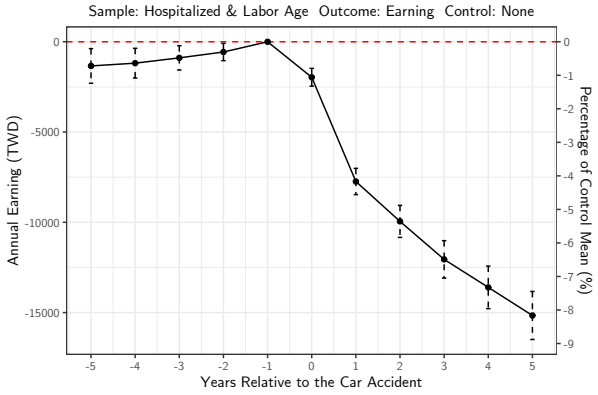


(F) NHI points

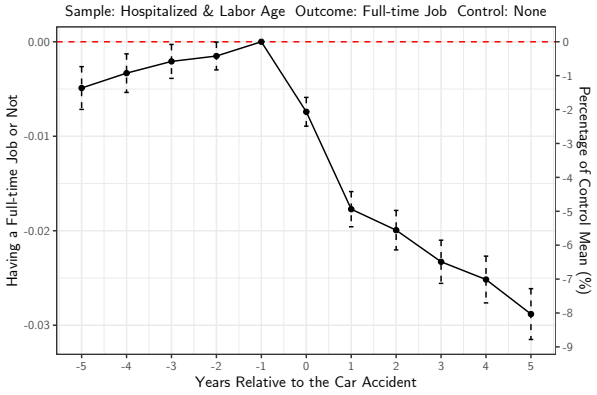
FIGURE 15: Robustness to Hospitalization within 24 Hours

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 $t = -1$. Whiskers show 95% confidence intervals.

H Main Results without Controlling Birth Year



(A) Earnings



(B) Full-time employment



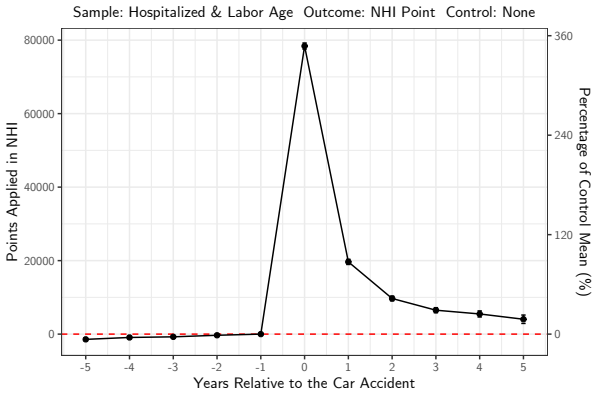
(C) Losing Job



(D) Inpatient days



(E) Outpatient visits



(F) NHI points

FIGURE 16: Dynamic Effects of Traffic Injuries without Covariates

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

I Conventional TWFE Event-Study Estimates

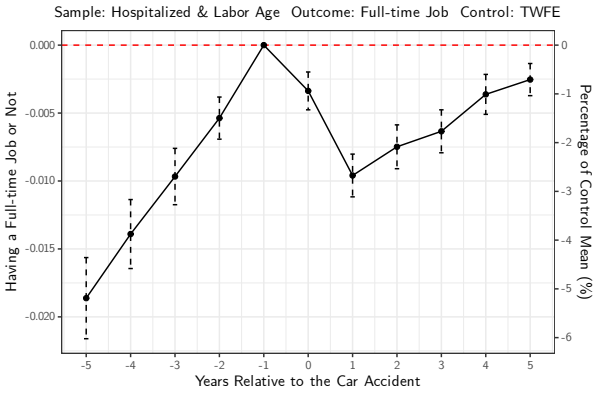
This appendix reports conventional TWFE event-study estimates for the main outcomes. The specification uses the same event-time normalization as the main dynamic analysis and includes birth-year controls to account for life-cycle differences across accident cohorts. These controls do not eliminate the 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 17 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



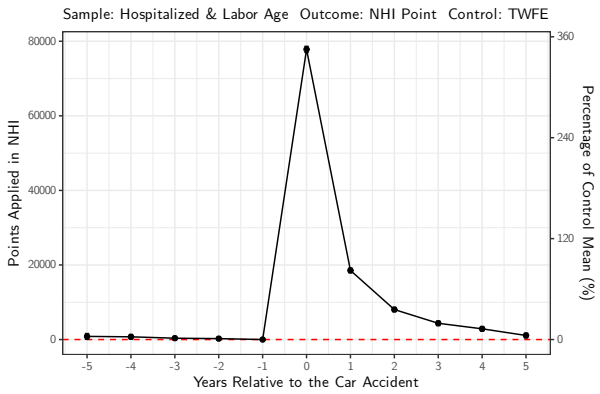
(B) Full-time employment



(C) Inpatient days



(D) Outpatient visits



(E) NHI points

FIGURE 17: 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. Whiskers show 95% confidence intervals. Earnings are expressed in 2021 TWD. NHI points are standardized reimbursement units under Taiwan's NHI.

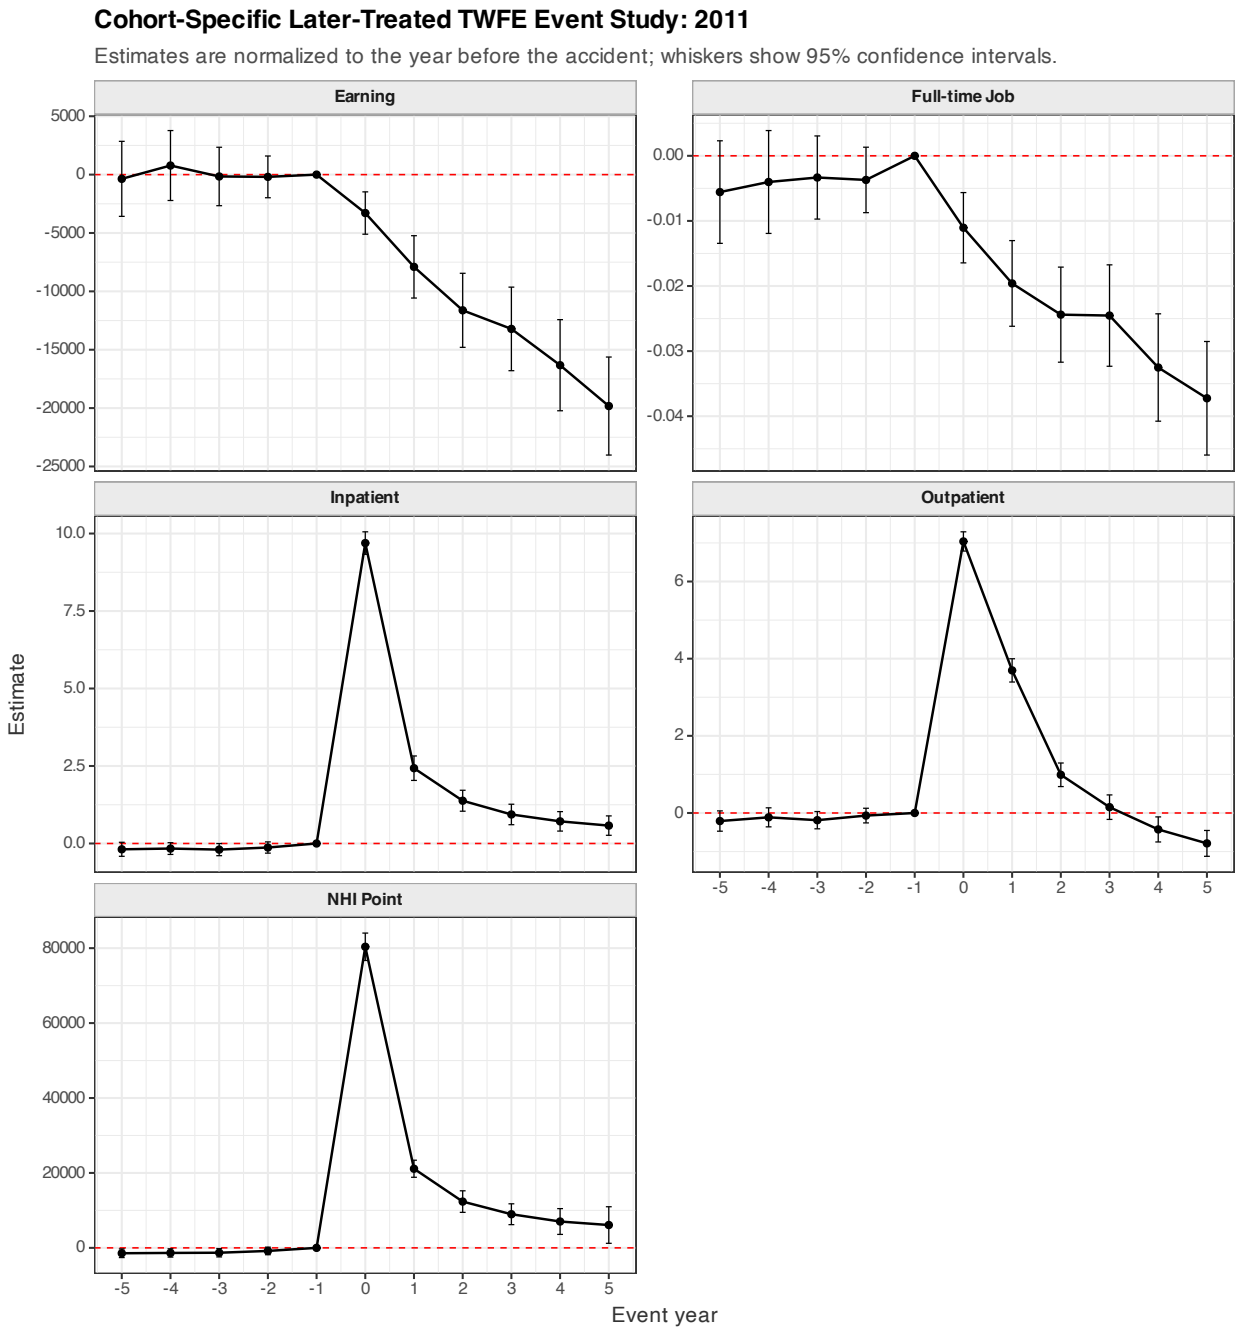


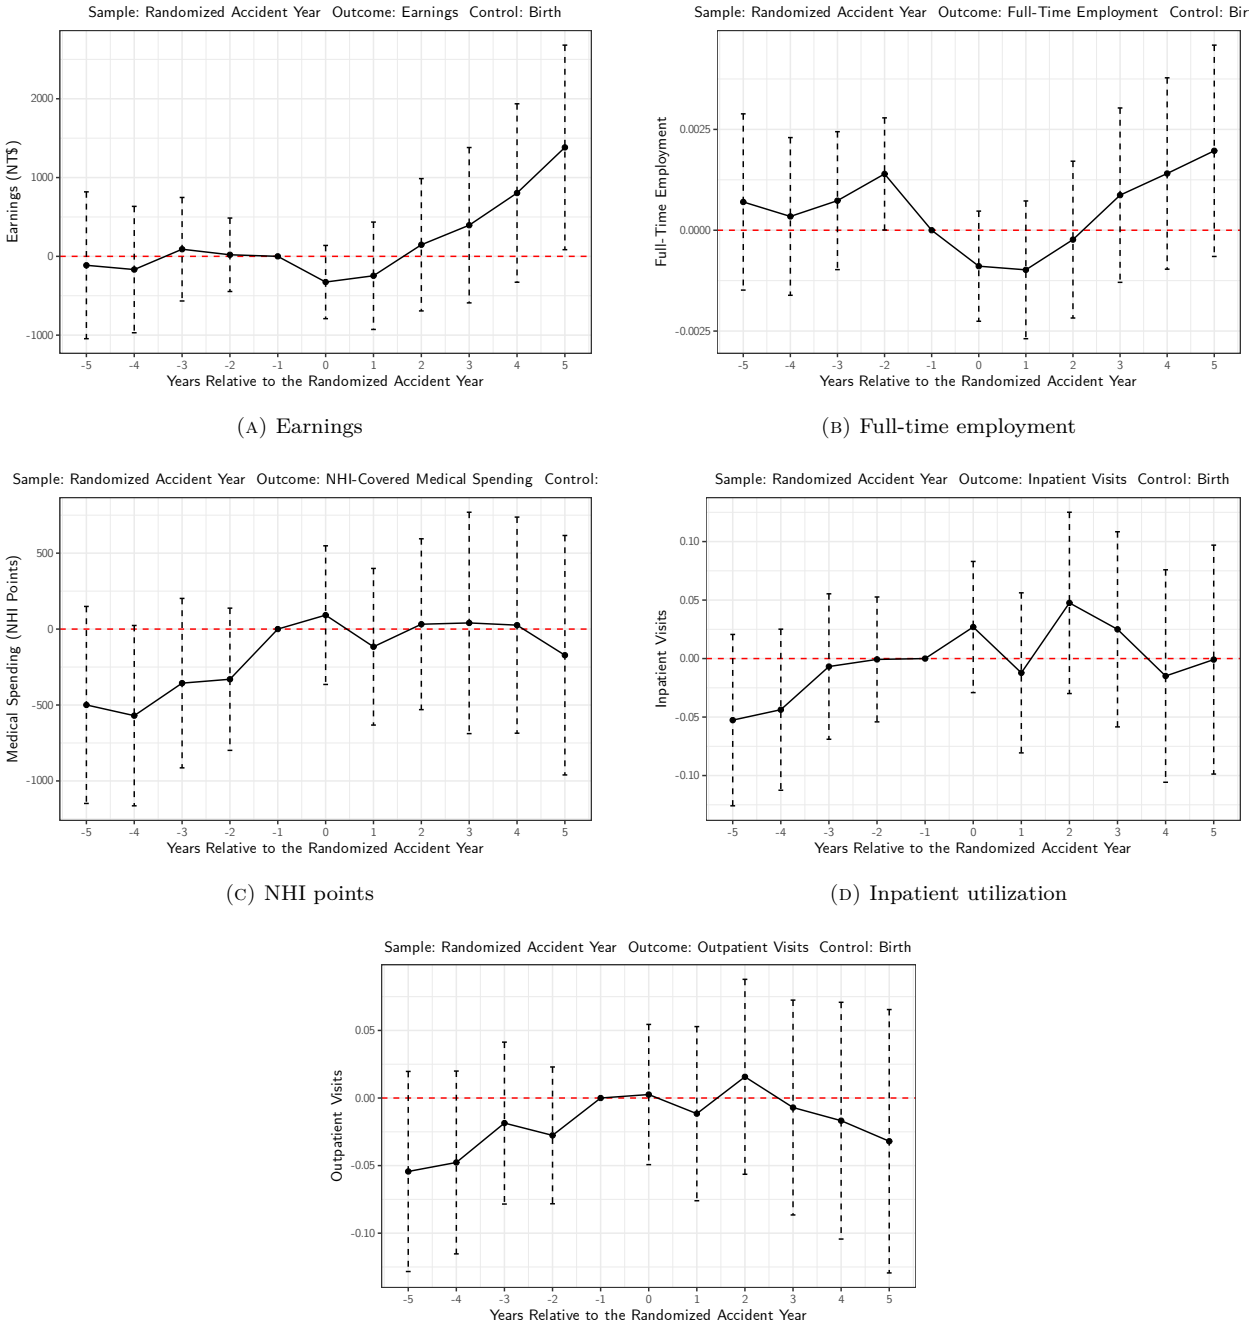
FIGURE 18: Cohort-Specific Later-Treated TWFE Event-Study Estimates

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. Whiskers show 95% confidence intervals.

J 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 19 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 19: Placebo Event Studies with Randomized Accident Years

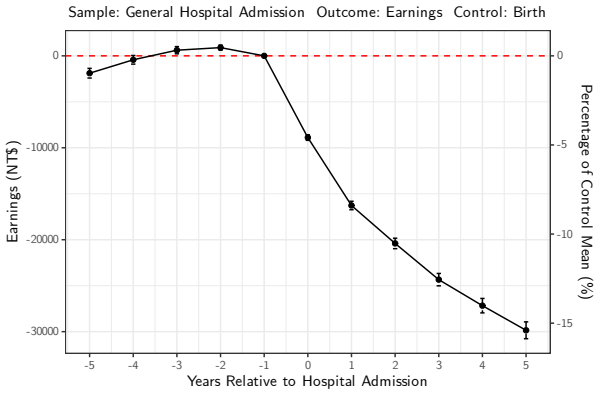
Notes: Each panel reports event-study estimates from the main specification after randomly re-assigning accident years within the main sample, preserving the overall distribution of accident timing. Estimates are normalized to the year before the placebo accident year. Whiskers show 95% confidence intervals.

K Benchmark: General Hospital Admissions

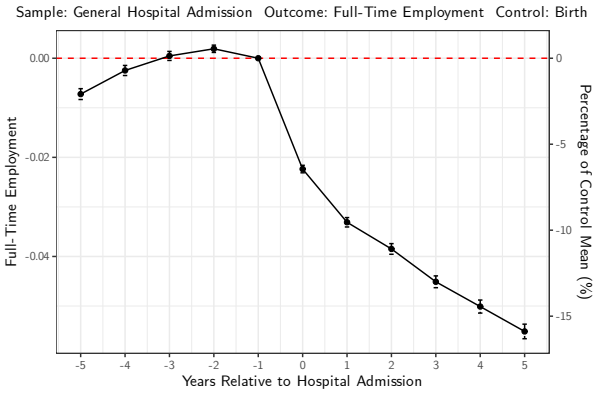
This appendix benchmarks the traffic-injury estimates against general (all-cause) hospital admissions, placing our shock in the same framework as the hospital-admissions literature (Dobkin et al., 2018). We apply the identical later-treated event-study specification to individuals experiencing a first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

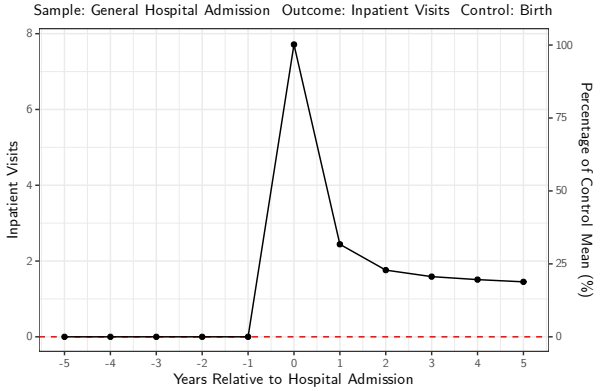
Health outcomes. At $t = 0$, a general admission increases inpatient utilization by about 7.7 days, outpatient visits by about 6.5, and NHI-covered spending by about 83,804 points, an admission-year spike similar in magnitude to the roughly 80,000-point spike after traffic injuries. Unlike traffic injuries, 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.



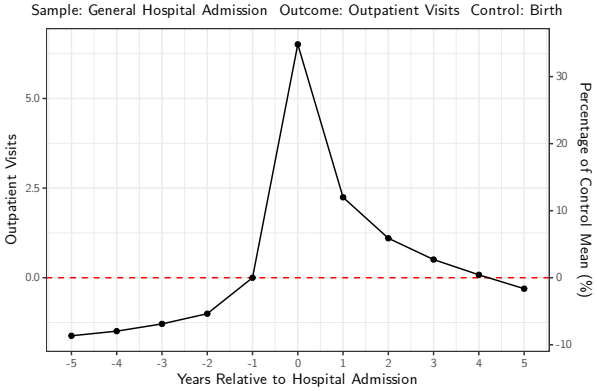
(A) Earnings



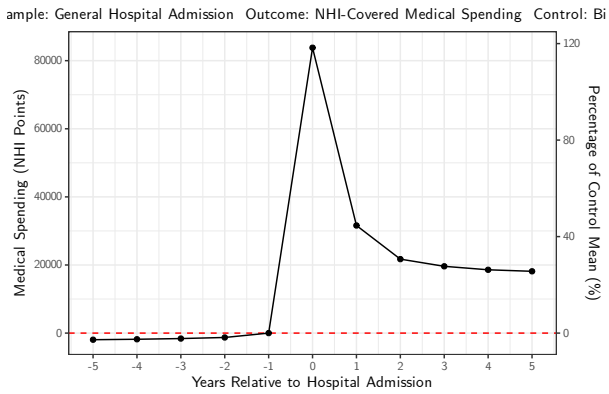
(B) Full-time employment



(C) Inpatient utilization



(D) Outpatient visits



(E) NHI points

FIGURE 20: 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. Whiskers show 95% confidence intervals. The general-admission sample displays non-negligible pre-trends in labor-market and outpatient outcomes; see the text for interpretation.

L Low-Utilization (Recovered) Subsample

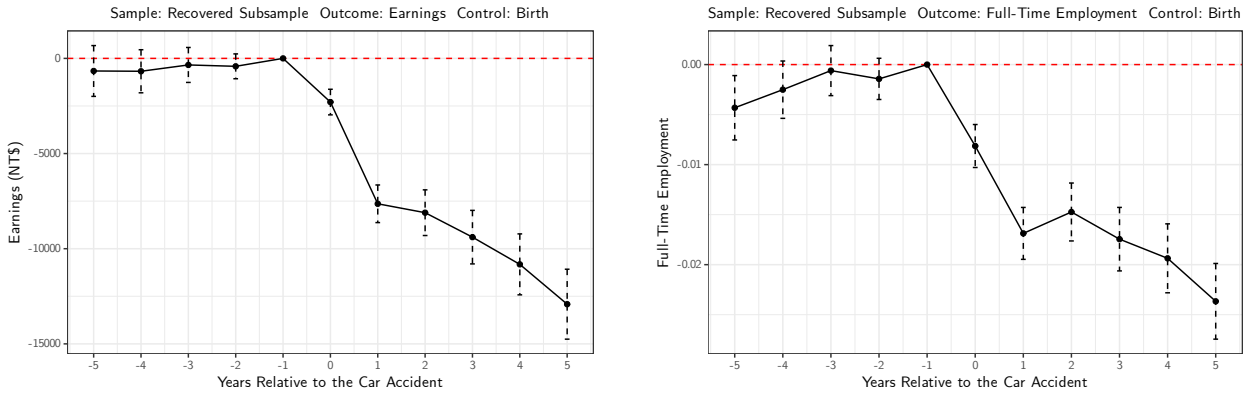
This appendix reports the low-utilization ("recovered") subsample analysis referenced in Section 5.2. The purpose is corroboration: if labor-market losses persist among workers whose medical utilization has returned to the range typical of uninjured individuals, then medical recovery and economic recovery 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 21 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.

[TODO: FILL: strengthen the recovery definition per the presubmission plan (utilization within the control band by year two *and* no catastrophic-illness certification, no sequelae spending, no rehabilitation codes) and report how the retained-loss shares change.]



(A) Earnings

(B) Full-time employment

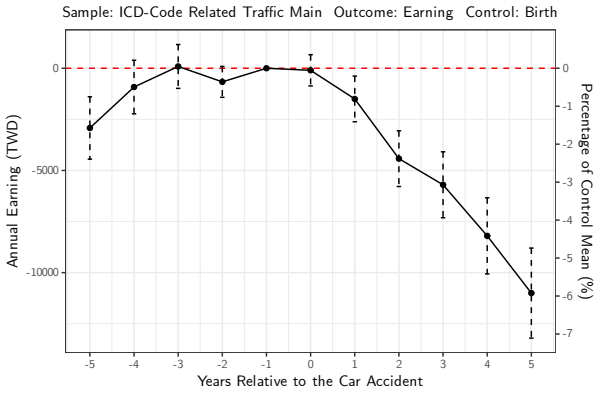
FIGURE 21: Event Studies for the Low-Utilization-at-Year-Two Subsample

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 $t = -1$. Whiskers show 95% confidence intervals.

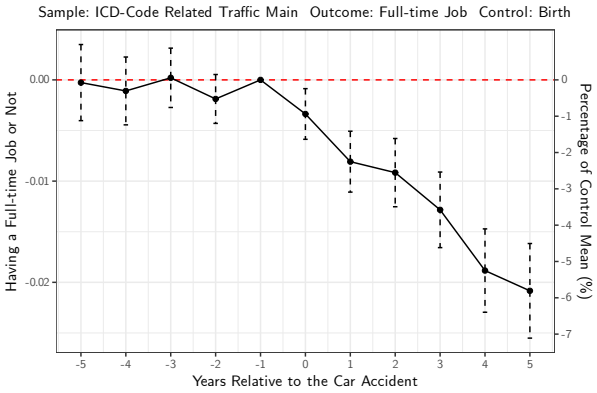
M Claims-Identifiable Traffic-Injury Subsample

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.

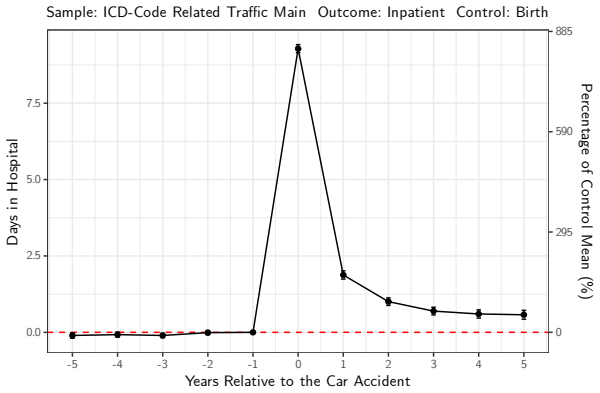
Conditioning on identification in medical claims does not select a subgroup with larger estimated responses in this hospitalized sample. The accident-year medical-spending effect is approximately 69,292 NHI points in the claims-identifiable subsample, compared with 78,405 points in the police-defined 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.



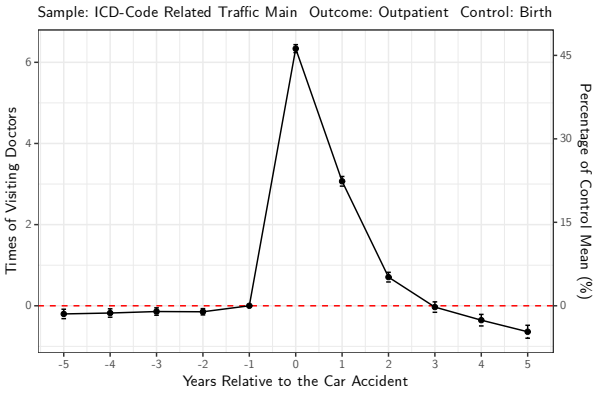
(A) Earnings



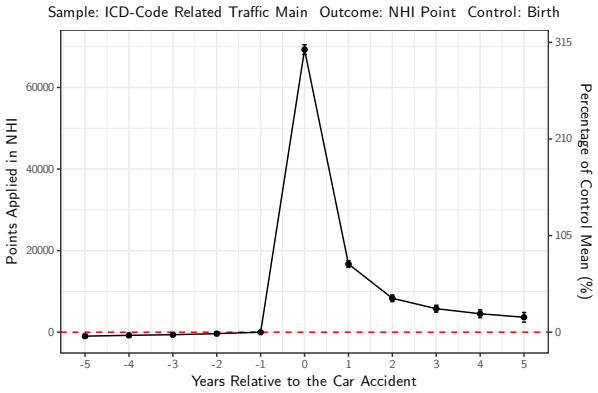
(B) Full-time employment



(C) Inpatient utilization



(D) Outpatient visits



(E) NHI points

FIGURE 22: 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 $t = -1$. Whiskers show 95% confidence intervals.

N Heterogeneity Analysis

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

N.1 Heterogeneity by Injury Severity Level

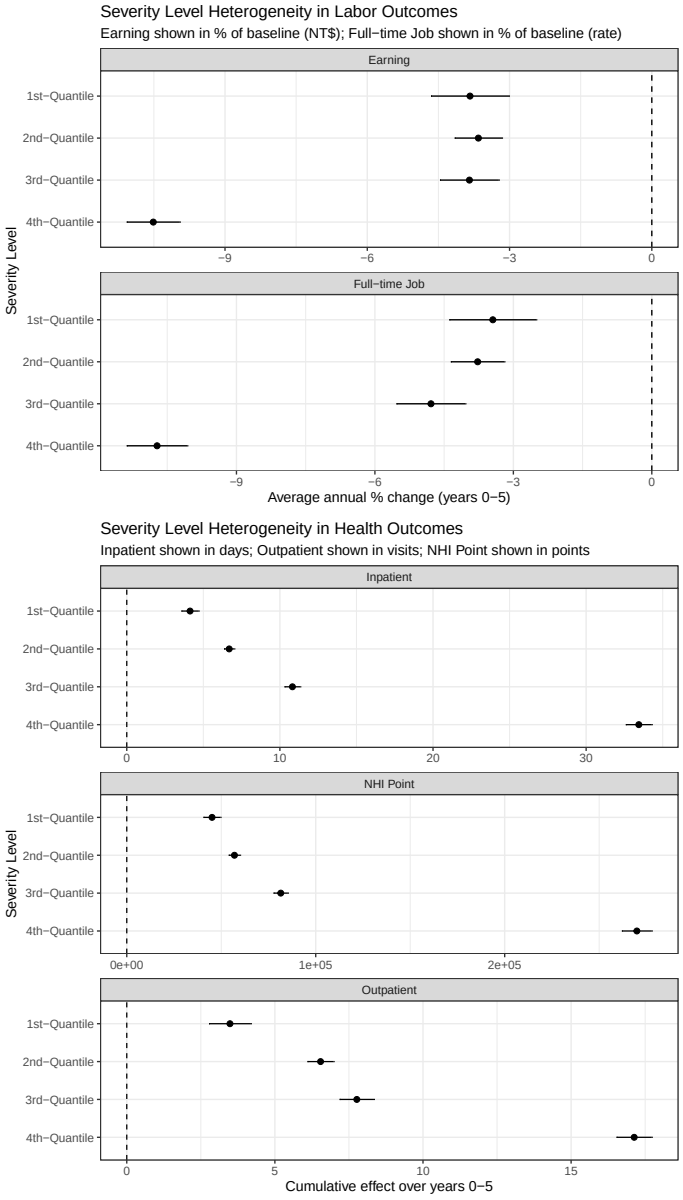


FIGURE 23: Heterogeneity by Injury Severity Level

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

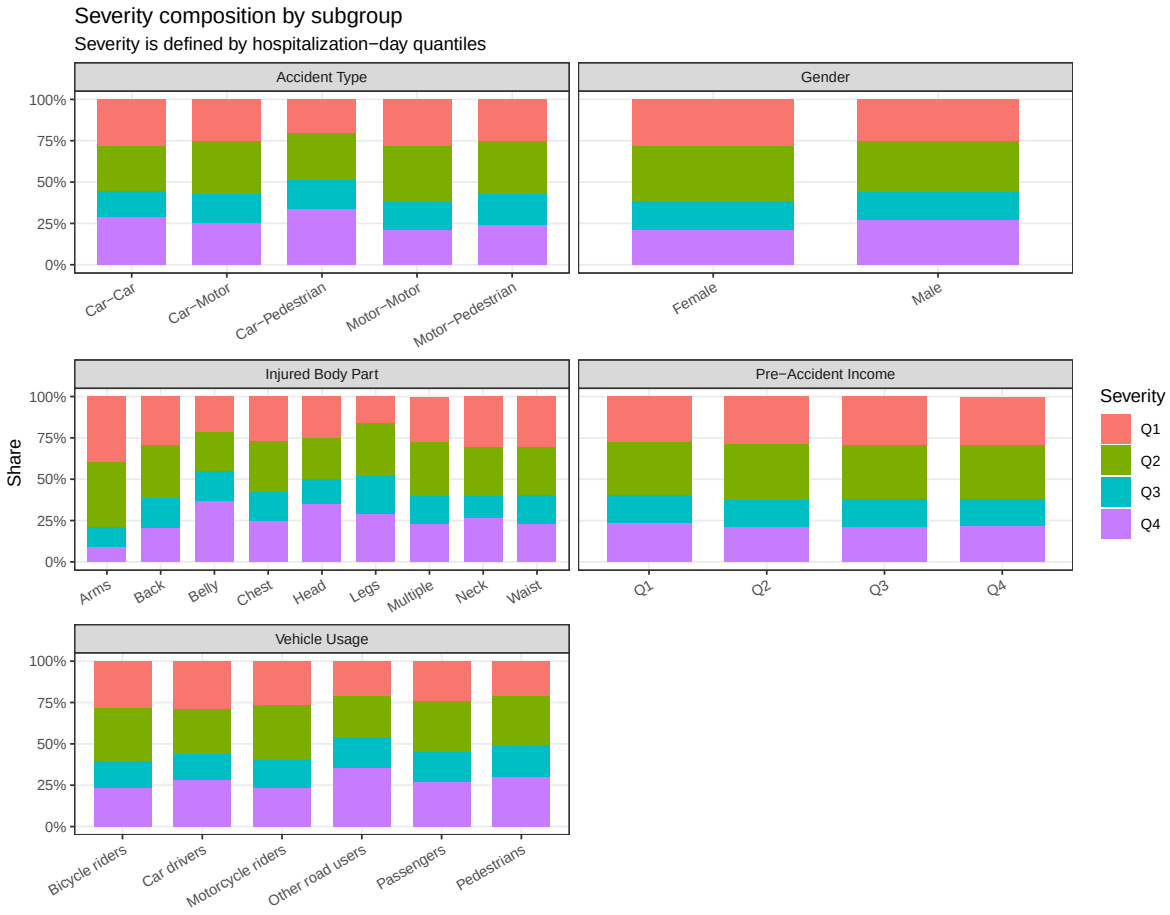


FIGURE 24: Severity Composition by Subgroup

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 8.

N.2 Heterogeneity by Pre-Accident Income Level

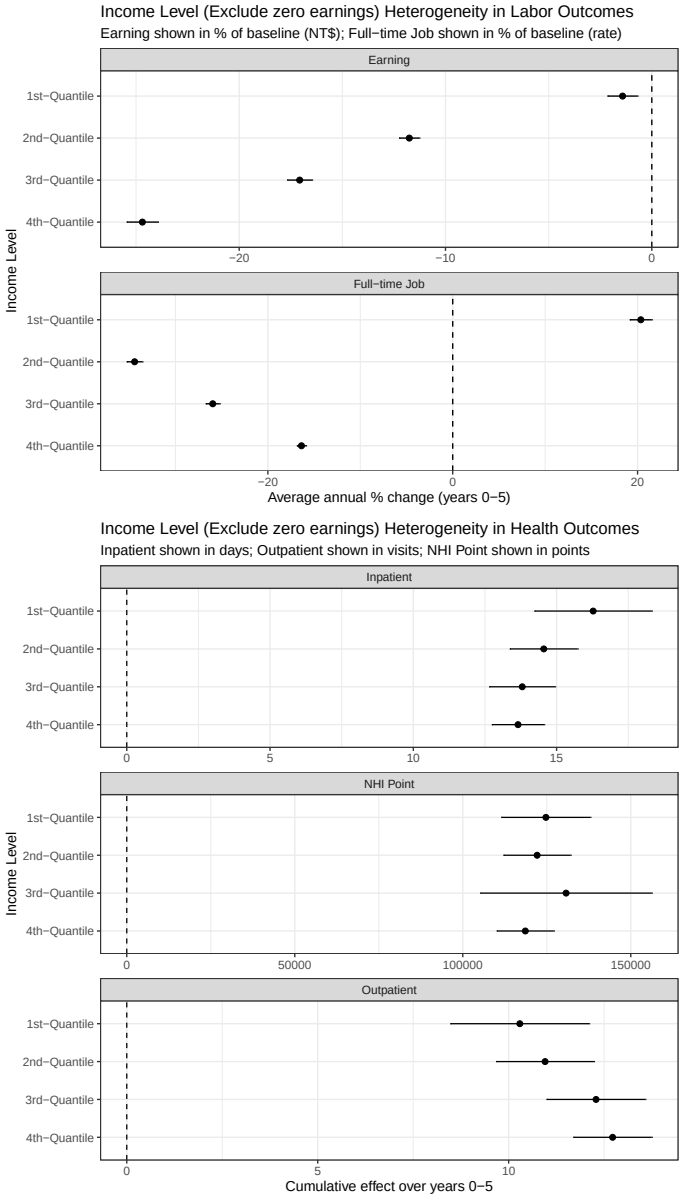


FIGURE 25: Heterogeneity by Pre-Accident Income Level

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

N.3 Heterogeneity by Injured Body Part

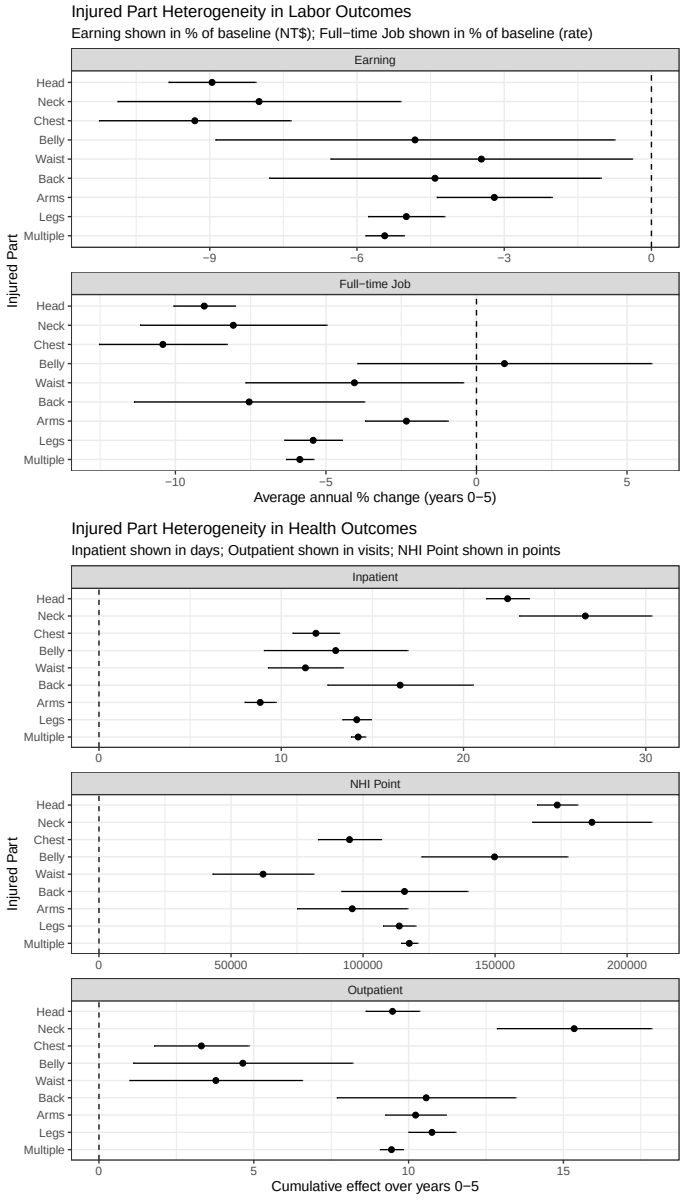


FIGURE 26: Heterogeneity by Injured Body Part

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

N.4 Heterogeneity by Vehicle Usage

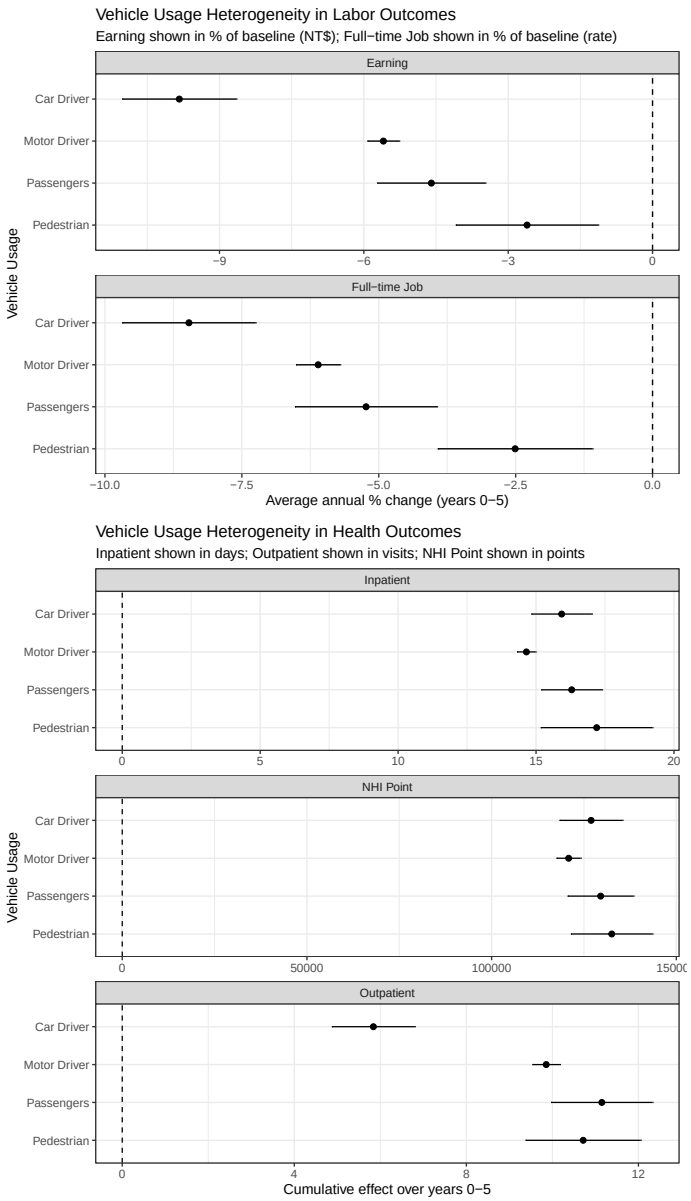


FIGURE 27: Heterogeneity by Vehicle Usage

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

N.5 Heterogeneity by Accident Case Type

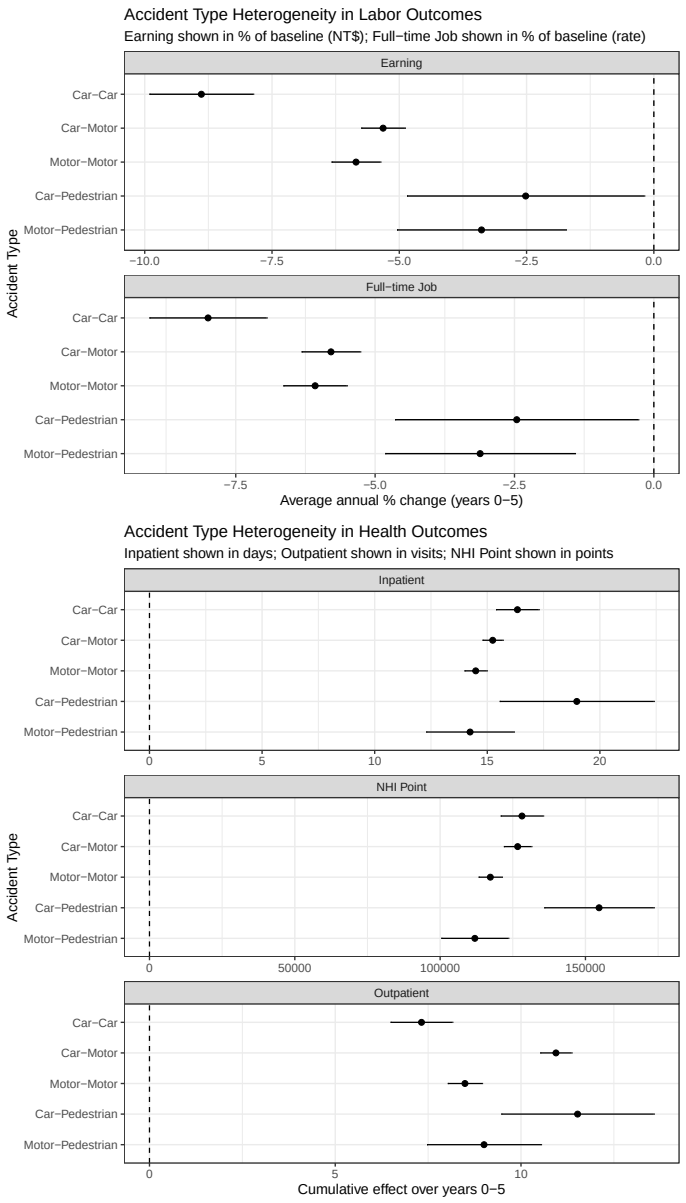


FIGURE 28: Heterogeneity by Accident Case Type

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

N.6 Heterogeneity by Gender

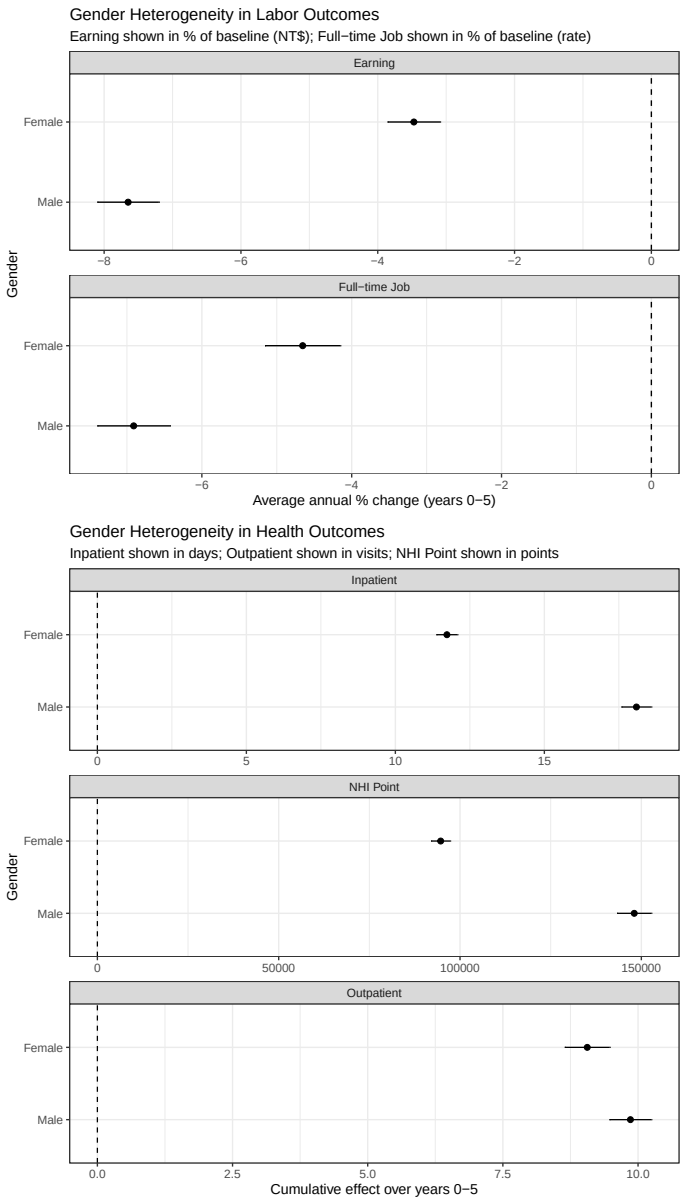


FIGURE 29: Heterogeneity by Gender

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

O 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 `FULL_TIME_WORK`=1. 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 8 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 8: 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 9 reports the implied bounds for the intensive contribution and the corresponding accounting total. These results use the non-balanced event-time Lee bounds.

Table 9: 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 10: 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.