Costs of Traffic Accidents: Accounting for the Survivors*

Kuan-Ming Chen† Chen-An Lien‡

March 12, 2026

Abstract

Evaluating investments in traffic safety requires an accurate accounting of costs and benefits. Conventional estimates typically rely on VSL-based scaling or recall exercises, which provide only coarse approximations and fail to capture the dynamic, realized consequences for survivors. We address this gap by constructing a 24-year population-level panel linking the universe of police-reported accidents in Taiwan to administrative health and employment registries. Using a staggered event-study design, we find that nonfatal injuries trigger persistent economic and health "scarring." Five years post-accident, annual earnings and employment probabilities remain 9% below baseline, while medical expenditures and cumulative mortality significantly increase. We estimate the average realized cost per injury at approximately USD 22,500. Because nonfatal injuries far outnumber fatalities, these long-run costs are a major component of the total economic burden that is often imprecisely estimated in traditional policy evaluations. Our findings suggest that incorporating these realized losses significantly increases the estimated returns to public safety infrastructure investments.

*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.

‡National Taiwan University. Email: ahn890209@gmail.com.

1 Introduction

Road traffic injuries impose a substantial and persistent burden on the global economy. Previous estimates suggest that economic losses from road traffic injuries will reach 0.12% of global GDP per year (Chen et al., 2019) or 3% when accounting for other losses such as medical expenditures (WHO, 2023). Among these costs, the foregone labor market earnings and medical expenditures due to injuries are the mostly cited costs. Governments rely on these estimates to decide investments in road safety and transport infrastructure, often around 1–2% of GDP in OECD economies (OECD International Transport Forum, 2023).

We argue that the basis behind these estimates can be much improved. Ideally, to estimate the labor market and medical costs associated with injuries, one would directly track the outcomes of injured individuals using administrative records, following changes in employment status, earnings, and medical utilization over time. In practice, however, even estimates of *injury* costs are largely derived from methods designed to value fatal risk reductions. Two methods dominate. The *human-capital* (HC) approach values a life by the present value of an individual's forgone market production, sometimes supplemented by direct medical costs. The *willingness-to-pay* (WTP) method infers a *Value of a Statistical Life* (VSL) from individuals' trade-offs between wealth and small changes in fatality risk, thereby capturing non-market losses such as pain and lost quality of life. WTP-based estimates are generally about twice as large as HC-based figures (Bougna et al., 2022); representative VSL values for high-income countries range between €1–4 million (CE Delft et al., 2019). Both approaches, however, are fundamentally centered on *fatality risk*.

Yet globally, road crashes cause about 1.19 million deaths annually but lead to 20–50 million nonfatal injuries (WHO, 2023), implying an injury-to-death ratio of roughly 17–42 to one.[1] Focusing primarily on fatalities therefore omits a substantial share of accident-related welfare losses.

Quantifying the long-run consequences of traffic injuries has been challenging. Many studies rely on cross-sections, short hospital records, or small surveys that cannot track individuals before and after accidents, making it difficult to isolate causal effects (Mohanan, 2013;

[1] This ratio is widely used in global road-safety discussions to emphasize the dominance of nonfatal injuries relative to fatalities.

Alam and Mahal, 2016). Others emphasize short-run medical costs or imputed productivity losses in estimating the economic burden of traffic injuries, which may understate persistent losses in employment, income, and health utilization (Blincoe et al., 2015). Although recent work has begun to exploit administrative data on traffic accidents, the research question differs from ours: Barr et al. (2025), for example, use minor car crashes as quasi-random financial shocks to study their broader economic consequences, rather than estimating the overall costs of injuries themselves.

This paper builds on and extends prior work by leveraging a newly constructed an administrative dataset linking the universe of police-reported traffic accidents in Taiwan to the National Health Insurance (NHI) and employment registries. The linkage yields a 24-year longitudinal panel covering the entire population, enabling us to track labor market outcomes and health utilization before and after accidents. We estimate the costs of injury by causally identify its impacts on labor market and medical utilization outcomes with a staggered-adoption event-study framework We use the "later-treated" as controls following Callaway and Sant'Anna (2021), comparing individuals injured in a given year to those who will be injured in the future. This design ensures comparability in both observed and unobserved characteristics while accommodating heterogeneous treatment timing. We focus on balanced event windows spanning five years before and after the accident to maintain consistent cohort composition.

We find that nonfatal traffic injuries impose large and persistent economic and health costs. On average, annual earnings and full-time employment probabilities decline by about 9% relative to pre-accident levels, with job separations spiking immediately after the accident and leaving lasting scars on employment and income trajectories. Health care utilization also rises substantially following injuries, including sizable increases in inpatient stays, outpatient visits, and NHI-covered medical spending. More severe injuries incur substantially larger economic and health costs. Labor-market losses are also greater for men than for women, while differences across income groups and road-user types primarily affect labor-market outcomes.

Our study makes three primary contributions to the literature on the economic and health consequences of physical shocks. First, we provide one of the first comprehensive

estimates of the total costs of nonfatal traffic injuries by integrating labor-market and health outcomes. While some research has begun to examine the impacts of traffic accidents on a smaller scale, these studies often focus in isolation on either financial medical expenses or foregone earnings using restricted samples. By tracking earnings, employment, medical utilization, and mortality simultaneously, we offer a more complete picture of the realized welfare losses borne by survivors. This approach ensures that we capture the full spectrum of post-accident life rather than just a single dimension of the clinical or economic shock.

Second, our analysis rests on a novel population-level data linkage spanning 24 years. We identify injured individuals from the universe of police-reported traffic accidents rather than from hospital discharge records. This distinction is critical because ICD codes in hospital data often document specific injuries or symptoms without recording the external cause. Furthermore, hospital records frequently miss injuries that do not result in immediate hospitalization. By linking police records to Taiwan's NHI and employment registries, we create a panel that captures the "missing middle" of injury severity. This data also allows us to identify cases designated as "catastrophic diseases," which is a formal status for severe, long-term conditions that serves as a proxy for disability.

Third, we enhance methodological rigor by employing a staggered event-study design (Callaway and Sant'Anna, 2021). This framework provides estimates that are robust to the heterogeneous treatment-timing and cohort-effect biases that can affect conventional two-way fixed effects models. By utilizing "not-yet-treated" individuals as a clean comparison group, we provide a robust causal interpretation of the post-injury path. This approach is consistent with the latest advancements in the econometrics of staggered treatment.

The studies closest to ours are Dano (2005) and Parro and Pohl (2021), who utilized administrative registers in Denmark and Chile to document lasting reductions in earnings and employment following road injuries. Our work extends this line of inquiry by expanding the scope of outcomes to include long-run health consequences and by utilizing a more representative identification strategy. Our approach avoids the selection issues and imcomplete data problems inherent in hospital-based datasets identified from ICD codes. Additionally, our use of modern event-study methods ensures that our results are robust to the timing and cohort biases that may have limited the causal reliability of earlier registry-based studies.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on Taiwan's road safety environment, health insurance, and compensation mechanisms. Section 3 describes the linked administrative datasets and sample construction. Section 4 outlines the empirical framework. Section 5 presents the main results and heterogeneity analysis. Section 6 conducts a back-of-the-envelope cost calculation and discusses policy implications. Section 7 concludes.

2 Background

2.1 Valuation Framework

Road safety policies rely on cost-benefit analysis to prioritize investments in infrastructure and enforcement. Central to these evaluations are estimates of the economic costs of traffic accidents, both fatal and nonfatal. Yet existing estimation methods face important conceptual and empirical limitations, leading to substantial variation in reported values across studies and institutional contexts. The most widely used approaches, VSL-based scaling and direct cost accounting, illustrate these challenges.

Public cost–benefit analyses often monetize safety improvements using the VSL, typically derived from hedonic wage studies or stated-preference surveys. VSL estimates are then scaled to value nonfatal injuries using severity weights, quality-adjusted life years (QALYs), or direct WTP elicitation for specific injury types. Alternatively, accounting-based methods aggregate observed medical costs and imputed productivity losses. Each approach faces well-documented limitations: revealed-preference methods suffer from risk endogeneity and worker self-selection; stated-preference surveys from hypothetical bias and limited information about long-run consequences; accounting approaches from incomplete cost coverage and uncertain productivity projections. Institutional differences across settings further generate substantial variation in estimates, limiting the usefulness of any single value as a universal shadow price (Ashenfelter, 2006).

These conceptual issues are reflected clearly in Taiwanese evidence. Hedonic-wage studies for Taiwan, beginning with Liu et al. (1997) and extended by subsequent studies, yield VSL estimates spanning roughly NT$5–50 million, depending on risk measurement, sample

composition, and model specification. Stated-preference and environmental-health studies produce similarly wide ranges due to differences in information structures and latency assumptions (Liu and Hammitt, 2004).

Transportation policy evaluations, however, typically rely on accident-cost estimates derived from compensation data and accounting approaches. Using administrative information on traffic accident cases, Chou et al. (2019) estimate per-fatality values of roughly NT$2.7–3.0 million for traffic accidents in Taiwan.

For nonfatal injuries, estimates of economic losses, defined as medical expenditures and productivity losses, also vary across valuation methods. Survey-based approaches measuring individuals' WTP to avoid traffic injuries imply costs of roughly NT$0.6 million per injured victim, while accounting-based estimates derived from court verdict data place the figure closer to NT$1.0 million. Although both approaches aim to measure the same economic consequences of injuries, the resulting estimates differ substantially.

Taken together, these estimates span a broad range across valuation approaches, highlighting how accident valuations depend strongly on measurement frameworks and institutional assumptions (see also Tsai et al., 2015).

These differences matter for policy. If shadow prices vary widely across domains, then the evaluation of road, safety programs becomes sensitive to the choice of valuation metric. In Taiwan, where fiscal resources devoted to safety programs have increased rapidly, relying on compensation-based benchmarks risks understating the social costs of injuries.

2.2 Research Context

2.2.1 Reporting and Classification of Traffic Accidents

When a traffic accident occurs in Taiwan, the involved parties may choose whether to report it to the police. While reporting is not mandatory, there is a strong incentive to do so: insurance companies generally require a police report to initiate the claims process. If 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.

This reporting system creates a standardized administrative trail for accidents involving personal injuries or deaths, particularly in A1 and A2 cases. 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. Established in 1995, it provides comprehensive access to medical services through a uniform benefits package and centralized administration.

In the aftermath of a traffic accident, injured individuals typically receive medical care without delay under the NHI. Even when legal liability has yet to be determined, the NHI covers most upfront medical costs. According to MOHW (2022), in 2022 the total medical expenditures associated with traffic accidents amounted to approximately 8.94 billion New Taiwan dollars. Of this amount, the NHI directly covered about 6.65 billion NT dollars (roughly 74%), while the remaining 26% was financed through other forms of insurance, both compulsory and private.

Most of this residual spending is borne by the CALI program, Taiwan's mandatory liability insurance for all vehicle owners. CALI provides standardized coverage for medical expenses and disability or death benefits up to statutory limits, and its payments complement NHI reimbursements. Although CALI also supports victims of hit-and-run accidents through

a supplementary compensation fund, the key point for our purposes is that NHI bears the majority of medical costs, with CALI functioning primarily as secondary payer.

Together, these mechanisms, universal public health insurance and mandatory liability coverage, form the financial infrastructure supporting accident victims in Taiwan. They ensure immediate access to care while also distributing costs between public and private payers.

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 victims whose impairments are severe enough to meet formal disability criteria.

For those who suffer complete loss of working capacity, CALI provides up to 2 million NT dollars 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. These compensation mechanisms provide partial income protection but leave substantial residual losses borne by injured individuals and their families.

3 Data

Our empirical analysis draws on rich administrative data that enable us to track individuals' labor market and health outcomes surrounding traffic accidents. We link the universe of police-reported traffic accidents to health and employment records from the national health insurance system to construct a 24-year longitudinal panel. The data provide person-level information on accident involvement, employment status, earnings, hospitalizations, and outpatient visits, observed both before and after the event. We describe each data source in detail below.

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 (deadly) and A2 (injured) accidents during this period. Minor A3 cases, which involve only property damage without bodily injury, are excluded. Since reporting is typically required for insurance claims and legal proceedings in A1 and A2 cases, coverage is near-universal for medically relevant accidents, as described in Section 2.2.

The NPA data are organized at two levels. The case-level file contains information for traffic accident. It includes a unique case identifier and detailed contextual information, including the date and location of the accident, the presumed cause as recorded by police officers and other details on accidents. The individual-level file contains information for every individual involved in the traffic accidents. It provides the case identifier, demographic, and role-specific information. This includes encrypted individual identifiers, gender, vehicle type, road-user role (e.g., driver, passenger, pedestrian), injury status, and the anatomical location of any injury.

These two datasets can be linked via unique case identifier, enabling 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

As our goal is to measure the aftermath of labor market outcomes and health-related information, we utilize 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. 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, out-

patient, 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 link the NPA accident records to administrative data from NHIRD using the unique individual identifier. By doing so, we construct an individual-level longitudinal dataset that allows us to examine the long-term health and labor outcomes associated with traffic-related injuries.

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 construct a matched comparison group never experienced any 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 (NTD/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 accident victims; *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.

3.4 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 accident victims and never-injured individuals, but a substantial gap in earnings. Among working-age individuals, the employment rate is 44.1% for future accident victims and 43.9% for their never-injured counterparts.[2] In contrast, average annual earnings are markedly lower for future accident victims (229,110 NTD) than for the never-injured group (248,344 NTD). 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 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 eventually become accident victims are not comparable to those who are never injured, even after controlling for basic demographics. Despite for the common belief that traffic accidents are a textbook example of exogenous shocks, it turns out that those who experience accidents are quite different from those who never experienced one.

[2]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.

4 Method

4.1 Identification Strategy

Rather than comparing injured individuals to those never injured, groups that differ systematically in unobserved risk exposure, we use a "later-treated" comparison group: individuals who will themselves be injured in the future but have not yet experienced an accident. This ensures closer similarity in both observed and unobserved characteristics.

Treatment is defined as an individual's first traffic accident injury, occurring in calendar year g, which defines the treatment cohort. Event time is measured relative to g, with $t < 0$ denoting pre-accident years and $t > 0$ post-accident years. The comparison group for cohort g at event time t consists of individuals not yet injured who will be in a later year.

This structure corresponds to a staggered adoption setting: different cohorts enter treatment in different calendar years, and treatment effects may vary both across cohorts and over time since the accident. Standard two-way fixed effects estimators are known to be biased in such settings when treatment effects are heterogeneous (Goodman-Bacon, 2021). We therefore adopt the framework proposed by Callaway and Sant'Anna (2021), which estimates treatment effects by comparing outcomes for each treated cohort to appropriate later-treated controls at the same event time, and then aggregates these cohort-specific effects to recover the overall dynamic treatment profile.

The key identifying assumption is that, in the absence of treatment, outcomes of treated and not-yet-treated individuals would have followed parallel trajectories over event time. A sufficient condition is that, conditional on observed covariates such as age, sex, and location, the *timing* of an individual's first accident is as good as random. While the probability of ever being injured may correlate with unobserved risk preferences or exposure, the precise year in which an accident occurs is largely idiosyncratic, depending on stochastic factors such as weatherand other traffic conditions. Under this quasi-random timing assumption, later-treated individuals provide a credible counterfactual.

4.2 Estimation Framework

Building on the identification strategy in Section 4.1, we estimate dynamic treatment effects in a setting where treatment timing is staggered and effects may vary across both cohorts and event time. We follow the framework of Callaway and Sant'Anna (2021), which accommodates these features and uses the "later-treated" group as the counterfactual for each cohort-year cell.

The key parameter of interest is the group-time average treatment effect, denoted $\text{ATT}(g, t)$, which measures the causal effect for individuals first treated in year g at relative time t, compared to individuals who have not yet been treated by that year. This formulation allows the comparison group to consist of individuals who will themselves receive treatment in the future, thereby improving comparability along both observed and unobserved dimensions.

Formally, the estimator for $\text{ATT}(g, t)$ is based on a doubly robust approach that combines inverse probability weighting (IPW) with outcome regression. For each group g and event time t, we estimate:

$$\text{ATT}(g,t) = \mathbb{E}\left[\left(\frac{G_g}{\hat{e}(X)} - \frac{C_{gt}}{1-\hat{e}(X)}\right)(Y_t - \hat{\mu}(X))\right],$$

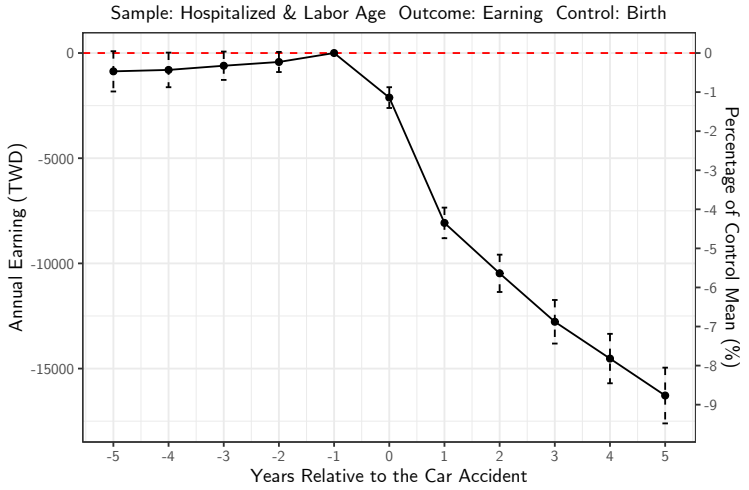
where G_g is an indicator for being first treated in period g, C_{gt} is an indicator for being not yet treated at time t, $\hat{e}(X)$ is the estimated propensity score conditional on covariates X, and $\hat{\mu}(X)$ is the predicted outcome from a pre-treatment regression. Covariates used for adjustment include sex, age, location, and calendar year fixed effects.

To illustrate, consider $\text{ATT}(2012, 5)$. This parameter measures the average difference in outcomes in year 2017 (five years after treatment) between individuals who first experienced a traffic accident in 2012 and those who would experience their first accident after 2017. In this setup, the latter group serves as the "later-treated" control, ensuring that both groups are subject to a similar underlying risk profile.

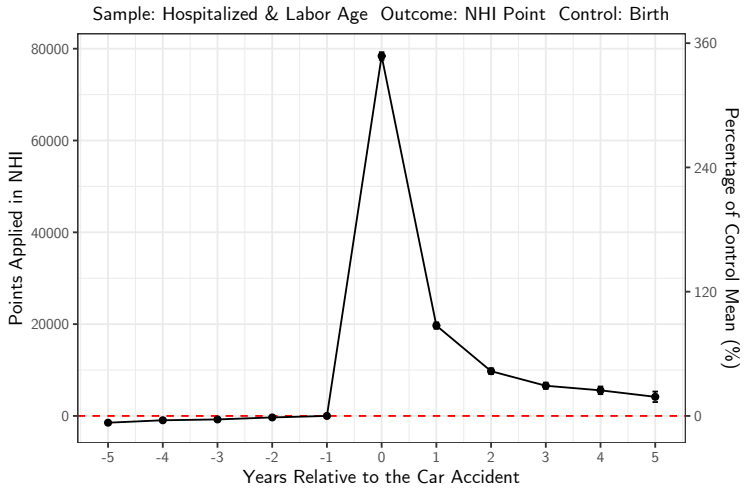
Given potential treatment effect heterogeneity, our main analysis focuses on a balanced event-time window from $t = -5$ to $t = +5$. This choice ensures that comparisons at each relative time are based on the same set of cohorts, thereby avoiding changes in sample composition that could arise if the 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.

5 Results



(A) Earnings (birth-year controls)



(B) NHI points (birth-year controls)

FIGURE 1: 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. Medical spending is measured in NHI points, a standardized reimbursement unit under Taiwan's NHI (NT$0.9 per point).

5.1 Dynamic Results

Before presenting the findings, we briefly summarize the empirical setup underlying the dynamic analysis. All estimates control for birth cohort fixed effects, which absorb life-cycle differences across individuals and prevent spurious pre-trends driven by cohort composition. Using later-treated individuals as controls, we estimate event-study specifications over a balanced window of $t \in [-5, 5]$, normalized to $t = -1$. We examine seven outcomes spanning labor market performance and health care utilization: earnings, full-time employment, job separation (*lost job*), firm attachment (*stay same*), inpatient days, outpatient visits, and total NHI points. Figure 1 presents the dynamic effects for earnings and medical expenditures; the left axis reports effects in natural units, while the right axis rescales the same estimates as percentages relative to the pre-accident control mean. Results for the remaining outcomes are presented in Appendix C.

Finding 1: Persistent and deepening earnings and employment losses. Traffic accident injuries lead to substantial and persistent declines in labor-market outcomes. Annual earnings remain stable prior to the accident but drop sharply in the accident year and continue to decline thereafter. Notably, the initial decline at $t = 0$ is modest compared to subsequent years: the losses deepen progressively over time, with annual earnings falling approximately 9 percent below the pre-accident control mean by year five, substantially larger than the immediate post-accident effect.

Employment attachment follows a similar deepening trajectory. The probability of holding a full-time job declines immediately after the accident and continues to fall throughout the five-year window, with no evidence of recovery. By year five, full-time employment is roughly 9 percent lower relative to the control mean.

Together, these results indicate that traffic injuries generate long-lasting labor-market scars, affecting both income and employment stability well beyond the initial health shock. The deepening trajectory suggests that the economic consequences of injury accumulate progressively rather than manifesting as a one-time level shift.

Finding 2: Short-run separations drive long-run employment gaps. To understand the mechanism underlying these persistent losses, we examine both job separations and firm attachment. The probability of losing a job (Figure 3b), a flow measure of employment separation, spikes sharply in the accident year. This increase is short-lived, with separation rates returning close to baseline in subsequent years.

In contrast, the probability of remaining in the same firm (Figure 3c) declines immediately after the accident and continues to fall over time. Five years after the accident, injured individuals are roughly 2 percentage points less likely to remain with their pre-accident employer, corresponding to a 6–7 percent reduction relative to the control mean.

Taken together, these patterns reveal a coherent adjustment process. Traffic accidents first induce an acute separation shock that disrupts existing worker–firm matches. Although the initial spike in separations is temporary, the destruction of firm attachment translates into persistent reductions in employment stability and earnings.

Finding 3: Elevated and sustained health care utilization. Medical utilization rises sharply at the time of the accident and remains elevated for several years thereafter. Inpatient days increase dramatically in the accident year, by roughly 10 days on average, reflecting the acute hospitalization episode. Although inpatient utilization declines after $t = 0$, it remains modestly above baseline for several subsequent years. Outpatient visits also rise substantially: utilization increases by approximately 42% in the first post-accident year and remains persistently elevated through $t = +5$. These patterns indicate that traffic injuries generate not only an acute hospitalization shock but also long-lasting demand for outpatient care and follow-up treatment.

Finding 4: Significant increases in medical expenditures. Consistent with heightened utilization, total NHI points—a comprehensive measure of reimbursed medical spending—jump sharply in the accident year. Expenditures increase by approximately 80,000 points at $t = 0$, corresponding to roughly a 360% increase relative to the pre-accident control mean. While spending declines substantially after the initial year, it remains persistently above baseline throughout the five-year horizon, indicating ongoing medical needs and sustained fiscal pressure on the public health insurance system.

To assess the sensitivity of these dynamics to the choice of event window, Appendix D extends the analysis to windows of $[-7, 7]$ and $[-10, 10]$. The qualitative patterns remain unchanged, though estimates become less precise at longer horizons. Appendix E further reports specifications without covariate adjustment, showing that birth-cohort controls primarily improve pre-trend fit rather than alter the substantive conclusions.

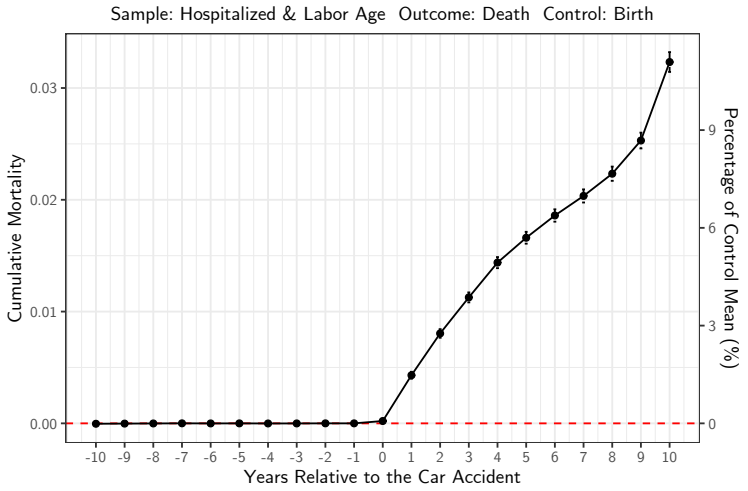


FIGURE 2: Dynamic Effects of Traffic Injuries on Mortality (Birth-Year Controls)

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.

Finding 5: Persistent increases in post-accident mortality risk. We next examine the effect of traffic accident injuries on mortality.Unlike the previous analyses that use later-treated individuals as controls, this panel uses a never-treated population as the comparison group. Although earlier sections documented observable differences between ever-injured and never-injured individuals, the estimates exhibit no differential pre-trends: coefficients in the pre-accident period are tightly centered around zero, suggesting similar baseline trends in cumulative mortality prior to the accident.

Figure 2 shows almost no immediate change in the accident year ($t = 0$). The divergence emerges in the first post-accident year, when the cumulative death indicator increases to approximately 0.005. Thereafter, the effect grows gradually and monotonically over time, consistent with an accumulating mortality gap between the injured and never-treated groups.

By five years after the accident, the cumulative increase reaches roughly 0.016. Relative to the control-group mean of 0.32, this corresponds to an increase of about 5 percent. The accumulation continues in the longer run, reaching approximately 0.03 ten years after the accident, or over 9 percent of the control mean.

Taken together, these dynamics indicate that traffic injuries are associated with a persistent and accumulating increase in mortality over time. The absence of pre-trends and the delayed, gradually widening post-treatment gap are consistent with long-run health deterioration or delayed complications following injury, rather than an effect driven solely by immediate fatal outcomes.

5.2 Heterogeneity Results

This section summarizes salient patterns of heterogeneity in post-accident impacts. The complete set of heterogeneity figures is reported in Appendix F.

Finding 6: Injury severity scales the magnitude of post accident losses. As shown in Figure 7, injury severity is the dominant source of heterogeneity across outcomes. More severe injuries are associated with substantially larger cumulative health care utilization and deeper labor-market losses, while mild injuries generate comparatively modest effects. Importantly, although the level of losses increases sharply with severity, the qualitative post-accident dynamics remain similar across groups, indicating that injury intensity primarily acts as a scale factor rather than altering the underlying adjustment process.

Finding 7: Income heterogeneity is concentrated in labor-market outcomes. Figure Figure 8 shows that, among individuals with positive earnings in the year prior to the accident, income-related heterogeneity is pronounced in labor-market impacts but limited in health outcomes. Higher-income groups experience substantially larger post-accident earnings and full-time employment losses, with the largest effects observed for those in the top income quartile. In contrast, health outcomes are broadly similar across income groups, with overlapping confidence intervals.

Finding 8: Injury location matters for both economic and health consequences. As illustrated in Figure 9, injuries to functionally critical or multiple body regions, such as the head, neck, or multi-site trauma, generate the largest and most persistent impacts on both health care utilization and labor-market outcomes. Limb injuries, by contrast, are associated with milder and more transient losses. This ranking highlights the role of functional impairment in shaping long-run economic and health consequences.

Finding 9: Health and economic impacts differ across road users. As shown in Figure 10 and Figure 11, health losses following traffic accidents are broadly similar across road users. While pedestrians and passengers exhibit slightly higher inpatient and outpatient utilization on average, the magnitude of health impacts does not differ dramatically across groups, and confidence intervals overlap substantially.

In contrast, economic losses display a markedly different pattern. Car drivers experience the largest declines in earnings and employment after accidents, exceeding those observed for other road users. This divergence indicates that heterogeneity in post-accident impacts is more pronounced along economic dimensions than health dimensions, and that physical exposure alone does not fully account for the distribution of economic losses.

Finding 10: Gender differences are more pronounced in labor outcomes. As shown in Figure 12, gender heterogeneity differs across outcome domains. Health care utilization and medical expenditures are broadly similar for male and female victims. In contrast, labor-market effects are systematically larger for men: average earnings and employment losses among male victims are approximately 1.5 to 2 times those observed for female victims. These patterns suggest that gender differences in post-accident impacts operate primarily through labor-market channels rather than medical utilization.

Finally, building on our event-study estimates of economic and medical losses, the next section conducts a simple back-of-the-envelope comparison between our implied costs and commonly used VSL.

6 Back-of-the-envelope costs and a VSL benchmark

Conventional approaches to valuing nonfatal traffic injuries often rely on stated-preference surveys or apply fixed ratios derived from VSL estimates for fatalities. These methods, while expedient, cannot capture the dynamic and heterogeneous nature of injury costs revealed by longitudinal administrative data. Our event-study design offers a more transparent alternative: measure realized outcomes across multiple domains over extended follow-up periods, and allow the data to reveal how costs evolve and vary across different victim profiles. Rather than assuming a single injury-severity weight or a uniform time path, this approach documents progressive deepening in labor market losses and rising medical needs, providing direct empirical evidence on the magnitude and trajectory of injury costs.

Aggregating cumulative medical and earnings losses within five years after hospitalization yields a per-injury cost of approximately TWD 176,000 (TWD 112,000 in excess medical spending plus TWD 64,000 in earnings losses). Extending this calculation to a twenty-year horizon, by mechanically projecting observed ten-year earnings losses, implies a total cost of about TWD 722,000 per injured individual. This estimate falls within the range of prior Taiwan-based valuations (TWD 600,000 to TWD 1 million). The gap relative to the upper bound primarily reflects our focus on tangible medical and labor market losses, omitting intangible components such as pain and suffering, psychological distress, and long-term quality-of-life reductions that figure prominently in legal damages.

A key advantage of our approach is that it reveals substantial heterogeneity in injury costs, variation that fixed-ratio VSL methods, by construction, cannot capture. Traditional approaches typically apply a single average weight to all nonfatal injuries, implicitly assuming similar long-run impacts across cases. The heterogeneity patterns documented in Appendix F, however, show that realized costs vary dramatically. Injury severity acts as a scale factor, with more severe cases generating losses that far exceed milder injuries. Vehicle-to-vehicle collisions impose labor-market losses roughly twice as large as other accident types, and car drivers face losses 1.5 to 3 times larger than other road users. Male victims experience labor losses 1.5 to 2 times larger than females. Income heterogeneity further reveals that job risk varies across earnings levels: top-quartile workers exhibit the largest propor-

tional earnings declines but the smallest employment losses, while lower-income workers face greater job instability.

This heterogeneity implies that uniform valuation methods systematically underestimate the economic burden for severe injuries, vehicle-to-vehicle collisions, car drivers, and male workers, while overestimating costs for milder cases. Where local injury profiles or demographic compositions diverge from conventional VSL reference populations, such biases can meaningfully distort policy priorities.

These estimates carry implications for infrastructure investment decisions. In Taiwan, nonfatal inpatient injuries outnumber traffic fatalities by a ratio of roughly eight to one. If cost-benefit analyses rely exclusively on fatality-avoidance benefits, typically valued at VSL estimates around TWD 2 million, they may systematically undervalue interventions that reduce injury severity or prevent nonfatal crashes. Incorporating the long-run costs of nonfatal injuries documented here would increase the estimated benefits of safety infrastructure. Our heterogeneity analysis further sharpen this point: more severe injuries generate disproportionately large long-run losses, and vulnerable road users, pedestrians, passengers, and motorcyclists, experience particularly severe health consequences when involved in crashes. These findings suggest that infrastructure investments targeting injury-severity mitigation and vulnerable-road-user protection may yield substantial social returns that conventional fatality-focused evaluations overlook.

7 Conclusion

This paper addresses these limitations by leveraging 24 years of longitudinal administrative data linking police-reported traffic accidents in Taiwan to national health insurance and employment records. Using a staggered event-study design with later-treated controls, we document the economic and health consequences of nonfatal traffic injuries for working-age individuals hospitalized after crashes.

We find that traffic injuries generate progressively deepening economic losses, annual earnings and employment decline by approximately 9 percent by year five, driven by initial job separations that permanently reduce firm attachment—and persistent medical burdens,

with outpatient visits rising 42 percent and medical expenditures increasing 360 percent in the first post-accident year, remaining elevated throughout the follow-up. Cumulative mortality rises progressively, reaching 5 percent above control levels by year five. Importantly, injury costs exhibit substantial heterogeneity: vehicle-to-vehicle collisions, car drivers, male victims, and severe injuries generate losses 1.5 to 3 times larger than other cases, while income heterogeneity reveals differential job-risk exposure across earnings levels. Fixed-ratio VSL methods systematically underestimate the burden for high-cost subgroups while overestimating costs for milder cases.

Aggregating cumulative medical and earnings losses over a twenty-year horizon yields a per-injury cost of approximately TWD 722,000, falling within the range of prior Taiwan-based valuations but derived from realized outcomes. In Taiwan, where nonfatal inpatient injuries outnumber traffic fatalities by roughly eight to one, incorporating these costs into infrastructure cost-benefit analyses would substantially increase the estimated returns to safety investments, particularly those targeting injury-severity mitigation and vulnerable-road-user protection. Evaluations that focus exclusively on fatality avoidance understate the true social burden of traffic accidents. Accounting for the enduring and heterogeneous economic and medical losses documented here is essential for infrastructure planning, insurance design, and the allocation of public safety resources. Improving road safety yields benefits not only through lives saved but also by preventing persistent losses among the far larger population of injured survivors.

References

Alam, K. and A. Mahal (2016). Economic impacts of road traffic injuries on households in south asia. *PLOS ONE 11*(10), e0164362.

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.

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.

Chen, S., M. Kuhn, K. Prettner, and D. E. Bloom (2019). The global macroeconomic burden of road injuries: estimates and projections for 166 countries. *The Lancet Planetary Health 3*(9), e390–e398.

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.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

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.

Mohanan, M. (2013). Causal effects of health shocks on consumption and debt: Quasi-experimental evidence from bus accident injuries. *Review of Economics and Statistics 95*(2), 673–681.

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.

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.

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

A Accident-Level Summary Statistics

Table 2: 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 2 (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 2 (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 3: T-test of labor outcome variables on victims before accident occurrence

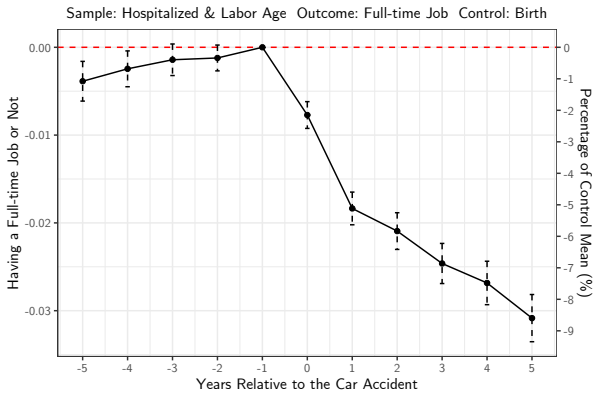
Dependent Variable	Earning	Full-time Job	Lost-Job
Constant	111,704.4	0.2030	0.058
	(96.97)	(0.0001)	(0.0004)
Injured	-7,382.2	-0.0011	-0.012
	(130.2)	(0.0002)	(0.0003)
S.E.: Clustered	by: ID	by: ID	by: ID
Observations	145,722,410	145,722,410	145,722,410
R2	0.18921	0.19705	0.10775

Table 4: T-test of health outcome variables on victims 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 data belongs to a traffic accident victim 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, the performance of these victims was already worse 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.

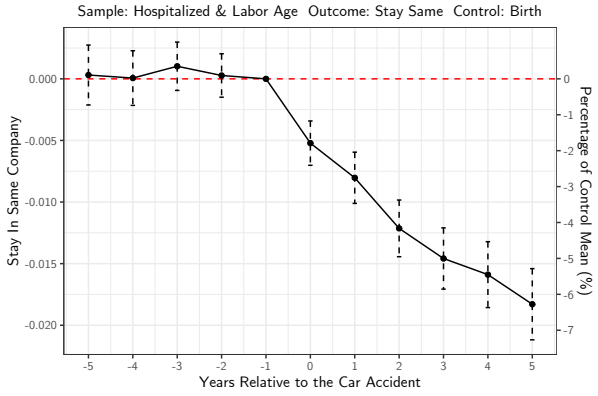
C More Main Results



(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)

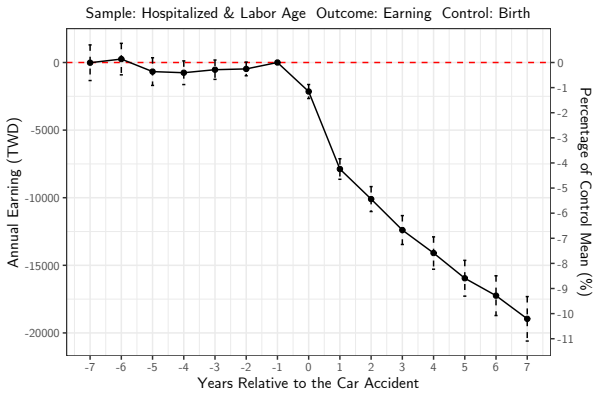


(E) Outpatient visits (birth-year controls)

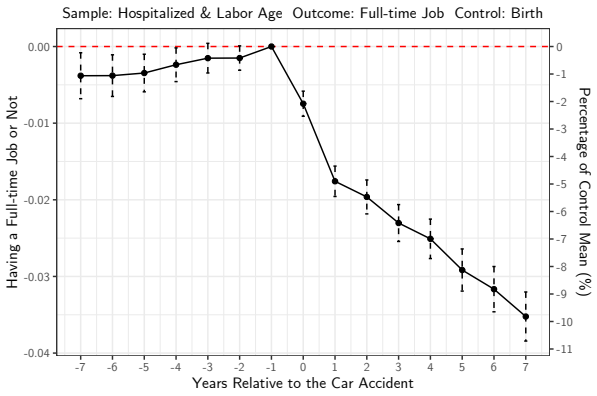
FIGURE 3: Dynamic Effects of Traffic Injuries On Labor and Health Outcomes

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

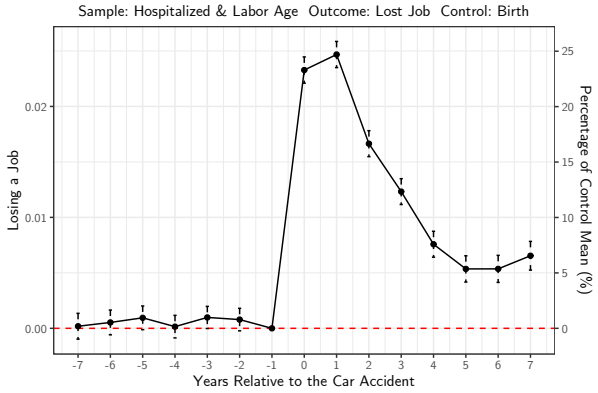
D Extended Window: 7 & 10 Years



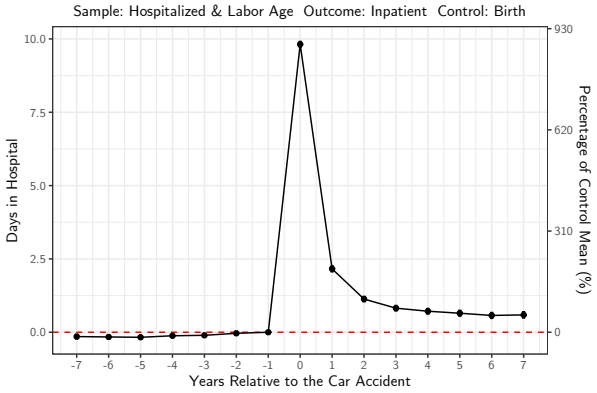
(A) Earnings (birth-year controls)



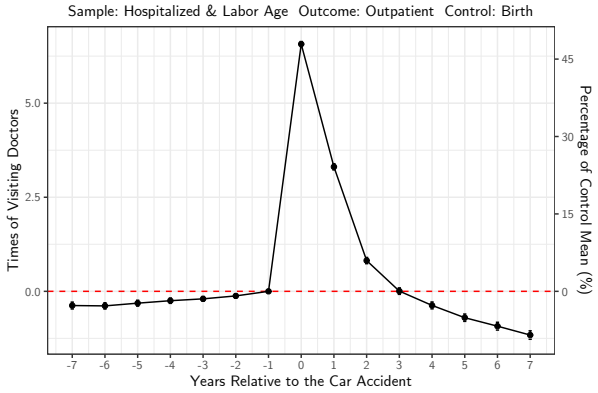
(B) Full-time employment (birth-year controls)



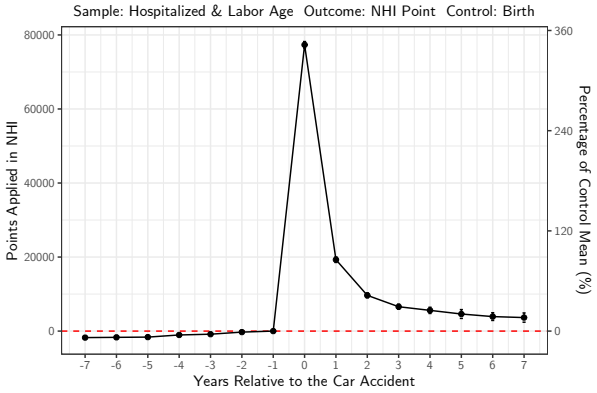
(C) Losing Job (birth-year controls)



(D) Inpatient days (birth-year controls)



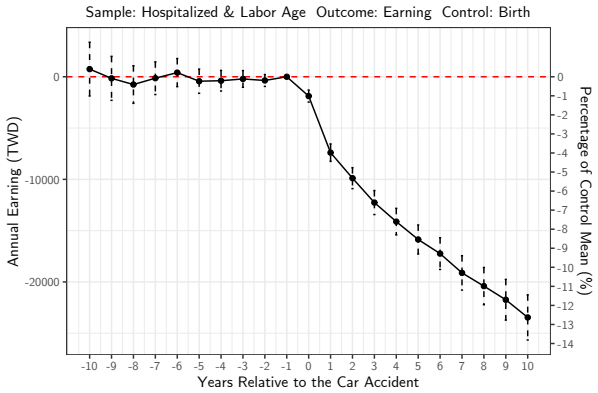
(E) Outpatient visits (birth-year controls)



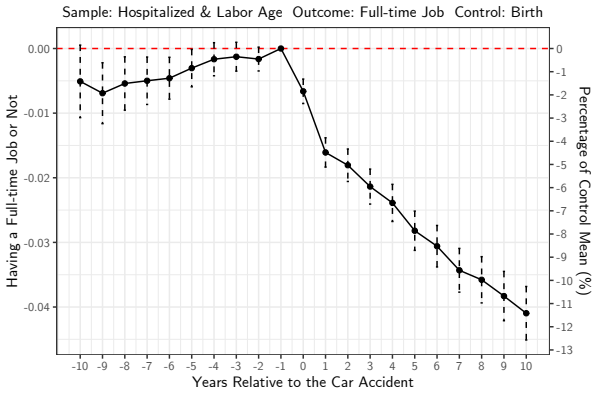
(F) NHI points (birth-year controls)

FIGURE 4: Dynamic Effects over a 7-Year Window

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



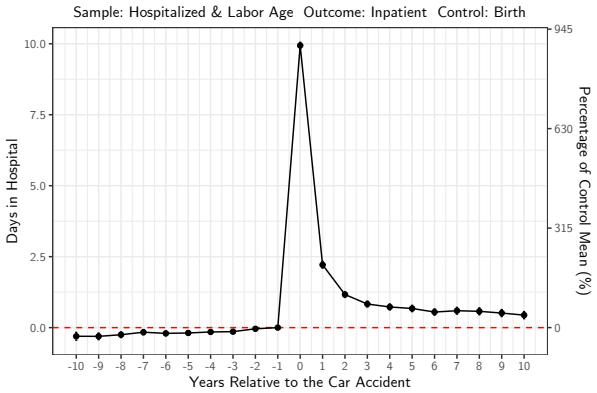
(A) Earnings (birth-year controls)



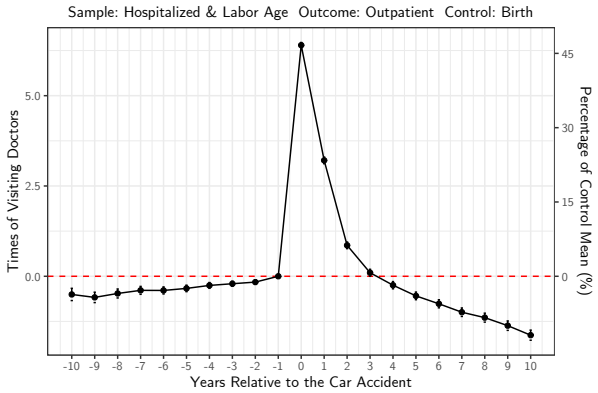
(B) Full-time employment (birth-year controls)



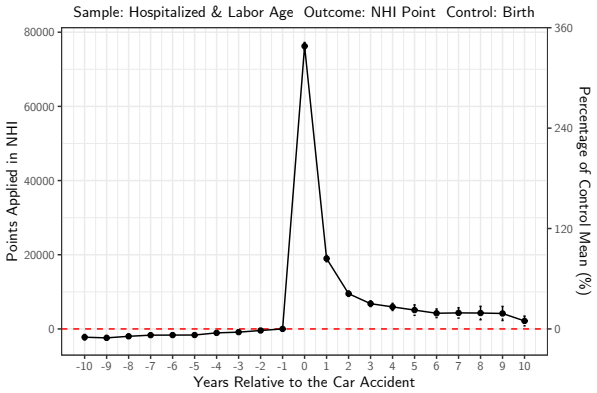
(C) Losing Job (birth-year controls)



(D) Inpatient days (birth-year controls)



(E) Outpatient visits (birth-year controls)

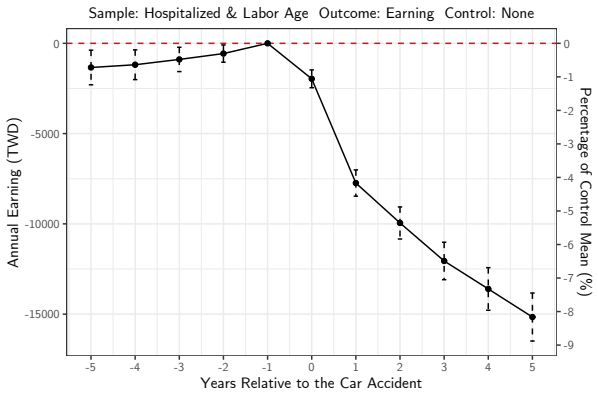


(F) NHI points (birth-year controls)

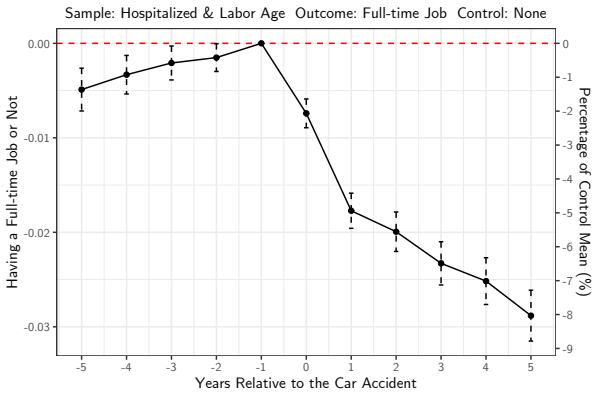
FIGURE 5: Dynamic Effects over a 10-Year Window

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

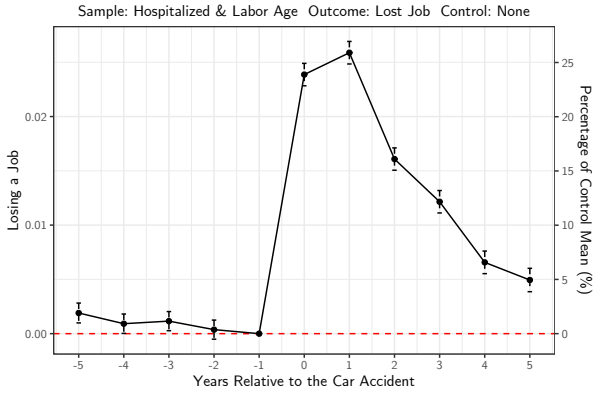
E Main Results without Controlling Birth Year



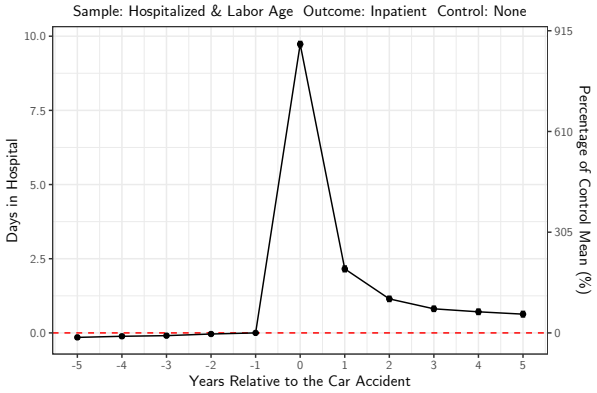
(A) Earnings



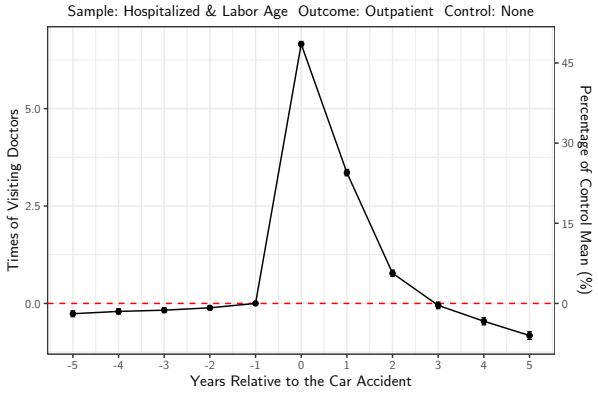
(B) Full-time employment



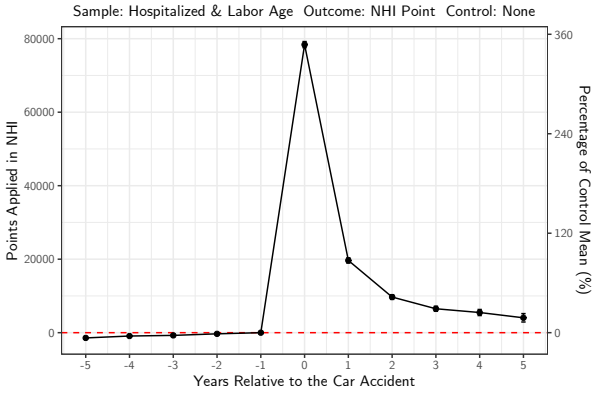
(C) Losing Job



(D) Inpatient days



(E) Outpatient visits



(F) NHI points

FIGURE 6: Dynamic Effects of Traffic Injuries without Covariates

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

F Heterogeneity Analysis

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

F.1 Heterogeneity by Injury Severity Level

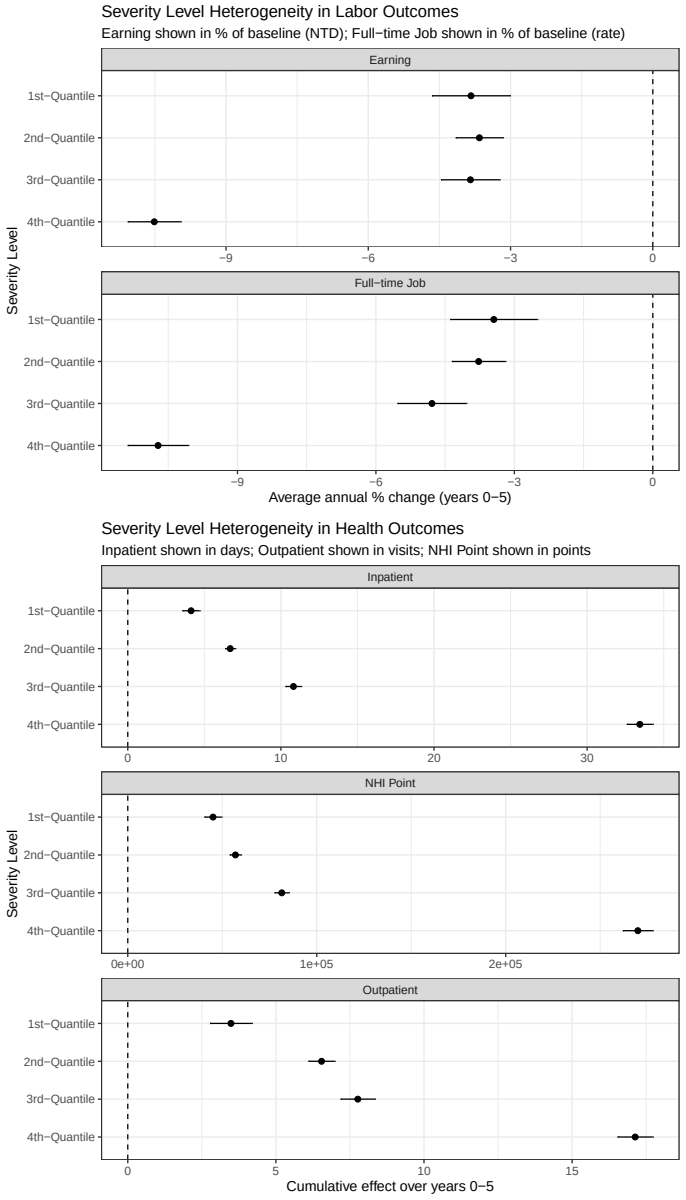


FIGURE 7: Heterogeneity by Injury Severity Level

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.

F.2 Heterogeneity by Pre-Accident Income Level

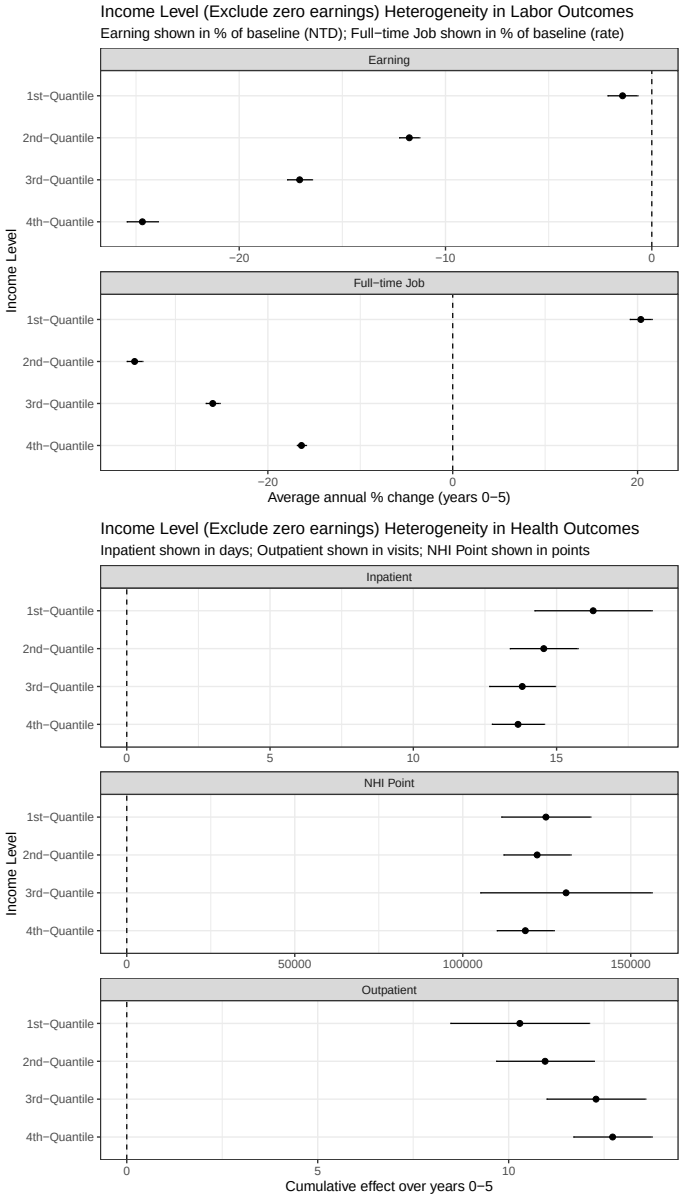


FIGURE 8: Heterogeneity by Pre-Accident Income Level

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.

F.3 Heterogeneity by Injured Body Part

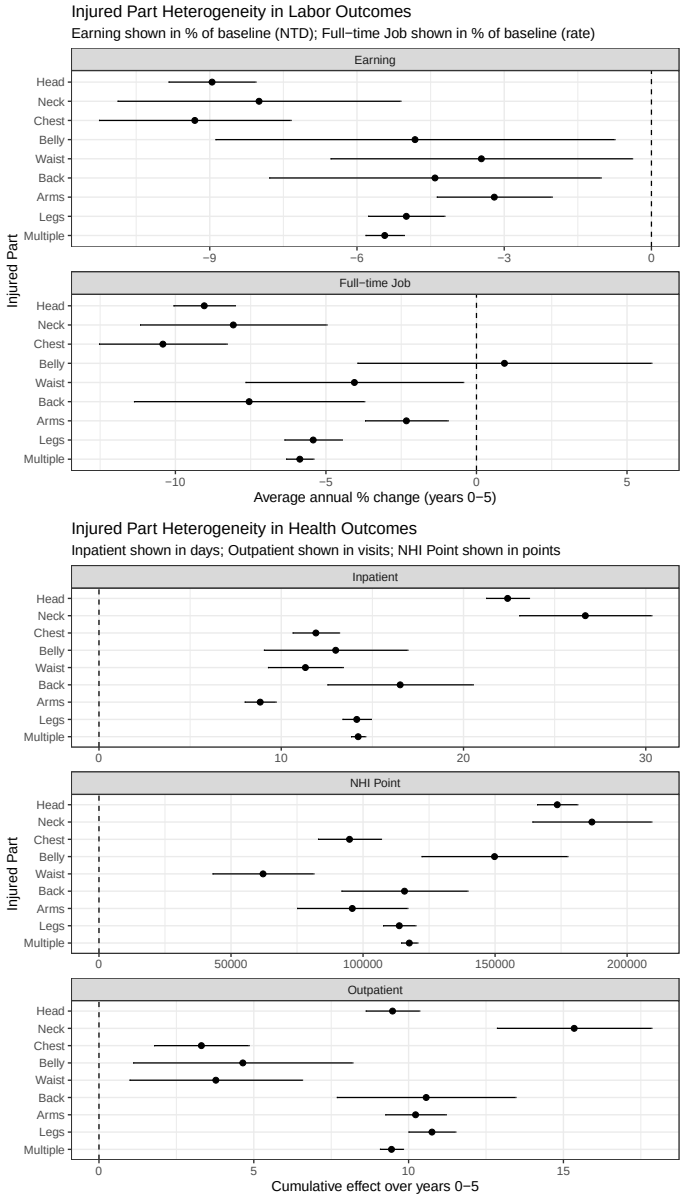


FIGURE 9: Heterogeneity by Injured Body Part

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.

F.4 Heterogeneity by Vehicle Usage

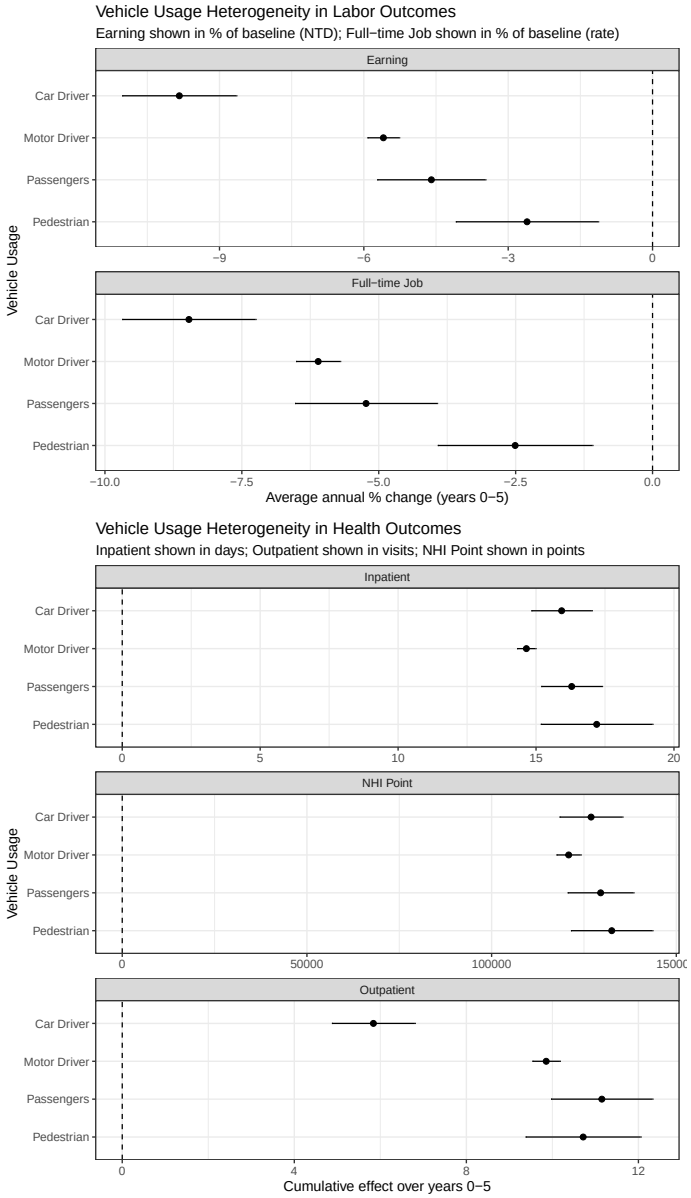


FIGURE 10: Heterogeneity by Vehicle Usage

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.

F.5 Heterogeneity by Accident Case Type

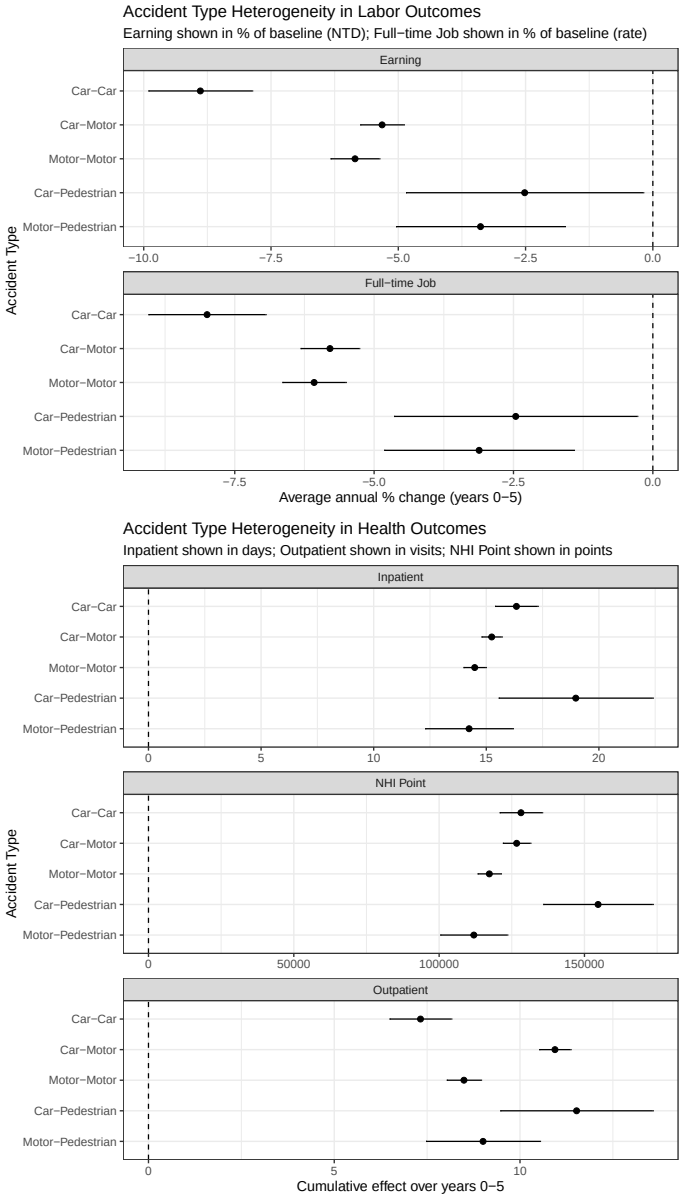


FIGURE 11: Heterogeneity by Accident Case Type

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.

F.6 Heterogeneity by Gender

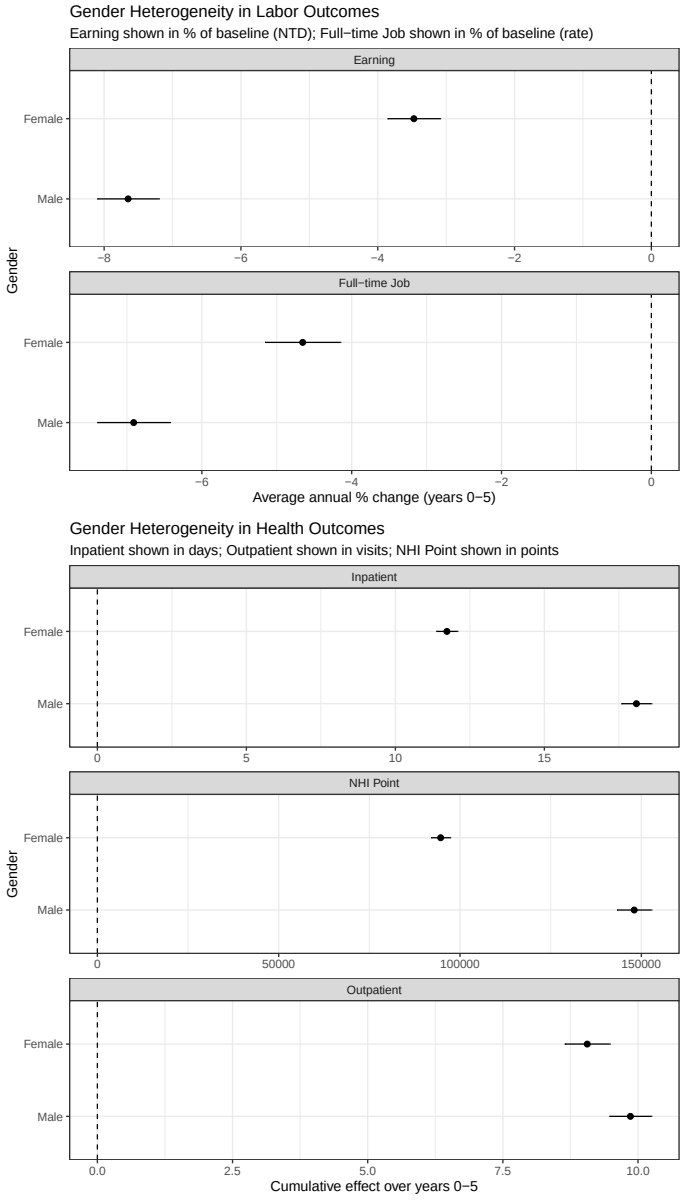


FIGURE 12: Heterogeneity by Gender

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.