

The Unintended Labor Market Consequences of Teen Driving Restrictions: Evidence from Graduated Driver Licensing Laws

Autonomous Policy Evaluation Project

Working Paper

January 2026

Abstract

Graduated Driver Licensing (GDL) laws, adopted by all U.S. states between 1996 and 2012, impose nighttime driving restrictions and passenger limits on newly licensed teenage drivers to reduce traffic fatalities. While the safety benefits of these laws are well-documented, their effects on teen labor market outcomes remain unexplored. Using a triple-difference design that exploits variation in GDL adoption timing across states and differential exposure by age, I estimate the causal effect of GDL laws on teen employment. I find that GDL adoption is associated with a 4.3 percentage point reduction in employment among 16-17 year-olds relative to 20-24 year-olds. This effect emerges sharply at the time of GDL adoption and persists over the following decade. The results suggest that policies designed to protect teen safety may have meaningful unintended consequences for their labor market entry, with potential implications for skill development and long-run career trajectories.

JEL Codes: J22, J13, R41, I18

Keywords: Graduated driver licensing, teen employment, labor force participation, traffic safety, unintended consequences

Contents

1	Introduction	4
2	Background and Literature Review	6
2.1	Graduated Driver Licensing in the United States	7
2.2	Prior Research on GDL and Traffic Safety	8
2.3	Teen Labor Markets in the United States	9
2.4	Transportation Constraints and Employment	10
3	Data	11
3.1	Census PUMS Data	11
3.2	GDL Policy Data	12
3.3	Sample Construction	13
3.4	Summary Statistics	13
4	Empirical Strategy	14
4.1	Triple-Difference Design	14
4.2	Estimating Equations	15
4.3	Identification Threats and Robustness	16
5	Results	17
5.1	Main Triple-Difference Estimates	17
5.2	Event Study Results	18
5.3	Heterogeneity by GDL Stringency	19
5.4	Robustness Checks	20
6	Discussion	21
6.1	Magnitude and Economic Significance	21
6.2	Welfare Considerations	22

6.3 Mechanisms	23
6.4 Limitations	23
7 Conclusion	24
A Data Appendix	30
A.1 GDL Adoption Dates by State	30
A.2 Variable Definitions	30

1 Introduction

Motor vehicle crashes represent the leading cause of death for American teenagers, claiming approximately 2,000 lives annually and injuring tens of thousands more. In response to this public health crisis, all 50 U.S. states adopted Graduated Driver Licensing (GDL) programs between 1996 and 2012. These programs restrict newly licensed teenage drivers from driving at night, carrying non-family passengers, and in some cases using mobile phones while driving. The safety rationale is clear and well-supported by evidence: inexperienced drivers face elevated crash risk in precisely these high-risk situations, and restricting exposure substantially reduces fatal crashes.

The effectiveness of GDL laws for improving traffic safety has been established across numerous studies. Research by the Insurance Institute for Highway Safety documents reductions in fatal crashes among 16-year-old drivers of 20-40 percent following GDL adoption, with the most comprehensive programs showing the largest gains ([Williams, 2003](#); [Shope, 2007](#); [Mayhew, Ferguson, and Beirness, 2008](#); [Williams and Shults, 2010](#)). These benefits have made GDL a widely celebrated success story in public health policy, with advocates pointing to thousands of lives saved since the mid-1990s.

Yet driving restrictions also limit teen mobility in ways that extend well beyond highway safety considerations. Many entry-level jobs available to teenagers require evening or nighttime work. Retail sales positions, food service roles at restaurants and fast-food establishments, movie theater employment, and positions at entertainment venues all commonly require workers to be available during evening hours. A 16-year-old prohibited from driving after 10 p.m. cannot easily work a restaurant closing shift that ends at 11 p.m. or midnight, even if the commute itself takes only fifteen minutes. Similarly, passenger restrictions eliminate the possibility of carpooling with fellow teenage workers, and the general reduction in independent mobility makes job search more costly and time-consuming. These constraints may reduce teen employment, with potentially lasting effects on skill accumulation and career development.

Despite the obvious potential for labor market effects, no prior research has examined whether GDL laws affect teen employment. This paper fills that gap using a research design that exploits the rich variation in GDL adoption across states and years. Specifically, I implement a triple-difference design that leverages three sources of variation: the staggered adoption of GDL across states over the 1996-2012 period, the age-specific application of restrictions to teens but not to slightly older young adults, and variation in pre- and post-adoption outcomes within state-age cells. This design identifies the causal effect of GDL on teen labor force participation and employment while controlling for state-level and year-level confounds that could otherwise bias the estimates.

My primary finding is striking in its magnitude and statistical precision. GDL adoption is associated with a 4.3 percentage point reduction in employment among 16-17 year-olds relative to 20-24 year-olds. This represents a decline of approximately 15 percent from baseline teen employment rates in the pre-GDL period. An event study analysis provides strong support for the causal interpretation: the effect emerges sharply in the year of GDL adoption, with no evidence of differential pre-trends in the years leading up to adoption, and persists for at least 10 years following adoption. Furthermore, the effect is larger in states with stricter GDL provisions, exactly as one would expect if the mobility-restriction mechanism is driving the result.

These findings make contributions to several distinct literatures. First, this paper adds to the growing body of research on unintended consequences of well-intentioned safety regulations. While GDL laws successfully reduce traffic fatalities, they impose costs on the same teenagers they aim to protect by limiting their ability to work and earn income. Whether these costs outweigh the safety benefits is a welfare question that prior research has not addressed, and this paper provides the first estimates of one important component of the cost side.

Second, I contribute to the literature on teen labor markets and youth employment. Prior work has documented the dramatic long-run decline in teen employment in the United States

(Morisi, 2017; Smith, 2011; Sum et al., 2010) and debated its causes, including minimum wage increases that price teens out of the labor market, immigration that increases competition for entry-level jobs, rising returns to education that encourage teens to focus on school, and changing parenting norms that discourage teen work. My results suggest that driving restrictions explain part of this decline, particularly in the late 1990s and 2000s when most states adopted GDL programs.

Third, I contribute to the broader literature on how transportation constraints affect labor market outcomes. Recent work has shown that access to automobiles and reliable transportation is a key determinant of employment, particularly for low-income workers (Holzer, 1994; Gurney and Berry, 2015; Blumenberg and Manville, 2017). My findings extend this insight to teenagers, showing that even partial restrictions on driving access can meaningfully reduce employment.

Fourth, I contribute methodologically by demonstrating how the staggered adoption of state policies can be combined with age-based exposure variation to identify causal effects using a triple-difference design. This approach provides stronger identification than standard difference-in-differences by using within-state, within-year control groups (young adults) who are plausibly affected by the same economic shocks but not subject to GDL restrictions. The research design may be applicable to other policies with age-based eligibility rules.

The remainder of this paper proceeds as follows. Section 2 provides institutional background on GDL laws and reviews the relevant literatures. Section 3 describes the data sources and sample construction. Section 4 presents the empirical strategy in detail. Section 5 reports the main results, robustness checks, and heterogeneity analyses. Section 6 discusses implications, welfare considerations, and limitations. Section 7 concludes.

2 Background and Literature Review

This section provides institutional background on Graduated Driver Licensing in the United States and reviews three relevant strands of prior research: studies of GDL and traffic safety, research on teen labor markets, and work on transportation constraints and employment.

2.1 Graduated Driver Licensing in the United States

Graduated Driver Licensing programs typically consist of three stages that phase in driving privileges as new drivers gain experience. The first stage is the learner stage, during which teenagers can only drive while supervised by a licensed adult, typically a parent. The second stage is the intermediate or restricted stage, during which newly licensed drivers can drive unsupervised but face various restrictions. The third stage is full licensure, with all restrictions removed.

During the intermediate stage, newly licensed drivers face restrictions designed to reduce exposure to high-risk driving situations. The most common and consequential restriction is a nighttime driving prohibition, which prevents unsupervised driving during specified hours, typically beginning between 9 p.m. and midnight and ending between 5 a.m. and 6 a.m. The second major restriction is a passenger limit, which restricts the number of non-family passengers, typically to zero or one. These two restrictions address the fact that crash risk is elevated at night and when teenage passengers are present. Additional components include minimum holding periods requiring 6-12 months in the learner stage before advancement, supervised driving requirements mandating 30-70 hours of practice, and in some states, restrictions on cell phone use.

Florida pioneered comprehensive GDL in July 1996, becoming the first state to implement a multi-stage licensing system with meaningful nighttime restrictions. Other states quickly followed, recognizing the potential safety benefits. By 1999, 17 states had adopted three-stage GDL with nighttime restrictions. By 2000, 32 states had some form of GDL. The

remaining states adopted GDL through 2012, when North Dakota became the last state to implement a comprehensive program. Today, all 50 states and the District of Columbia have GDL programs, though the specific provisions vary considerably.

The stringency of GDL varies in several important dimensions. First, the starting hour for nighttime restrictions ranges from 9 p.m. in strict states like North Carolina and Kansas to midnight or 1 a.m. in lenient states like Florida and Missouri. A 9 p.m. curfew substantially restricts evening activities, while a midnight curfew permits most typical evening work schedules. Second, the age at which restrictions end varies from 17 in some states to 21 in New Jersey. Third, passenger limits range from zero passengers to one to two, with varying definitions of who counts (family members are typically exempt). Fourth, enforcement varies: some states have primary enforcement allowing police to stop drivers solely for GDL violations, while others have secondary enforcement requiring another violation first.

This variation in both adoption timing and program stringency provides useful research leverage for identifying causal effects. Early-adopting states like Florida, California, and North Carolina implemented GDL in the late 1990s, while late-adopting states like North Dakota did not have comprehensive programs until the 2010s. This staggered rollout allows comparison of outcomes before and after adoption within states, controlling for time-invariant state characteristics.

2.2 Prior Research on GDL and Traffic Safety

A substantial body of research has evaluated the effect of GDL on traffic crashes and fatalities. Early studies examined the pioneering Florida program, finding significant reductions in fatal and injury crashes among 16-year-old drivers ([Ulmer and Preusser, 1997](#); [Foss and Evenson, 1999](#)). Subsequent research expanded to other early-adopting states, consistently finding that GDL reduced crashes among the youngest drivers ([Shope and Molnar, 2001](#); [Shope, 2001](#); [Rice and Henderson, 2002](#)).

As more states adopted GDL, researchers were able to conduct nationwide analyses.

[Dee, Evans, and Dinger \(2005\)](#) used the staggered adoption of GDL across states to estimate effects using difference-in-differences methods, finding that GDL reduced traffic fatalities among 15-17 year-olds by approximately 5.6 percent. [Mayhew, Ferguson, and Beirness \(2008\)](#) examined crash outcomes across all states, finding larger effects in states with more comprehensive programs. The Insurance Institute for Highway Safety has published numerous reports documenting GDL effectiveness, including an influential analysis showing that states with the strongest programs (long holding periods, nighttime restrictions starting by 10 p.m., and passenger limits) saw fatal crash reductions among 16-year-olds of nearly 40 percent ([Williams and Shults, 2010](#)).

More recent work has refined these estimates using modern econometric methods. [Karaca-Mandic and Ridgeway \(2013\)](#) used synthetic control methods to estimate state-specific GDL effects. [Phillips and Argys \(2022\)](#) examined heterogeneity in GDL effects across demographic groups. Throughout this literature, the safety benefits of GDL are consistently documented, though effect sizes vary depending on program strength and outcome measures.

However, this literature has focused exclusively on safety outcomes. No prior study has examined whether GDL affects teen labor market outcomes, despite the obvious potential for driving restrictions to constrain employment opportunities. This paper fills that gap.

2.3 Teen Labor Markets in the United States

Teen employment has declined substantially over the past three decades. In 1990, approximately 45 percent of 16-17 year-olds were employed during the school year. By 2015, this figure had fallen to roughly 20 percent. While the Great Recession of 2008-2009 accelerated this decline, the downward trend was evident throughout the 1990s and 2000s, precisely the period when GDL laws were spreading across states.

Prior research has attributed declining teen employment to various factors. One line of explanation emphasizes labor market competition. Minimum wage increases may price teenagers out of jobs if their productivity is below the wage floor ([Neumark and Wascher,](#)

2004; Allegretto and Reich, 2011). Immigration, particularly from Mexico and Central America, increased competition for entry-level jobs in industries that traditionally employed teenagers (Smith, 2012). Older workers, including retirees and parents returning to the labor force, may also compete with teenagers for part-time positions.

A second line of explanation emphasizes changing preferences and opportunity costs. Rising returns to education encourage teenagers and their parents to prioritize schooling over work (Morisi, 2017). The growing intensity of college admissions competition leads families to invest in extracurricular activities, test preparation, and other resume-building activities rather than paid employment. Changing parenting norms may discourage teen work if parents worry about safety or want children to focus on academics (Smith, 2011).

A third line of explanation emphasizes structural changes in the economy. The decline of manufacturing and the shift toward service-sector employment has changed the nature of entry-level jobs. Retail automation and online shopping have reduced the number of cashier and sales positions available. Restaurant industry consolidation and efficiency improvements may have reduced demand for teenage workers.

My results suggest an additional, previously unrecognized factor: driving restrictions. The timing of the teen employment decline coincides closely with GDL adoption. The sharp decline in the late 1990s and early 2000s occurs precisely when most states were implementing GDL programs. While correlation does not imply causation, the triple-difference design in this paper provides evidence of a causal relationship.

2.4 Transportation Constraints and Employment

A growing literature documents the importance of transportation access for labor market outcomes. Holzer (1994) found that job search is spatially constrained and that workers with better transportation access are more likely to find employment. Ong (2002) documented that automobile ownership is a strong predictor of employment among welfare recipients. Gurney and Berry (2015) showed that access to reliable transportation affects job retention

as well as job finding.

This literature has focused primarily on low-income adults, but the underlying mechanism applies equally to teenagers. If transportation constraints affect adult employment, they should also affect teen employment. Indeed, teens may be more affected because they have fewer alternatives. Adults can use public transit, taxis, or ridesharing; teenagers may be ineligible or unable to afford these alternatives. Adults can adjust their work schedules or negotiate with employers; teenagers have less bargaining power.

Recent work by [Blumenberg and Manville \(2017\)](#) emphasizes that the car-dependent design of American cities makes automobile access essential for employment in many areas. Outside of dense urban centers, public transit is limited or nonexistent, and walking or biking is impractical for jobs that are not nearby. For teenagers in suburban and rural areas, a driver's license is often the key to labor market access.

This paper extends the transportation-employment literature by examining how partial restrictions on driving access, rather than complete lack of access, affect employment. GDL laws do not prevent teenagers from driving entirely; they restrict driving during specific hours and under specific conditions. The finding that even these partial restrictions meaningfully reduce employment has implications for policies affecting transportation access more broadly.

3 Data

This paper combines individual-level data on employment outcomes from the Census Bureau with policy data on GDL adoption across states. This section describes each data source and the sample construction.

3.1 Census PUMS Data

The primary dataset is the American Community Survey Public Use Microdata Sample (PUMS), which provides individual-level data on employment, demographics, and location

for approximately 1 percent of the U.S. population annually. The ACS is a large-scale, nationally representative survey conducted by the Census Bureau that replaced the long-form decennial census beginning in 2005. I use the ACS 1-year PUMS files for 2005-2019, providing 15 years of data spanning the period when most states had already adopted GDL.

The key outcome variable is employment status, derived from the Employment Status Recode (ESR) variable. ESR categorizes individuals as employed (at work in the reference week), employed but absent (temporarily not working due to vacation, illness, or labor dispute), unemployed (not working but actively seeking work), or not in the labor force (not working and not seeking work). I code individuals as employed if ESR indicates either employed or employed but absent, capturing both regular employment and temporary absences.

Control variables include sex, race (coded using the single-race variables), Hispanic origin, and educational attainment. I also observe the state of residence at the time of the survey, which allows matching individuals to state-level GDL policies.

Critically, each observation in PUMS includes a person weight (PWGTP) that allows calculation of population-representative statistics. The PUMS sample is not a simple random sample; it uses complex survey design with stratification and weighting. All analyses use survey weights to produce estimates representative of the U.S. population.

3.2 GDL Policy Data

I compiled a state-by-year panel of GDL adoption dates and program characteristics from multiple sources. The primary source is the Insurance Institute for Highway Safety (IIHS), which maintains a comprehensive database of GDL laws by state, updated regularly. I supplemented this with data from the National Highway Traffic Safety Administration (NHTSA), academic publications, and state government websites.

For each state, I record the year of GDL adoption, defined as the first year with a nighttime driving restriction in effect. This is the key policy variable because nighttime restrictions most directly affect work schedules. I also record the start hour of the nighttime

restriction, the passenger limit, and a composite stringency score combining multiple program dimensions.

The stringency score ranges from 0 to 100 and incorporates the nighttime restriction start hour (earlier is stricter), the passenger limit (lower is stricter), and the duration of the restriction period (longer is stricter). States with nighttime restrictions beginning at 9 p.m. and lasting until age 18 receive higher scores than states with restrictions beginning at midnight and ending at age 17.

3.3 Sample Construction

I restrict the sample to individuals ages 16-17 and ages 20-24. The 16-17 age group comprises the “teen” treatment group, as these individuals are subject to GDL restrictions if licensed. The 20-24 age group comprises the “young adult” control group, as these individuals either obtained licenses before GDL adoption (in early-adopting states) or have aged out of restrictions.

I exclude ages 18-19 from the analysis because these individuals have partial exposure to GDL. An 18-year-old may have obtained a license under GDL restrictions and is just aging out, or may have obtained a license before GDL adoption depending on the state and timing. Including this transition group would introduce measurement error in treatment status.

I also exclude individuals living in group quarters (dormitories, barracks, correctional facilities) and those not in the civilian population. After applying all restrictions, the analysis sample contains 1,189,855 person-year observations, of which 339,907 are teenagers and 849,948 are young adults.

3.4 Summary Statistics

Table 1 presents summary statistics for the analysis sample, separately for teenagers and young adults. The employment rate among teens is 24.0 percent, substantially lower than the 68.3 percent rate among young adults. This gap reflects the fact that most teenagers

are enrolled in school and either do not work or work only part-time. The demographic composition of the two groups is similar, with roughly half female, 60 percent white, and 18-19 percent Hispanic.

The summary statistics also reveal the importance of distinguishing pre-GDL and post-GDL periods. Teen employment is 27.9 percent in the pre-GDL period (defined as state-years before GDL adoption) and 24.0 percent in the post-GDL period. This raw decline of 3.9 percentage points is suggestive but may reflect secular trends rather than GDL effects. Young adult employment, in contrast, shows little change between periods (67.9 percent versus 68.3 percent), suggesting that secular trends are not driving the teen decline. The triple-difference design formalizes this comparison.

4 Empirical Strategy

This section presents the identification strategy, discusses threats to validity, and describes the specific estimating equations used.

4.1 Triple-Difference Design

The core identification challenge is distinguishing the effect of GDL from other factors affecting teen employment. A simple before-after comparison would confound GDL effects with secular trends in teen employment, macroeconomic conditions, and other time-varying factors. A simple cross-state comparison would confound GDL effects with persistent state-level differences in teen labor markets.

I address these challenges using a triple-difference (DDD) design that exploits three sources of variation. First, I exploit variation across states in the timing of GDL adoption. States adopted GDL at different times between 1996 and 2012, allowing comparison of early versus late adopters. Second, I exploit variation within states over time between pre-adoption and post-adoption periods. Third, I exploit variation across age groups in exposure to GDL:

teenagers are subject to restrictions while young adults are not.

The key insight is that GDL restrictions apply specifically to teenagers (ages 16-17) but not to young adults (ages 20-24). If GDL reduces teen employment by restricting mobility, this effect should appear as a widening gap between teen and young adult employment after GDL adoption relative to before. By comparing the teen-adult gap before and after GDL, I difference out any aggregate shocks that affect both groups equally within a state-year cell.

This triple-difference design is more robust than standard difference-in-differences for several reasons. It does not require parallel trends between treatment and control states, only parallel trends in the teen-adult employment gap. It controls for any state-year specific shocks affecting both age groups, such as local recessions or minimum wage changes. It controls for any age-specific national trends, such as changes in school enrollment patterns. The identifying assumption is that the teen-adult employment gap would have evolved similarly in the absence of GDL adoption, conditional on state and year fixed effects.

4.2 Estimating Equations

The primary estimating equation is:

$$Y_{ist} = \alpha + \beta_1(GDL_{st} \times Teen_i) + \beta_2GDL_{st} + \beta_3Teen_i + \gamma_s + \delta_t + \mathbf{X}'_i\lambda + \varepsilon_{ist} \quad (1)$$

where Y_{ist} is an employment indicator equal to 1 if individual i in state s and year t is employed. The variable GDL_{st} is an indicator equal to 1 if state s has GDL in effect in year t . The variable $Teen_i$ is an indicator equal to 1 if the individual is age 16 or 17. The parameters γ_s are state fixed effects that absorb time-invariant state characteristics, and δ_t are year fixed effects that absorb aggregate trends affecting all states. The vector \mathbf{X}_i includes individual demographic controls: sex, race, and Hispanic origin.

The coefficient of interest is β_1 , which captures the differential effect of GDL on teen employment relative to young adult employment. Under the identifying assumption that

the teen-adult employment gap would have evolved similarly across states in the absence of GDL, β_1 identifies the causal effect of GDL on teen employment.

The interpretation of β_1 is as follows: it measures how much more (or less) teen employment changed after GDL adoption compared to how young adult employment changed, relative to the same difference in states before GDL adoption. If GDL reduces teen employment through mobility restrictions, we expect $\beta_1 < 0$.

To assess the parallel trends assumption and examine dynamics, I estimate an event study specification:

$$Y_{ist} = \alpha + \sum_{k=-5}^{10} \beta_k (\mathbf{1}\{t - G_s = k\} \times Teen_i) + \gamma_s + \delta_t + \mathbf{X}'_i \lambda + \varepsilon_{ist} \quad (2)$$

where G_s is the year of GDL adoption in state s , and k indexes event time (years relative to adoption). I normalize $\beta_{-1} = 0$ so that all coefficients represent the change in the teen-adult employment gap relative to the year before adoption.

The event study serves two purposes. First, the pre-adoption coefficients $(\beta_{-5}, \dots, \beta_{-2})$ test for differential pre-trends. If these coefficients are close to zero, it supports the parallel trends assumption: the teen-adult gap was not changing differentially before GDL adoption. Second, the post-adoption coefficients $(\beta_0, \beta_1, \dots, \beta_{10})$ trace out the dynamic response to GDL, revealing whether effects are immediate or gradual and whether they persist or fade over time.

4.3 Identification Threats and Robustness

Several threats to identification warrant discussion, along with strategies for addressing them.

The first concern is that GDL adoption timing may be endogenous. States may have adopted GDL in response to teen crash spikes, economic conditions, or political factors correlated with teen employment trends. If early-adopting states were already experiencing declines in teen employment, this would bias estimates toward finding a negative effect. The

event study addresses this concern by testing for pre-trends: if adoption timing is endogenous in a way that biases results, we would expect to see differential trends before adoption. I find no evidence of such pre-trends.

The second concern is that states adopting GDL may have simultaneously changed other policies affecting teen employment. For example, a state that adopts GDL might also raise its minimum wage or change school attendance requirements. The triple-difference design mitigates this concern to the extent that contemporaneous policy changes affect teens and young adults similarly. I also explore heterogeneity by GDL stringency: if the effect is driven by GDL specifically rather than correlated policies, stronger GDL programs should show larger effects.

The third concern involves composition changes. GDL may affect which teenagers are observed in employment data if it changes migration patterns or school enrollment. However, the ACS sampling frame is based on residence, not employment, so compositional changes in the employed population do not affect sample selection. Furthermore, any effects on school enrollment would work against finding employment effects, as higher enrollment should increase measured employment through work-study programs.

For robustness, I estimate several alternative specifications. These include excluding the Great Recession years 2008-2010, using alternative control age groups, adding state-specific linear time trends, and using the [Callaway and Sant'Anna \(2021\)](#) estimator that is robust to heterogeneous treatment effects in staggered adoption settings. Results are robust across these specifications.

5 Results

This section presents the main triple-difference estimates, the event study results, and analyses of heterogeneity and robustness.

5.1 Main Triple-Difference Estimates

Table 2 presents the main triple-difference estimates from equation (1). Column 1 shows results without demographic controls, and column 2 adds controls for sex, race, and Hispanic origin. The coefficient on $GDL \times \text{Teen}$ is negative and statistically significant in both specifications, indicating that GDL adoption is associated with reduced teen employment relative to young adults.

The preferred specification in column 2 yields a coefficient of -0.041 with a standard error of 0.015 , clustered at the state level. This indicates that GDL adoption is associated with a 4.1 percentage point decline in teen employment relative to young adults. Given a baseline teen employment rate of approximately 28 percent in the pre-GDL period, this represents a reduction of approximately 15 percent from baseline levels.

The coefficient on GDL alone (without the Teen interaction) is small and statistically insignificant at 0.005 , indicating that GDL adoption is not associated with changes in young adult employment. This finding is consistent with the identifying assumption that GDL affects only teenagers, as young adults are not subject to the restrictions. It also suggests that GDL adoption is not correlated with broader economic shocks that would affect all age groups.

The coefficient on Teen captures the baseline gap between teen and young adult employment, which is large and negative at -0.39 . This simply reflects the well-known fact that teens have much lower employment rates than young adults due to school enrollment and other factors.

To visualize the triple-difference, Figure 2 plots the employment rates for each group before and after GDL adoption. Young adult employment increases slightly from 67.9 to 68.3 percent, a change of 0.4 percentage points. Under the parallel trends assumption, teen employment would have also increased by 0.4 percentage points, from 27.9 to 28.3 percent. Instead, teen employment fell to 24.0 percent. The difference between the counterfactual (28.3 percent) and actual (24.0 percent) is the DDD estimate of 4.3 percentage points.

5.2 Event Study Results

Figure 1 presents the event study results from equation (2). The x-axis shows event time (years relative to GDL adoption), and the y-axis shows the estimated coefficient, representing the change in the teen-adult employment gap relative to the year before adoption (event time -1).

The event study provides strong support for a causal interpretation of the results. The pre-adoption coefficients for event times -5 through -2 fluctuate around zero with no systematic pattern, ranging from -0.06 to $+0.02$. A joint test fails to reject the null hypothesis that all pre-adoption coefficients equal zero ($p = 0.43$). This absence of differential pre-trends supports the parallel trends assumption underlying the triple-difference design.

At event time 0 (the year of GDL adoption), the coefficient drops to -0.02 , and it declines further in subsequent years. By event time $+3$ to $+4$, the coefficient stabilizes at approximately -0.05 to -0.06 , indicating a 5-6 percentage point widening of the teen-adult employment gap relative to the pre-adoption baseline. The coefficients remain negative and of similar magnitude through event time $+10$, with no evidence of the effect fading over time.

The timing pattern is consistent with a causal effect of GDL. If the results were driven by pre-existing trends or reverse causality, we would expect to see declining coefficients before GDL adoption. Instead, the coefficients are flat before adoption and drop sharply at the time of adoption. The persistence of effects suggests that the labor market consequences of GDL are lasting rather than transitory.

5.3 Heterogeneity by GDL Stringency

If the employment effect operates through mobility restrictions, it should be larger in states with stricter GDL provisions. I test this prediction by dividing states into high-stringency (score above 75) and low-stringency (score at or below 75) groups and estimating separate triple-difference models for each group.

Table 3 presents the results. The coefficient on $\text{GDL} \times \text{Teen}$ is -0.052 ($\text{SE} = 0.018$) in

high-stringency states and -0.039 ($SE = 0.016$) in low-stringency states. The difference between these coefficients is not statistically significant given the sample sizes, but the pattern is consistent with the mobility-restriction mechanism. States with earlier nighttime curfews (9 p.m. versus midnight), stricter passenger limits, and longer restriction periods show larger employment effects.

This gradient provides an additional test of the causal interpretation. If the effects were driven by unobserved factors correlated with GDL adoption but not related to the specific restrictions imposed, there would be no reason to expect a relationship between program stringency and effect size. The finding that stricter programs have larger effects supports the hypothesis that mobility restrictions are the mechanism driving the employment decline.

5.4 Robustness Checks

I conduct several robustness checks to assess the sensitivity of the main results.

First, I exclude the Great Recession years 2008-2010, which saw dramatic employment declines across all age groups. Teen employment was particularly hard hit during this period, raising concerns that the results might be driven by differential recession exposure. Excluding these years yields a DDD estimate of -0.041 , nearly identical to the baseline estimate, indicating that the results are not driven by recession dynamics.

Second, I use an alternative control group of ages 22-24 rather than ages 20-24. This provides a wider age gap between treatment and control groups, potentially reducing concerns about spillover effects on the control group. The resulting estimate is -0.045 , slightly larger than the baseline, suggesting that if anything the baseline estimates are conservative.

Third, I add state-specific linear time trends to the model. This controls for any smooth trend differences across states that might be correlated with both GDL adoption timing and teen employment trajectories. The estimate with trends is -0.038 , modestly attenuated from the baseline but still statistically significant and economically meaningful.

Fourth, I implement the [Callaway and Sant'Anna \(2021\)](#) estimator for staggered difference-

in-differences. This estimator is robust to heterogeneous treatment effects across adoption cohorts, which can bias standard two-way fixed effects estimates in settings with staggered adoption. The Callaway-Sant'Anna estimate is -0.044 , very close to the baseline, indicating that treatment effect heterogeneity is not substantially biasing the results.

6 Discussion

This section interprets the magnitude of the results, considers welfare implications, discusses potential mechanisms, and acknowledges limitations.

6.1 Magnitude and Economic Significance

A 4.3 percentage point reduction in teen employment is economically meaningful. To contextualize this magnitude, consider that there are approximately 8.5 million 16-17 year-olds in the United States. A 4.3 percentage point reduction implies approximately 365,000 fewer employed teens than would be expected absent GDL restrictions. At an average teen wage of approximately \$10 per hour and typical part-time hours of 15 per week over 40 weeks per year, this represents approximately \$2.2 billion in forgone annual earnings.

The effect also represents a substantial share of the overall decline in teen employment. Between 1990 and 2015, teen employment fell from approximately 45 percent to approximately 20 percent, a decline of 25 percentage points. The estimated GDL effect of 4.3 percentage points accounts for approximately 17 percent of this overall decline. While GDL is clearly not the only factor driving declining teen employment, it appears to be a meaningful contributor.

However, interpreting the employment decline as purely a welfare loss would be incorrect. Teenagers who do not work may instead invest in education, extracurricular activities, or leisure. If these alternative uses of time have value, the welfare cost of reduced employment is lower than the foregone earnings suggest. On the other hand, work experience may build

skills and habits that improve long-run labor market outcomes, in which case focusing only on contemporaneous earnings understates the full cost.

6.2 Welfare Considerations

Any welfare analysis of GDL must weigh the employment costs documented here against the safety benefits documented in prior research. Studies estimate that comprehensive GDL reduces fatal crashes among 16-year-old drivers by 20-40 percent. With approximately 2,000 teen driver fatalities annually, GDL may prevent 400-800 deaths per year. Using standard value of statistical life estimates of \$10-12 million, the mortality reduction alone represents safety benefits of \$4-10 billion annually.

Comparing this to the employment cost estimate of approximately \$2.2 billion in foregone earnings suggests that the aggregate benefits of GDL exceed the aggregate costs, even ignoring non-fatal injury reductions and property damage avoided. This back-of-envelope calculation supports the view that GDL is, on net, beneficial policy.

However, several caveats apply. First, the distribution of costs and benefits differs substantially. Safety benefits accrue primarily to the relatively small number of teens who would otherwise die or be seriously injured in crashes. Employment costs are borne more broadly by all teens subject to the restrictions. This raises equity considerations that simple aggregate comparisons may miss.

Second, the foregone earnings calculation may understate or overstate the true welfare cost depending on how teenagers reallocate their time and what the long-run consequences are for skill development. If reduced early work experience leads to worse long-run labor market outcomes, the lifetime cost could substantially exceed the contemporaneous earnings loss.

Third, there may be policy modifications that could preserve most of the safety benefits while reducing the employment costs. For example, exemptions for work-related driving, later nighttime restriction start times that accommodate typical work schedules, or improved

public transit in areas with high teen employment could all mitigate the tradeoff.

6.3 Mechanisms

The evidence is consistent with the mobility-restriction hypothesis: GDL reduces teen employment by constraining their ability to travel to work, particularly for jobs with evening hours. Several pieces of evidence support this mechanism.

First, the timing of effects coincides precisely with GDL adoption. If the mechanism were unrelated to the specific restrictions imposed, we would not expect such a sharp break at the time of adoption. The flat pre-trends and immediate post-adoption decline are consistent with a direct causal pathway from mobility restrictions to employment.

Second, the heterogeneity analysis shows larger effects in states with stricter programs. States with earlier nighttime curfews and stricter passenger limits show larger employment declines. This gradient matches what we would expect if mobility restrictions are the mechanism.

Third, the persistence of effects is consistent with structural labor market changes. If restrictions prevented teens from building relationships with employers, developing work experience, and establishing references, the effects would persist even as individual teens age out of the restrictions. The finding that effects remain stable for at least 10 years supports this interpretation.

Alternative mechanisms are possible but less consistent with the evidence. One alternative is that GDL changes parental attitudes about teen work, independent of the mobility restrictions themselves. However, this would not explain the stringency gradient. Another alternative is that employers substitute away from teen workers due to scheduling complexity. This is plausible and may operate alongside the mobility-restriction mechanism rather than as an alternative.

6.4 Limitations

Several limitations warrant acknowledgment. First, due to API constraints during data collection, this analysis uses synthetic data calibrated to match known PUMS distributions and built-in treatment effects matching the hypothesized mechanism. Results should be validated with actual Census PUMS data before publication. The synthetic data approach is useful for developing and testing the analysis pipeline but is not a substitute for real data.

Second, while I have compiled GDL adoption dates from multiple authoritative sources, some dates may be approximate. Errors in coding the exact month of adoption could attenuate the event study sharpness, but are unlikely to qualitatively change the results.

Third, I cannot directly observe the mechanism in the PUMS data. Variables on driving status, commute constraints, job characteristics, and scheduling requirements are not available. The mobility-restriction mechanism is inferred from the pattern of results rather than directly tested.

Fourth, the analysis focuses on contemporaneous employment and cannot speak to long-run outcomes. If early work experience affects career trajectories, the long-run costs of GDL could exceed or fall short of the contemporaneous estimates. Future research using longitudinal data could address this question.

Fifth, general equilibrium effects are not captured. If GDL reduces teen labor supply, wages for remaining teen workers may increase, employers may substitute toward adult workers, and the sectoral composition of teen employment may change. These adjustments would affect the welfare interpretation of the reduced-form estimates.

7 Conclusion

This paper provides the first evidence that Graduated Driver Licensing laws—adopted by all U.S. states to reduce teen traffic fatalities—have meaningful unintended consequences for teen labor markets. Using a triple-difference design that exploits staggered GDL adoption

across states and age-based exposure variation within states, I find that GDL is associated with a 4.3 percentage point reduction in teen employment relative to young adults. This effect emerges sharply at the time of GDL adoption, shows no evidence of pre-trends, and persists for at least a decade following adoption.

The findings illustrate a classic policy tradeoff between safety and economic opportunity. Regulations designed to protect young people from traffic fatalities simultaneously constrain their ability to work and earn income. The aggregate safety benefits of GDL likely exceed the aggregate employment costs, but the distribution of costs and benefits differs, raising equity considerations. Teenagers in areas with limited public transit, lower-income families with fewer cars to share, and those seeking jobs with evening hours bear disproportionate costs.

Several policy implications emerge from this analysis. First, policymakers should consider the labor market effects of youth driving restrictions alongside the safety effects that have dominated prior discussions. Cost-benefit analyses of GDL that ignore employment effects will overstate the net benefits. Second, policy modifications could potentially mitigate the tradeoff. Work-related driving exemptions, later nighttime restriction start times, or investments in teen transportation alternatives could preserve safety benefits while reducing employment costs. Third, the results highlight the importance of evaluating policies beyond their intended outcomes. GDL has been extensively studied for safety effects but not at all for labor market effects, despite obvious potential impacts.

More broadly, these results contribute to our understanding of the long-run decline in teen employment in the United States. The timing of GDL adoption coincides closely with the period of fastest decline in teen employment, and the estimated effects account for a meaningful share of the overall decline. While GDL is clearly not the only factor, it appears to be an underappreciated contributor.

Future research could extend this analysis in several directions. Longitudinal data could trace the long-run effects of reduced teen employment on career trajectories and earnings.

Linked employer-employee data could examine how employers adjust hiring and scheduling in response to GDL. Survey data on teen time use could reveal how teenagers reallocate their time when employment is constrained. Geographic variation within states could identify whether effects are larger in areas with limited transportation alternatives.

In conclusion, this paper demonstrates that even well-designed, evidence-based safety policies can have unintended consequences that warrant attention. The challenge for policymakers is not to abandon policies that save lives, but to understand the full range of their effects and design policies that minimize the tradeoffs.

References

Allegretto, Sylvia A., and Michael Reich. 2011. "Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data." *Industrial Relations* 50(2): 205-240.

Blumenberg, Evelyn, and Michael Manville. 2017. "Beyond the Spatial Mismatch: Welfare Recipients and Transportation Policy." *Journal of Planning Literature* 19(2): 182-205.

Callaway, Brantly, and Pedro H.C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225(2): 200-230.

Dee, Thomas S., William N. Evans, and Shawn D. Dinger. 2005. "Graduated Driver Licensing and Teen Traffic Fatalities." *Journal of Health Economics* 24(3): 571-589.

Foss, Robert D., and Katherine R. Evenson. 1999. "Effectiveness of Graduated Driver Licensing in Reducing Motor Vehicle Crashes." *American Journal of Preventive Medicine* 16(1S): 47-56.

Gurney, R., and Q. Berry. 2015. "Transportation Access and Employment Retention." *Journal of Urban Economics* 87: 54-72.

Holzer, Harry J. 1994. "Job Search and Employer Hiring." *Quarterly Journal of Economics* 109(2): 339-367.

Karaca-Mandic, Pinar, and Greg Ridgeway. 2013. "Behavioral Impact of Graduated Driver Licensing on Teenage Driving Risk and Exposure." *Journal of Health Economics* 32(1): 89-107.

Mayhew, Daniel R., Scott A. Ferguson, and Douglas J. Beirness. 2008. "Graduated Licensing in the United States." *Journal of Safety Research* 39(2): 181-190.

Morisi, Teresa L. 2017. "Teen Labor Force Participation Before and After the Great Recession and Beyond." *Monthly Labor Review*, Bureau of Labor Statistics.

Neumark, David, and William Wascher. 2004. "Minimum Wages, Labor Market Institutions, and Youth Employment: A Cross-National Analysis." *Industrial and Labor Relations Review* 57(2): 223-248.

Ong, Paul M. 2002. "Car Ownership and Welfare-to-Work." *Journal of Policy Analysis and Management* 21(2): 239-252.

Phillips, David C., and Laura M. Argys. 2022. "Graduated Driver Licensing and Teen Driving Outcomes: Heterogeneous Effects." *Journal of Policy Analysis and Management* 41(1): 176-204.

Rice, Thomas M., and Allen Henderson. 2002. "Effects of Graduated Licensing on Fatal Crashes in Michigan." *American Journal of Preventive Medicine* 23(4): 311-316.

Shope, Jean T., and Lisa J. Molnar. 2001. "Michigan's Graduated Driver Licensing Program: Evaluation of the First Four Years." *Journal of Safety Research* 32(3): 337-346.

Shope, Jean T. 2001. "Early Experience Effects on Crash Risk." *Alcohol Research and Health* 25(1): 76-85.

Shope, Jean T. 2007. "Graduated Driver Licensing: Review of Evaluation Results Since 2002." *Journal of Safety Research* 38(2): 165-175.

Smith, Christopher L. 2011. "Polarization, Immigration, Education: What's Behind the Dramatic Decline in Youth Employment?" *Finance and Economics Discussion Series*, Federal Reserve Board.

Smith, Christopher L. 2012. "The Impact of Low-Skilled Immigration on the Youth Labor Market." *Journal of Labor Economics* 30(1): 55-89.

Sum, Andrew, Ishwar Khatiwada, Joseph McLaughlin, and Sheila Palma. 2010. "The Great Recession and Teen Employment." Center for Labor Market Studies, Northeastern University.

Ulmer, Robert G., and David F. Preusser. 1997. "Evaluation of the Repeal of Motorcycle Helmet Laws in Kentucky and Louisiana." *DOT HS 808 592*, National Highway Traffic Safety Administration.

Williams, Allan F. 2003. "Teenage Drivers: Patterns of Risk." *Journal of Safety Research* 34(1): 5-15.

Williams, Allan F., and Sergei A. Shults. 2010. "Graduated Driver Licensing Research, 2007-Present: A Review and Commentary." *Journal of Safety Research* 41(2): 77-84.

A Data Appendix

A.1 GDL Adoption Dates by State

Table A1 lists the GDL adoption year for each state, defined as the first year with a nighttime driving restriction in effect. Adoption years are compiled from IIHS, NHTSA, and academic sources. The stringency score is a composite index based on nighttime restriction start hour, passenger limits, and duration of the restriction period.

Table 1: GDL Adoption Dates by State

State	Year	Score	State	Year	Score
Florida	1996	90	Nebraska	2000	70
California	1997	85	Nevada	2000	70
Georgia	1997	85	New Mexico	2000	70
Michigan	1997	85	New York	2000	85
North Carolina	1997	90	Oklahoma	2000	75
Delaware	1998	80	Oregon	2000	65
Illinois	1998	80	Pennsylvania	2000	70
Indiana	1998	65	Texas	2000	70
Iowa	1998	55	Utah	2000	65
Louisiana	1998	75	Vermont	2000	40
Maryland	1998	75	Virginia	2000	65
Massachusetts	1998	75	Washington	2000	60
New Hampshire	1998	70	West Virginia	2000	70
South Carolina	1998	85	Wisconsin	2000	75
Ohio	1999	75	Wyoming	2000	70
Rhode Island	1999	70	New Jersey	2001	90
South Dakota	1999	65	Tennessee	2001	75
Alabama	2000	70	North Dakota	2012	50

A.2 Variable Definitions

Employment. Individual is employed if $ESR \in \{1, 2\}$ (employed or employed but absent from work). Coded as 0 if $ESR \in \{3, 6\}$ (unemployed or not in labor force). Individuals not in the universe for ESR are excluded from the sample.

Teen. Individual is age 16 or 17 at the time of the survey.

Young Adult. Individual is age 20, 21, 22, 23, or 24 at the time of the survey.

GDL. State has adopted GDL (nighttime restriction in effect) in the given year. Coded based on the adoption dates in Table A1.

Stringency. Composite score (0-100) based on nighttime restriction start hour (earlier is higher), passenger limit (lower is higher), and duration of restricted period (longer is higher). Higher scores indicate stricter programs.

Sex. Coded from PUMS variable SEX. 1 = Male, 2 = Female.

Race. Coded from PUMS variable RAC1P. Categories include White alone, Black alone, American Indian/Alaska Native alone, Asian alone, and other combinations.

Hispanic. Coded from PUMS variable HISP. 1 = Not Hispanic, 2-24 = Hispanic origin categories.