Shaping the Habits of Young Drivers*

Timothy J. Moore

Purdue University and NBER

Todd Morris

Max Planck Institute for Social Law and Social Policy

February 28, 2020

Abstract

Young, inexperienced drivers have the highest risk of traffic accidents. In this paper, we examine a targeted intervention in Australia that banned first-year drivers from carrying multiple passengers between 11pm and 5am. Using daytime outcomes as controls, we find that the reform reduced late-night accidents with multiple passengers by 54% and deaths by 75%, with passengers and other drivers accounting for most of the lives saved. We do not find offsetting effects for other types of accidents. The effects persist beyond the first year, suggesting that the intervention changed the habits of young drivers throughout their highest-risk years.

Keywords: traffic fatalities, driving restrictions, habit formation

JEL Classification: I18, K32

*__Acknowledgments:__ We thank the NSW Centre for Road Safety and Roads and Maritime Services NSW, especially Phil Sparkes, Emma Shearer and Padma Uppuluri, for assistance and access to data. We are grateful for comments and suggestions from John Denew, Gregory Gilpin, Ben Hansen, Peter Siminski, Stefan Staubli and Jenny Williams. **Correspondence:** Moore: timothyjmoore@gmail.com; Morris: toddstuartmorris@gmail.com.

1 Introduction

Crashes and fatalities on the road continue to damage the economy and society. For the year of 2010, the National Highway Traffic Safety Administration (NHTSA) estimated the total cost of crashes to American society at $836 billion, a cost equivalent to 5.5% of GDP (Blincoe et al., 2015). This cost reflected 32,999 fatalities, 3.9 million injuries, and 24 million damaged vehicles. Despite consistent progress in reducing traffic fatality rates, the U.S. lags behind other developed countries. Figure 1 compares the trend in U.S. traffic fatality rates from 1970 to 2018 with the trend for New South Wales, Australia's most populous state and the setting for our analysis. While the two jurisdictions had similar traffic fatality rates between 1970 and 1985, New South Wales has since seen a much larger reduction in fatality rates. Following a series of major policy interventions and awareness campaigns, the fatality rate in New South Wales is now almost two-thirds lower than in the U.S.

Despite these long-running improvements in road safety, there is substantial heterogeneity in the risk of crashes and fatalities. Young, inexperienced drivers have the highest crash and fatality risk (Mayhew, Simpson and Pak, 2003; McCartt, Shabanova and Leaf, 2003; Insurance Institute for Highway Safety, 2017), especially when driving at night and with passengers (e.g., see Chen et al., 2000; Tefft, Williams and Grabowski, 2013). Worldwide, 370,000 people aged 15–29 die each year from traffic accidents (World Health Organisation, 2018), and traffic fatalities were the leading cause of death for 16–24-year-old Americans from 2012 to 2014 (Webb, 2016). With these statistics in mind, it is not surprising that many road safety policies have targeted young drivers. So far, the most common policy response has been to adopt or refine graduated licensing systems. Graduated licensing systems now operate in every U.S. state and many other jurisdictions around the world. These systems impose specific restrictions on novice drivers, which are progressively relaxed as drivers gain experience and demonstrate a clean driving record. Most systems have three stages. The first stage, a learner period, allows learners to drive under the supervision of an experienced adult. The second stage, a probation period, allows unsupervised driving but imposes targeted restrictions on drivers, such as the time of day they can drive or the number of passengers they can carry. Typically, these restrictions are lifted when drivers reach the final stage, an unrestricted period.

In the probation period, most U.S. states currently enforce night-time curfews and passenger restrictions. The NHTSA (1998, p.16) sets out the rationale for such restrictions as follows: "By reducing the risk exposure of teenage drivers and allowing them time to mature before we give them the keys and unlimited use of the car, we will increase the likelihood that they will safely make it through their early driving years. And by creating safer teen drivers today, we also are helping them become safer, more responsible young adult drivers tomorrow."

Although it makes sense to target high-risk drivers in high-risk settings, the efficacy of these restrictions depends on several understudied factors. These factors include (i) the degree of compliance and which drivers comply, (ii) any spillover effects of these restrictions, and (iii) any long-term impacts. From a theoretical perspective, the direction of any long-term effects is ambiguous. On the one hand, any short-term improvements in road safety may be fully or partially offset in the long term if learning by doing is present, since drivers may graduate from the probation period with limited experience in risky driving environments. Indeed, Dee and Evans (2001) demonstrate that learning by doing has important implications on traffic fatalities in the context of legal access to alcohol.[1] On the other hand, if habit formation occurs — because individuals derive disutility from changes in their driving behavior over time — driving restrictions may improve road safety in the long term by causing a persistent reduction in the amount of driving in high-risk settings.

In this paper, we shed light on the importance of these factors in the context of a targeted intervention in New South Wales, Australia that banned first-year drivers from carrying multiple passengers between 11:00 pm and 5:00 am. Using linked administrative data on the universe of drivers' license records and serious (fatal and non-fatal) crashes, we exploit the fact that the restriction only applies between 11:00 pm and 5:00 am and only restricts multi-passenger driving. These features of the restriction and our data allow us to credibly estimate the short- and long-term effect of the restriction on crash rates using a differences-in-differences approach. Unlike previous studies, which exploit the staggered rollout of graduated licensing systems across U.S. states (Dee, Grabowski and Morrisey, 2005; Karaca-Mandic and Ridgeway, 2010; Gilpin, 2019), our estimates are immune to the concerns raised by several recent studies about the interpretation of two-way fixed effect estimates (e.g., Borusyak and Jaravel, 2017; De Chaisemartin and D'HaultfŒuille, 2017; Goodman-Bacon, 2018; Abraham and Sun, 2018), and it is straightfor-

[1] Exploiting variation in the minimum legal drinking age across U.S. states and over time, Dee and Evans find a spike in traffic fatalities at the drinking age and lower traffic fatalities among drivers who have had legal access to alcohol for longer.

ward to verify that our parallel-trends assumption held before the reform was introduced in 2007. Moreover, by focusing on the types of crashes targeted by the restriction, late-night crashes with multiple passengers, our estimates not only have more statistical power than previous estimates but are also robust to general trends in crash rates among young people.

We estimate that the restriction caused a 54% reduction in the late-night multi-passenger crash rates of first-year drivers. This large effect is precisely estimated and robust to a wide range of specification checks. We show that the reduction in crashes led to statistically and economically significant reductions in injuries and fatalities. For example, our estimates imply that the restriction reduced the associated number of fatalities from late-night multi-passenger crashes by 75% and overall fatalities involving first-year drivers by 13.5%. This 13.5% reduction in fatalities is impressive given that, before the restriction, less than 3% of crashes involving first-year drivers occurred during the restricted time-window with multiple passengers.

To examine if there were any spillovers from the restriction, we start by estimating how the crash reduction among first-year drivers affected other road users. We find that the restriction had large positive externalities on other road users. Given the types of crashes that were most affected — crashes involving multiple vehicles and multiple passengers — first-year drivers account for just 30% of the total reduction in injuries and fatalities, with the remaining 70% explained by their passengers (40%) and other road users (30%). Then, we examine if the restriction caused offsetting increases in other types of crashes, such as crashes just outside the restricted time-window or crashes during the restricted time-window with fewer passengers. We find no evidence of such effects.

To examine the long-term effects of the restriction, we estimate the impact on drivers who were previously subject to the restriction. We find that the road safety benefits persist in an attenuated manner in the five years after the probation period — that is, previously restricted drivers continue to have fewer late-night crashes with passengers because of the restriction. The estimated long-term effect on crash rates is around one-third of the size of the effect on first-year drivers in percentage terms, and the estimates imply that previously restricted drivers explain around one-quarter of the total reduction in crashes.

To study the key mechanism behind these effects, we draw on detailed information collected at the scene of the crash by police to examine whether the decrease in crash rates is more likely to result from a decrease in exposure — the amount of late-night driving with passengers (in hours/miles) — or crash risk (i.e., crashes per mile/hour driven). Our estimates indicate that the short- and long-term effects are predominantly caused by reduced exposure; that is, the short-term effect on crash rates appears to result from high rates of compliance with the restriction and the long-term effects from a persistent reduction in late-night driving with passengers. This analysis, combined with several other properties of the long-term effect on crash rates, indicates that graduated licensing restrictions can shape the long-term driving habits of young people.

This change in driving habits occurs in years in which young drivers continue to have high crash and fatality rates. Despite the progress in reducing traffic fatalities in New South Wales, newly licensed drivers still have much higher crash rates than more experienced drivers. Figure 2 shows how crash rates varied with experience for newly licensed drivers in the three years prior to the reform. Evidently, crash rates are highest for newly licensed drivers, with around 2% of newly licensed drivers having a serious crash within their first three months of independent driving. These rates decline considerably within the first year of driving; after one year of driving, crash rates are 50% lower. However, crash rates remain high for second- and third-year drivers and continue to decline beyond the first year. These declines are especially large for the type of crashes targeted by the restriction: late-night crashes with multiple passengers (Figure 2b). This demonstrates the promise of targeted interventions in changing the habits of young drivers throughout their highest-risk years.

Related literature. This paper contributes to a large literature in economics on the determinants of road crashes and fatalities. Economists have made many important contributions here, examining the importance of seat belts and seat belt regulation (Levitt and Porter, 2001*b*; Dee and Evans, 2001; Cohen and Einav, 2003; Carpenter and Stehr, 2008); air bags (Levitt and Porter, 2001*b*); police enforcement (DeAngelo and Hansen, 2014); beer taxes (Dee, 1999); and the minimum legal drinking age and drink-driving (Dee, 1999; Dee and Evans, 2001; Levitt and Porter, 2001*a*; Carpenter, 2004, 2006; Carpenter and Dobkin, 2009; Lovenheim and Slemrod, 2010; Lindo, Siminski and Yerokhin, 2015). However, economists have shown surprisingly little interest — with three key exceptions (Dee, Grabowski and Morrisey, 2005; Karaca-Mandic and Ridgeway, 2010; Gilpin, 2019) — in understanding the efficacy of graduated licensing systems.[2]

[2] There has been considerably more interest in the efficacy of these policies outside of economics. Most of these studies find improvements in road safety from graduated licensing systems overall and their individual components

Dee, Grabowski and Morrisey (2005) use variation across states in the U.S. in the adoption of — and type of restrictions applied by — graduated licensing systems. They find that graduated licensing reduced traffic fatalities among 15–17-year-olds by at least 5.6%. Karaca-Mandic and Ridgeway (2010) use a structural model to examine whether this reduction is caused by a reduction in crash risk or a reduction in driving. Consistent with our results, they find that graduated licensing reduces the amount teenagers drive but not their crash risk. Gilpin (2019) further decomposes the effect of graduated licensing systems on fatalities and finds that, after accounting for the reduction in license acquisition, there is no consistent effect on fatalities per licensee. Interestingly, Gilpin also disaggregates the effects of graduated licensing systems into the effects of individual components and does not find a statistically significant impact of either night-time curfews or passenger restrictions.

We build on these studies in several ways. First, we show that night-time passenger restrictions can be a very effective component of graduated licensing systems, not only reducing the total number of crashes and fatalities but also the rates per licensee. This finding contrasts with the statistically insignificant estimates of Gilpin (2019) for the U.S. of the impact of night-time curfews and passenger restrictions. This discrepancy is likely to reflect, at least in part, the fact that our estimates focus on the types of crashes that are targeted by the restriction, which considerably increases the power and robustness of our estimates. Second, we show that graduated licensing regulations can cause large positive spillovers on other road users that should be considered in any calculations of the welfare benefits of these restrictions. Third, we explicitly examine whether behavioral responses to these restrictions may result in an increase in other types of crashes and, somewhat surprisingly, we find no evidence to validate this concern. Fourth, we show interesting heterogeneity in the effects on crash rates. This heterogeneity indicates much higher rates of compliance among drivers who complied with speed and alcohol restrictions, indicating that graduated licensing restrictions may be relatively ineffective in changing the behavior of the most dangerous drivers. Finally, we show that graduated licensing regulations can cause persistent improvements in road safety; while Dee, Grabowski and Morrisey (2005) and Karaca-Mandic and Ridgeway (2010) examine the long-term effects of graduated licensing regulations, neither study finds a significant effect.

This paper also contributes to the literature on habit formation. Studies discussing the

(e.g., see Foss, Feaganes and Rodgman, 2001; Shope et al., 2001; Eisenberg, 2003; Fell, Todd and Voas, 2011).

importance of habits in consumer behavior have a long history in economics, dating back to Marshall (1920) and Duesenberry (1949). More recently, economists have used empirical methods to examine whether habit formation occurs in important types of behaviors, such as voting (Shachar, 2003; Gerber, Green and Shachar, 2003; Bechtel, Hangartner and Schmid, 2018; Fujiwara, Meng and Vogl, 2016); exercise (Charness and Gneezy, 2009; Royer, Stehr and Sydnor, 2015); nutrition (Naik and Moore, 1996; List and Samek, 2015; Loewenstein, Price and Volpp, 2016); alcohol consumption (Grossman, Chaloupka and Sirtalan, 1998; Williams, 2005); and electricity consumption (Ito, Ida and Tanaka, 2018). To our knowledge, though, this paper provides the first evidence of habit formation in driving behavior. This result is important; it suggests that policymakers can use short-term interventions to persistently reduce the amount people drive (and potentially improve driving), which may be desirable in many circumstances due to the negative externalities produced by car travel.

The rest of this paper is organized as follows. Section 2 explains the policy environment for young drivers in New South Wales and the details of the late-night passenger restriction. Section 3 describes the data and empirical approach. Section 4 presents the results. Section 5 concludes and discusses the implications for policy.

2 Policy background

2.1 Graduated licensing system in New South Wales

The current three-stage graduated licensing system was introduced in 2000. First, a permit to learn to drive can be obtained from age 16 by passing a written test. This allows people to drive when accompanied by a fully licensed driver. Second, an initial probationary ("P1") license can be obtained from age 17 by passing both written and driving tests. A P1 license permits unsupervised driving, although drivers are prohibited from traveling faster than 90 km/h (56 mph) and required to have a zero blood alcohol concentration. After 2007, they were also subject to the passenger restrictions that we detail below. After holding a P1 license for one year, individuals can obtain a P2 license by passing a written exam. The P2 license increases the speed restriction to 100 km/h (63 mph) and removes the passenger restriction. After holding a P2 license for two years, drivers can apply for an unrestricted license.

The system in New South Wales (NSW) is similar to the systems in other Australian states.

Like NSW, most states allow supervised driving from the age of 16 and unsupervised driving from the age of 17. In addition, drink-driving laws are uniform across Australia, with learner and probationary drivers required to have a blood alcohol concentration (BAC) of zero regardless of their age.[3] Furthermore, peer-passenger restrictions are enforced not only in NSW but also in Victoria and Queensland, the two most populous states in Australia after NSW. Queensland's restriction is almost identical to the one in NSW, while the Victorian restriction applies all day.

Internationally, the NSW system is more similar to systems in the U.S. than the less restrictive systems in Europe, where night-time and passenger restrictions are rare. However, there are a few important differences to note. First, in most U.S. states, teenagers can drive — both supervised and unsupervised — from a younger age. In most U.S. states, teenagers can drive with supervision from the age of 15 and without supervision from the age of 16. Second, the probation period is shorter in most U.S. states. Hence, all states in the U.S. allow teenagers to progress to an unrestricted license by the age of 18, while drivers in NSW must be at least 20. Finally, most U.S. states enforce night-time curfews and all-day passenger restrictions on drivers in the probation period, while NSW has a hybrid policy, with passenger restrictions between 11:00 pm and 5:00 am.

2.2 Introduction of passenger restrictions

In January 2007, the NSW Government announced several changes to the graduated licensing system in an attempt to improve the road safety of young people. Besides the introduction of the late-night passenger restriction, these changes — which all became effective on July 1, 2007 — included an immediate three-month license suspension for any P1 driver caught speeding and a ban on all mobile phone use for P1 and Learner drivers. In addition, Learners under the age of 25 who obtained their license on or after July 1, 2007 were required to hold their license for at least 12 months (previously six months) and log 120 hours of supervised driving (previously 50 hours), including 20 hours of night driving.

In this paper, we focus on the impact of the late-night passenger restriction, which prohibits P1 drivers under 25 from carrying more than one passenger under the age of 21 between 11:00 pm and 5:00 am.[4] The restriction is backed by a significant penalty. As of July 2017, the penalty is

[3]Prior to May 3, 2004, the BAC limit for probationary drivers in NSW was 0.02.

[4]Drivers can be granted an exemption under special circumstances, e.g., for employment reasons or family responsibilities.

a fine of \$549 (around \$384 US) and three demerit points — the latter is a significant penalty in itself as P1 drivers have their licenses suspended after receiving four or more demerit points. Moreover, the restriction is enforced. According to the *The Daily Telegraph* (Hildebrand and Kaye, 2008), about one per cent of all P1 drivers were caught breaking the passenger restriction by NSW Police in the year after the restriction was introduced.

As the passenger restriction only applies late at night, we can isolate its impact from the impact of other changes to the graduated licensing system under the assumption that these changes had the same effect, in percentage terms, on late-night crash rates as crash rates at other times of the day. Later, we show support for this assumption using the sample of crashes involving P1 drivers under 25 with zero or one passenger. As these drivers are not subject to the passenger restriction at any hour of the day, they provide a useful placebo group.

Expected effects on road safety. The passenger restriction may have short- and long-term effects on road crashes and fatalities. In the short term, the restriction effectively increases the marginal cost of a late-night trip with multiple passengers for P1 drivers because of the risk of a fine (and license suspension for drivers with prior infringements). As such, we would expect a reduction in these types of trips and an associated reduction in crashes. The restriction may also affect other road users and other types of crashes. Positive spillovers may occur on other road users if they are less likely to be involved in a crash with a P1 driver during the restricted time-window. In contrast, negative spillovers may arise from behavioral responses to the restriction by P1 drivers and their passengers. For example, an increase in crashes could result from P1 drivers (i) shifting late-night trips with multiple passengers to times just outside the restricted window, (ii) making more trips during the restricted window with zero or one passenger, or (iii) traveling more frequently during the restricted window with other drivers or as pedestrians.

In the long term, once drivers move onto a P2 license and are no longer subject to the passenger restriction, there is no effect on the marginal cost of a late-night trip with multiple passengers. However, the restriction may still affect crash rates, through long-term effects on exposure — the amount people drive late at night with multiple passengers (in hours/miles) — or crash risk (i.e., crashes per mile/hour driven). For example, if learning by doing is present, we would expect a long-term increase in crash risk — and an associated increase in crashes —

because the restriction limits drivers from gaining night-time driving experience with passengers. In contrast, if habit formation occurs — because individuals derive disutility from changes in their driving behavior over time — we would expect a persistent decrease in exposure and an associated decrease in crashes. As these mechanisms imply opposing effects on crash rates, the magnitude and direction of any long-term effects will shed light on the relative importance of each mechanism.[5]

3 Data and empirical approach

This section describes our data and empirical strategy. Section 3.1 describes our administrative crash and license data. Section 3.2 describes our baseline differences-in-differences approach, which uses crash rates during the unrestricted window (5:00 am to 10:59 pm) as the comparison group for crash rates during the restricted window (11:00 pm to 4:59 am). Section 3.3 presents graphical evidence in support of the key identifying assumptions and shows preliminary evidence of a large impact on the late-night multi-passenger crash rates of P1 drivers.

3.1 Data and descriptives

We use two linked sources of administrative data provided by the NSW Department of Transport. The first dataset, CrashLink, contains detailed information on road crashes collected at the scene of the crash by NSW police. Crashes are included in CrashLink if they satisfy four criteria: (i) the crash was reported to the police; (ii) the crash occurred on a road open to the public in NSW; (iii) the crash involved at least one moving road vehicle; and (iv) the crash involved either at least one vehicle being towed or at least one person being injured or killed. Crashes that satisfy criteria (ii)–(iv) are legally required to be reported to the police.

CrashLink contains detailed information at the crash, vehicle and person level. At the crash level, CrashLink includes information on the time of the crash in one-hour intervals; the month, year and day of the week of the crash; and the location and speed limit of the road where the crash occurred. At the vehicle level, CrashLink includes information on the type of traffic unit, e.g., car, motorbike, or pedestrian; the number of people traveling in the vehicle; and whether the vehicle was considered to be the main contributor to the crash.[6] At the person

[5]Of course, it is also possible to come up with plausible explanations for a long-term decrease in crash risk or a long-term increase in exposure. We examine these possibilities in Section 4.4.

[6]A vehicle is considered to be the key vehicle if it performed the maneuver that was the main contributor to

level, information is available on the drivers involved in the crash and anyone who was injured or killed in the crash. This information includes the age (in years) and gender of the person; the role of the person in the crash, e.g., driver, passenger, or pedestrian; the current license status of each driver at the time of the crash; and indicators for whether the person was injured or killed, including information on the severity of any injuries.

The second dataset contains the complete driver-license history (for car licenses) from 1996 of everyone in NSW who obtained their first license under the age of 30. The data includes the license class, e.g., Learner, P1, P2, or Unrestricted; the date the license became effective; the date the license expired; the person's age in months on the date the license became effective; and a unique de-identified customer ID that allows us to merge these records with CrashLink. This data allows us to determine (i) whether P1 drivers are subject to the restriction in each month and (ii) whether P2 and Unrestricted drivers were previously subject to the passenger restriction and, if so, how long ago.

We restrict the sample to the period between June 2004 and September 2014. We do not consider crashes prior to June 2004 because the BAC limit for probationary drivers was reduced from 0.02 to zero in May 2004. We make this restriction because we are concerned that any resulting change in drink-driving behavior would disproportionately affect night-time crash rates (and thus confound our estimates). We do not include crashes after September 2014 because of a reporting change in October 2014, which meant that the police were no longer required to attend the scene of the crash if no one was injured or killed.

Table 1 summarizes the crashes of P1 drivers under 25 over the sample period.[7] On average, there were 4,776 serious crashes per annum, or around four crashes per 100 drivers. In 44% of crashes, there was at least one casualty (injury or fatality), most of which were injuries. On average, P1 drivers were involved in 23.1 fatal crashes per annum (leading to 26.4 fatalities). Most crashes occurred during the unrestricted window, and P1 drivers usually had at most one passenger. However, multi-passenger crashes during the restricted window were much more serious in terms of the average number of fatalities per crash.

---

the impact. We use this variable as a proxy for which vehicle caused the crash.

[7]We exclude 0.1% of crashes in which P1 drivers are not classified as driving a car, car derivative or light truck.

3.2 Empirical strategy

We use a differences-in-differences approach to identify the causal effect of the passenger restriction on the late-night multi-passenger crash rates of P1 drivers. Specifically, we compare the change after the reform in the multi-passenger crash rates of P1 drivers during the restricted window (11:00 pm to 4:59 am) to the corresponding change during the unrestricted window (5:00 am to 10:59 pm). We compare these changes in relative terms rather than in levels because crash rates are higher during the unrestricted window. Hence, the key identifying assumption is that, if not for the restriction, the percentage change in multi-passenger crash rates for P1 drivers would have been the same during the restricted and unrestricted windows.

To implement this approach, we construct a monthly panel on the crash outcomes of each P1 driver. The main outcome is the number of crashes involving each driver, but we also consider other outcomes, such as the number of crash-related injuries and fatalities. We disaggregate all crash outcomes based on (i) whether the crash occurred during the restricted window (11:00 pm to 4:59 am) or the unrestricted window (5:00 am to 10:59 pm) and (ii) whether P1 drivers were carrying multiple passengers or 0–1 passenger.[8]

We implement our approach by estimating the following Poisson regression models:

$$y_{itw} = \exp\{\beta_{DID}(Rest._w \times After_t) + \beta_1 Rest._w + \beta_2 After_t + \mathbf{X'_{itw}}\gamma\}\varepsilon_{itw} \tag{1}$$

where i denotes an individual who had a P1 license and was under the age of 25 in month-year t; w denotes the time-window (11:00 pm to 4:59 am or 5:00 am to 10:59 pm); y_{itw} is the crash outcome of interest, such as the number of multi-passenger crashes involving driver i in month-year t and time-window w; $Rest._w$ is a dummy variable equal to one for the 11:00 pm to 4:59 am window to control for the lower crash rates during this window; $After_t$ is a dummy variable equal to one for months after the restriction was introduced, i.e., after June 2007; $\mathbf{X'_{itw}}$ includes (i) individual-level controls for driver age and gender, and their interactions with an indicator for the restricted time-window, to control for changes in the composition of drivers

[8]Unfortunately, we are unable to disaggregate crash outcomes based on whether the passengers were under the age of 21 as information on the age of passengers is only available for passengers who were injured or killed in the crash. As such, the estimated impact on late-night multi-passenger crash rates is likely to be a lower bound of the intended impact of the policy on the late-night crash rates of P1 drivers with multiple passengers *under the age of 21.*

that could disproportionately affect late-night crash rates and (ii) fixed effects for both calender month by time window, to control for any seasonality in the temporal crash patterns of drivers, and calender year, to control for changes in the crash rates of drivers over time; and ε_{itw} is a multiplicative error term. In all regressions, we cluster standard errors at the level of the individual driver to allow for an arbitrary correlation between each driver's errors, though standard errors are similar without clustering.

The key term is $Rest._w \times After_t$. Its coefficient, β_{DID}, estimates the change in multi-passenger crash rates during the restricted window after the passenger restriction was introduced divided by the corresponding change during the unrestricted window. Namely, we can derive the following expression for the exponential of β_{DID}:

$$\exp(\beta_{DID}) = \frac{\left(\dfrac{\mathbb{E}[y_{itw}|Rest._w = 1, After_t = 1, \mathbf{X'_{itw}}]}{\mathbb{E}[y_{itw}|Rest._w = 1, After_t = 0, \mathbf{X'_{itw}}]}\right)}{\left(\dfrac{\mathbb{E}[y_{itw}|Rest._w = 0, After_t = 1, \mathbf{X'_{itw}}]}{\mathbb{E}[y_{itw}|Rest._w = 0, After_t = 0, \mathbf{X'_{itw}}]}\right)}$$

Thus, the estimated percentage change in the relevant crash outcome (e.g., late-night multi-passenger crash rates) is equal to $100 \times (\exp(\beta_{DID}) - 1)$. There are two main threats to interpreting this expression as the causal effect of the passenger restriction. First, multi-passenger crash rates during the unrestricted window may follow a different trend to multi-passenger crash rates during the restricted window. Second, the comparison group may be contaminated if drivers respond to the restriction by shifting multi-passenger trips to the unrestricted window and this causes an associated increase in crash rates. In the next section, we examine these threats empirically and find no evidence to validate either concern. Rather, the graphical evidence provides strong support for the identification strategy.[9]

[9]Note that while it is more common to implement identification strategies of this nature using a log-linear specification, Poisson regression allows us to reliably estimate the effects at the individual-level. Although the log-linear model produces a near-identical estimate when crashes are aggregated to the half-year–year–time-window level, the estimates are not robust at the individual level due to the high proportion of zero counts (see Table A1). The high proportion of zeros, and low counts in general, also creates problems for the inverse-hyperbolic-sine (IHS) transformation, which is non-logarithmic for values close to zero. Estimating regressions at the individual level has two key advantages here. First, it allows us to cluster standard errors at the individual level and add individual-level controls. Second, for examining the long-term effects of the restriction, it eliminates measurement error in whether drivers were previously subject to the passenger restriction and allows us to examine how the long-term effects vary with (i) time since the restriction and (ii) time subject to the restriction.

3.3 Graphical evidence and checks of identifying assumptions

Figure 3 shows how crash rates have changed over time during the restricted and unrestricted windows for P1 drivers under 25. Given that our parallel-trends assumption is that changes in crash rates would have been similar in relative terms, we present the natural log of crash rates.[10] Figure 3a adds credibility to the parallel-trends assumption, showing similar trends in multi-passenger crash rates during the restricted and unrestricted windows before the passenger restriction was introduced. Immediately after the restriction was introduced, though, there was a large and persistent decrease in multi-passenger crash rates during the restricted window, with little change during the unrestricted window. Meanwhile, Figure 3b shows little change in 0–1-passenger crash rates at the time of the reform during either the restricted or unrestricted windows. Hence, the drop in multi-passenger crash rates during the restricted window after the reform does not appear to be driven by general trends in late-night crash rates.

Figure 4 validates the identification strategy further. It shows, for P1 drivers under 25, the average monthly crash rates in each one-hour window before and after the reform (under a log transformation). The vertical distance between markers shows how crash rates have changed in each one-hour window. Figure 4a shows a decrease in multi-passenger crash rates in every hour after the reform, but the decrease is considerably larger for hours during the restricted window. There is also no evidence of an increase in crash rates just outside the restricted window. This indicates that the comparison group is unlikely to be contaminated by P1 drivers shifting multi-passenger trips to the unrestricted window, assuming that such trips are most likely to be shifted to hours just outside the restricted window. Figure 4b shows the corresponding graph for 0–1-passenger crash rates. While crash rates have decreased in every hour of the day after the reform, the decrease is similar in all hours of the day. Overall, Figures 3 and 4 support the key identifying assumptions and suggest that the restriction has considerably reduced the late-night multi-passenger crash rates of P1 drivers.

4 Results

This section presents the results of the regression analysis. In Section 4.1, we present the baseline estimates, which show a large and precisely estimated decrease in the late-night multi-

[10] Before taking logs, we divide crash rates by the number of hours in each time-window (i.e., 18 or 6), which makes it easier to compare the two lines.

passenger crash rates of P1 drivers. We also verify the robustness of these estimates and examine heterogeneity in the effects. In Section 4.2, we examine the spillover effects of the reform. We find positive spillovers on other road users — from a reduced likelihood of being injured or killed in a crash with a P1 driver — but no evidence of an increase in other types of crashes resulting from behavioral responses to the restriction. In Section 4.3, we modify the baseline regressions to examine the long-term effects of the restriction. We find that the benefits of the restriction persist in an attenuated manner after the probation period, with previously restricted drivers having fewer late-night crashes with multiple passengers. In Section 4.4, we examine the key mechanism behind the short- and long-term effects of the restriction. We find that fewer late-night trips with passengers is likely to explain most, if not all, of the short- and long-term effects. As a final step, we explore several properties of the long-term effects to examine whether habit formation can explain the persistent effect on crash rates. Overall, this analysis reinforces habit formation as the likely explanation for the persistent effect on crash rates.

4.1 Short-term effects

Baseline estimates. Table 2 presents the baseline estimates. The estimates indicate that the passenger restriction has considerably reduced the number of multi-passenger crashes during the restricted window involving P1 drivers under 25. The differences-in-differences estimate of -0.784 ($p < 0.01$) indicates a 54.4% decrease in such crashes,[11] implying that the passenger restriction has prevented 59.5 serious crashes per annum since it was introduced. Table 2 shows that the restriction has also reduced the number of injuries and fatalities from these crashes. The estimates show statistically significant reductions in the number of crashes involving at least one casualty (52.7%, $p < 0.01$); the total number of casualties (48.2%, $p < 0.01$), the number of crashes involving at least one hospital admission or fatality (43.6% $p < 0.01$), the total number of hospital admissions and fatalities (45.4% $p < 0.01$), the number of fatal crashes (69.7% $p < 0.1$), and the total number of fatalities (75.3% $p < 0.05$). Overall, the points estimates imply that the restriction has prevented 48 casualties per annum, 12.1 hospital admissions per annum and 3.4 fatalities per annum.

Robustness. We verify the robustness of these results in several ways. First, we confirm that the results are not driven by a general decrease among P1 drivers in the proportion of

[11] $54.4 = 100 \times (1 - \exp\{-0.784\})$

crashes that occurred late at night. Specifically, we estimate equation (1) using the sample of crashes involving P1 drivers under 25 with 0–1 passenger. This regression tests whether any other factors, such as concurrent changes to the graduated licensing system, disproportionately affected the late-night crash rates of these drivers. The estimate in column 8 of Table 2 is close to zero and statistically insignificant (-0.023, $\sigma = 0.036$). Therefore, the estimated reduction in late-night multi-passenger crashes cannot be explained by a general trend among P1 drivers in the share of crashes that occurred late at night.

Second, we show that the parallel-trends assumption was satisfied before the reform. We estimate regressions that show the change over time in the late-night crash rates of P1 drivers under 25, relative to the change in their crash rates during the unrestricted window. Specifically, we estimate the following Poisson regressions:

$$y_{itw} = \exp\left(\sum_{\substack{j=2004:2 \\ j\neq 2007:1}}^{2014:1} \beta_{DID,j}\left(Rest._w \times \mathbf{1}(t \in j)\right) + \beta_1 Rest._w + hy_t + \mathbf{X'_{itw}}\gamma\right)\varepsilon_{itw} \tag{2}$$

where y_{itw} is the number of crashes involving driver i in month-year t and time-window w; hy_t are fixed effects for each half-year time period; and the $\beta_{DID,j}$ coefficients estimate the change in multi-passenger crash rates in the restricted window in each half-year period relative to the reference period, January–June 2007, the six months before the restriction was introduced.[12] Figure 5 presents the regression estimates. To aid interpretation, we translate the estimates into the implied percentage change in crash rates. Figure 5a shows that while there was no trend in multi-passenger crash rates during the restricted window before the restriction was introduced, there was an immediate, persistent and large decrease of around 50% after the restriction was introduced. In contrast, Figure 5b shows no change in 0–1-passenger crash rates.

Third, we confirm that the effects cannot be explained by P1 drivers responding to the reform by shifting multi-passenger trips from the restricted window to the unrestricted window. We estimate regressions that show how multi-passenger crash rates have changed after the reform

[12] We exclude crashes in June 2004 and July–September 2014 to implement this regression.

in each one-hour window. Specifically, we estimate the following regressions:

$$y_{ith} = \exp\left(\sum_{\substack{k=00:00\text{–}00:59 \\ k\neq 12:00\text{–}12:59}}^{23:00\text{–}23:59} \beta_{DID,k}\left(After_t \times \mathbf{1}(h=k)\right) + v_h + u_t + \mathbf{X'_{ith}}\right)\varepsilon_{ith} \tag{3}$$

where y_{ith} is the number of crashes involving driver i in month-year t and hour h; v_h is a set of fixed effects for each one-hour window; $\mathbf{X'_{ith}}$ includes individual-level controls for driver age and gender, and their interactions with indicators for each one-hour window; and the $\beta_{DID,k}$ coefficients estimate the change in crash rates in each one-hour window after the reform relative to the change at 12:00–12:59 pm. The idea behind these regressions is that while some trips during the restricted window may be substitutable to other times of the day, we would expect any spillovers to occur mainly in hours just outside the restricted window. Figure 6 presents the regression estimates. To aid interpretation, we translate the estimates into the implied percentage change in crash rates. Compared to the change at 12:00–12:59 pm, Figure 6a shows (i) a large decrease in multi-passenger crash rates during each of the restricted hours (35–65%) and (ii) no evidence of any increase in crash rates just outside the restricted window.[13] In contrast, Figure 6b shows similar changes in 0–1-passenger crash rates for all hours of the day.

Fourth, we show that the estimated effects are significantly larger among drivers who were more likely to be constrained by the restriction. To start, we compare the effects between P1 drivers aged 17–20 and 21–24. As the passenger restriction only applies to passengers under the age of 21, we would expect the second group of drivers to be less constrained by the passenger restriction, assuming that drivers tend to travel with people of a similar age. Table A2 confirms this prediction, showing that the restriction caused a significantly larger reduction ($p < 0.05$) in multi-passenger crash rates during the restricted window for P1 drivers aged 17–20 than those aged 21–24. In Table A2, we also examine how the effects varied based on the number of passengers. All else equal, we would expect the probability that drivers are carrying more than one passenger under the age of 21 to be increasing in the total number of passengers. Therefore, we would expect the absolute value of the estimates to be increasing with the number of passengers. Table A2 confirms this prediction. The estimates show a significantly larger

[13]In fact, the estimates for hours leading up to the restricted window show weak evidence of the opposite effect, a decrease in multi-passenger crash rates, which could occur if P1 drivers chose not to drive with multiple passengers before the restriction starts because they would be restricted on their return trip. However, none of the estimates prior to the restricted window are statistically significant at the 10% level.

reduction ($p < 0.01$) in the crash rates of drivers who were carrying three or more passengers than those with only two passengers.

Fifth, we show that the estimates are unlikely to be explained by P1 drivers under 25 choosing not to report crashes to police — in which these drivers were illegally carrying multiple passengers under the age of 21 — after the restriction was introduced. Note that this behavior is illegal and potentially difficult to get away with given the serious nature of the crashes reported in CrashLink. Nonetheless, we would expect selective non-reporting to be less of a concern for crashes involving multiple vehicles, especially crashes that were caused by P1 drivers under 25. Among the relevant crashes, Table A3 shows a 66.0% reduction in crashes involving multiple vehicles and a 62.7% reduction in crashes involving multiple vehicles where the P1 driver was likely to have caused the crash.[14] As these estimates are larger than the baseline estimates in Table 2 — and around three-quarters of crashes involve multiple vehicles — the estimates are unlikely to be strongly affected by selective non-reporting.

Finally, we examine the temporal crash patterns of other drivers. Given that most crashes involve multiple vehicles, changes in the temporal crash patterns of other drivers could bias the estimates in Table 2. To implement this test, we aggregate crashes involving drivers aged 25 and above (75% of all crashes) to the month–year–time-window level and estimate equation (1).[15] This regression tests whether, for other drivers, there was a disproportionate change in late-night crashes after the passenger restriction was introduced.[16] Table A4 shows that the differences-in-differences estimate from this large sample of crashes is a precisely estimated zero (0.0001, $\sigma = 0.0220$). Thus, the temporal crash patterns of other drivers did not change after the restriction was introduced, at least not in a way that would bias the estimates in Table 2.

Heterogeneity. Table 3 shows how the estimated effect on late-night multi-passenger crash rates varies based on the characteristics of the driver and the crash. In terms of the characteristics of the driver, while two-thirds of the *absolute* reduction in crashes is explained by male drivers, the point estimate implies that the percentage change in crash rates is larger,

[14] For each crash in the data, exactly one vehicle is classified as the "key traffic-unit". A vehicle is considered to be the key traffic-unit if it performed the maneuver that was the main contributor to the impact. We use this variable as a proxy for which driver caused the crash.

[15] We cannot estimate this regression at the individual level, as we do not have license information on (most) older drivers unless they were involved in crashes.

[16] We do not restrict the sample based on the number of passengers, as all crashes involving these drivers could potentially involve, in another vehicle, a P1 driver under 25 with multiple passengers.

though not statistically different, for female drivers ($p = 0.133$). Similarly, while 17–18-year-old drivers explain most of the absolute reduction in crashes, the percentage change in crash rates is statistically indistinguishable for 17–18-year-olds and 19–20-year-olds. In terms of the characteristics of the crash, most of the *absolute* reduction in crashes is explained by crashes involving multiple vehicles; crashes in Sydney, where most of the NSW population lives; crashes on Friday and Saturday nights; and crashes on major roads with a speed limit of 60 km/h or higher (37.5 mph). In *relative* terms, the estimated effects are statistically larger for crashes involving multiple vehicles ($p < 0.01$) and crashes on main roads with a speed limit of 60 km/h or higher ($p = 0.058$).

4.2 Spillovers

Direct effects on P1 passengers and other road users. In this section, we show how different types of road users were affected. Given the type of crashes that were most affected by the restriction — crashes with multiple passengers and multiple traffic units — it is important to understand not only how P1 drivers under 25 were affected but also the effects on their passengers and other road users. In Table 4, we quantify the effects of the restriction on the number of casualties for different types of road users. At the 5% level, Table 4 shows that the restriction has decreased the number of casualties for P1 drivers, their passengers, and other road users (i.e., pedestrians and people in other vehicles). The point estimates imply that the restriction has prevented 13.7 casualties per annum for P1 drivers, 17.8 casualties for their passengers,[17] and 13.7 casualties for other road users.[18] Thus, the restriction has had large positive spillovers on P1 passengers and other road users, with P1 drivers accounting for less than 30% of the total reduction in casualties.

Indirect effects from behavioral responses. To understand the overall effects of the restriction, it is also important to consider the extent of any effects on other types of crashes stemming from behavioral responses to the policy. In this section, we discuss and examine different responses that could cause an increase in other types of crashes. Before we begin, note

[17] All of the estimated impact on P1 passengers is explained by passengers under the age of 21, consistent with the rules of the policy (see Table A5).

[18] Interestingly, the percentage reduction in casualties for other road users (from crashes involving P1 drivers with multiple passengers) is significantly larger than the reduction in casualties for people traveling with a P1 driver. This stems from the fact that the restriction caused a larger percentage reduction in crashes involving multiple vehicles.

that we have already shown that there is no evidence of any increase in crashes resulting from P1 drivers under 25 shifting multi-passenger trips from the restricted window to the unrestricted window.

Another way that P1 drivers may respond is by driving with fewer passengers during the restricted window. For example, a group of four teenagers may choose to go in two cars, instead of one, which could result in an increase in the number of late-night crashes involving P1 drivers with zero or one passenger. However, the negative and statistically insignificant estimate in column 8 of Table 2 showed no evidence of such an increase. Moreover, the estimates remain negative and statistically insignificant if we restrict the sample further to P1 drivers aged 17–20, the group who are most likely to be constrained by the restriction, or drivers with exactly one passenger (see Table A6).

We may also expect P1 drivers and their passengers to substitute to other forms of transport, which could increase their chance of being involved in other types of late-night crashes. For example, instead of driving with friends, young people may choose to walk or travel with drivers who are not subject to the restriction. To examine these responses, we estimate whether there was any effect on the number of traffic casualties during the restricted window among people aged 16–20 from crashes *not* involving P1 drivers under 25. The estimates in columns 4 and 5 of Table A6 are negative and statistically insignificant, providing no evidence of any increase in casualties from young people traveling more as passengers in other vehicles or as pedestrians.

In summary, we find no evidence that behavioral responses to the restriction resulted in an increase in other types of crashes.

4.3 Long-term effects

In this section, we examine the long-term effects of the passenger restriction. Specifically, we examine whether the restriction had any long-term effect on drivers' propensity to crash during the restricted window with multiple passengers. To start, we consider the impact on drivers who graduated from a P1 license within the last five years.[19] We estimate the following regressions:

$$y_{itw} = \exp\{\beta_{DID,LT}(Rest._w \times Prev\_rest._i) + \beta_1 Rest._w + \beta_2 Prev\_rest._i + \mathbf{X'}_{itw}\gamma\}\varepsilon_{itw} \qquad (4)$$

[19] We exclude drivers who were over the age of 25 when they graduated from a P1 license.

where $crashes_{itw}$ is the number of multi-passenger crashes involving driver i in month-year t and time-window w.[20] These regressions modify equation (1) by replacing $After_t$, an indicator for whether the restrictions apply in month t, with $Prev\_Rest._i$, an indicator for whether driver i was previously subject to the P1 passenger restriction. The key identifying assumption for interpreting $\beta_{DID,LT}$ as a causal effect is as follows: if not for the passenger restriction, the proportion of crashes that occurred during the restricted window would have been the same for drivers who were previously subject to the passenger as those who were not.

Table 5 presents the estimates of equation (4). The estimates indicate that the passenger restriction has significantly reduced the late-night multi-passenger crash rates of drivers who were previously subject to the restriction. The estimate of $\beta_{DID,LT}$ in column 1 of -0.200 ($p < 0.01$) implies a 18.1% reduction in such crashes. As we would expect, this estimate is smaller than the estimated impact on drivers who were subject to the restriction. However, this estimate indicates that the effects of the restriction persist in a meaningful way. The estimated long-term reduction in late-night multi-passenger crash rates is around one-third of the size of the direct effect on P1 drivers, and the estimates imply that previously restricted drivers explain around one-quarter of the total reduction in crashes.

Robustness. We verify the robustness of the estimated long-term effect on crashes in two ways. First, we show in Table 5 that the estimated effect on 0–1-passenger crash rates is smaller and statistically insignificant at the 10% level (-0.023, $\sigma = 0.029$). Thus, the long-term effects on late-night multi-passenger crash rates cannot be explained by a general decrease in late-night crash rates.

Second, we examine the validity of the parallel-trends assumption by estimating equation (2) for drivers in their first five years after graduating from a P1 licence. Figure A1 displays the estimates. Prior to the reform, there were similar trends for multi-passenger crashes during the restricted and unrestricted windows. After the restriction was introduced, though, and previously restricted drivers flow into the sample, crash rates fall gradually and remain around 20% lower on average. In contrast, there is little change for 0–1-passenger crash rates.

[20] To avoid any bias from the decrease in late-night crashes among P1 drivers, we exclude the 6% of crashes that involved a P1 driver under 25 with multiple passengers. The estimates are similar, but slightly larger, without this restriction.

4.4 Key mechanism: Fewer trips or a decrease in crash risk?

Since a reduction in crashes can result from a reduction in exposure (miles/hours driven) or a reduction in crash risk (i.e., crashes per mile/hour driven), we examine the key mechanism behind the short- and long-term effects of the restriction.

Short-term effects. As discussed in Section 2.2, theory suggests that the restriction should reduce the number of late-night trips with passengers by P1 drivers under 25, as these drivers would otherwise risk a fine (and license suspension for drivers with prior infringements). However, it is not clear if and how the restriction would affect their crash risk.

Unfortunately, it is not possible to directly observe the importance of each mechanism. Without comprehensive data on when P1 drivers are driving and with how many passengers, we cannot estimate whether there has been a reduction in late-night driving with passengers. Thus, the importance of each mechanism has to be inferred. We do so by analyzing heterogeneity in the effects of the restriction.

Primarily, we focus on heterogeneity based on whether the P1 driver operated the vehicle that was the main contributor to the crash. This variable is available for all crashes and provides a useful proxy for which driver caused the crash. The key idea here is as follows: if there was a decrease in the average crash risk of P1 drivers, we would expect to find a larger decrease, in percentage terms, for crashes that were caused by P1 drivers than crashes that simply involved — but were not caused by — these drivers.

However, the estimates in Table 6 show the opposite pattern. The reduction in crash rates is statistically larger ($p < 0.01$) for crashes that were likely to have been caused by other drivers. This indicates that, if anything, the average crash risk of P1 drivers on the road may have increased because of the restriction.[21] By extension, this result suggests that the reduction in crashes among P1 drivers results entirely from fewer late-night trips with passengers.

This conclusion is further supported by the heterogeneity in the estimates based on whether P1 drivers were speeding or drinking alcohol at the time of the crash. If there was a decrease in the crash risk of P1 drivers, we may expect a larger decrease, in percentage terms, in crashes

[21] Note that this does not necessarily imply an increase in the crash risk of individual drivers. Rather, the increase in crash risk may reflect a more dangerous composition of P1 drivers on the road, which is supported by the statistically smaller reductions in crash rates among speeding and drinking drivers (discussed below).

where P1 drivers were speeding or had consumed alcohol (i.e., situations where their crash risk may be above average). However, the estimates in Table 6 show statistically larger reductions in crashes where P1 drivers were not speeding ($p < 0.01$) and not drinking ($p < 0.01$).

As a final step, we quantify the effect on late-night crashes involving — but not caused by — P1 drivers with multiple passengers. Crash rates of this type provide a useful proxy for the amount P1 drivers are driving late at night with multiple passengers. This assumes that the temporal crash patterns of other drivers are stable over time, which as discussed in Section 4.1, is strongly supported by the data. Table 6 shows that the restriction decreased the number of late-night crashes involving — but not caused by — P1 drivers with multiple passengers by 72.7% ($p < 0.01$).[22] If we assume that this reduction is roughly proportional to the change in the amount P1 drivers are driving late at night with passengers, this estimate implies a very high rate of compliance with the restriction.

Overall, while we do not explicitly observe driving behavior, the estimates suggest that the restriction reduced crashes by reducing the amount P1 drivers drove late at night with multiple passengers rather than reducing their risk of a crash during such trips.

Long-term effects. Given the short-term effects discussed above, we would expect the long-term effect on crash rates to result mainly from a persistent decrease in exposure — that is, less driving with multiple passengers during the restricted window — rather than a decrease in crash risk. To examine the relative importance of these mechanisms, we use the same approach as above and consider heterogeneity in the long-term effects based on (i) which driver caused the crash and (ii) whether previously restricted drivers were speeding or had consumed alcohol above the legal limit. By the same logic as above, if there was a decrease in the average crash risk of previously restricted drivers, we would expect to find a larger decrease, in percentage terms, for crashes that were caused by previously restricted drivers and crashes where these drivers were defying speed or alcohol limits. However, the estimates in Table 6 show little evidence of such heterogeneity; the estimates are slightly larger, though not statistically different, for crashes that involved — but were not caused by — previously restricted drivers and crashes where these drivers were complying with speed and alcohol limits. As such, fewer late-night trips with passengers is likely to explain most of the long-term impact on crash rates.

[22] $72.7 = 100 \times (1 - \exp\{-1.300\})$

This conclusion is further supported by the estimated reduction in late-night crashes involving — but not caused by — previously restricted drivers of 23.2% ($p < 0.1$).[23] Again, if we assume that this reduction is roughly proportional to the change in late-night driving with multiple passengers, this estimate indicates that the restriction had a substantial long-term impact on the amount of late-night driving with multiple passengers.

Habit formation is a plausible explanation for why the passenger restriction may cause a persistent reduction in late-night driving with passengers. Because P1 drivers were prohibited from driving during the restricted window with multiple passengers, we would expect these individuals to make fewer late-night trips with passengers even after transitioning to a P2 license if they derive disutility from changes in their driving behavior over time. We examine whether the estimates support this explanation in three ways.

1. Heterogeneity by years of exposure. Our first hypothesis for habit formation to explain the long-term impact on crash rates is that the effect on driving habits, and thus crash rates, would be larger for individuals who were subject to the restriction for longer. In our sample, the time individuals were subject to the restriction varies because some drivers had already started their P1 license period when the restriction was introduced. For example, an individual who held a P1 license from September 2006 to September 2007 was subject to the restriction for just 3–4 months, while someone who obtained their P1 license after July 2007 was subject to the restriction for at least one year. We use this variation to examine how the long-term effects vary with the number of years drivers were exposed to the restriction.[24] Specifically, we estimate the following Poisson regressions:

$$\begin{aligned} y_{itw} = &\exp\{\beta_{DID,LT}(Rest._w \times Prev\_rest._i) + \beta_1 Rest._w + \beta_2 Prev\_rest._i \\ &\beta_{exp}(Years\_restricted_i \times Rest._w) + \beta_3 Years\_restricted_i + \mathbf{X'}_{itw}\gamma\}\varepsilon_{itw} \end{aligned} \tag{5}$$

where $Years\_restricted_i$ is the number of years individual i was subject to the restriction. We use this regression to show how the estimated effect on crash rates varies with years subject to

[23] $23.2 = 100 \times (1 - \exp\{-0.264\})$

[24] We calculate years of exposure using the dates individuals moved onto a P2 license. For example, someone who moved onto a P2 license before July 2007 was not exposed to the restriction (years restricted = 0); someone who moved onto a P2 license in January 2008 was exposed to the restriction for 6.5 months (years restricted = 0.54); and someone who moved onto a P2 license after July 2008 was exposed to the restriction for at least one year (years restricted = 1). The estimates are similar if we calculate exposure based on the dates individuals obtained a P1 license (see Figure A2).

the restriction in Figure 7a. Specifically, we present the estimates and 95% confidence intervals of $\beta_{DID,LT} + (\beta_{exp} \times Years\_restricted)$, which shows how the long-term effect varies with the number of years individuals were subject to the restriction. Figure 7a shows that the estimated reduction in crash rates is larger and more precisely estimated for drivers who were subject to the restriction for longer. While this pattern may indicate stronger habit formation among individuals who were subject to the restriction for longer, the effects are not conclusive as the estimate of β_{exp} is statistically indistinguishable from zero (-0.196, $\sigma = 0.196$).

2. Heterogeneity by recency of exposure. Our second hypothesis for habit formation to explain the long-term effect on crash rates is that the impact on crash rates would (i) fall immediately after drivers are no longer subject to the restriction and (ii) weaken further over time. We examine these hypotheses in Figure 8. In Figure 8a, we allow the long-term effects to vary over time by estimating the following Poisson regressions:

$$
\begin{aligned}
y_{itw} = &\exp\{\beta_{LT}(Rest._w \times Prev\_rest._i) + \beta_1 Rest._w + \beta_2 Prev\_rest._i + \\
&\beta_{trend}\left(f(Yrs\_since\_P1_{it}) \times Rest._w \times Prev\_rest._i\right) + \beta_3\left(f(Yrs\_since\_P1_{it}) \times Rest._w\right) + \\
&\beta_4\left(f(Yrs\_since\_P1_{it}) \times Prev\_rest._i\right) + \beta_5 f(Yrs\_since\_P1_{it}) + \mathbf{X'}_{itw}\gamma\}\varepsilon_{itw} \qquad (6)
\end{aligned}
$$

where $f(Yrs\_since\_P1_{it})$ is a function of the number of years since driver i became subject to the restriction.[25] Figure 8a plots the estimates and 95% confidence intervals of $\beta_{LT} + \beta_{trend} \times \ln(Yrs\_since\_P1)$, which shows how the long-term effect varies with the number of years since the restriction under the assumption of logarithmic decay/growth. Figure 8 shows several interesting results. First, the estimated effect on crash rates is smaller in the second year after the start of the probation period, when drivers are no longer subject to the restriction. Second, the effect in the second year remains large (around 40%) and statistically significant at the 1% level. Third, the effect weakens over time, with a positive and statistically significant ($p < 0.05$) estimate of β_{trend}. Fourth, the estimated impact on late-night multi-passenger crash rates remains negative and statistically significant for several years, eventually returning to zero around six years after

[25] In these regressions, we define previously restricted drivers as drivers who obtain their P1 license after June 2007. Drivers who obtain their license prior to June 2006 form the control group. Drivers who obtain their P1 license in between these dates are omitted, as Figure 7a shows that the long-term effects are concentrated among drivers who were subject to the restriction for one year. The sample for these regressions consists of drivers who obtained their P1 license between one and six years ago — that is, we estimate the effect on drivers in the first five years after they were subject to the restriction.

the start of the probation period. Finally, we find a similar trend in the estimated effects in Figure 8b using a more flexible choice of $f(\cdot)$, where we use indicator variables to divide drivers into four groups based on the number of years since they started the probation period.

While the long-term effects decay over time, Figure 9 demonstrates the importance of these effects on the crash rates of drivers in their first six years of independent driving. Figure 9 shows how late-night multi-passenger crash rates changed with years of driving experience before and after the reform. We also plot the estimated counterfactual for post-reform drivers from non-parametric estimates of equation (6).[26] Before the reform, crash rates were much higher among inexperienced drivers. Drivers in their first year had crash rates nearly twice as high as drivers in their second year, three times as high as drivers in their third year, and around six times as high as drivers in their sixth year. After the reform, crash rates are nearly 70% lower in the first year and continue to decline with experience but in a more gradual manner. The dashed counterfactual line shows that while we would have expected a modest decrease in crash rates if the reform had not occurred, the observed decrease in crash rates is considerably larger. Comparing the post-reform line with the estimated counterfactual demonstrates that most of the crash reduction occurs in the first year. However, we estimate meaningful reductions in the next four years, years in which the crash risk of young drivers, while lower, remains high compared to older drivers. This demonstrates the importance of the restriction in developing safer driving habits among young drivers in their highest-risk years.

3. Correlation between short- and long-term effects. Our third and final hypothesis for habit formation to explain the long-term effects is that the long-term impact on driving habits would be largest for individuals whose driving behavior was most affected by the restriction in the short term. As such, we would expect the long-term impact of the restriction on driving behavior to resemble the short-term impact, in terms of the types of trips and crashes that are most strongly affected. This approach draws on other studies that document evidence of habit formation, which show strong correlations between the short- and long-term effects either under different treatments (Charness and Gneezy, 2009; Ito, Ida and Tanaka, 2018) or for different subgroups (Fujiwara, Meng and Vogl, 2016).[27] Figure 10 presents the short- and long-term estimates for 14 types of crashes, with the estimates translated into the implied

[26]We estimate a separate treatment effect for each one-year period after the start of the probation period.

[27]In a different context, Fitzpatrick and Moore (2018) use a similar idea to examine the link between the increases in retirement and mortality at age 62 in the United States.

percentage change in crash rates. Evidently, there is a strong positive correlation between the short- and long-term effects. For example, the short- and long-term estimates are largest for crashes involving multiple vehicles; crashes *not* caused by restricted/previously restricted drivers; crashes where these drivers were complying with speed and alcohol limits; and crashes involving female drivers. This suggests that the restriction affected the driving behavior of similar individuals in the short and long term.

Overall, the results in this section reinforce habit formation as the likely explanation for the estimated long-term effect on late-night multi-passenger crash rates. First, we showed that a persistent reduction in late-night trips with passengers is more likely to explain the long-term effect on crash rates than a reduction in crash risk. Second, we showed that the effect on crash rates falls immediately once drivers are no longer subject to the restriction. Third, we showed evidence of larger effects on drivers who were subject to the restriction for (i) a longer period of time and (ii) more recently. Finally, we showed a positive correlation between the types of crashes that are most strongly affected in the short and long term, consistent with a larger impact on the driving habits of individuals who had a stronger short-term response to the restriction.

5 Conclusion

We show that a late-night passenger restriction on first-year drivers has substantially improved road safety in New South Wales (NSW), Australia. We estimate that this targeted, light-touch intervention has reduced late-night crashes with multiple passengers among first-year drivers by 54% and associated fatalities by 75%. Our estimates imply that the restriction has directly prevented 32.5 people from minor injuries each year, 12.1 people from being admitted to hospital and 3.4 people from being killed. These reductions are economically significant overall: a 1.6% decrease in the total number of minor injuries from crashes involving first-year drivers, a 2.4% decrease in the total number of hospital admissions, and a 13.5% decrease in the total number of fatalities. Moreover, we find no evidence that these benefits were attenuated either by an increase in other types of crashes or by an increase in crashes after the first year. In fact, we find that the road safety benefits persist in a weaker form for several years.

Our estimates imply large economic benefits. As cited in Naude, Makwasha and McGeehan (2015), the Australian Road Research Board estimates the average value of a statistical life at

$7,573,412 (in 2013 dollars), the average value of an injury resulting in a hospital admission at $100,431, and the average value of a minor injury at $31,739. Combining these values with the annual injury and fatality reductions implied by our estimates, we estimate that the restriction has reduced the economic costs of road crashes involving first-year drivers by $28.0M per annum, a 9.1% reduction — equivalent to an average annual economic saving of $236 per first-year driver ($165 US).

Like other studies evaluating the efficacy of graduated licensing systems, this paper cannot evaluate the welfare costs of the passenger restriction. While such costs may include additional work for police, the largest cost may be the disutility imposed on young people. Placing a value on this disutility is difficult, although we would expect that at least some drivers may be willing to pay $236 to avoid the restriction. Nonetheless, as noted by Dee, Grabowski and Morrisey (2005), many parents, policymakers and other citizens may reject such cost-benefit appraisals in favor of policies that are effective in improving the road safety of young people in their highest-risk years.

Overall, our estimates are larger than Dee, Grabowski and Morrisey's (2005) estimates of the impact of graduated licensing systems in the U.S. on traffic fatalities among 15–17-year-olds. While Dee, Grabowski and Morrisey (2005) estimate that graduated licensing reduced traffic fatalities among 15–17-year-olds by at least 5.6%, our estimates imply that, by itself, the late-night passenger restriction caused a 13.5% reduction in the number of fatalities from crashes involving first-year drivers. Moreover, this comparison understates just how large our estimates are for several reasons.

First, unlike Dee, Grabowski and Morrisey (2005), our estimated effects on fatalities are per licensee, rather than per capita. This distinction is important as, based on the recent work of Gilpin (2019) in the U.S. context, a reduction in licenses appears to be a key mechanism behind the impact of graduated licensing restrictions on fatalities.

Second, the road-safety benefits of the passenger restriction are especially impressive considering that it only applies between 11:00 pm and 5:00 am. Back-of-the-envelope calculations suggest that an all-day passenger restriction, as is present in most U.S. states, would result in much larger benefits. If we assume the same percentage reduction in injuries and fatalities from multi-passenger crashes during daytime hours, all-day restrictions would reduce the total number of minor injuries from crashes involving first-year drivers by 10.7%, the total number of hospital admissions by 9.4% and the total number of fatalities by 24.4%, producing an average annual economic saving of $553 per driver.

Third, the late-night passenger restriction is just one of many components of the graduated licensing system in NSW. As such, the overall benefits of the NSW graduated licensing system may be considerably larger if other aspects of the NSW system — such as mandatory minimum levels of supervised driving experience and a minimum three-year probation period — also prevent road crashes and fatalities. Moreover, the peer-passenger restriction was just one of several changes to the NSW system in July 2007. Although a careful evaluation of the concurrent changes to the NSW system is beyond the scope of this paper, there is suggestive evidence of their effectiveness that warrants further research: over the sample period, overall crash rates were 21.4% lower for first-year drivers after these changes.

Finally, unlike previous studies examining the efficacy of graduated licensing systems by Dee, Grabowski and Morrisey (2005) and Karaca-Mandic and Ridgeway (2010), we show convincing evidence that the passenger restriction caused persistent improvements in road safety. These effects are around one-third of the size of the direct effects of the restriction and are important in extending the improvements in road safety to years in which crash rates remain high.

Overall, this paper shows that passenger restrictions can be a very effective component of graduated licensing systems. While passenger restrictions have become common in Australia and the U.S., most jurisdictions in Europe do not enforce such restrictions. Our results indicate that European governments could save lives and significantly reduce the economic and social costs of road crashes by introducing passenger restrictions on novice drivers.

References

Abraham, Sarah, and Liyang Sun. 2018. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects."

Bechtel, Michael M, Dominik Hangartner, and Lukas Schmid. 2018. "Compulsory voting, habit formation, and political participation." *Review of Economics and Statistics*, 100(3): 467–476.

Blincoe, LJ, TR Miller, Eduard Zaloshnja, and Bruce A Lawrence. 2015. "The economic and societal impact of motor vehicle crashes, 2010 (revised)." Washington, DC: National Highway Traffic Safety Administration (No. DOT HS 812 013).

Borusyak, Kirill, and Xavier Jaravel. 2017. "Revisiting event study designs."

Carpenter, Christopher. 2004. "How do zero tolerance drunk driving laws work?" *Journal of Health Economics*, 23(1): 61–83.

Carpenter, Christopher. 2006. "Did Ontario's zero tolerance & graduated licensing law reduce youth drunk driving?" *Journal of Policy Analysis and Management*, 25(1): 183–195.

Carpenter, Christopher, and Carlos Dobkin. 2009. "The effect of alcohol consumption on mortality: Regression discontinuity evidence from the minimum drinking age." *American Economic Journal: Applied Economics*, 1(1): 164–82.

Carpenter, Christopher S, and Mark Stehr. 2008. "The effects of mandatory seatbelt laws on seatbelt use, motor vehicle fatalities, and crash-related injuries among youths." *Journal of Health Economics*, 27(3): 642–662.

Charness, Gary, and Uri Gneezy. 2009. "Incentives to exercise." *Econometrica*, 77(3): 909–931.

Chen, Li-Hui, Susan P Baker, Elisa R Braver, and Guohua Li. 2000. "Carrying passengers as a risk factor for crashes fatal to 16-and 17-year-old drivers." *Journal of the American Medical Association*, 283(12): 1578–1582.

Cohen, Alma, and Liran Einav. 2003. "The effects of mandatory seat belt laws on driving behavior and traffic fatalities." *Review of Economics and Statistics*, 85(4): 828–843.

DeAngelo, Gregory, and Benjamin Hansen. 2014. "Life and death in the fast lane: Police enforcement and traffic fatalities." *American Economic Journal: Economic Policy*, 6(2): 231–257.

De Chaisemartin, Clément, and Xavier D'HaultfŒuille. 2017. "Fuzzy differences-in-differences." *Review of Economic Studies*, 85(2): 999–1028.

Dee, Thomas S. 1999. "State alcohol policies, teen drinking and traffic fatalities." *Journal of Public Economics*, 72(2): 289–315.

Dee, Thomas S, and William N Evans. 2001. "Behavioral policies and teen traffic safety." *American Economic Review*, 91(2): 91–96.

Dee, Thomas S, David C Grabowski, and Michael A Morrisey. 2005. "Graduated driver licensing and teen traffic fatalities." *Journal of Health Economics*, 24(3): 571–589.

Duesenberry, James. 1949. *Income, Saving, and the Theory of Consumer Behavior.* Cambridge, Massachusetts: Harvard University Press.

Eisenberg, Daniel. 2003. "Evaluating the effectiveness of policies related to drunk driving." *Journal of Policy Analysis and Management*, 22(2): 249–274.

Fell, James C, Michael Todd, and Robert B Voas. 2011. "A national evaluation of the nighttime and passenger restriction components of graduated driver licensing." *Journal of Safety Research*, 42(4): 283–290.

Fitzpatrick, Maria D, and Timothy J Moore. 2018. "The mortality effects of retirement: Evidence from Social Security eligibility at age 62." *Journal of Public Economics*, 157: 121–137.

Foss, Robert D, John R Feaganes, and Eric A Rodgman. 2001. "Initial effects of graduated driver licensing on 16-year-old driver crashes in North Carolina." *Journal of the American Medical Association*, 286(13): 1588–1592.

Fujiwara, Thomas, Kyle Meng, and Tom Vogl. 2016. "Habit formation in voting: Evidence from rainy elections." *American Economic Journal: Applied Economics*, 8(4): 160–88.

Gerber, Alan S, Donald P Green, and Ron Shachar. 2003. "Voting may be habit-forming: Evidence from a randomized field experiment." *American Journal of Political Science*, 47(3): 540–550.

Gilpin, Gregory. 2019. "Teen driver licensure provisions, licensing, and vehicular fatalities." *Journal of Health Economics*, 66: 54–70.

Goodman-Bacon, Andrew. 2018. "Difference-in-differences with variation in treatment timing."

Grossman, Michael, Frank J Chaloupka, and Ismail Sirtalan. 1998. "An empirical analysis of alcohol addiction: Results from the monitoring the future panels." *Economic Inquiry*, 36(1): 39–48.

Hildebrand, Joe, and Byron Kaye. 2008. "Lifesaving P-plate laws: One year on and the rules are making a big impact." *The Daily Telegraph*, July 4, Sydney, Australia.

Insurance Institute for Highway Safety. 2017. "Fatality facts 2016." http://www.iihs.org/iihs/topics/t/general-statistics/fatalityfacts/overview-of-fatality-facts#Age-and-gender (accessed on 27/02/2018).

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. "Moral suasion and economic incentives: Field experimental evidence from energy demand." *American Economic Journal: Economic Policy*, 10(1): 240–67.

Karaca-Mandic, Pinar, and Greg Ridgeway. 2010. "Behavioral impact of graduated driver licensing on teenage driving risk and exposure." *Journal of Health Economics*, 29(1): 48–61.

Levitt, Steven D, and Jack Porter. 2001*a*. "How dangerous are drinking drivers?" *Journal of Political Economy*, 109(6): 1198–1237.

Levitt, Steven D, and Jack Porter. 2001*b*. "Sample selection in the estimation of air bag and seat belt effectiveness." *Review of Economics and Statistics*, 83(4): 603–615.

Lindo, Jason M, Peter Siminski, and Oleg Yerokhin. 2015. "Breaking the link between legal access to alcohol and motor vehicle accidents: Evidence from New South Wales." *Health Economics*.

List, John A, and Anya Savikhin Samek. 2015. "The behavioralist as nutritionist: Leveraging behavioral economics to improve child food choice and consumption." *Journal of Health Economics*, 39: 135–146.

Loewenstein, George, Joseph Price, and Kevin Volpp. 2016. "Habit formation in children: Evidence from incentives for healthy eating." *Journal of Health Economics*, 45: 47–54.

Lovenheim, Michael F, and Joel Slemrod. 2010. "The fatal toll of driving to drink: The effect of minimum legal drinking age evasion on traffic fatalities." *Journal of Health Economics*, 29(1): 62–77.

Marshall, Alfred. 1920. *Principles of Economics: An Introductory Volume.* 8th ed. London: Macmillan.

Mayhew, Daniel R, Herbert M Simpson, and Anita Pak. 2003. "Changes in collision rates among novice drivers during the first months of driving." *Accident Analysis & Prevention*, 35(5): 683–691.

McCartt, Anne T, Veronika I Shabanova, and William A Leaf. 2003. "Driving experience, crashes and traffic citations of teenage beginning drivers." *Accident Analysis & Prevention*, 35(3): 311–320.

Naik, Narayan Y, and Michael J Moore. 1996. "Habit formation and intertemporal substitution in individual food consumption." *Review of Economics and Statistics*, 321–328.

National Highway Traffic Safety Administration. 1998. "Saving teenage lives: The case for graduated driver licensing." Washington DC: Department of Transportation.

Naude, Clifford, Tariro Makwasha, and Ed McGeehan. 2015. "Methodologies for updating parameter values for economic evaluation of road transport projects in Australia."

Royer, Heather, Mark Stehr, and Justin Sydnor. 2015. "Incentives, commitments, and habit formation in exercise: Evidence from a field experiment with workers at a Fortune-500 company." *American Economic Journal: Applied Economics*, 7(3): 51–84.

Shachar, Ron. 2003. "Party loyalty as habit formation." *Journal of Applied Econometrics*, 18(3): 251–269.

Shope, Jean T, Lisa J Molnar, Michael R Elliott, and Patricia F Waller. 2001. "Graduated driver licensing in Michigan: Early impact on motor vehicle crashes among 16-year-old drivers." *Journal of the American Medical Association*, 286(13): 1593–1598.

Tefft, Brian C, Allan F Williams, and Jurek G Grabowski. 2013. "Teen driver risk in relation to age and number of passengers, United States, 2007–2010." *Traffic Injury Prevention*, 14(3): 283–292.

Webb, Caitlin. 2016. "Motor vehicle traffic crashes as a leading cause of death in the United States, 2012–2014." Washington, DC: National Highway Traffic Safety Administration (No. DOT HS 812 297).

Williams, Jenny. 2005. "Habit formation and college students' demand for alcohol." *Health Economics*, 14(2): 119–134.

World Health Organisation. 2018. "Global health estimates 2016: Deaths by cause, age, sex, by country and by region, 2000–2016." http://www.who.int/healthinfo/global_burden_disease/estimates/en/ (accessed on 26/05/2018).

Figure 1: Traffic fatality rates in New South Wales and the United States (1970–2018)

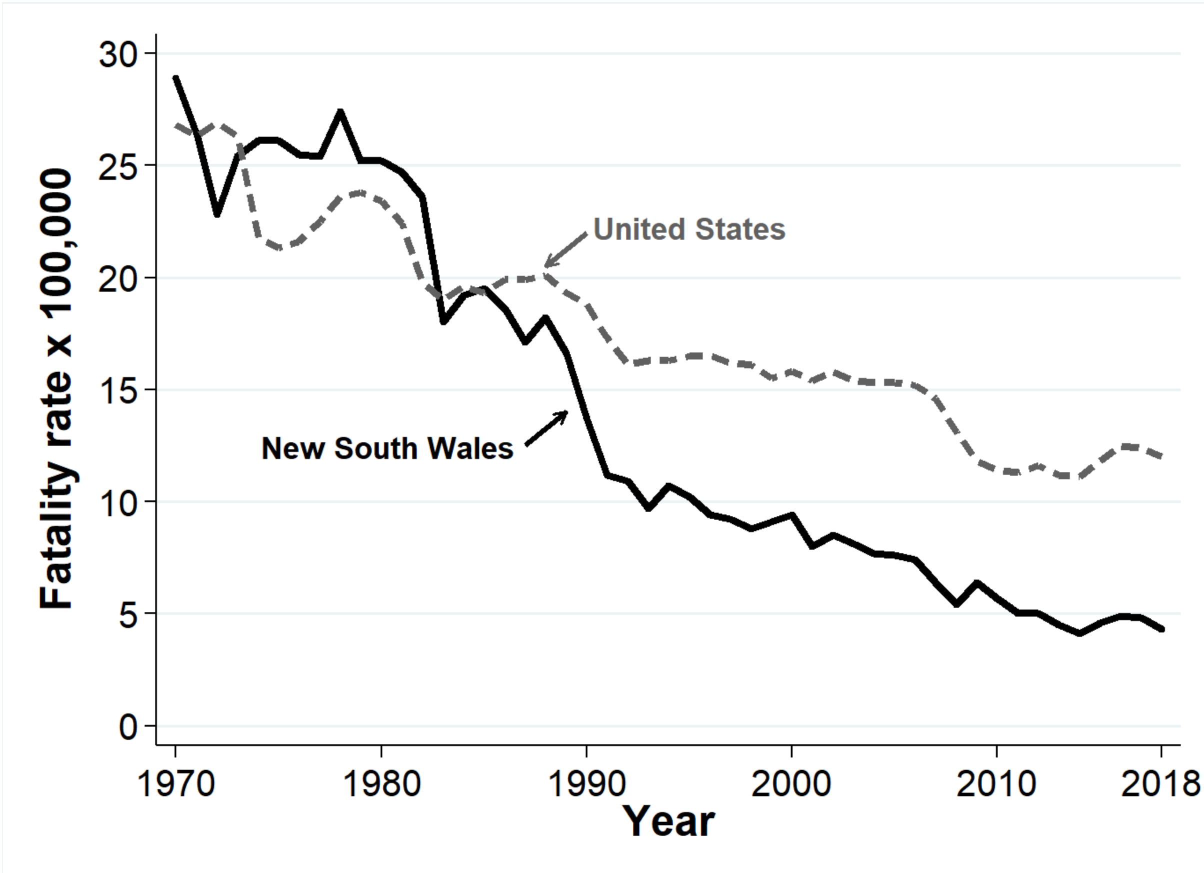

Notes: This figure shows the trends in traffic fatality rates (per 100,000 population) in New South Wales and United States since 1970. Source: New South Wales Centre for Road Safety and National Highway Traffic Safety Administration.

Figure 2: Pre-reform crash rates by years since the start of the probation period

(a) All crashes

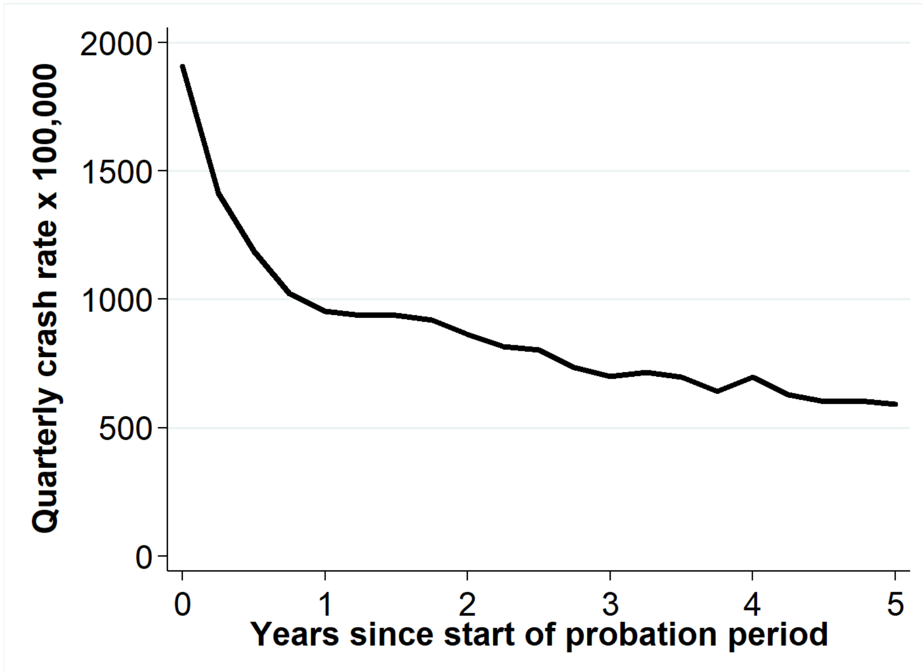

(b) Late-night multi-passenger crashes

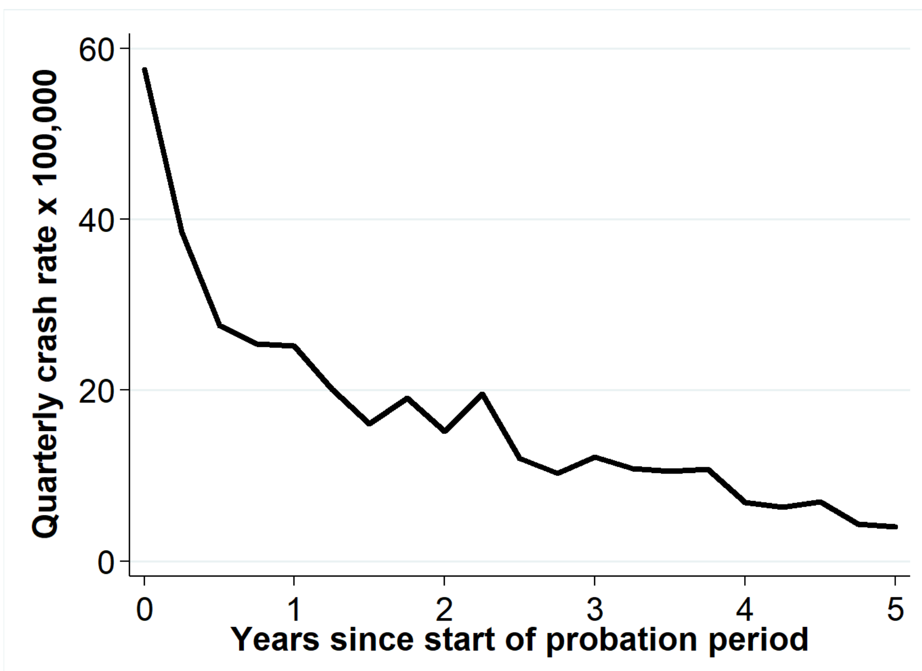

Notes: These figures show crash rates before the reform by the number of years since the start of the probation period. The sample comes from linked administrative driver-license and crash records from June 2004 to June 2007. The sample includes all drivers who obtained a P1 license (i) before the age of 25 and (ii) within the last five years.

Figure 3: Crash rate trends of probationary drivers by no. of passengers and time of day

(a) Carrying two or more passengers

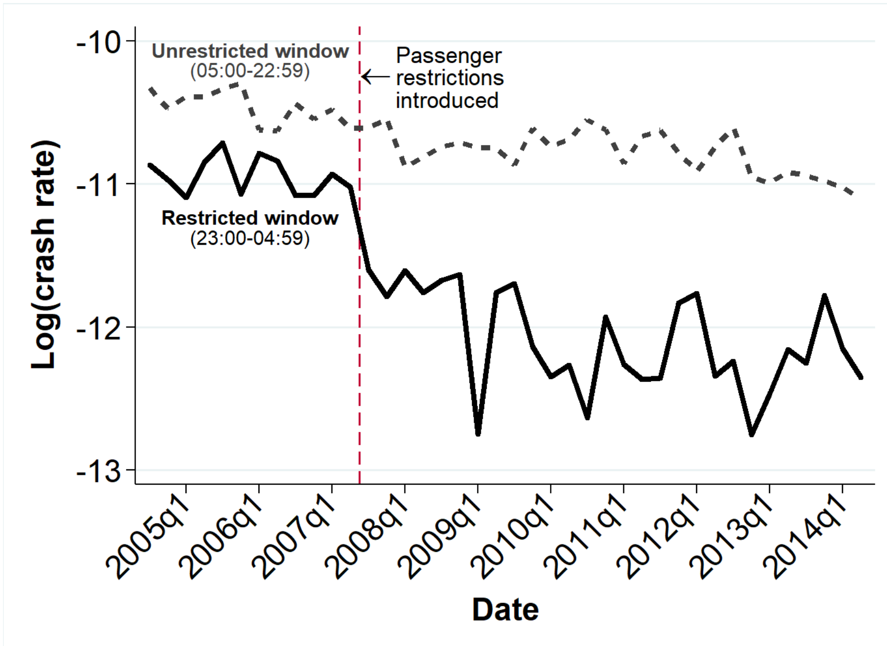

(b) Carrying zero or one passenger

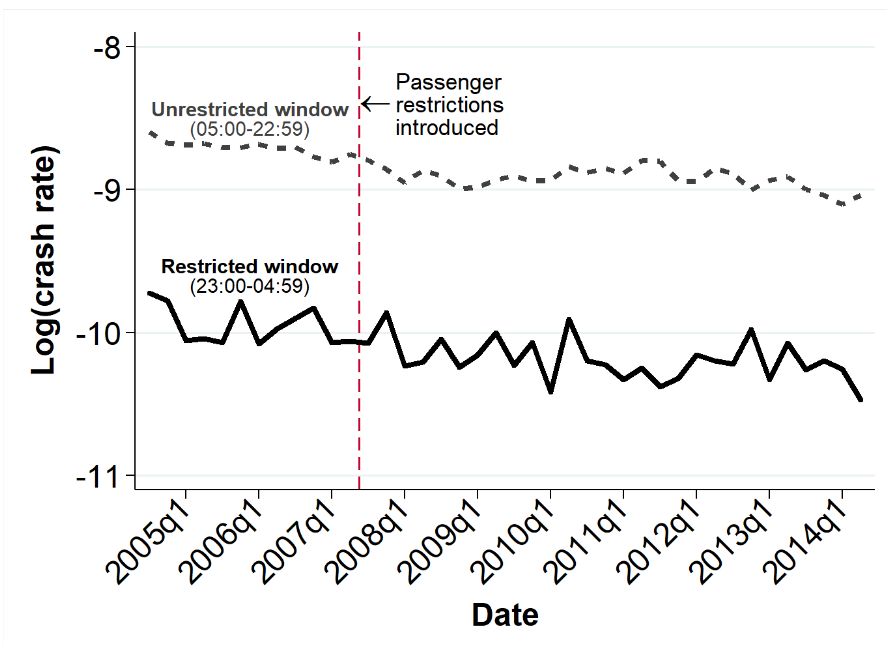

Notes: These figures show the natural log of the average crash rates of P1 drivers under 25 in each quarter-year period. We divide crash rates by the number of hours in the relevant window (18 or 6) before taking logs to make it easier to compare crash trends in each window. The figures show (i) similar trends in crash rates during the restricted and unrestricted windows before the reform; (ii) an immediate and persistent decrease in late-night multi-passenger crash rates after the passenger restriction was introduced; and (iii) little change in other crash rates. The sample comes from linked administrative driver-license and crash records from 2004 Q3 to 2014 Q3.

Figure 4: Average crash rates of probationary drivers by hour before and after reform

(a) Carrying two or more passengers

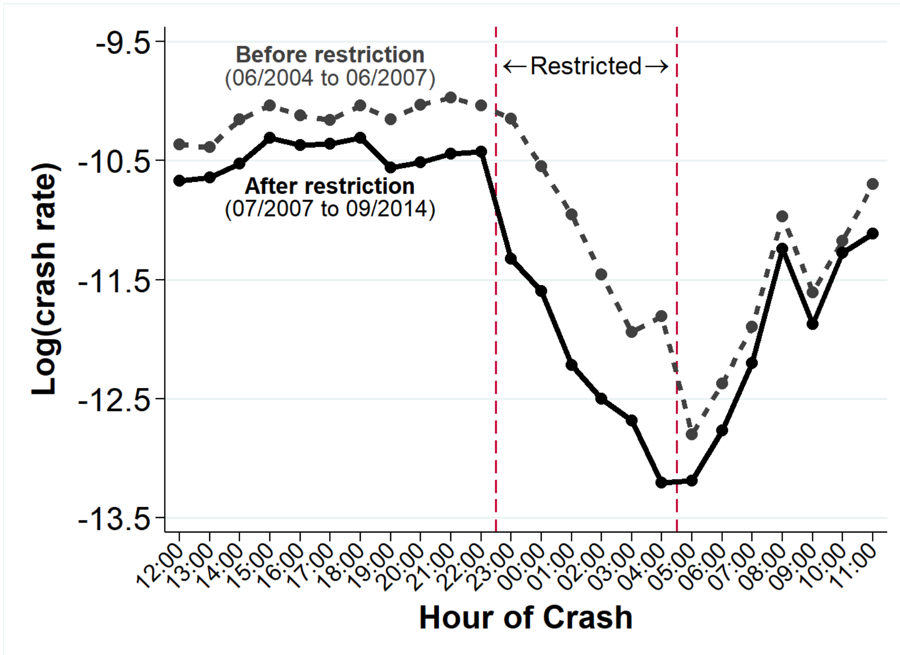

(b) Carrying zero or one passenger

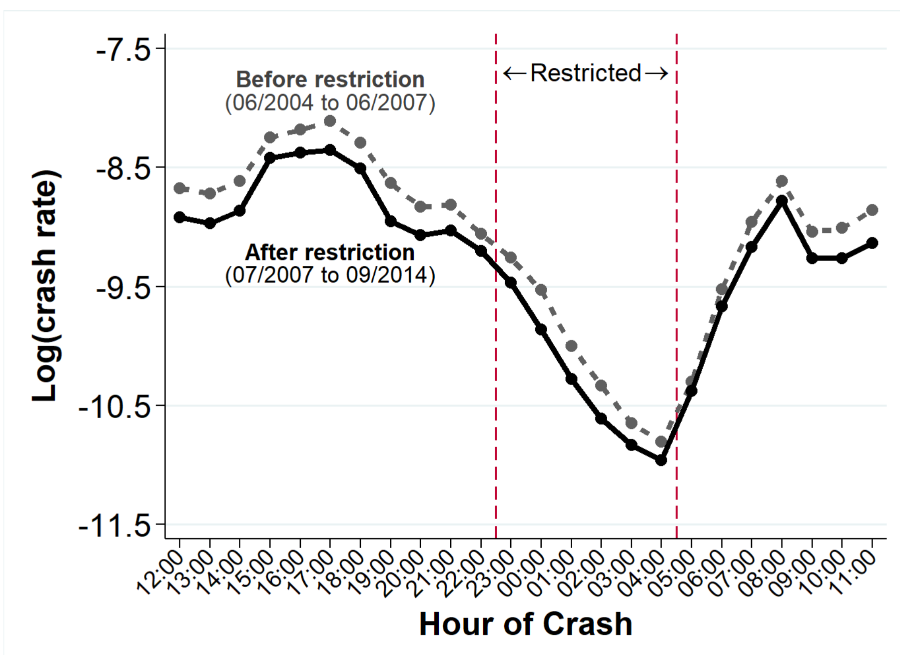

Notes: These figures show the natural log of the average monthly crash rates of P1 drivers under 25 in each one-hour window before and after the reform. The figures show a decrease in crash rates in every hour after the passenger restriction was introduced, but the decrease is particularly large for multi-passenger crash rates during the restricted window. There is also no sign of an increase in multi-passenger crash rates just outside the restricted window after the restriction was introduced. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. See the text in Section 3.3 for more details.

Figure 5: Estimated percentage change in the late-night crash rates of probationary drivers compared to January–June 2007, the six months before the reform

(a) Carrying two or more passengers

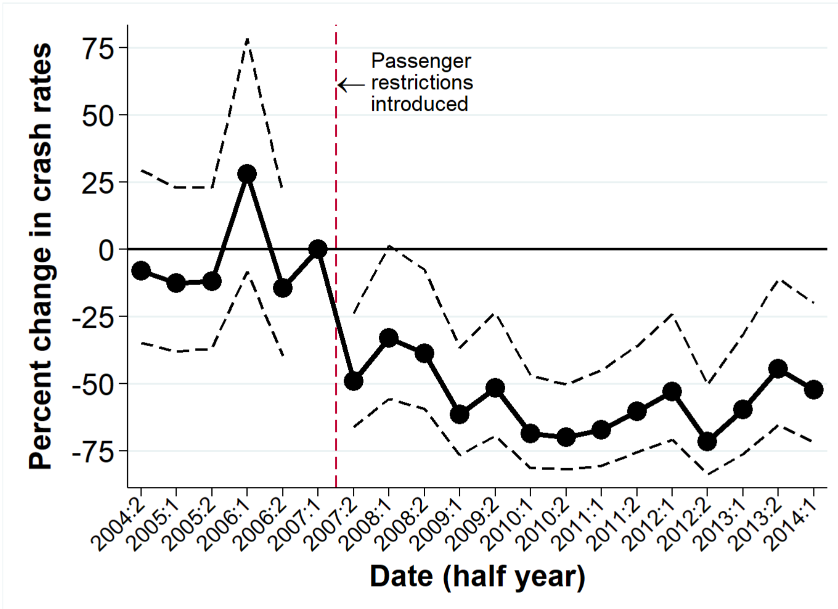

(b) Carrying zero or one passenger

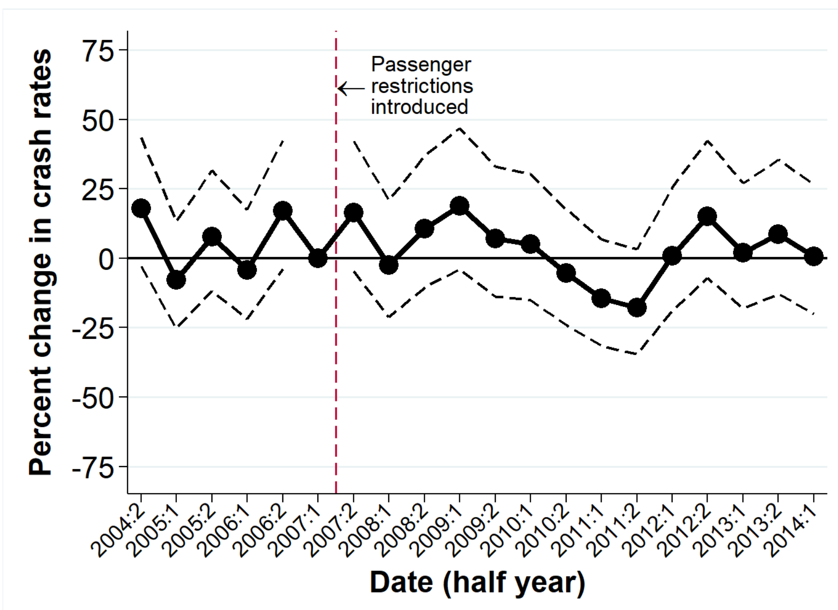

Notes: These figures show the estimates and 95% confidence intervals of the $\beta_{DID,j}$ terms in equation (2). The estimates show the change in the late-night crash rates of P1 drivers under 25 in each six-month period compared to January–June 2007, the six months before the reform (using crash rates during the unrestricted window as the comparison group). The estimates show (i) support for the parallel-trends assumption, with all of the estimates close to zero prior to the reform, and (ii) a large, immediate and persistent decrease in late-night multi-passenger crash rates after the passenger restriction was introduced (but no change for 0–1-passenger crash rates). The sample comes from linked administrative driver-license and crash records from July 2004 to June 2014. For more details, see the text under "Robustness" in Section 4.1.

Figure 6: Estimated percentage change in the hourly crash rates of probationary drivers after the reform compared to the change at 12:00–12:59

(a) Carrying two or more passengers

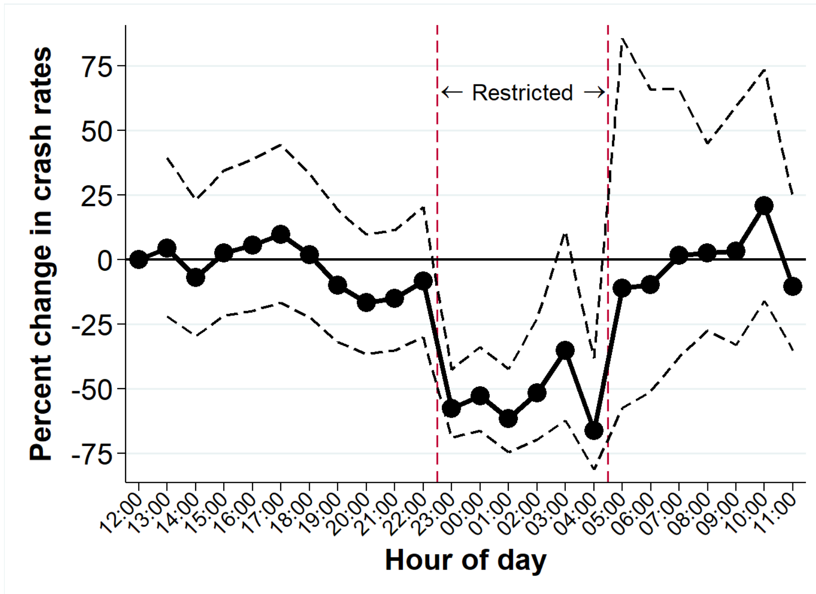

(b) Carrying zero or one passenger

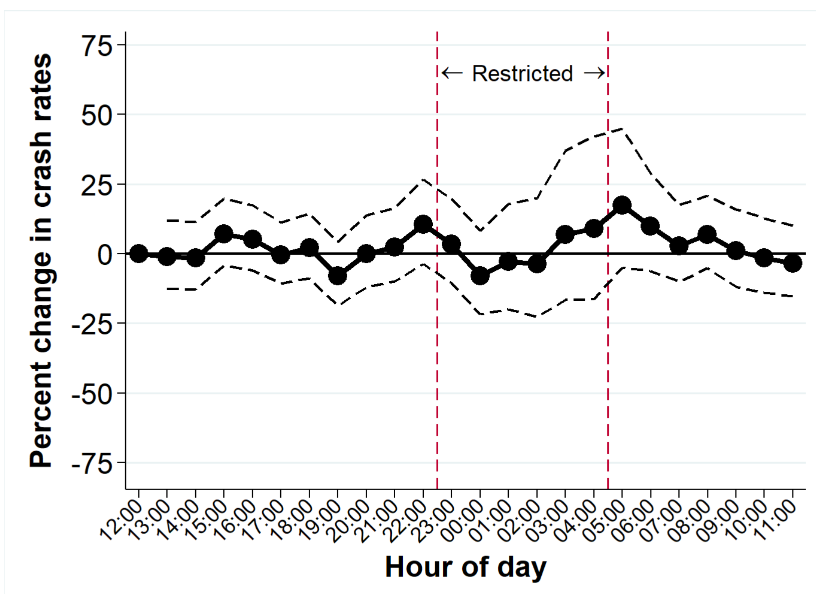

Notes: These figures show the estimates and 95% confidence intervals of the $\beta_{DID,k}$ terms in equation (3). The estimates show the change in the crash rates of P1 drivers under 25 in each one-hour window after the reform relative to the change at 12:00–12:59. To aid interpretation, we translate the estimates into the implied percentage change in crashes. The estimates show (i) little change in crash rates in any of the unrestricted hours compared to the change at 12:00–12:59 and (ii) a large decrease in multi-passenger crash rates during restricted hours (but no change in 0–1-passenger crash rates). The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see the text under "Robustness" in Section 4.1.

Figure 7: Estimated long-term impact on late-night crash rates by years exposed to restriction

(a) Carrying two or more passengers

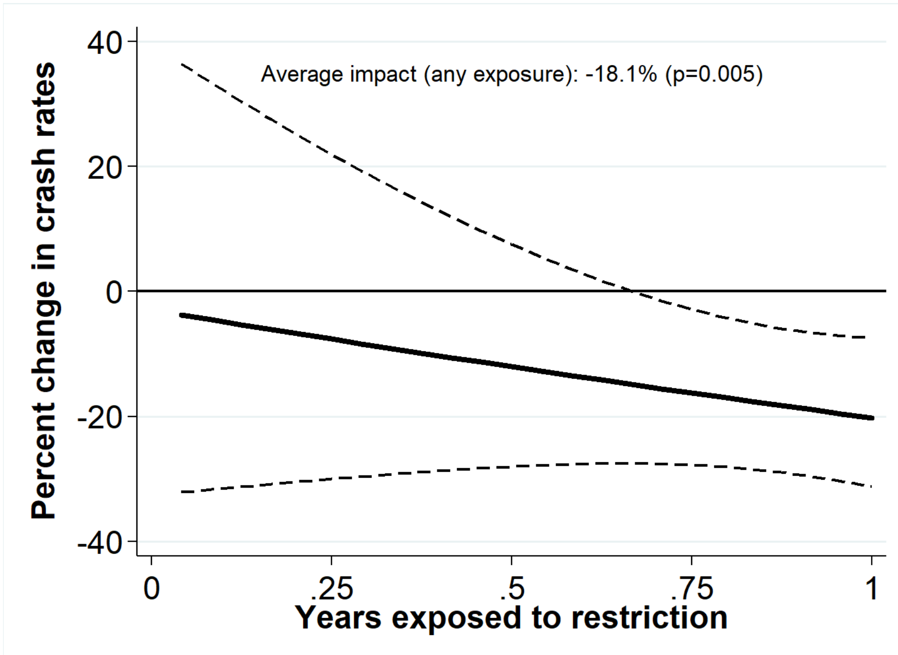

(b) Carrying zero or one passenger

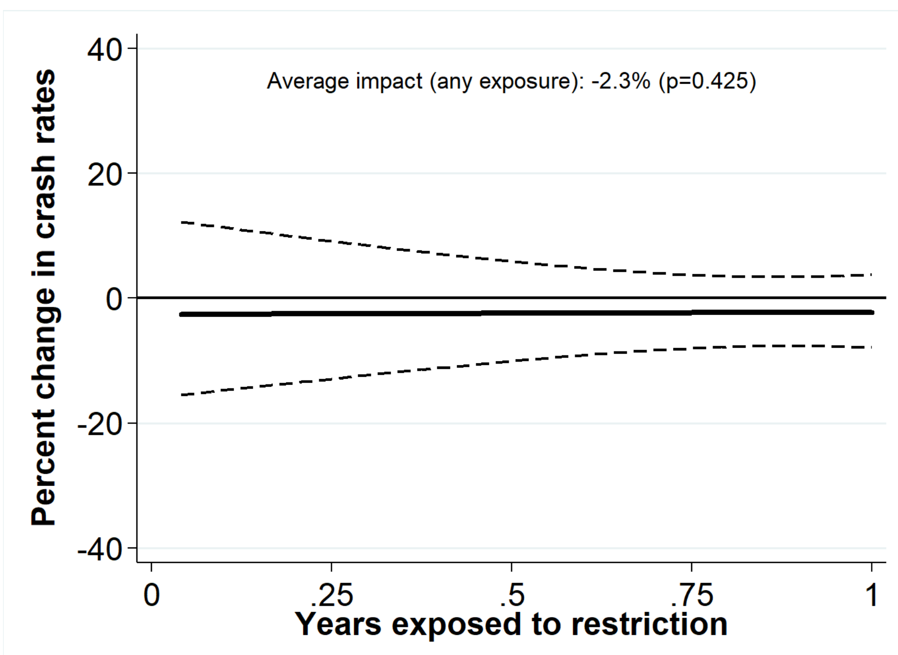

Notes: This figure shows the estimated impact, with 95% confidence intervals, of the long-term impact of the restriction on late-night multi-passenger crash rates by the number of years drivers were exposed to the restriction. The estimates, which come from equation (5), show a larger long-term impact on the crash rates of drivers who were subject to the restriction for longer. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see Section 4.4.

Figure 8: Estimated long-term impact on late-night multi-passenger crash rates

(a) Logarithmic decay

(b) One-year bins (last bin is two years)

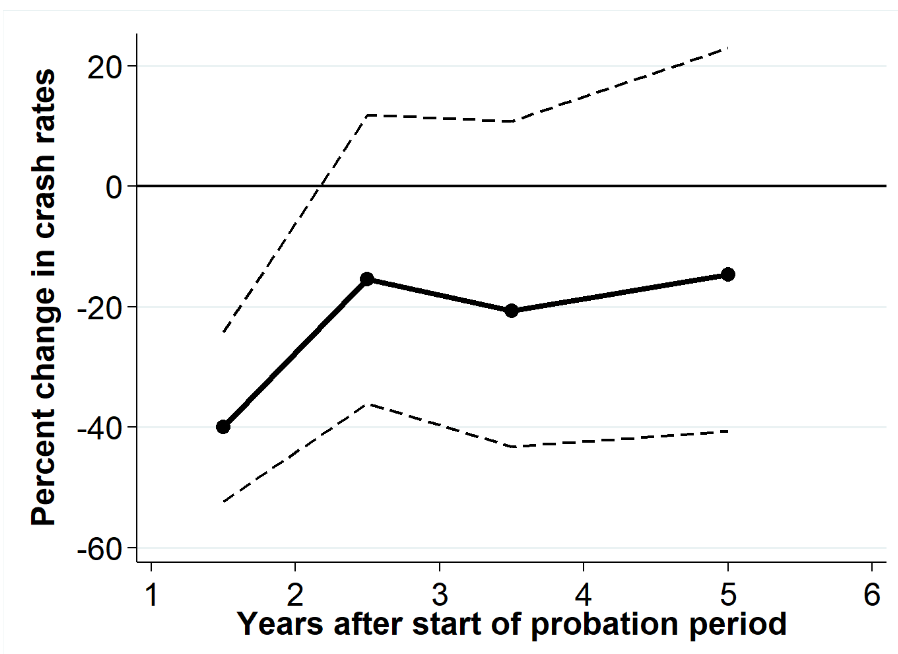

Notes: This figure shows the estimated impact, with 95% confidence intervals, of the restriction on the late-night multi-passenger crash rates of drivers in the long-term. The estimates, which come from equation (6), are translated into the implied percentage change in crash rates. The estimates show a persistent but smaller effect on the crash rates of previously restricted drivers, which diminishes over time. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see Section 4.4.

Figure 9: Change in late-night multi-passenger crash rates by years of driving experience

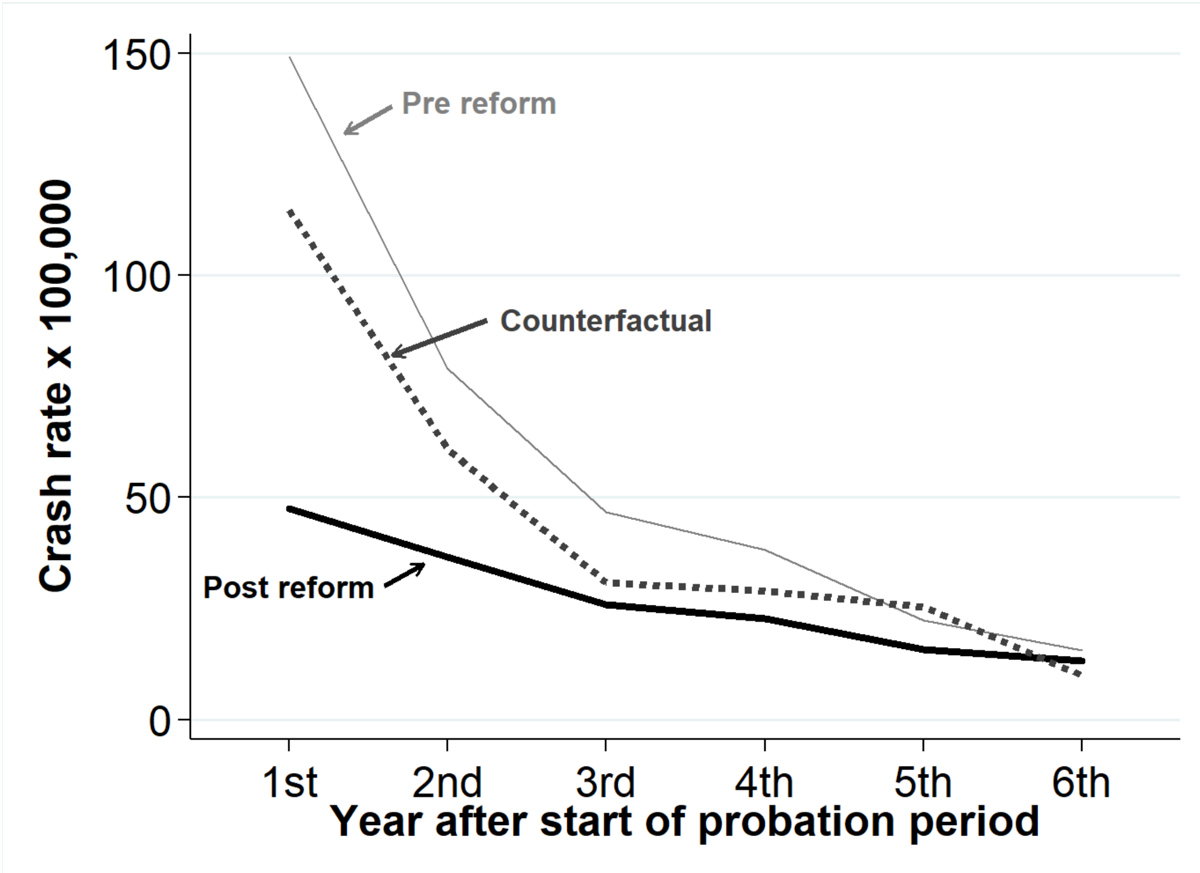

Notes: These figures show the annual multi-passenger crash rates of two groups: (i) drivers who obtained their P1 license before June 30, 2006 ("pre reform") and (ii) drivers who obtained their P1 license after June 30, 2007 ("post reform"). We also plot the estimated counterfactual for post-reform drivers using non-parametric estimates of equation (6), in which we estimate treatment effects for each of the six years above. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. The sample includes drivers who obtained a P1 license (i) before the age of 25 and (ii) within the last six years. Drivers who obtained their P1 license between July 2006 and June 2007 are excluded as they became subject to the restriction after the start of their probation period.

Figure 10: Correlation between estimated short- and long-term change in late-night multi-passenger crash rates for different crash types

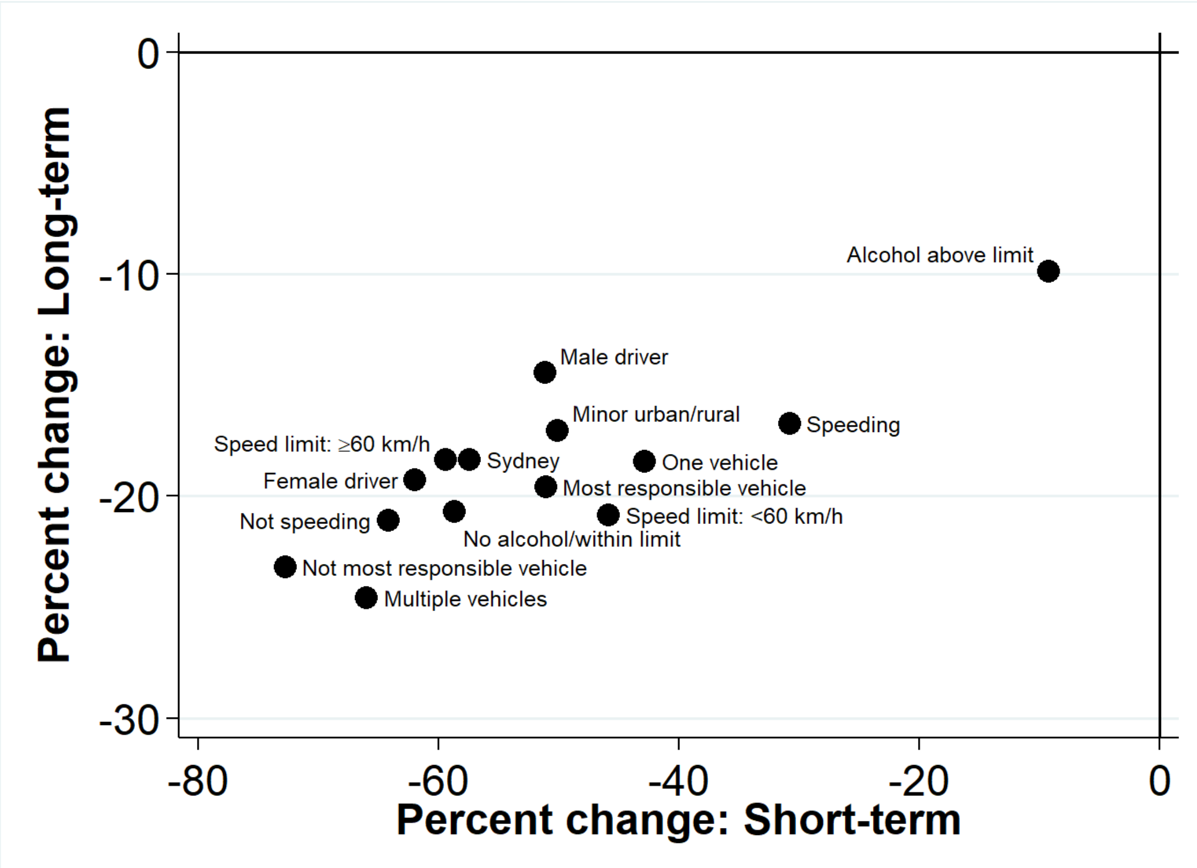

Notes: This figure shows the correlation between the estimates from equation (1) and equation (4) for 14 crash types. To aid interpretation, we translate the estimates into the implied percentage change in crash rates. The figure shows a positive correlation between the types of crashes most affected by the restriction in the short and long term. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see Section 4.4.

Table 1: Descriptive statistics of crashes involving probationary drivers

	Total crashes (1)	Casualty crashes (2)	No. of casualties (3)	Fatal crashes (4)	No. of fatalities (5)	Fatalities per 1,000 crashes (6)
All crashes						
Annual mean	4,776	2,093	2,923	23.1	26.4	5.5
(Std. dev.)	(892)	(366)	(561)	(17.7)	(21.2)	
Annual mean per 1,000 drivers	39.6	17.3	24.2	0.19	0.22	
Annual means by crash type						
Unrestricted window: 5:00 am to 10:59 pm	4,339	1,914	2,662	18.9	21.1	4.9
Restricted window: 11:00 pm to 4:59 am	437	179	262	4.3	5.3	12.1
Zero or one passenger	4,090	1,776	2,323	17.5	19.3	4.7
Multiple passengers	687	316	600	5.6	7.2	10.5
Restricted window & multiple passengers	83	38	83	1.5	2.2	**26.5**

Notes: This table presents descriptive statistics of the average number of crashes per quarter involving P1 drivers under 25. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. See the text in Section 3.1 for more details.

Table 2: Baseline estimates: The causal effect of the passenger restriction on the late-night crash rates of restricted drivers

	Main estimates (Carrying 2+ passengers)							Placebo (0–1 passenger)
	All crashes (1)	Casualty crashes (2)	No. of casualties (3)	Hospital admission/ fatality (4)	No. admitted/ killed (5)	Fatal crashes (6)	No. killed (7)	All crashes (8)
Effect of restriction	-0.784***	-0.748***	-0.658***	-0.573***	-0.605***	-1.194*	-1.399**	-0.023
	(0.074)	(0.109)	(0.142)	(0.180)	(0.217)	(0.611)	(0.701)	(0.036)
Implied percentage reduction	54.4%	52.7%	48.2%	43.6%	45.4%	69.7%	75.3%	2.3%
Implied annual reduction	59.5	27.0	48.0	9.1	15.5	1.9	3.4	7.5

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents Poisson regression estimates from equation (1) of the causal effects of the passenger restriction on the multi-passenger crash rates of P1 drivers under 25 during the restricted window (11:00 pm to 4:59 am). Column 8 presents placebo estimates where equation (1) is estimated on the sample of crashes where P1 drivers under 25 are carrying zero or one passenger. All regressions include individual-level controls for driver age and gender, and their interactions with an indicator for the restricted time-window, and fixed effects for (i) calender month by time window and (ii) calender year. The number of observations in each regression is 29,830,442. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. See the text in Section 4.1 for more details.

Table 3: Heterogeneity in the effects of the passenger restriction on crash rates

	Coefficient estimate (1)	Pre-reform annual mean (2)	Implied annual reduction (3)	p-value on test of equal treatment effects (4)
Driver gender				
Male	-0.716***	115.5	39.3	
	(0.087)			
Female	-0.966***	45.7	20.2	0.133
	(0.142)			
Driver age				
17–18	-0.884***	119.4	46.8	
	(0.089)			
19–20	-0.666***	31.1	9.7	0.252
	(0.167)			
Number of vehicles				
One	-0.559***	79.8	23.3	
	(0.105)			
Multiple	-1.078***	81.4	36.4	$<$0.001
	(0.111)			
Crash location				
Sydney	-0.853***	93.7	34.2	
	(0.101)			
Minor urban/rural	-0.694***	67.5	24.6	0.287
	(0.109)			
Day of crash				
Friday & Saturday nights	-0.787***	104.8	37.3	
	(0.097)			
Weeknights	-0.743***	56.4	20.7	0.775
	(0.119)			
Road type				
Residential: Speed limit $<$ 60 km/h	-0.613***	56.8	19.6	
	(0.116)			
Main: Speed limit $\geq$ 60 km/h	-0.901***	103.8	39.1	0.058
	(0.097)			

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents Poisson regression estimates from equation (1). The estimates show heterogeneity in the effects of the passenger restriction on different types of crash rates during the restricted window (11:00 pm to 4:59 am) involving P1 drivers under 25 with multiple passengers. All regressions include individual-level controls for driver age and gender, and their interactions with an indicator for the restricted time-window, and fixed effects for (i) calender month by time window and (ii) calender year. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see the text under "Heterogeneity" in Section 4.1.

Table 4: The estimated effect of the passenger restriction on casualties for different road users

		Restricted driver's car			
	All involved (1)	All in car (2)	Driver (3)	Passengers (4)	Other road users (5)
Effect of restriction	-0.658***	-0.534***	-0.607***	-0.469**	-1.267***
	(0.142)	(0.156)	(0.143)	(0.189)	(0.264)
Implied percentage reduction	48.2%	41.4%	45.5%	37.4%	71.8%
Implied annual reduction	48.0	32.6	13.7	17.8	13.7

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents Poisson regression estimates from equation (1) of the causal effects of the passenger restriction on the number of casualties — by the type of road user — from crashes during the restricted window (11:00 pm to 4:59 am) involving P1 drivers under 25 with multiple passengers. All regressions include individual-level controls for driver age and gender, and their interactions with an indicator for the restricted time-window, and fixed effects for (i) calender month by time window and (ii) calender year. The number of observations in each regression is 29,830,442. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see Section 4.2.

Table 5: Estimated long-term impact on late-night crash rates

	Main estimate (2+ passengers)	Placebo (0–1 passenger)
	All crashes (1)	All crashes (2)
Previously restricted	-0.200***	-0.023
	(0.072)	(0.029)
Implied percentage reduction	18.1%	2.3%
Implied annual reduction	20.1	13.5

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents Poisson regression estimates of the causal effects of the passenger restriction on the late-night crash rates of previously restricted drivers (11:00 pm to 4:59 am) from equation (4). All regressions include individual-level controls for driver age and gender, and their interactions with an indicator for the restricted time-window, and fixed effects for (i) calender month by time window and (ii) calender year. The number of observations in each regression is 76,366,594. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. See the text in Section 4.3 for more details.

Table 6: Using effect heterogeneity to examine the key mechanism: A reduction in driving or a reduction in crash risk?

	Coefficient estimate (1)	Pre-reform annual mean (2)	Implied annual reduction (3)	p-value on test of equal treatment effects (4)
Panel A: Short-term effects on restricted drivers				
Caused crash				
Other driver	-1.300***	28.5	13.3	
	(0.205)			
Restricted driver	-0.715***	132.6	47.0	0.008
	(0.080)			
Speeding				
Restricted driver not speeding	-1.026***	107.7	48.2	
	(0.095)			
Restricted driver speeding	-0.368***	53.5	10.3	<0.001
	(0.124)			
Alcohol				
Restricted driver not drinking (BAC=0)	-0.883***	144.6	57.7	
	(0.080)			
Restricted driver over legal limit (BAC>0)	-0.097	16.5	0.9	0.006
	(0.274)			
Panel B: Long-term effects on previously restricted drivers				
Caused crash				
Other driver	-0.264*	20.4	3.5	
	(0.154)			
Previously restricted driver	-0.218***	83.0	12.3	0.794
	(0.082)			
Speeding				
Previously restricted driver not speeding	-0.237***	70.4	11.0	
	(0.085)			
Previously restricted driver speeding	-0.183	33.1	4.2	0.743
	(0.140)			
Alcohol				
Previously restricted within legal limit	-0.232***	87.6	13.5	
	(0.077)			
Previously restricted driver over legal limit	-0.104	15.9	1.1	0.622
	(0.248)			

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents Poisson regression estimates from equation (1) in Panel A and equation (4) in Panel B. The estimates show heterogeneity in the effects of the passenger restriction on different types of late-night multi-passenger crash rates involving drivers subject to the restriction in Panel A and previously subject to the restriction in Panel B. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. See the text in Section 4.4 for more details.

For online publication

Web Appendix for "Shaping the Habits of Young Drivers: The Persistent Effect of Passenger Restrictions"

Timothy J. Moore

Purdue University and NBER

Todd Morris

Max Planck Institute for Social Law and Social Policy

List of Figures

List of Tables

Figure A1: Estimated change in late-night crash rates compared to January–June 2007, the six months before the reform: Drivers 0–5 years after P1 period

(a) Carrying two or more passengers

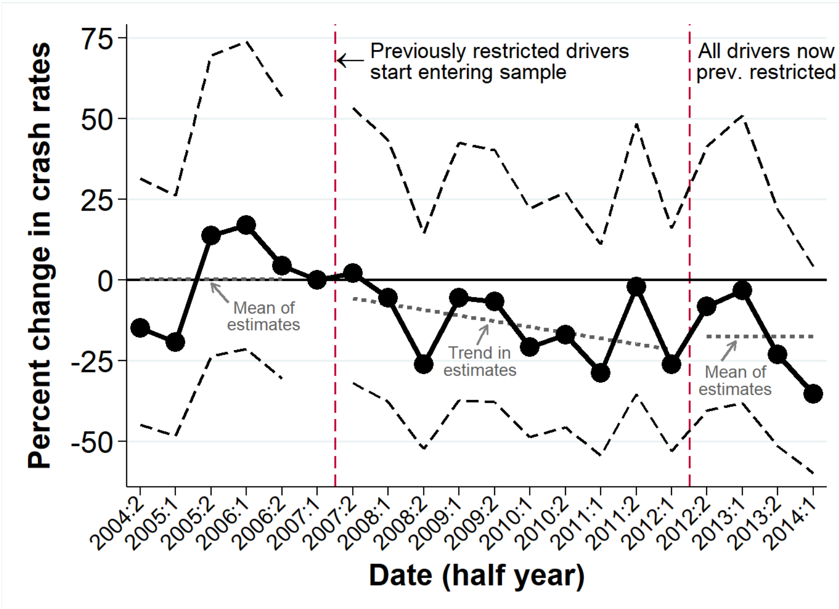

(b) Carrying zero or one passenger

Notes: These figures show the estimates and 95% confidence intervals of the $\beta_{DID,j}$ terms in equation (2) for drivers who graduated from a P1 license in the last five years. The estimates show the change in the late-night crash rates of these drivers in each six-month period compared to January–June 2007, the six months before the reform (using crash rates during the unrestricted window as the comparison group). The estimates show (i) support for the parallel-trends assumption, with all of the estimates hovering around zero prior to the reform, and (ii) a gradual decline in late-night multi-passenger crash rates after the passenger restriction was introduced and previously restricted drivers flow into the sample. The sample comes from linked administrative driver-license and crash records from July 2004 to June 2014. For more details, see the text under "Robustness" in Section 4.3.

Figure A2: Estimated long-term impact on late-night crash rates by years exposed to restriction (using P1 start dates to calculate exposure)

(a) Carrying two or more passengers

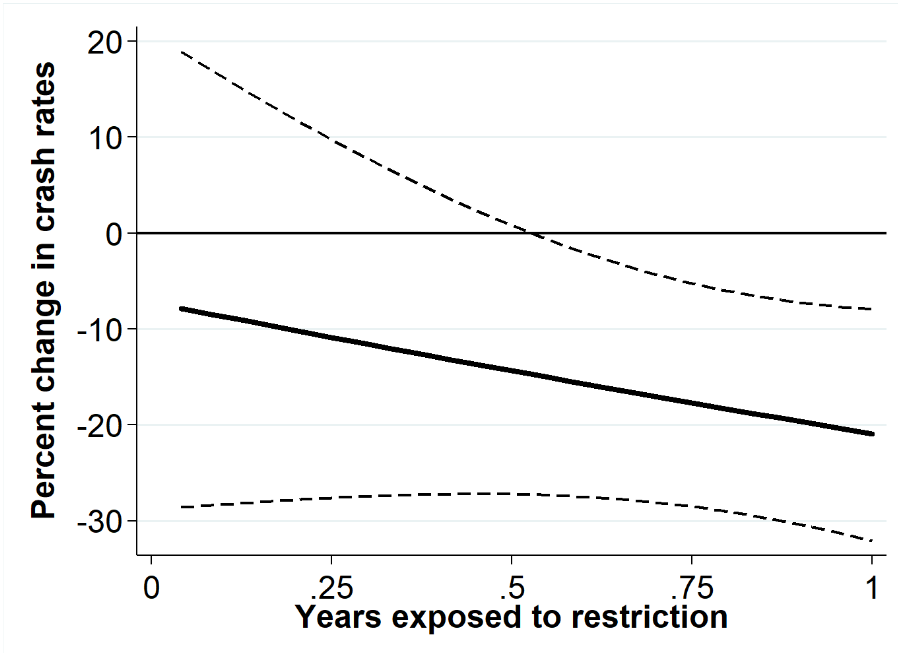

(b) Carrying zero or one passenger

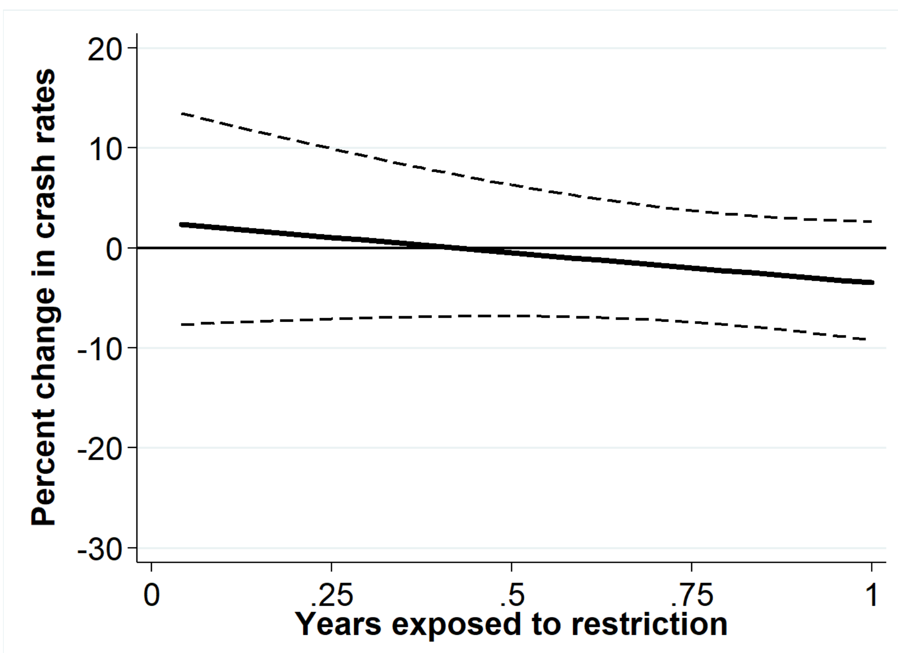

Notes: This figure shows the estimated impact, with 95% confidence intervals, of the long-term impact of the restriction on late-night multi-passenger crash rates by the number of years drivers were exposed to the restriction. The estimates, which come from equation (5), show a larger impact on the crash rates of previously restricted drivers who were subject to the restriction for longer. For more details, see Section 4.4.

Table A1: Comparison of Poisson estimates to Negative Binomial, log-linear and inverse-hyperbolic-sine estimates

	Level of aggregation		
	Aggregate half year (N=40)	Aggregate month (N=240)	Individual month (N=28,905,680)
Poisson			
Effect of restriction	-0.779***	-0.779***	-0.779***
	(0.080)	(0.073)	(0.075)
Negative Binomial			
Effect of restriction	-0.779***	-0.779***	-0.779***
	(0.080)	(0.073)	(0.075)
Log (crashes + 1)			
Effect of restriction	-0.783***	-0.713***	0.00005***
	(0.086)	(0.073)	(0.00001)
Inverse hyperbolic sine (crashes)			
Effect of restriction	-0.811***	-0.878***	0.00006***
	(0.088)	(0.086)	(0.00001)

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: This table shows how the baseline Poisson regression estimates from equation (1) compare to Negative Binomial estimates and OLS estimates when crashes are transformed using a logarithmic or inverse-hyperbolic-sine (IHS) transformation. We take the log of the number of crashes plus one to avoid zero counts. The four approaches show near-identical estimates when crashes are aggregated to the half-year–year–time-window level. However, unlike the Poisson and Negative Binomial estimates, the log + 1 and IHS estimates are sensitive to the level of aggregation, as both functions are non-logarithmic for counts close to zero. The sample comes from linked administrative driver-license and crash records from July 2004 to June 2014. (We exclude crashes in June 2004 and July–September 2014 so that the sample is consistent across columns). See the text in Section 3.2 for more details.

Table A2: Robustness: Estimates by age of driver and number of passengers

	Coefficient estimate (1)	Pre-reform annual mean (2)	Implied annual reduction (3)	p-value on test of equal treatment effects (4)
Driver age				
17–20	-0.835***	150.5	56.3	
	(0.078)			
21–24	-0.274	10.7	2.1	0.027
	(0.240)			
Number of passengers				
Two	-0.553***	74.6	21.7	
	(0.101)			
Three or more	-1.038***	86.6	37.5	0.001
	(0.110)			

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.
Notes: This table presents regression estimates of β_{DID} from equation (1). The estimates show that the passenger restriction has caused larger reductions in late-night multi-passenger crash rates in situations where drivers were more likely to be constrained by the restriction. For more details, see the text under "Robustness" in Section 4.1.

Table A3: Robustness of estimates to selective non-reporting

	All crashes (1)	2+ vehicle crashes (2)	2+ vehicles & P1 u25 caused crash (3)
Effect of restriction	-0.784***	-1.078***	-0.985***
	(0.074)	(0.111)	(0.132)
Implied percentage reduction	54.4%	66.0%	62.7%

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents estimates of equation (1). The estimates examine the robustness of the baseline results to selective non-reporting. Compared to the baseline estimates, the estimates are larger for crashes involving multiple vehicles and crashes involving multiple vehicles that were caused by a P1 driver under 25, which suggests that selective non-reporting is unlikely to significantly affect the results. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see the text under "Robustness" in Section 4.1.

Table A4: Placebo effect of restriction on drivers aged 25 and over

	All crashes (1)
Placebo effect of restriction	0.0001
	(0.0220)
Observations	248

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Notes: This table presents estimates of equation (1) for the sample of drivers aged 25 and over. We include all crashes, regardless of the number of passengers, and aggregate crashes to the month–year–time-window level, as for this population we only have license information on drivers who were involved in crashes. The estimate shows that, in percentage terms, there is no disproportionate change in the number of late-night crashes involving drivers over 25 relative to the change in crashes at other times of the day. The sample comes from administrative crash records over the period from June 2004 to September 2014. See the text in Section 4.4 for more details.

Table A5: Robustness: The estimated effect of the passenger restriction on P1 passenger casualties by passenger age

	Passengers under 21 (1)	Passengers 21 or older (2)
Effect of restriction	-0.580***	0.087
	(0.206)	(0.343)
Implied percentage reduction	44.0%	-9.1%
Implied annual reduction	17.9	-0.6

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver.

Notes: This table presents Poisson regression estimates from equation (1) of the causal effects of the passenger restriction on the number of casualties among P1 passengers from crashes during the restricted window (11:00 pm to 4:59 am) involving P1 drivers under 25 with multiple passengers. All regressions include individual-level controls for driver age and gender, and their interactions with an indicator for the restricted time-window, and fixed effects for (i) calender month by time window and (ii) calender year. The number of observations in each regression is 29,830,442. The estimate shows that, in terms of casualties, the impact on P1 passengers appears to be confined to those under the age of 21. This is consistent with the rules of the policy, which only limits the number of passengers below this age. The sample comes from linked administrative driver-license and crash records from June 2004 to September 2014. For more details, see Section 4.2.

Table A6: Spillovers from behavioral responses to the passenger restriction on crash and casualty rates from other types of crashes

	Crashes			Casualties	
	0–1 passenger	1 passenger		No P1 driver u25 involved	
	P1 driver aged 17–20	P1 driver aged 17–24	P1 driver aged 17–20	Passengers aged 16–20	Pedestrians aged 16–20
Effect of restriction	-0.025	-0.056	-0.037	-0.108	-0.045
	(0.039)	(0.061)	(0.064)	(0.076)	(0.103)
Observations	23,789,720	29,830,442	23,789,720	248	248

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are clustered by driver in columns 1–3. Robust standard errors in parentheses in columns 4 and 5.

Notes: This table presents estimates of β_{DID} from equation (1). The estimates examine whether there was any increase in other types of crashes resulting from behavioral responses by P1 drivers under 25 and their passengers. In columns 4 and 5, casualties are aggregated to the month–year–time-window level. The sample comes from linked administrative driver-license and crash records in columns 1–3, and administrative crash records in columns 4 and 5, and covers the period from June 2004 to September 2014. For more details, see the text under "Indirect effects from behavioral responses" in Section 4.2.