The Increased Risk of High Drivers

By Daniel Palau and Mitchell VanVuren*

We build and estimate a multinomial model of fatal car accidents involving two drivers using data from the Fatal Accident Reporting System. We estimate that high drivers are 2.83 times more likely to cause a fatal accident than sober drivers. We also estimate that, on average, 1.6% of nighttime drivers have cannabis in their system. These estimates differ substantially from recent estimates based on roadside surveys. Our model also estimates the relative risk of drunk drivers as 5.94, a value consistent with previous literature.

As of 2016, twenty-four states and the District of Columbia have laws legalizing the consumption of marijuana in some form, medical or recreational. The recreational legislation movement is more recent and relevant, as adults 21 and over can purchase, possess, and grow marijuana for personal use in four states and Washington D.C. with even more states considering future legalization. The current political environment and trends indicate that an increasing amount states will legalize in the future. Naturally, increased access to cannabis will lead to an increase of Americans consuming marijuana and driving. As the number of high drivers on the road looks to go up, it is increasingly important to explore the ensuing risks and consequences of high driving for public knowledge and to inform future legislation and regulation.

While marijuana and cannabinoid products are possible sources of tax revenue for affected states, there are significant negative effects involved. While some argue that drivers under the influence of marijuana compensate well for impairment on the road by decreasing speed and following distance, others believe that the inherent risks present while operating a vehicle on a mind-altering drug are extremely dangerous. Marijuana legislation currently in place clearly assumes that driving under the influence is dangerous to the driver and others on the road, setting legal limits of THC in the bloodstream to warrant arrest. No matter the set legal limits, further study must be done to fully understand the nuances of marijuana's varying effects on different users. Finally, marijuana and driving has rapidly entered the public eye, as prominent and political organizations like Mothers Against Drunk Driving (MADD) begin to take policy stands. For example, in 2015 MADD changed its mission to include that it, "helps [to] fight drugged driving," stating that while alcohol and drugs like cannabis, "are different, the results are the same - needless death and injuries." (MADD, 2016) When notable organizations decry the idea of legalizing another dangerous substance like

* Palau: University of Oregon. VanVuren: University of Oregon.

alcohol, it is certainly in the public interest to quantify the risks involved while driving under the influence of marijuana.

Though some marijuana users may stop themselves from driving under the influence, many others take the chance and drive after consumption. It is therefore necessary to understand the risk that high drivers impose to others and themselves on the road, compared to their sober counterparts. As most individuals fail to account for the full social costs of their actions, it is paramount to discover the relative risk high drivers impose on the road to calculate such costs and provide policy makers accurate information with which to form policy.

It may seem quite easy to determine relative risk. One would simply divide the number of fatal accidents involving high drivers by the number of high drivers on the road at a given time. The first metric can be found using the Fatality Analysis Reporting System provided by the National Highway Traffic Safety Administration, which reports national data on fatal accidents each year. While not perfect, the FARS allows an as accurate as possible estimation of relative risk using the most current data. Unfortunately the latter metric is nearly impossible to accurately measure. While some studies use roadside survey data to circumvent the issue, we believe that the error in using such data leads to less accurate reports of relative risk than possible. Measuring alcohol impairment on the side of the road across the United States is relatively easy when making use of widespread implied consent laws and breathalyzer tests. Because marijuana measurement requires blood tests and/or urinalysis, the collection of reliable measure of marijuana-related impairment in roadside surveys is nearly impossible. Levitt and Porter (2001) resolve these data issues by taking advantage of FARS data on two-car crashes involving varying driver types, and the fact that the distribution of different two-car crashes are determined by a multinomial distribution. Under the same conditions, the relative risk of high drivers can be estimated.

We expand the model of Levitt and Porter (2001) to include high drivers and estimate that on average high drivers are 2.83 times more likely to cause a fatal crash. For reasons described later in the paper, this number is best interpreted as a lower bound on the true increase in risk. Our model specification does not allow us to easily control for observable confounding factors such as gender or age, but we are able to bound the bias in our estimate under certain assumptions. We also estimate that, from 10 p.m. to 3 a.m., approximately 1.6% of drivers are high.

I. Literature Review

The crux of this study centers around the relative risk that drivers under the influence of cannabis impose on the road. Not limited to economic analysis, researchers in a broad scope of fields (pharmacology, psychology, biology, etc.) have approached the question using a variety of data sets and methodologies. The following review gives focus to research beginning in scientific and simulator studies, followed by traditional economic approaches, ending with an evaluation of a relevant methodology.

For decades the scientific community has grappled with whether or not cannabis consumption truly impairs driving — typically in relation to alcohol, a known impairment device. A meta-analysis and review by Sewell, Poling and Sofuoglu (2009) discusses cognitive, epidemiological, and experimental studies published in seeking to find marijuana-caused driving impairment. Citing over 60 cognitive studies, the authors adjudge that marijuana causes impairment in every performance area that can reasonably be connected with safe driving of a vehicle, such as tracking, motor coordination, visual functions, and particularly complex tasks that require divided attention, and conclude that an increased risk of being in a fatal traffic accident must follow cannabis consumption. Experimental studies, using driving simulators or closed driving courses, have found mixed results, with some researchers finding distinct impairment and others not. Most authors dismissing increased marijuana-related fatality risk claim that drivers' awareness of impairment leads to risk neutralizing behaviors like decreased speed and increased following distance. On the other hand, some argue that behavioral strategies cannot make up for impairment due to marijuana. For example, high drivers typically cannot stay within the bounds of their lane, erratically swerving with increased doses of THC. In general, a meta-analysis of over 120 experimental studies (Sewell, Poling and Sofuoglu, 2009) shows that a higher concentration of THC in a drivers blood will lead to greater the impairment and that more frequent users of marijuana show less impairment than infrequent users at the same dose, either because of physiological tolerance or learned compensatory behavior.

In a review of culpability studies analyzing driver responsibility for crashes, Sewell, Poling and Sofuoglu (2009) again reveals mixed conclusions. For example, one study (Drummer, 1995) found, using blood samples of traffic fatalities in Australia, that drivers testing positive for cannabis were less likely to have been judged as the responsible party. Conversely, Terhune (1986) finds that cannabis users had a responsibility rate of 76% versus 42.5% for a control group. Sewell concludes that though culpability studies have been contradictory, all clearly find "that the combination of alcohol and cannabis has worse consequences than use of cannabis alone."

Regardless of fatal crash risk, clinical research has been conducted on marijuana's effects on subjects' general risk taking behavior, as most people anecdotally assume that marijuana smoking itself exhibits risky behavior and disposition. In a risk-taking laboratory experiment (Lane et al., 2005), the highest administered THC dose led to increased selection of risky responses, while at lower doses, riskiness increases were not as pronounced or significant. The authors conclude that altered sensitivity to the consequences of their actions may be the key instrument in marijuana smoker's altered decision making and risk taking.

Beyond the traditional hard sciences, economists have studied traffic fatalities and alcohol extensively, including surrounding the legalization of medical marijuana in several states. Anderson, Hansen and Rees (2013) find that in the years following medical marijuana legalization laws, legalization is associated with an 8 to 11 percent decrease in traffic fatalities. Though not implying anything about the relative safety or risk of driving high, the authors suggest that alcohol and marijuana are substitutes, meaning a significant amount of the decrease in fatalities is due to drivers shifting away from dangerous alcohol consumption following cannabis legalization. The authors state that this is possible even if driving under the influence of marijuana is just as dangerous as driving drunk. This could be because marijuana is typically consumed at home, typically taking the option of driving while high away, or it could be that marijuana smokers choose to drive less than alcohol drinkers overall, among other explanations.

Scientists and economists alike have made prolific use of odds ratios in estimating fatal crash risk. An odds ratio is a measure of association between an exposure and an outcome, for example THC and a fatal car crash respectively (Szumilas, 2010). Most commonly applied in case control studies, odds ratios are a consistent tool in determining the relative risk of drugged drivers. Relevant odds ratios are calculated by solving the ratio of high drivers and sober drivers involved in fatal crashes, divided by the ratio of high drivers and sober drivers on the road not involved in fatal crashes. An odds ratio of 1.0 indicates no relationship between the two factors, while a ratio of over 1.0 indicates a positive relationship. To solve the problem of not knowing how many drugged drivers are on the road at a given time, most studies use the National Highway Traffic Safety Administrations (NHTSA) National Roadside Survey of Alcohol and Drug Use by Drivers (NRS) to determine the odds ratio denominator, which includes self-reported and voluntary driver data stratified at a random sample of weekend nighttime drivers in the lower 48 United States (Berning, Compton and Wochinger, 2015). Most studies concurrently use the NHTSA Fatality Analysis Reporting System (FARS), which contains a multiplicity of fatal crash information, to determine the odds ratio numerator. Using a 95% confidence interval, odds ratios of being involved in fatal traffic accidents under the influence of marijuana have been reported from 0.2 with any dosage, to 6.6 at significantly high doses of THC (Ramaekers et al., 2004). According to the NHTSA's own estimates, an alarming 12.6% of drivers on the road test positive for THC during the weekend nighttime. Though, "Drivers testing positive for THC were overrepresented in the crash-involved (case) population," the NHTSA reports an unadjusted odds ratio of 1.25 for THC-positive drivers, and a ratio of 6.75 for drivers with a BAC greater than 0.05 (Compton and Berning, 2015). It is additionally worth noting that the average BAC is more than double the 0,08 legal limit, with the NHTSA reporting that the most frequently recorded BAC among drinking drivers in fatal crashes in 2010 was 0.18 g/dL (NHTSA, 2012). It remains unclear if lowering the legal limit on alcohol would reduce fatal crashes.

A novel take on predicting the number of drunk drivers on the road comes from Levitt and Porter (2001). The authors estimate the effects of driving while intoxicated on fatal car crashes, and how various policy decisions affect measures associated with fatal collisions involving alcohol. Wanting to the compute relative risk of drunk driving while facing the fundamental problem of note knowing the amount of drunk and sober drivers on the road and inherent problems with roadside testing, Levitt and Porter take advantage of the fact that two-car collisions follow a binomial distribution. While the number of crashes involving a single sober driver and a single drunk driver will vary linearly with the number of drunk drivers on the road, the number of crashes involving two drunk drivers will vary with the square of the number of drunk drivers on the road. By exploiting this difference, the authors are able to compute an odds ratio without relying on survey data. Using fatality data only from FARS, and identifying two types of drivers, drunk and sober, the researchers make several assumptions to simplify the study: that driver type is independent of interactions on the road, that fatal crashes are the error of a single driver, that the composition of driver types in one fatal crash is independent of the composition of driver types in other fatal crashes, and that drinking increases the likelihood that a driver makes an error resulting in a fatal two-car crash. The papers most important contribution is the methodology used to assess the relative risk of drinking drivers, making use of maximum likelihood estimation. Under this model, the authors find that legally drunk drivers are 13 times more likely to be involved in a fatal crash than a sober driver. The authors conclude by estimating an appropriate punitive sum to account for a drunk drivers external cost to society (30 cents per mile), and an appropriate fine ($8,000 per arrest). The assumptions, methodology, and data can be reasonably transposed to address driving under the influence of marijuana, with a few new wrinkles (including increased possible driver types).[1][2]

II. Model

We build our model of fatal accidents following Levitt and Porter (2001) and classify each driver as either sober, drunk, or high. This model is notable because it allows us to estimate the relative risk of high drivers without requiring knowledge of the relative number of high drivers on the road. Although there are attempts at establishing the approximate number of high drivers using roadside surveys (Berning, Compton and Wochinger, 2015), using these surveys is prob-

[1] One notable study to take advantage of the methodology of Levitt and Porter (2001) is a technical report written by Martinelli and Maria-Paulina Diosdado-De-La-Pena (2008) in the field of civil engineering that examines the relative risk that Sport Utility Vehicles (SUVs) pose on other passenger cars. Facing a similar problem of not knowing how many SUVs were on the road at a given time, the group used Levitt and Porters model to estimate the odds ratio. Interestingly, the authors separate the United States into six independent areas in order to combat the requirement of space homogeneity. The study, using only FARS data, finds that SUVs are 2.7 times as likely to be involved in a fatal crash as their smaller counterparts.

[2] Loughran and Seabury (2007) conduct a similar study estimating the risk of older drivers on the road using the same methodology. Separating all drivers into either younger (25-64 years old) or older (over 65), the researchers, using FARS data from 1973 to 2003, find that older individuals are 16% more likely to cause a fatal crash than their younger peers. The authors explain that the riskiest older drivers self-regulate, meaning that many physically and mentally deteriorating older individuals reduce the hours that they drive, or pull themselves off the road permanently. Though finding an unadjusted relative crash risk odds ratio of 6.73, the authors adjust the final number to 1.16, after accounting for relative crash fatality rates, as older drivers are simply more likely to die in crashes than younger adults.

lematic due to potential bias in the sampling procedure, as well as the fact that these methods do not allow the proportion of high drivers to vary over time or across locations. Our method allows us to both avoid any survey bias and allow variation in high driving patterns across geography and time. We will briefly state the assumptions of the model.[3]

- Assumption 1 — There are three driver types: Sober (S), Drunk (D), and High (H), and all drivers can be placed into one of these categories. Note that this assumes that there are no drivers who are both Drunk and High. This lets us write N_S, N_D, N_H, and N_{tot} for the number of Sober, Drunk, High, and total drivers and say $N_S + N_D + N_H = N_{tot}$.
- Assumption 2 — Define I to be an indicator variable equal to 1 if two drivers interact. We assume that $P(i|I=1) = \frac{N_i}{N_{tot}}$ and $P(i,j|I=1) = P(i|I=1)P(j|I=1)$. In words, we assume that the probability of an interaction is independent of driver type (i.e. a Drunk driver is equally likely to interact with a Sober driver and a Drunk driver). This assumption is unlikely to hold in samples that are aggregated over long periods of time or across large regions of geography. Violations of this assumption will bias our estimates downwards (Levitt and Porter, 2001). We can mitigate or eliminate this bias by allowing variation across units of observation over which this assumption is plausible.
- Assumption 3 — A fatal crash occurs due to the error of a single driver. Levitt and Porter (2001) shows that violations of this assumption will leads to downwards bias in the estimate of relative crash risk.
- Assumption 4 — Driver types are independent across crashes. This assumption essentially amounts to assuming that one drunk driver being in a fatal accident does not influence that probability that another drunk driver is in an unrelated accident.
- Assumption 5 — We write θ_i for the probability that an i-type driver causes a crash and assume that $\theta_D, \theta_H \geq \theta_S$. The assumption that drunk and high drivers are more dangerous that sober drivers is well-supported by the literature (Zador, Krawchuk and Voas, 2000; Hall and Solowij, 1998).

From these assumptions, we can derive the probability distribution of driver types conditional on a fatal crash occurring. The second assumption gives the probability of an interaction between two drivers.

$$P(i,j|I=1) = \frac{N_i N_J}{N_{tot}^2} \tag{1}$$

[3] Levitt and Porter (2001) discusses relaxing these assumptions.

We need the probability of an accident occurring given an interaction between two drivers. Define A to be an indicator equal to one if an accident occurs between two drivers.

$$P(A=1|i,j,I=1) = \theta_i + \theta_j - \theta_i\theta_j \approx \theta_i + \theta_j \tag{2}$$

Multiplying (1) and (2) gives the probability of a fatal accident occurring between drivers i and j conditional on an interaction.

$$P(i,j,A=1|I=1) = \frac{N_i N_j(\theta_i + \theta_j)}{N_{tot}^2} \tag{3}$$

The final piece we need, $P(i,j|A=1)$, comes from summing (3) across all combinations of driver types.

$$P(A=1|I=1) = \frac{2(\theta_S N_S^2 + (\theta_S+\theta_D)N_S N_D + \theta_D N_D^2 + (\theta_S+\theta_H)N_S N_H + \theta_H N_H^2 + (\theta_D+\theta_H)N_D N_H}{N_{tot}^2} \tag{4}$$

Dividing (3) by (4) yields the equation we desire.

$$P(i,j|A=1) = \frac{P(i,j,A=1|I=1)}{P(A=1|I=1)} \tag{5}$$
$$= \frac{N_i N_j(\theta_i+\theta_j)}{2(\theta_S N_S^2 + (\theta_S+\theta_D)N_S N_D + \theta_D N_D^2 + (\theta_S+\theta_H)N_S N_H + \theta_H N_H^2 + (\theta_D+\theta_H)N_D N_H}$$

Let P_{IJ} denote $P(i=I, j=J|A=1)$. Then equation 5 gives us a system of six equations in six unknowns. Unfortunately, as probabilities, these equations necessarily sum to 1 meaning that we cannot identify all 6 parameters. Instead, we multiply all the equations by $\frac{1/\theta_S N_S^2}{1/\theta_S N_S^2}$ and rewrite $\gamma_D = \frac{\theta_D}{\theta_S}$, $\gamma_H = \frac{\theta_H}{\theta_S}$, $\nu_D = \frac{N_D}{N_S}$, and $\nu_H = \frac{N_H}{N_S}$. This reduces our system of equations to six equations in four unknowns, allowing us to identify the risk that Drunk and High drivers pose relative to Sober drivers. The resulting system of equations is:

$$\delta = 1 + (1+\gamma_D)\nu_D + \gamma_D\nu_D^2 + (1+\gamma_H)\nu_H + \gamma_H\nu_H^2 + (\gamma_D+\gamma_H)\nu_D\nu_H \tag{6}$$

$$P_{SS} = \frac{1}{\delta} \tag{7}$$

$$P_{SD} = \frac{(1+\gamma_D)\nu_D}{\delta} \tag{8}$$

$$P_{DD} = \frac{\gamma_D \nu_D^2}{\delta} \tag{9}$$

$$P_{SH} = \frac{(1+\gamma_H)\nu_H}{\delta} \tag{10}$$

$$P_{HH} = \frac{\gamma_H \nu_H^2}{\delta} \tag{11}$$

$$P_{DH} = \frac{(\gamma_D+\gamma_H)\nu_D\nu_H}{\delta} \tag{12}$$

The final step is to derive the likelihood function of our model, allowing us to estimate the parameters using observational data. Let A_{ij} denote the observed number of fatal crashes between driver type i and driver type j. We start with the likelihood function for a multinomial model.

$$L(\vec{A}) = \frac{(A_{SS}+A_{SD}+A_{DD}+A_{SH}+A_{HH}+A_{DH})!}{A_{SS}!\,A_{SD}!\,A_{DD}!\,A_{SH}!\,A_{HH}!\,A_{DH}!} P_{SS}^{A_{SS}} P_{SD}^{A_{SD}} P_{DD}^{A_{DD}} P_{SH}^{A_{SH}} P_{HH}^{A_{HH}} P_{DH}^{A_{DH}} \tag{13}$$

Combining (6)-(13) provides a likelihood function for $\gamma_D, \gamma_H, \nu_D, \nu_H$, allowing us to estimate them using our data. Our empirical strategy involves direct maximum likelihood estimation of (13). In our estimation, we allow ν_H and ν_D to vary over different units of observation, making our likelihood function the product of (13) over each unit of observation with γ_H, γ_D constrained to be constant.

In the likely situation of heterogeneity of risk within driver groups, θ_S, θ_D, and θ_H can be thought of as the mean risk of causing a fatal accident for each group. If this is the case, our estimates should not be interpreted as causal[4] as they will reflect differences in the compositions of each group.

III. Data

Our data source is the Fatality Analysis Reporting System which contains rich micro-level data for every accident in the United States. Beginning in 1983, local police departments were required to report any fatal crashes across the United States, making the FARS a complete database of fatal crashes. We look exclusively at accidents involving exactly two drivers. Due to changing patterns in marijuana use over time, the legalization of medical and recreational marijuana in some states, and low marijuana testing rates in the 80s and 90s, we restrict our

[4]For example, a $\hat{\gamma_H}$ of 3 would not imply that an individual is 3 times more likely to cause a fatal accident when high.

sample to accidents occurring between January 1, 2006 and December 31, 2014.

We also restrict our sample to fall between 10 p.m. and 3 a.m.[5] Although the number of accidents involving a high driver peaks at around 5 p.m. (see Figure 1) and follow a distribution very similar to that of sober drivers, we restrict the hours in our sample to maximize the number of accidents involving drunk drivers. Due to the relative rarity of high driving, particularly of accidents involving two high drivers, we rely heavily on collisions involving exactly one drunk and one high driver to estimate the relative number of high drivers on the road. Hence restricting our sample to nighttime hours minimizes the standard errors of our model. This has an added effect of making our sample directly comparable to the NHTSA's roadside survey.

Applying the above restrictions reduces our sample to slightly over 15,000 accidents. Of these, we also drop any accidents for which time of day is not reported (<1% of the remaining data) and any accidents involving a driver who is both drunk and high[6] (about 4% of the remaining data). After dropping these data, our final sample contains 14,480 two-driver accidents.

We may be worried about compositional differences between sober drivers and drunk or high drivers, particularly for traits that could be highly correlated with driving risk such as gender or age. Table 1 compares the percentage of drivers who are male, younger than 26, and both between the entire sample, the sample of high drivers, and the sample of drunk drivers. The composition of the sample varies across sober, high, and drunk drivers. In particular, a disproportionately large number of high drivers are also young. While our model can be used to estimate the relative risk of young drivers, it provides no way of distinguishing between an increase in risk for young drivers due to their higher propensity to drive while high or an increase in the risk of high drivers due to the inherent riskiness of young drivers. We address this concern in a later section.

We rely on blood and urine tests to determine marijuana involvement; any driver who tests positive for cannabinoids is classified as a high driver. To determine alcohol involvement, we use law enforcement officers' judgment as well as the values of reported BAC tests. We classify any driver who the police report as being involved with alcohol or who has a positive BAC test as drunk. The FARS provides an accident-level statistic on the number of drunk drivers involved in a crash; however, these data were incorrectly derived for the years 1999-2007. Ignoring these years, our measure of alcohol involvement agrees with the FARS-reported number of drunk drivers for more than 98% of the data so we are confident that our measure is accurate.

Both these methods of classification can cause bias in our estimates if drivers' observable characteristics (e.g. age, race, or gender) are correlated with both that risk of causing a fatal accident and the probability that an officer chooses

[5] Some restrictions on the data are necessary to make the maximum likelihood model computationally tractable.

[6] We could include these drivers as a fourth type in our model, but doing so would make the model significantly more computationally difficult while providing very little additional information.

to conduct a drug test or declare alcohol involvement. Assuming there is very little systematic variation in the fatal crash risk of sober drivers, this bias will be small. To help mitigate this bias, we may wish to restrict our sample to states that meet a set threshold for the percentage of drivers tested, but making such a restriction does not significantly change our estimate of γ_H and modestly increases our standard errors.[7]

IV. Estimation

To estimate our model, we perform direct maximum likelihood estimation of equation 13. We wish to allow ν_H and ν_D to vary over the smallest unit of observation possible to minimize downwards bias due to violations of Assumption 2. Unfortunately, computational tractability limits this possibility; allowing variation over hour, year, and state would result our likelihood function containing 4592 parameters over which to maximize. To help bring this number down to something reasonable, we group the years 2006-2008, 2009-2011, and 2012-2014 together, which we will refer to as YearG. We also group states together by census region. This reduces the number of parameters in our likelihood function to 326. Additionally, we estimate the model twice more, allowing variation only over different hours and over Hour-YearG combinations. For details on how we estimate the model, refer to the technical appendix.

Table 2 presents the results of our estimation. Each column contains estimates from one of our three model specifications, allowing for variations in ν_H and ν_D over smaller units of observation from left to right. As we expect, allowing ν_H and ν_D to vary over consecutively smaller units of observation increases our estimates, indicating that the amount of downwards bias due to violations of Assumption 2 is shrinking. Due to the remaining downwards bias, our estimates are best thought of as lower bounds for the true risk of driving while high. In each model specification, we are able to reject the null hypothesis of a relative risk factor equal to 1 at the 5% level.

Our preferred estimate of γ_H is 2.83, indicating that the average high driver is 2.83 times more likely to cause a fatal accident than the average sober driver; the corresponding estimate of γ_D is 5.47. We report the computed standard errors using the Hessian approximation of the variance-covariance matrix as well as bootstrapped 95% confidence intervals. For each model specification, the traditional confidence intervals for γ_D are very similar to the bootstrapped intervals. The traditional confidence intervals for γ_H are moderately larger than the bootstrapped intervals. Figure 2 provides a potential explanation for this discrepancy. The collection of bootstrapped estimates of γ_D in the first and second models looks approximately normal, while bootstrap estimates of γ_H appear to be slightly skewed. This indicates our sample size was sufficiently large to use the

[7] Our analysis used a minimum testing rate of .7, where an accident was considered tested if at least one driver involved was given a drug test.

Hessian approximation of the variance-covariance matrix for γ_D, but not large enough to use the approximation for γ_H. This is likely due to the relative rarity of high driving and a lack of High-High collisions in particular.

Our estimate of γ_H is directly comparable to the odds ratios estimated by previous studies. Most notable, Compton and Berning (2015) estimates an odds ratio of 1.25 using the NHTSA's National Roadside Survey and FARS. This is significantly lower than our preferred estimate of 2.83. Their estimation procedure makes use of the NHTSA's roadside survey, which estimates that 12.6% of nighttime weekend drivers are high. While our implied percentage of high drivers on the road are not directly comparable, as our model doesn't distinguish between weekend and non-weekend accidents, our average implied percentage of high drivers on the road is 1.6% with a maximum of 9.0% [8] (see Figure 3a for a histogram of the estimates). Violations of Assumption 2 will bias these estimates upwards, making them most easily interpretable as upper bounds. Thus our estimates are inconsistent with the claim that 12.6% of nighttime drivers are high. This explains the difference between our estimate of γ_H and that of Compton and Berning (2015).

Our estimate of γ_D can be easily compared to the relative risk estimated in Levitt and Porter (2001). Our estimate in column 1 of table 2 is 5.46; using the same unit of observation, they estimate a relative risk parameter of 4.87. These estimates are fairly similar, although they are different enough to warrant further investigation. To examine the trend in drunk driving risk over time, we attempt to reproduce the sample selection in Levitt and Porter (2001) and estimate their model over time, beginning with the years 1983-1993 and incrementing by a single year up to 2004-2014. Our person-level measure of drunk driving matches the FARS accident-level measure of drunk driving up until 1999 (when the FARS self-reports that the accident-level drunk driving statistic is improperly computed) with the exception of the year 1989 where the person-level measurement disagrees with the accident-level measurement for over 50% of the data. Because of this, we drop any accident from the year 1989 in our analysis.[9] We can see from Figure 4 that the relative riskiness of drunk driving has been trending upwards over time, consistent with the observation that our estimate of γ_D is higher than that of Levitt and Porter (2001).

This increase in relative risk could be due to an increase in the risk that a drunk driver will cause a fatal accident (i.e. an increase in θ_D) or a decrease in the risk that a sober driver will cause a fatal accident (i.e. a decrease in θ_S). Figure 5 compares the BAC distribution of drinking drivers involved in fatal accidents for the years 1983-1993 and 2006-2014. The mean of the distribution are 16.29 and 16.69 respectively, indicating that the increase in relative risk cannot be explained by an increase in the BAC of the average drunk driver. One possible

[8] We compute the percentage of drivers who are high as $\frac{1}{\frac{1}{\nu_H}+\frac{\nu_D}{\nu_H}+1}$

[9] By including this data and using the accident-level measurement of drunk driving, we are able to reproduce column 2 of table 2 in Levitt and Porter (2001).

alternative explanation for this increase is that the increasing severity of drunk driving punishments has been more effective at deterring potential drunk drivers with relatively low propensity for risk taking behavior, leaving only the riskiest drivers on the road. It is also possible that new automobile safety features that require attentive drivers, such as back-up cameras and collision warning systems, have decreased θ_S while having no impact on θ_D, resulting in an increase of γ_D over time. Additionally, it is possible that data quality has improved over time, allowing us to distinguish between drunk and sober drivers with less noise.

As shown in Table 1, we may be concerned that our estimate of γ_H is biased due to the fact that male and young drivers are disproportionately represented among high drivers. To examine the magnitude of this bias, we estimate the relative risk factor for males, drivers younger than 26, and drivers who are both male and younger than 26. As Table 1 demonstrates, drivers in these categories are more likely to be impaired. Thus we estimate the risk factor twice, once including our entire sample of drivers and once excluding all accidents involving an impaired driver.[10] The results of this estimation are shown in Table 3.

We estimate that male drivers are 1.88 times more likely to cause a fatal accident than female drivers. We drop any accidents involving a drunk or high driving when computing this estimate, so we are confident that drunk and high driver are not responsible for this increase in risk. Because there are more male drivers who are high than we would expect if being male was independent of driving while high, this will bias our estimates upwards. It is fairly simple to estimate the extent of the upward bias. From Table 1, we know the overall percentage of male drivers as well as the percentage of high drivers who are male. Using these percentages and our estimate of 1.88, we can compute an expected value of γ_H due to the compositional change if high drivers were no more dangerous than sober drivers. To compute this value, define θ_F to be the probability that a female driver causes a fatal accident. Then the probability that a male driver causes an accident is $1.88\theta_F$. We know that 77.7% of drivers in fatal accidents are male, so the average risk is $(.777(1.88) + .223)\theta_F$. We perform the same computation using 81.5%, the percentage of high drivers that are male, and divide the latter by the former to get a value of 1.02. This value is very small relative to our estimated value of γ_H, so we are confident that the upward bias due to the increased percentage of males who are high is negligible.

If we make an additional assumption, we can compute a second estimate of γ_H while controlling for the difference in the percentage of males between the two populations. Define θ_{SM}, θ_{SF}, θ_{HM}, and θ_{HF} to be the probability that a sober male, sober female, high male, and high female cause an accident respectively. Our new assumption is that $c\theta_{SM} = \theta_{HM}$ and $c\theta_{SF} = \theta_{HF}$ for some c. Informally, this means that being high multiplies the probability of causing an accident by a constant for both males and females. Our estimated relative risk ratio of 2.83 is the ratio between the two average risks. Using

[10] We exclude accidents involving drivers for whom age and gender are not known

the data from Table 1 for the percentage of males in fatal crashes, we have $\frac{.815\theta_{HM}+.185\theta_{HF}}{.777\theta_{SM}+.223\theta_{SF}} = 2.83 \Rightarrow c\frac{.815(1.88)\theta_{SF}+.185\theta_{SF}}{.777(1.88)\theta_{SF}+.223\theta_{SF}} = 2.83 \Rightarrow c = 2.77$. Note that this is equivalent to dividing 2.83 by the 1.04 we found as the expected value of γ_H due to the compositional change assuming that high drivers were no more likely to cause fatal accidents than sober drivers.

Beginning with the second and third column, we estimate that young drivers, whether they are male or not, are no more likely to cause a fatal accident than other drivers. This observation does not hold up when looking at other literature (Jonah, 1986; Williams, 2003) and is almost certainly due to the fact that we would need to use a smaller unit of observation to expect Assumption 2 to hold. Young drivers are likely concentrated around high school and college campuses, as well as in cities. This geographic concentration means that young drivers are more likely to collide with young drivers than with older drivers. These strong violations of Assumption 2 bias our estimates downwards to the extent that young drivers appear as safe as older drivers. Levitt and Porter (2001) estimate that young drivers are 2.78 times more likely to cause a fatal accident. If we assume that $c\theta_{SY} = \theta_{HY}$ and $c\theta_{SO} = \theta_{HO}$ where θ_{SY}, θ_{HY}, θ_{SO}, and θ_{HO} are the probabilities of causing an accident for sober young drivers, high young drivers, sober older drivers, and high older drivers respectively (note the similarity to the assumption made in the previous paragraph), then we can adjust our estimate of γ_H from 2.83 to 2.25. Other studies estimate relative risk parameters for young drivers of 1.44 and 1.17 (Zador, Krawchuk and Voas, 2000; Mao et al., 1997).

Figure 6 plots the adjusted estimates of γ_H as a function of the relative risk of male and young drivers using the method of the previous two paragraphs. Our estimate of γ_H is only slightly effected even if the hypothetical relative risk parameter for males is large. The relative risk parameter of young drivers has much more influence on our adjusted estimate of γ_H. This is because the change in composition between the sober and high driver populations is more severe for young drivers than it is for high drivers.

V. Limitations and Suggestions for Further Study

In this section, we discuss the limitations of our results and suggest potential areas for future study. Most of the limitations of our results are due to various assumptions made in our model or limitations of the FARS data. These limitations raise interesting and policy-relevant questions that should be the subject of future study.

As a limitation, it's important to remember that our estimates are non-causal, meaning the estimates cannot suggest that consuming cannabis and driving will cause increased fatal accidents. Because we don't employ any form of exogenous variation, our model cannot account for unobservable factors, such as individual risk-taking propensity, that may be correlated with both the probability that an individual chooses to drive while high, and the chance that an individual causes a fatal accident. We also do not allow heterogeneity in the relative risk between individuals; as a concrete example, it's possible that a small number of individuals who are experienced, chronic marijuana users may properly compensate for their impairment without increasing their probability of causing a fatal crash.

We are additionally limited by the data provided by the FARS. Although BAC is reported for the vast majority of drunk drivers, there is no measure of THC levels for drivers who are reported to be high. As THC can remain in an individual's bloodstream at residual levels for days (Cary, 2006), it's possible that some drivers that we identify as high are actually experiencing very little of the effects of THC. To maintain computational tractability, we also choose not to make use of the data on BAC other than to identify drunk drivers. It is unlikely that this impacts our estimates of γ_H; the estimate of γ_D is the mean relative risk parameter for drunk drivers and fully controls for the influence of drunk drivers. For a thorough examination of how risk varies across different BACs, see Levitt and Porter (2001).

As mentioned above, we are also limited in our ability to disaggregate the data into acceptably small units of observation. Because Assumption 2 is unlikely to hold at high levels of aggregation, our estimates are likely biased downwards. Thus, our estimates are best interpreted as lower bounds of the true value of relative risk.

With many states in the process of legislating policy for recreational or medical marijuana use, understanding the externalities of high driving is extremely policy-relevant. While our study demonstrates that there are likely fairly large externalities involved, more study is certainly needed. In particular, with a better understanding of how driving risk varies with a driver's level of THC and with information on the distribution of THC levels of high drivers, policy makers can make better-informed, evidence-based policies regarding the legal threshold for high driving.

VI. Conclusion

As policy makers are rapidly writing marijuana legislation, it is important that they have a solid understanding of the potential risks that marijuana poses, such as high driving. These risks can be hard to estimate due to the speed at which marijuana has become mainstream and the inherent difficulties in collecting reliable data on high driving. Our model of fatal crashes allows us to estimate the relative risk of high drivers without relying on potentially unreliable data from roadside surveys. Our preferred specification estimates that high drivers are 2.83 times more likely to cause a fatal accident than sober drivers. We find very little evidence that this estimate is biased by observable differences in the populations of high and sober drivers, namely by gender and age. We also find convincing evidence that high driving is safer than drunk driving, supporting claims that the legalization of marijuana may reduce traffic fatalities due to substitution effects.

Our results are inconsistent with the NHTSA's National Roadside Survey of Alcohol and Drug Use by Driver's claim that 12.6% of weekend nighttime drivers are high. Because of this, our preferred estimate is sufficiently different from previous odds-ratio estimates that rely on this survey data. In particular, we find an average implied percentage of high drivers on the road of 1.6%. When compared to the implied percentages of drunk drivers on the road (see Figure 3), it is clear that our estimates imply that high driving is a significantly rarer phenomenon than drunk driving.

VII. Tables

Table 1—: Summary of Sample

	Total	High	Drunk
Male	77.7%	81.5%	82.0%
Younger than 26	31.3%	54.3%	33.8%
Male and below 26	23.2%	44.7%	27.1%

Table 2—: Model Results using FARS data

	(1)	(2)	(3)
Estimate of γ_D	5.47 (0.33)	5.49 (0.33)	5.94
Estimate of γ_H	2.60 (0.63)	2.62 (0.64)	2.83
Bootstrap 95% Interval for γ_D	[4.82,6.17]	[4.89,6.29]	[5.59,6.94]
Bootstrap 95% Interval for γ_H	[1.69,4.42]	[1.62,4.34]	[1.79,5.13]
Unit of Observation	Hour	Hour x YearG	Hour x YearG x Region

Standard errors in parentheses where available.
All bootstrap intervals computed using 500 samples except where indicated.

VIII. Figures

Table 3—: Estimated Relative Risk of Potential Risk Factors

	Male	Younger than 26	Both
Risk: including impaired drivers	1.94	1.00	1.00
Risk: excluding impaired drivers	1.88	1.00	1.00

In estimating this table, we drop any accidents involving a driver for which age or gender is unknown.
We use Hour x YearG x Region as our unit of observation, making these estimates directly comparable to those of table 2.

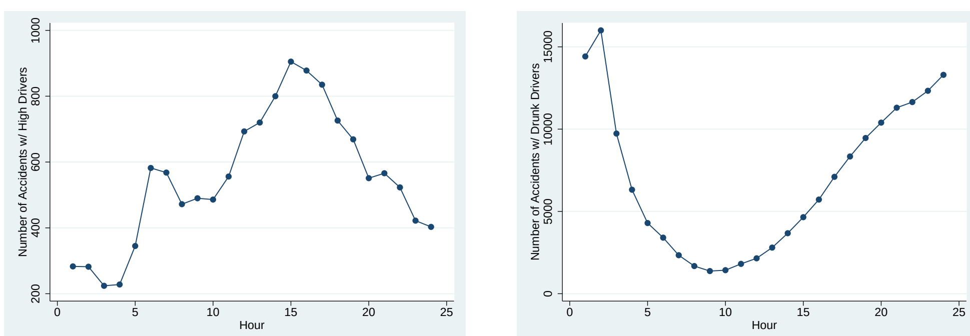

a: Accidents Involving a High Driver by Hour

b: Accidents Involving a Drunk Driver by Hour

Figure 1. : Accidents by Hour

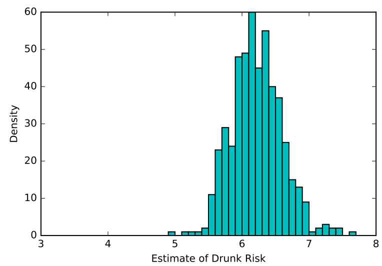

a: Bootstraps of γ_D: third specification

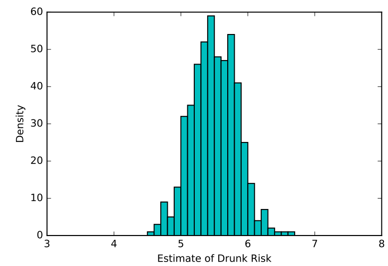

b: Bootstraps of γ_D: second specification

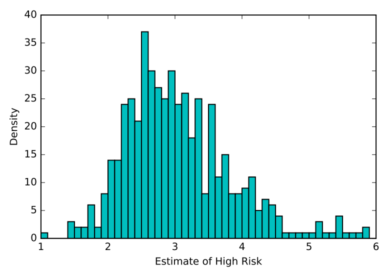

c: Bootstraps of γ_H: third specification

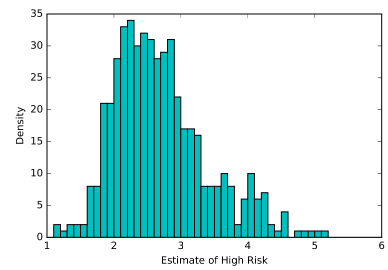

d: Bootstraps of γ_H: second specification

Figure 2. : Histogram of bootstrapped estimates of γ_H and γ_D

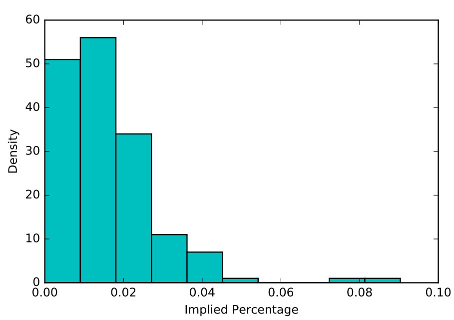

a: Implied percentage of drivers who are high

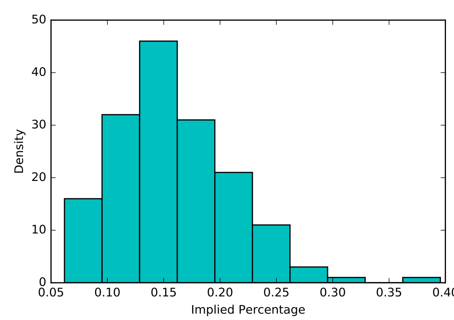

b: Implied percentage of drivers who are drunk

Figure 3. : Histograms of implied percentages

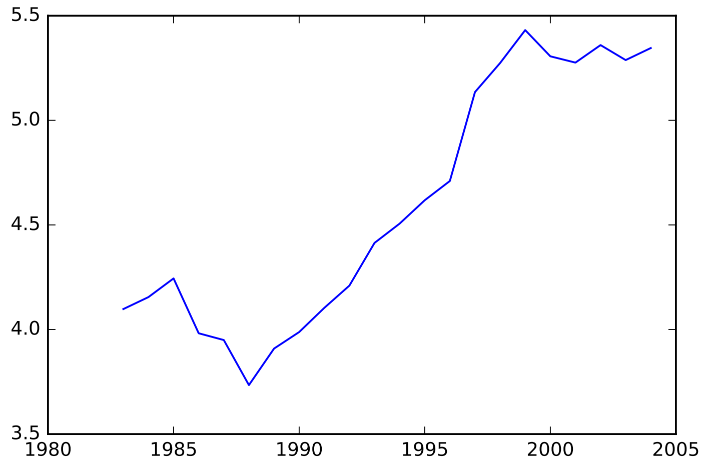

Figure 4. : Estimated Relative Risk of Drunk Drivers over Time

These estimates were computed independently from any estimates involving high driving. A label of "1983" on the x-axis indicates that the estimate was computing using data from 1983-1993.

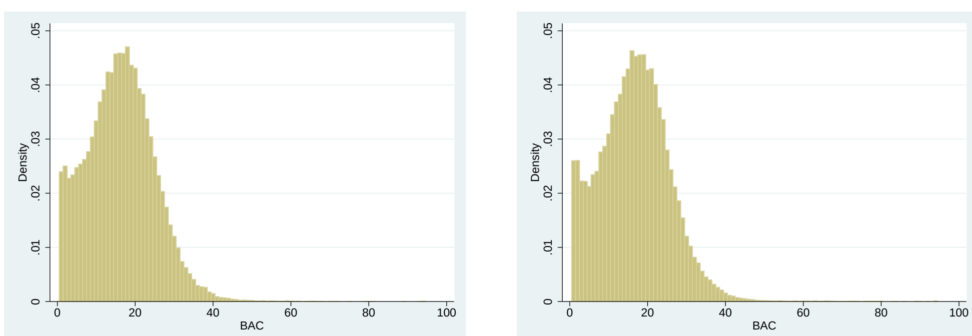

a: 1983-1993: mean of 16.34

b: 2006-2014: mean of 16.43

Figure 5. : Histograms of Non-zero BAC of Drivers in Fatal Accidents

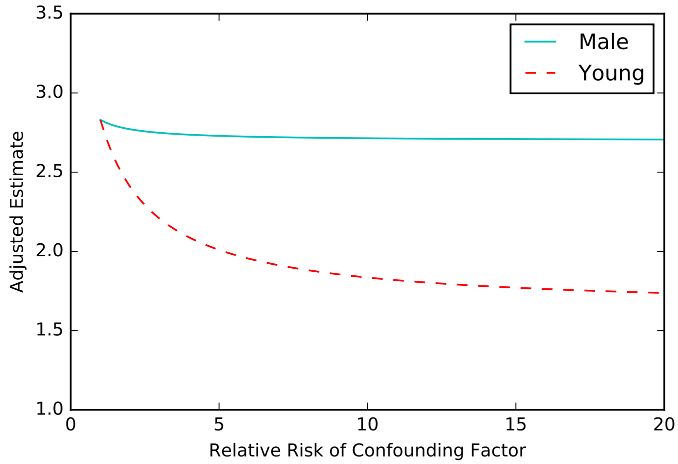

Figure 6. : Adjusted Estimates of γ_H for Different Risk Parameters of Confounding Factors

REFERENCES

Anderson, D Mark, Benjamin Hansen, and Daniel I Rees. 2013. "Medical marijuana laws, traffic fatalities, and alcohol consumption." *Journal of law and economics*, 56(2): 333–369.

Berning, By Amy, Richard Compton, and Kathryn Wochinger. 2015. "Research Note Results of the 2013 2014 National Roadside Survey of Alcohol and Drug Use by Drivers." February.

Cary, Paul. 2006. "The Marijuana Detection Window: Determining the length of time cannabinoids will remain detectable in urine following smoking: A Critical Review of Relevant Research and Cannabinoid Detection Guidance for Drug Courts." 2.

Compton, Richard P, and Amy Berning. 2015. "Drug and Alcohol Crash Risk." February.

Drummer, O. H. 1995. "A review of the contribution of drugs in drivers to road accidents." , (May).

Hall, W, and N Solowij. 1998. "Adverse effects of cannabis." *Lancet*, 352(9140): 1611–1616.

Jonah, B a. 1986. "Accident risk and risk-taking behaviour among young drivers." *Accident; analysis and prevention*, 18(4): 255–271.

Lane, Scott D, Don R Cherek, Oleg V Tcheremissine, Lori M Lieving, and Cythia J Pietras. 2005. "Acute marijuana effects on human risk taking." *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology*, 30(4): 800–9.

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

Loughran, David S., and Seth A. Seabury. 2007. "Estimating the Accident Risk of Older Drivers." RAND Institute for Civil Justice.

MADD. 2016. "Drugged Driving — MADD,." *http://www.madd.org/drugged-driving/*, [Online; accessed 18-May-2016].

Mao, Y, J Zhang, G Robbins, K Clarke, M Lam, and W Pickett. 1997. "Factors affecting the severity of motor vehicle traffic crashes involving young drivers in Ontario." *Injury prevention : journal of the International Society for Child and Adolescent Injury Prevention*, 3(3): 183–9.

Martinelli, David R, and Maria-Paulina Diosdado-De-La-Pena. 2008. "Safety Externalities of SUVs on Passenger Cars : An Analysis Of the Peltzman Effect Using FARS Data." West Virginia University.

NHTSA. 2012. "Prevalence of High BAC in Alcohol-Impaired Driving Fatal Crashes." August.

Ramaekers, J. G., G. Berghaus, M. Van Laar, and O. H. Drummer. 2004. "Dose related risk of motor vehicle crashes after cannabis use."

Sewell, R. Andrew, James Poling, and Mehmet Sofuoglu. 2009. "The Effect of Cannabis Compared with Alcohol on Driving." *American Journal on Addictions*, 18(3): 185–193.

Szumilas, Magdalena. 2010. "Explaining Odds Ratios." *Journal of the Canadian Academy of Child and Adolescent Psychiatry*, 341(August): c4414.

Terhune, Kenneth W. 1986. "Problems and methods in studying drug crash effects." *Alcohol, Drugs & Driving*, 2(3-4): 1–13.

Williams, Allan F. 2003. "Teenage drivers: Patterns of risk." Vol. 34, 5–15.

Zador, Paul L., Sheila A. Krawchuk, and Robert B. Voas. 2000. "Alcohol-related relative risk of driver fatalities and driver involvement in fatal crashes in relation to driver age and gender: An update using 1996 data." *Journal of Studies on Alcohol*, 61(3): 387–395.

Technical Appendix

We use Python to compute the maximum of our likelihood function. In particular, we compute our likelihood function using the numpy package. Numpy computes array operations in C, which allows us to compute the likelihood function much more quickly than any Python-only operation could. We then use the scipy.optimize package for black-box minimization of our likelihood function. The sequential least squares quadratic programming algorithm gives us the quickest and most stable convergence. To compute our bootstrap intervals, we fix the number of accidents per unit of observation and resample within each unit from its own probability distribution. We then feed this new sample directly into our maximization method to compute a bootstrapped estimate.