

NBER WORKING PAPER SERIES

ESTIMATING THE EFFECT OF ALCOHOL
ON DRIVER RISK USING ONLY
FATAL ACCIDENT STATISTICS

Steven D. Levitt
Jack Porter

Working Paper 6944
<http://www.nber.org/papers/w6944>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 1999

We would like to thank Joe Altonji, Gary Chamberlain, Austan Goolsbee, Lars Hansen, James Heckman, Guido Imbens, James Poterba, Michael Thaddeus, and numerous seminar participants for helpful comments and advice, as well as Christopher Ruhm for generously providing us with data. Financial support of the Milton Fund, administered by Harvard Medical School, and the National Science Foundation, administered by the NBER, is gratefully acknowledged. The views expressed here are those of the author and do not reflect those of the National Bureau of Economic Research.

© 1999 by Steven D. Levitt and Jack Porter. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Estimating the Effect of Alcohol on Driver
Risk Using Only Fatal Accident Statistics
Steven D. Levitt and Jack Porter
NBER Working Paper No. 6944
February 1999
JEL No. C13, I12, K42

ABSTRACT

Measuring the relative likelihood of fatal crash involvement for different types of drivers would seem to require information on both the number of fatal crashes by driver type and the fraction of drivers on the road falling into each category. In this paper, however, we present a methodology for measuring fatal crash likelihood that relies solely on fatal crash data. The key to our identification strategy is the hidden richness inherent to two-car crashes. Crashes involving two drinking drivers are proportional to the square of the number of drinking drivers on the road; crashes with one drinking and one sober driver increase linearly in the number of drinking drivers. Imposing a limited set of assumptions (e.g. independence across crashes, equal mixing on the roads), we are able to estimate both the likelihood of causing a fatal crash and the fraction of drivers of each type on the road. Our estimates suggest that drivers with alcohol in their blood are at least eight times more likely to cause a fatal crash; legally drunk drivers pose a risk at least 15 times greater than sober drivers. Males, young drivers, and drivers with bad past driving records are all more dangerous, but the impact of these other factors is far less than that of alcohol.

Steven D. Levitt
Department of Economics
University of Chicago
1126 E. 59th St.
Chicago, IL 60637
and NBER
slevitt@midway.uchicago.edu

Jack Porter
Department of Economics
Harvard University
Cambridge, MA 02138
jporter@harvard.edu

Motor vehicle crashes claim over 40,000 lives a year in the United States, approximately the same number of Americans killed over the course of either the Korean or Vietnam wars. The death toll in motor vehicle accidents roughly equals the combined number of suicides and homicides, and motor vehicle deaths are thirty times as frequent as accidental deaths due to firearms. Motor vehicle accidents are the leading cause of death for Americans aged 6-27.

Alcohol is often implicated in automobile deaths. According to police reports, at least one driver has been drinking (although not necessarily over the legal blood-alcohol limit) in over 40 percent of fatal crashes. During the time periods where alcohol usage is greatest, that proportion rises to more than 60 percent.¹

Without knowing the fraction of drivers on the road who have been drinking, however, it is impossible to draw conclusions about the relative fatal crash risk of drinking versus sober drivers. Past research has attempted to measure this fraction through the use of random roadblocks and driver stops (Lehman et. al. 1975, Lund and Wolfe 1991, Hurst et al. 1994).² While these studies are extremely valuable, they suffer from a number of important limitations. First, they are costly to undertake and consequently are performed only rarely. Second, in such experiments, drivers who have been stopped cannot be compelled to submit to alcohol tests. In

¹ Because many fatal crashes involve more than one vehicle, the actual fraction of drinking drivers involved in these crashes is lower than the values cited above. Overall, roughly 30 percent of drivers in fatal crashes have been drinking, with that percentage rising to 50 percent during peak alcohol usage times.

² There is also survey data asking drivers whether they have driven when they have “had too much to drink” (Liu et al 1997). In addition to any question about the accuracy of the responses given, these surveys have not attempted to ask drivers to report the percentage of miles driven with and without the influence of alcohol. Without that number, accurate estimates of the elevated risk of drinking drivers cannot be computed.

practice, roughly ten percent of drivers refuse to participate -- presumably those most likely to have been drinking (Lund and Wolfe 1991). The assumptions adopted for dealing with this sample selection are critical to the interpretation of the data. Third, even if the estimates obtained are reliable, they reflect the specific circumstances at a particular time and place and the extent to which the conclusions are broadly generalizable is unknown.³

In this paper, we adopt a radically different strategy for estimating the fraction of drinking drivers and the extent to which their fatal crash likelihood is elevated. Specifically, we rely exclusively on data from fatal crashes. A priori, it would seem that such an exercise, if even feasible, would require an extremely restrictive set of assumptions and the imposition of an arbitrary functional form. Separately identifying the fraction of drinking drivers on the road and their relative fatal crash risk using only the fraction of drinking drivers in fatal crashes is ostensibly equivalent, for instance, to determining the relative free-throw shooting ability of two basketball players based *only* on the number of free throws successfully completed by each. Without knowing how many free-throw attempts each player had, such an exercise would appear futile. In the realm of economics, our task is equivalent to separately identifying per capita income and population, based only on aggregate income data.

Despite the apparent difficulty of this exercise, the set of assumptions required for identification of the model are in actuality quite natural and do not even require the imposition of arbitrary distributional assumptions. The complete set of assumptions required are as follows.

³ A separate literature examines the impact of public policies on fatal car crashes (Cook and Tauchen 1982, Asch and Levy 1987, Saffer and Grossman 1987, Homel 1990, Chaloupka et al. 1993, Grossman et al. 1993, Ruhm 1996). As discussed in the concluding section, our work is complementary to this policy-related literature.

First, there are two types of drivers, drinking and sober, with different likelihoods of causing fatal crashes. Second, drinking and sober drivers mix equally on the roads within the relevant geographic area and time period under consideration. Third, one driver's error causes a fatal crash. Fourth, fatal crashes are independent of one another. Finally, drinking drivers are assumed to have a weakly greater likelihood of causing a fatal crash than sober drivers. This set of assumptions provides a natural baseline for analyzing the problem. Possible violations of the assumptions are discussed in Section II, where it is argued that the likely direction of bias is towards under-estimating how dangerous drinking drivers are. Thus, the most reasonable interpretation of our results is as a lower bound on this value.

The ability to identify the parameters arises from a hidden richness in the data due to the fact that crashes often involve multiple drivers. For two-car crashes, the relative frequency of accidents involving two drinking drivers, two sober drivers, or one of each provides extremely useful information. Indeed, given the set of assumptions outlined above, this information alone is sufficient to separately identify both the relative likelihood of causing a fatal crash on the part of drinking and sober drivers, and the fraction of the drivers on the road who have been drinking. The intuition underlying the identification of the model is extremely simple. The number of two car fatal crash opportunities is dictated by the binomial distribution. Consequently, the number of fatal two-car crash opportunities involving two drinking (sober) drivers is proportional to the *square* of the number of drinking (sober) drivers on the road. The number of fatal crash opportunities involving exactly one drinking and one sober driver is *linearly* related to both the number of drinking and sober drivers. Identification of the model arises from these intrinsic non-linearities. These non-linearities are not artificially imposed on the problem via arbitrary

functional form assumptions, but rather, are the immediate implication of the binomial distribution, which relies only on the assumptions stated above concerning independence of crashes and equal mixing of the different types on the road.

Applying the model to fatal accident data in the United States over the period 1983-93, we obtain a number of interesting results. Drivers identified by police as having been drinking (but not necessarily legally drunk) are at least eight times more likely to cause a fatal crash than are drivers with no reported alcohol involvement. Drivers above the blood-alcohol limit of 0.10 are at least 14 times more likely to be the cause of fatal crashes. When we apply the model to other observable traits, males, young drivers, and those with bad previous driving records are also more likely to cause crashes. Drinking, however, is far more important than these other characteristics and much of the apparent impact of gender and past driving record actually reflects differential propensities to drink and drive across groups. The exception is young drivers: sober, young drivers are almost three times as likely to cause a fatal crash as are other sober drivers. The peak hours for drinking and driving are between 1 a.m. and 3 a.m., when as many as 25 percent of drivers are estimated to have been drinking. The proportion of drinking drivers appears to have fallen by about one-quarter over the course of our sample. The relative fatal crash risk of drinking drivers, in contrast, appears to have been stable.

The remainder of the paper is structured as follows. Section II derives the basic model and discusses the sensitivity of the results to alternative modeling assumptions. Section III describes and summarizes the data used. Section IV presents the empirical estimates of the relative crash risk and number of drinking and sober drivers, as well as a number of extensions to the basic model. Section V considers possible uses of our measure in analyzing public policy

towards road safety, as well as potential applications of the methodology in other settings where interactions occur.

Section II: A Model of Fatal Crashes

In this section, we present a simple model of fatal accidents, demonstrating how identification of the underlying structural parameters (the fraction of drivers on the road who have been drinking and the relative likelihood of causing a fatal crash risk by drinking and sober drivers) naturally emerges from the model. A number of features of the model are worth noting. First, identification relies only on the distribution of crashes in a particular geographic area over a given period of time. Consequently, the model does not necessitate comparisons across times and places that may differ in systematic, yet unobservable ways, leading to biased estimates. Second, although the model is identified off non-linearities, these functional forms are not arbitrarily imposed. Rather, the structural equations that will be estimated follow directly from the restrictions dictated by nature in the form of the binomial distribution. Third, the approach we outline provides a previously unattainable flexibility in measuring drinking and driving. The solution to the model depends only on fatal crash tallies, data that are already collected and widely available. Consequently, parameter estimates can be obtained almost without cost for any geographic area or time period of interest to the researcher, e.g. the Chicago metro area, on weekends between 10 p.m. and 2 a.m. (although standard errors increase as the number of fatal crashes on which the estimate is based shrinks).

Assumptions of the Model

We begin by outlining the five assumptions underlying the model. The first assumption is

(A1) *There are two driver types, D and S*

where D and S correspond to drinking and sober drivers respectively, although other categories of driver types could also be used. Restricting the analysis to two types is done primarily to ease exposition of the model, which readily generalizes to multiple types. In some of our empirical estimation we allow four types. In theory, any number of types could be incorporated if enough data existed. As we demonstrate later, the parameter estimates from a model assuming exactly two types have a straightforward interpretation when there is heterogeneity in driver risk within these two categories of drivers.

The second assumption of the model is that of “equal mixing” of drinking and sober drivers on the roads. By equal mixing, we mean two things: (1) that the number of interactions that a driver has with other cars is independent of the driver’s type, and (2) a driver’s type does not affect the composition of the driver types with which he or she interacts. A formal statement of this assumption requires some notation. Let the total number of drivers be N , and let the total number of drivers of type i be N_i . Thus, by assumption A1 that there are only two types, $N_D + N_S = N$. Define I to be an indicator variable equal to 1 if two cars interact (i.e. pass each other on the road) and equal to zero otherwise. Denoting two drivers as having types i and j , then,

$$(A2) \text{ (i) } \Pr(i|I = 1) = \frac{N_i}{N_D + N_S}$$
$$\text{(ii) } \Pr(i, j|I = 1) = \Pr(i|I) \Pr(j|I)$$

Assumption A2 is essentially a homogeneity requirement. Over a small enough geographic range and time period, A2 is certainly reasonable. For example, on a particular stretch of highway over a fifteen minute period, there may be little reason to think that drinking and sober drivers are not equally mixed. As the unit of observation expands either spatially or with respect to time, this homogeneity condition clearly becomes suspect. We devote a great deal of attention to possible violations of A2 and their impact on the results in the empirical section of the paper.

The third assumption of the model is

(A3) A fatal car crash results from a single driver's error

A3 rules out the possibility that each of the drivers shares some of the blame for a crash. As discussed later in this section, while assumption A3 is critical to the identification of the model, it is possible to sign the direction of bias introduced by A3 if, in fact, both drivers contribute to fatal crashes.

The fourth modeling assumption is

(A4) The composition of driver types in one fatal crash is independent of the composition of driver types in other fatal crashes.

This assumption allows us to move from individual crash probabilities to probabilities involving multiple crashes. Given the level of aggregation used in the empirical analysis (e.g. weekend nights between the hours of midnight and 1 a.m. in a given state and year), there is little reason that this assumption should fail, although for very localized observations (a short stretch of road

over a fifteen minute time period) it may be less applicable.

The final assumption required to solve the model is that drinking (weakly) increases the likelihood that a driver makes an error resulting in a fatal two-car crash. Denote the probability that a driver of type i makes a mistake that causes a fatal two-car crash as θ_i . Formally,

$$(A5) \quad q_D \geq q_S$$

the existing evidence concerning the relative crash risk of drinking and sober drivers overwhelmingly supports this assumption (e.g. Linnoila and Mattila 1973, Borkenstein et al. 1974, Dunbar et al. 1987, Zador 1991).

Fatal crash data and the parameters of interest

Having laid out the assumptions of the model, we derive the link between fatal crash data and the parameters of interest in three steps. First, we derive the formulas corresponding to the likelihood of two cars interacting with one another. Second, we determine the likelihood of a crash conditional both on the drivers' types and an interaction taking place between two cars. We then use Bayes' Theorem to back out the probability that a given pair of driver types will be involved, conditional on a fatal crash occurring. Third, we derive the likelihood function and discuss the identification issues involved in its estimation.

Assumption A2 gives the joint distribution for a pair of driver types, conditional on an interaction between two drivers:

$$\Pr(i, j|I = 1) = \frac{N_i N_j}{(N_D + N_S)^2} \quad (1)$$

where i and j are drivers of a particular type, i.e. either drinking or sober. So, for example, given that an interaction occurs between two cars, the probability that both are sober drivers is $(N_S)^2/(N_D+N_S)^2$. Interactions between drivers in this model, as reflected in equation (1), are logically equivalent to randomly drawing balls labeled either “S” or “D” out of an urn.

Define A to be an indicator variable equal to one if there is a fatal accident and equal to zero otherwise. Assumption A3 implies that the conditional probability of a fatal two-car crash given that two drivers of type i and j pass on the road is

$$\Pr(A = 1|I = 1, i, j) = q_i + q_j - q_i q_j \approx q_i + q_j \quad (2)$$

The likelihood of a fatal crash is the sum of the probabilities that either driver makes a fatal error, minus the probability that both drivers make a mistake. Given that the chance of either driver making a fatal mistake is *extremely* small, the chance of both drivers making an error is vanishingly small and can be ignored.⁴ So, for example, given that two sober drivers pass each other on the road, the probability of a fatal accident is $2\theta_s$. More generally we could allow for heterogeneity in driver risk within each driver type, as discussed in the extensions to the basic model. In that case, θ_i and θ_j in equation (2) represent the *mean* driver risk for the population of

⁴ There are roughly 13,000 fatal two-car crashes in the United States annually. The total number of vehicle miles driven is approximately 2 trillion. If every car interacted with an average of five other cars per mile, then the implied θ_i is on the order of 10^{-9} and the interaction term is on the order of 10^{-18} .

drivers of type i and j on the road.

Multiplying equations (1) and (2), the joint probability of driver types and a fatal crash conditional on an interaction between two drivers is as follows

$$\Pr(i, j, A = 1 | I = 1) = \frac{N_i N_j (q_i + q_j)}{(N_D + N_S)^2} \quad (3)$$

In words, given that two random drivers interact, the probability that a fatal crash occurs and that the drivers involved are of the specified types, is simply equal to the likelihood that two drivers passing on the road are of the specified types multiplied by the probability that a fatal crash occurs when these drivers interact.

The key relationship that we seek is the probability of driver types conditional on a fatal accident occurring, rather than conditional on an interaction taking place. That value can be obtained from equation (3) through an application of Bayes' Theorem:

$$\Pr(i, j | A = 1) = \frac{\Pr(i, j, A = 1 | I = 1)}{\Pr(A = 1 | I = 1)} = \frac{N_i N_j (q_i + q_j)}{2[q_D(N_D)^2 + (q_D + q_S)N_D N_S + q_S(N_S)^2]} \quad (4)$$

Although the expression in equation (4) looks somewhat complicated, it is in fact quite straightforward. For each combination of driver types, the numerator is proportional to the number of fatal crashes involving those two types. The denominator is a scaling factor assuring that the sum of the probabilities add to one.

Let P_{ij} represent the probability that the drivers are of type i and j given that a fatal crash

occurs. We can explicitly state the values of P_{ij} by simply substituting for i and j in equation (4):

$$P_{DD} = \Pr(i = D, j = D | A = 1) = \frac{q_D(N_D)^2}{q_D(N_D)^2 + (q_D + q_S)N_D N_S + q_S(N_S)^2} \quad (5)$$

$$\begin{aligned} P_{DS} &= \Pr(i = D, j = S | A = 1) + \Pr(i = S, j = D | A = 1) \\ &= \frac{(q_D + q_S)N_D N_S}{q_D(N_D)^2 + (q_D + q_S)N_D N_S + q_S(N_S)^2} \end{aligned} \quad (6)$$

$$P_{SS} = \Pr(i = S, j = S | A = 1) = \frac{q_S(N_S)^2}{q_D(N_D)^2 + (q_D + q_S)N_D N_S + q_S(N_S)^2} \quad (7)$$

Note that the ordering of the driver types does not matter. Consequently in equation (6), the probability of a mixed drinking-sober crash is the sum of the probability that i is sober and j is drinking plus the probability that j is sober and i is drinking.

Examination of equations (5)-(7) reveals that there are only three equations, but four unknown parameters (θ_D , θ_S , N_D , N_S). Consequently, all four parameters cannot be separately identified. Closer examination of equations (5)-(7) reveals that only the ratios of the parameters could possibly be identified. Therefore, let $\theta = \theta_D / \theta_S$ and $N = N_S / N_D$. θ is the relative likelihood that a drinking driver will cause a fatal two-car crash compared to a sober driver, and N is the ratio of sober to drinking drivers on the road at a particular place and time.⁵ Expressing equations

⁵Since we have three observable pieces of data (the number of drinking-drinking, drinking-sober, and sober-sober crashes), one might expect that we might be able to do better than to identify only two parameters, the ratios θ and N . In fact, although there are three equations, the

5-7 in terms of θ and N yields

$$P_{DD}(q, N|A) = \frac{q}{q + (q + 1)N + N^2} \quad (8)$$

$$P_{DS}(q, N|A) = \frac{(q + 1)N}{q + (q + 1)N + N^2} \quad (9)$$

$$P_{SS}(q, N|A) = \frac{N^2}{q + (q + 1)N + N^2} \quad (10)$$

The final step is deriving the likelihood function. The values in equations (8)-(10) provide the likelihoods of observing the various combinations of driver types conditional on a crash occurring. Using assumption A4, which provides independence across fatal crashes, and given the total number of fatal crashes, the joint distribution of driver types involved in fatal accidents is given by the multinomial distribution. Defining A_{ij} as the number of fatal crashes involving one driver of type i and one driver of type j and A_{total} as the total number of fatal crashes

$$\Pr(A_{DD}, A_{DS}, A_{SS}|A_{total}) = \frac{(A_{DD} + A_{DS} + A_{SS})!}{A_{DD}! A_{DS}! A_{SS}!} (P_{DD})^{A_{DD}} (P_{DS})^{A_{DS}} (P_{SS})^{A_{SS}} \quad (11)$$

Substituting the solutions to P_{DD} , P_{DS} , and P_{SS} into equation (11) yields the likelihood function for

three equations are linearly dependent (i.e. the equations sum to one), so in practice only two parameters can be identified.

the model. In the empirical work that follows, we perform maximum likelihood estimation of equations (8)-(11), directly linking our empirical estimation strategy to the model. Note that maximization of equation (11) with respect to P_{DD} , P_{DS} , and P_{SS} yields the following intuitive result:

$$\hat{P}_{DD} = \frac{A_{DD}}{A_{total}}, \hat{P}_{DS} = \frac{A_{DS}}{A_{total}}, \text{ and } \hat{P}_{SS} = \frac{A_{SS}}{A_{total}} \quad (12)$$

In other words, the maximum likelihood estimate of the fraction of crashes involving two drinking drivers is simply the observed fraction of such crashes in the data.

In order to solve the model for θ , we take the following ratio in which N , the the ratio of sober to drinking drivers cancels out:

$$\frac{(A_{DS})^2}{A_{DD}A_{SS}} = \frac{(\hat{P}_{DS})^2}{\hat{P}_{DD}\hat{P}_{SS}} = \frac{(q+1)^2 N^2}{qN^2} = 2 + q + \frac{1}{q} \quad (13)$$

It is worth pausing here to note the significance of equation (13), which says that it is possible to determine the relative crash risk of drinking and sober drivers (θ) based *solely* on the observed distribution of fatal crashes. The key to this result is the cancellation of the N term. N disappears from equation (13) because, by the binomial distribution, the squared number of interactions between drinking and sober drivers is in fixed proportion to the product of drinking-drinking and sober-sober interactions. Therefore, information on the relative number of drinking and sober drivers on the road is not needed to identify the model.

Defining the value in the left-hand-side of equation (13) as R ,

$$R \equiv \frac{(A_{DS})^2}{A_{DD}A_{SS}} \quad (14)$$

then rearranging equation (13) and multiplying through both sides by θ yields an equation that is quadratic in θ :

$$q^2 + (2 - R)q + 1 = 0 \quad (15)$$

the solutions to which are

$$q = \frac{(R - 2) \pm \sqrt{R^2 - 4R}}{2} \quad (16)$$

Ignore for the time being values of $R < 4$, which do not yield a real-valued solution for θ . When $R=4$, the only solution is $\theta=1$, a limiting case where drinking drivers pose no greater risk of causing two-car crashes than sober drivers. Note that for $R=4$, the distribution of fatal crashes precisely matches the distribution of crash opportunities as given by the binomial distribution. The only circumstance under which this occurs is when the crash likelihoods are equal. For all $R > 4$, there are two real solutions: one with $\theta > 1$ and the other with $\theta < 1$. By assumption A5, which requires drinking drivers to be at least as dangerous as sober drivers, we select the first of those two solutions. Having obtained a solution for θ , it is straightforward to back out the relative number of drinking and sober drivers. Standard errors for both θ and N are readily attainable from the Hessian of the likelihood function.

Having obtained a solution for θ , it is straightforward to back out the relative number of drinking and sober drivers. In particular, the first order condition from maximum likelihood estimation of equation (11) provides a solution for N in terms of θ :

$$\frac{N}{q} = \frac{A_{DS} \left(\frac{1}{1+q} \right) + A_{SS}}{A_{DS} \left(\frac{q}{1+q} \right) + A_{DD}} \quad (17)$$

Using the definitions of N and θ , the left-hand side of (17) is $\frac{N_s q_s}{N_D q_D}$. This ratio is the relative

number of two-car crashes caused by sober drivers relative to the number caused by drinking drivers. The right-hand side of equation (17) expresses that ratio solely in terms of θ and the observed distribution of crashes. Sober drivers cause $1/(1+\theta)$ of the crashes between drinking and sober drivers and all crashes involving two sober drivers; drinking drivers cause $\theta/(1+\theta)$ of the crashes between drinking and sober drivers and all crashes involving two drinking drivers.

Now consider the case where $R < 4$. A non-real valued solution to equation (16) emerges. Values of $R < 4$ are not consistent with the binomial distribution, i.e. there is no combination of θ and N that can generate this outcome. Looking at equation (14), low values of R result when there are too few drinking-sober crashes. Such values of R may arise in practice either because of small numbers of observed crashes, or due to a violation of the equal mixing assumption, as will be discussed below. It is important to note, however, that observed values of $R < 4$ do not invalidate the maximum likelihood estimation. When $R < 4$, the maximum likelihood estimate of θ is one and the maximum likelihood estimate of N is the observed ratio of sober and drinking drivers involved in two car crashes. Note that regardless of the value of R , standard errors for both θ and N can be computed.

Figure 1 provides a sense of how the estimates of θ and N vary with the distribution of two-car crashes. The y-axis in Figure 1 reflects θ , the relative crash likelihood of drinking drivers. The x-axis is the number of mixed drinking-sober crashes (A_{DS}). The two curves plotted in the figure correspond to alternative numbers of drinking-drinking and sober-sober crashes. In the top line is the case where A_{DD} and A_{SS} (the number of fatal crashes involving two drinking and two sober drivers) are 20 and 80 respectively. In the bottom line, both A_{DD} and A_{SS} are held constant at 50. Holding constant the number of drinking-drinking and sober-sober crashes, the estimated θ increases roughly in proportion to the square of the number of crashes involving one sober and one drinking driver, although over the relevant range the relationship appears almost linear since the constant of proportionality ($1 / A_{DD}A_{SS}$) is so small. The flat portion of the curves correspond to values of $R < 4$ described in the preceding paragraph. While the shape of the curves are similar, note that for a given number of mixed drinking-sober crashes, the implied θ is much higher when the number of drinking-drinking crashes is lower. The intuition for this result is that holding constant drinking-sober crashes, a low number of drinking-drinking crashes implies fewer drinking drivers on the road since the latter are a function of the square of the number of drinking drivers. In order for a small number of drinking drivers to crash into sober drivers so frequently, the drinking drivers must be at high risk to cause a crash.

Incorporating one-car crashes

The model above is derived without reference to one-car crashes. Intuition might suggest that there would be useful information provided by such crashes. In fact, that intuition turns out to be correct only to a limited degree. Since one-car crashes lack the interactive nature of two-

car crashes which provides identification of the model, there is relatively little to be gained by adding one-car crashes to the model.

Denote λ_D and λ_S as the probabilities that drinking and sober drivers make an error resulting in a fatal one-car crash, paralleling the θ_D and θ_S terms in the two-car case.

Let C be an indicator for the presence of a one-car crash, i.e. C equals one if a one-car crash occurs and zero otherwise. Then, letting Q_D and Q_S denote the respective probabilities that a drinking or sober driver is involved in a given one-car crash:

$$Q_D = \Pr(i = D|C = 1) = \frac{I_D N_D}{I_D N_D + I_S N_S} \quad (17)$$

$$Q_S = \Pr(i = S|C = 1) = \frac{I_S N_S}{I_D N_D + I_S N_S} \quad (18)$$

Defining $\lambda = \lambda_D / \lambda_S$, then the ratio of equations (17) and (18) can be rewritten as

$$\frac{O_D}{O_S} = I / N \quad (19)$$

Thus, unlike the two-car case, it is impossible to separately identify the parameters of the model using only one-car crashes. Algebraically, adding equation (19) to the two-car crash model provides one additional equation and one extra unknown. Since N is identified from two-car crashes, it is possible nonetheless to back out estimates of λ . Note, however, that it is the identification coming off of the two-car crashes that is critical to obtaining that parameter, and

adding the one-car crashes does not affect the solution to the two-car case.⁶

Relaxing the Assumptions

Before proceeding to the data, it is worth considering how each of the assumptions made influence the solution to the model, possible alternative assumptions, as well as the likely direction of bias in the estimation induced by violation of the assumptions.

The assumption of exactly two types is relatively innocuous. The introduction of more types allows for greater differentiation of estimates for particular sub-groups of the population (e.g drinking teenagers or sober motorists with clean driving records). From a practical perspective, however, the number of two-car crashes is not great enough to support more than a small number of categories. In the presence of heterogeneity within categories (i.e. there is a distribution of driving abilities among drinking and sober drivers), the estimates obtained from the two-type model are nonetheless readily interpretable. The estimates obtained are weighted averages across drivers in the category, with weights determined by the number of drivers of each ability *on the road*.⁷ This is a surprising and useful result. Given that we only observe drivers in

⁶ In the empirical estimation, we include one-car crashes for two reasons. First, we are interested in estimating λ , even if identification hinges on two-car crashes. Secondly, in our empirical estimates, we will generally impose equality restrictions across θ for different geographic areas or time periods. Once such restrictions are imposed, one-car crashes are useful in estimating the parameters of the model.

⁷Suppose driver types D and S have distributions of driver risk given by $f_D(q)$ and $f_S(q)$. Then the probability of a fatal accident given only the types i and j of two interacting drivers is

$$\begin{aligned} \Pr(A = 1 | i, j, I = 1) &= \int (q_i + q_j - q_i q_j) f_{i,j}(q_i, q_j) dq_i dq_j \\ &= E(q_i) + E(q_j) - E(q_i q_j) \approx E(q_i) + E(q_j) \end{aligned}$$

fatal crashes, intuition might suggest that the weights would be based on the distribution of abilities among crashers rather than drivers as a whole. If our coefficients reflected only the abilities of those in crashes, then the results would be much less useful for public policy since the distribution of crashers is likely to be very different than the underlying distribution of drivers.⁸

The second assumption, equal mixing/homogeneity on the road, is much more critical to the results. The likely manner in which this assumption will be violated is through spatial or temporal clumping of similar types of drivers. For instance, if the unit of observation were all crashes in the United States in a given year, then it is clear that drinking drivers would not be randomly distributed, but rather, concentrated during nighttime hours and especially weekend nights. Even within a smaller unit of analysis (e.g. weekends between midnight and 1 a.m. in a particular state and year) there may be clumping. Roads near bars may contain a higher fraction of drinking drivers, or the proportion of drinking drivers may rise sharply immediately following bar closings.⁹ Such a non-random distribution of drivers will result in a greater number of drinking-drinking and sober-sober interactions than predicted by the binomial distribution, with correspondingly fewer drinking-sober interactions. Looking at equation (14), non-random mixing will lead to smaller values of R and consequently a downwardly biased estimate of θ ,

Thus with heterogeneity we can simply reinterpret θ as the mean driver risk for the given types.

⁸ In the presence of heterogeneity, our simple model can do no better than to identify the ratio of the means of the distributions. Identification of higher moments of the distributions would require imposing arbitrary parametric assumptions.

⁹ It is at least theoretically possible that this assumption could be violated in the opposite way, i.e. drinking drivers being less likely to interact with other drinking drivers. For instance, if all drinking drivers are traveling North on a two-lane highway and all sober drivers are traveling South, then drinking and sober drivers will disproportionately interact with the opposite type.

which is an increasing function of R . Conversely, N , the ratio of sober to drinking drivers, will be biased upward. Violations of this assumption are likely to be less extreme as the geographic and temporal units of analysis shrink. In fact, that is precisely the pattern revealed by the empirical estimates in Section IV.

The third assumption requires that one driver be wholly at fault in a fatal crash, rather than both drivers sharing some fraction of the blame. A simple alternative model would allow both drivers to play a role in the crash: if one driver makes a fatal error, the second driver has an opportunity to take an action to avoid the crash. Let μ denote the relative inability of drinking drivers to avoid a fatal crash that another driver initiates, with $\mu \geq 1$. It is straightforward to demonstrate that this more general model yields a solution $R = 2 + \theta/\mu + \mu/\theta$, paralleling equation (13), but with θ/μ replacing θ . In our basic model μ is implicitly set equal to one because there is no scope for crash avoidance. If in actuality drinking drivers are less proficient in avoiding crashes initiated by the other driver, then the parameter identified in our model is actually $\theta/\mu < \theta$. Allowing for sober drivers to be more skilled in averting potential fatal crashes yields a larger estimated value of θ for any observed distribution of crashes, implying that the estimates obtained in this paper for θ are lower bounds on the true value.¹⁰

¹⁰ An even more general model would also allow the seriousness of mistakes made by drivers to vary by type. Let $\delta > 1$ capture the additional difficulty of avoiding a mistake made by a drinking driver. The parameters that we identify in two-car and one-car crashes respectively in this expanded model are $\theta\delta/\mu$ and λ/μ respectively. $\theta\delta$ is the seriousness-weighted ratio of drinking to sober driver mistakes – precisely what we are attempting to capture in the model, thus this further generalization does not pose any problems to our estimation. It is worth noting that δ does not appear in the solution for one-car crashes. If δ were large, one would expect to see big differences in the empirical estimates for one-car and two-car crashes. In practice, the two sets of estimates are close, perhaps suggesting that $\delta \approx 1$.

One could also imagine other possible models of fatal crashes. For instance the probability of a crash (conditional on an interaction occurring) might be modeled as $\theta_1\theta_j$. In other

The final two assumptions of the model, independence of crashes and that drinking drivers pose a higher risk, are unlikely to impose any important biases on the empirical estimates. The existing evidence overwhelmingly confirms the assumption that drinking drivers have greater likelihoods of fatal crash involvement. Given the unit of analysis of the paper, the independence assumption appears quite reasonable. Even if this assumption were to fail (i.e. the presence or absence of one crash influences other potential crashes), there is no reason to expect that violation of this assumption should systematically bias the estimates in one direction. For bias to occur, the presence of one crash has to differentially affect drinking-sober crashes relative to drinking-drinking or sober-sober crashes.

Section III: Data on Fatal Crashes

The primary source of data on fatal motor vehicle accidents in the United States is the Fatality Analysis Reporting System (FARS) administered by the National Highway Transportation Safety Administration. Local police departments are required by federal law to submit detailed information on each automobile crash involving a fatality. Compliance with this law was uneven until 1983, thus we restrict the analysis of the paper to the years 1983-1993. Because our primary interest is the impact of alcohol on driver risk, we limit our sample to those hours (8 p.m.-5 a.m.) when drinking and driving is most common. Our sample includes over 100,000 one-car crashes and over 40,000 two-car crashes. During these hours, almost 60 percent of drivers involved in fatal crashes have been drinking, compared to less than 20 percent of drivers at all other times of

words, each driver has to make a mistake in order for a crash to occur. In such a model, the solution to equation (13) is $R=4$ for all observed crash distributions (a result strongly rejected by the data) and only the composite parameter θ/N is identified.

the day. Crashes involving three or more drivers, which represent less than 6 percent of fatal crashes, are dropped from the sample. All crashes from a handful of state-year pairs with obvious data problems are also eliminated.

Among the variables collected in each fatal crash is information on the time and location of the accident and whether the drivers involved were under the influence of alcohol. In extensions to the basic model, we also utilize information on the age, sex, and past driving record of those involved in fatal crashes. Therefore, we exclude from the sample any crash where one or more of the drivers are missing information about the time or location of the accident, police reported drinking status, age, sex, or past driving record. Combined, these missing data lead to the exclusion of roughly 8 percent of all crashes.

Two measures of alcohol involvement are included in FARS. The first of these is the police officer's evaluation of whether or not a driver had been drinking. The officer's assessment may be based on formal breath, blood, or urine test results, or other available evidence such as a driver's behavior (for those drivers not killed in the crash) or alcohol on the driver's breath. The primary advantage of this measure is that it is available for virtually every driver involved in a fatal crash. There are at least two drawbacks of this variable. First, it does not differentiate between varying levels of alcohol involvement. In particular, no distinction is made between those drivers who are legally drunk and those who have been drinking but are below the legal limit. Second, the measure is often subjective and relies on the discretion of the police at the scene of the accident.

In spite of these shortcomings, the police officer's assessment of whether or not a driver has been drinking serves as our primary measure of alcohol involvement. As a consequence, the

coefficients we obtain with respect to the elevated risk associated with drinking and driving are based on the entire population of drinking drivers, not just the sub-group of legally drunk drivers. As a check on the results obtained, we also examine a second measure of alcohol involvement which is measured blood alcohol content (BAC). The major problem with this variable is the frequent failure to conduct such tests, despite federal law mandating that all drivers involved in fatal crashes be tested. Evidence suggests that the likelihood of BAC testing is an increasing function of actual blood alcohol levels, suggesting that this measure will be most flawed for low BAC motorists. Thus, for the purpose of analysis, we compare two groups of drivers: (1) those with measured BAC greater than 0.10 percent (the legal limit in most states for most of the sample period) and (2) those who test free of alcohol or who are not tested, but who police describe as not having been drinking. We eliminate those who test positive for alcohol, but who are below the 0.10 threshold. Furthermore, because sample selection in the pool of drivers who are tested for BAC is a major concern when using this measure, we exclude all crashes occurring in states that do not test at least 95 percent of those judged to have been drinking by the police in our sample in that year (regardless of whether the motorist in question was tested). This requirement excludes more than 80 percent of the fatal crashes in the sample.

The presence of classification error in each of our measures of drinking, like the violations of the model's assumptions in the previous section, will almost certainly lead to an underestimate of the relative fatal crash risk of drinking drivers, making our estimates conservative. This point is discussed in detail in the next section.

The time series of fatalities in all motor-vehicle crashes rises from an initial value of 42,589 in 1983 to a peak of just over 47,000 in 1988, and then declines to roughly 40,000 by

1993. The percentage of deaths occurring in crashes with at least one drinking driver steadily falls over the sample from 55.5 percent in the beginning to 43.5 percent in the end.

Table 1 presents means for the data in our sample.¹¹ Slightly more than half of all drivers involved in fatal crashes are reported to be drinking by police. Drivers in fatal crashes are mostly male (83.0 percent), somewhat less than half are under the age of 25, and about one-third qualify as having bad previous driving records under our definition: two or more minor blemishes in the last three years (any combination of moving violations and reported accidents), or at least one major blemish (driving while intoxicated conviction, license suspension/revocation). The fraction of drivers who are both male and drinking (49 percent) is higher than would be expected based on the marginal distributions if gender and drinking were uncorrelated (i.e., $.569 * .830 = .472$). This implies that males involved in fatal crashes are more likely to have been drinking than are females. Young drivers and those with bad previous driving records are also more likely to have been drinking.

In one-car crashes, 66.2 percent of drivers are classified as drinking by the police. 13.7 percent of two-car crashes involve two drinking drivers, 53.4 percent have exactly one drinking driver, and in the remaining 32.9 percent of cases neither driver was drinking. When we restrict our sample to states with good BAC reporting practices, 56 percent of drivers in one-car fatal crashes have a BAC over .10. In two-car crashes, there are two drivers with BACs over the legal limit 5.4 percent of the time, and exactly one driver over the limit in 45.6 percent of the crashes.

¹¹Because the degree of aggregation in our analysis varies from hour-year to hour-year-state-weekend, standard deviations, minimums, and maximums are not particularly meaningful.

Section IV: Estimation of the Model

Table 2 presents maximum likelihood estimates of equation (11), focusing on how the relative fatal crash risk for drinking drivers in two-car crashes (θ) and one-car crashes (λ) are affected as we increasingly disaggregate the data. Each column represents a different specification, with the distinction between columns being the unit of observation over which “equal mixing” of drivers is assumed. Column 1 is the most restrictive, imposing equal mixing of all drivers in all years, locations, and hours of the day. Equal mixing is unlikely to hold at such a high level of aggregation. This restriction is continually relaxed moving from left to right in the table. In Column 3, for instance, equal mixing is imposed for crashes in each hour-year pair in the sample (e.g. it is assumed that between 2 a.m. and 3 a.m. in 1991, drinking and sober drivers are equally mixed across all locations in the United States). By column 8, equal mixing is assumed only within a given hour and weekend status (equal to one on Friday or Saturday night and zero otherwise) for a particular state and year.¹² Theory predicts that all of the estimates presented in the table are likely to be lower bounds on the true parameters, but that the downward bias will be mitigated moving from columns 1 to 8. In fact, that is precisely the pattern observed in the table. The two-car fatal crash risk of drinking drivers rises monotonically from four times greater than sober drivers to eight times greater.¹³ For one-car crashes, the value goes from almost six to

¹² Stated more formally, in maximum likelihood estimation of equation (11) in column 8, we restrict θ and λ to be the same for all observations, but allow N to vary by state x year x hour x weekend.

¹³ To the extent that the pool of drinking drivers, when they are themselves sober, are systematically safer/more dangerous than other sober drivers, this coefficient gives a biased estimate of the incremental effect of alcohol on driver risk. Previous research (Hurst et al. 1994) suggests that drinking drivers, when sober, are actually safer than the typical sober driver. If that is the case, then the coefficients in Table 2 understate the true impact of alcohol on driver risk.

more than nine times larger. In all cases, the parameters are precisely estimated and the null hypothesis of equality between drinking and sober drivers (i.e. $\theta=1$, $\lambda=1$) is resoundingly rejected. The restrictions implied by the specifications in columns 1-7 relative to column 8 are each rejected using the relevant likelihood ratio test. For this reason, and because theory predicts that less restrictive specifications should minimize the downward bias, we take column 8 as our preferred specification.¹⁴ All results presented in the remaining tables correspond to column 8.

Although the parameters are not directly comparable to past estimates in the literature, it is nonetheless useful to consider relative magnitudes. Zylman (1973) finds that the relative crash risk of those with positive levels of alcohol is only 2.2 times that of sober drivers, much smaller than our estimates. Borkenstein et al. (1964) finds the relative likelihood of causing a fatal crash to be two times higher for drivers with BAC between .05-.099, 10.1 times higher for BACs between .10-.149, and 30-40 times greater for BACs over .15. For reasonable distributions of BACs across drinking drivers, the results of Table 2 are consistent with the Borkenstein et al (1964) estimates. Zador (1991), comparing FARS data to the results of the national roadside testing survey of Lund and Wolfe (1991), estimates that .05-.09 BACs are associated with nine times greater risk, .10-.15 BACs are associated with 50 times greater risk, and BACs above .15 have 300-600 times greater risk. These estimates are far greater than those that we obtain.

Table 3 presents separate estimates of relative crash risk across the years of our sample.

Estimates that further disaggregate drivers by type, presented below, will also shed light on this issue.

¹⁴ One could imagine further disaggregating the data. The main stumbling block is lack of data. By column 8, we are allowing over 6,000 different cells in which there are approximately 43,000 two-car crashes.

For two-car crashes, the parameter estimates are fairly stable. A test of the null hypothesis of equality across all of the years, reported in the bottom of the table, cannot be rejected. The one-car crash estimates show a puzzling break between 1986 and 1987, but are otherwise stable. Given that each year's estimates are derived independently, the stability of the parameters suggests that the estimation approach is robust.

The final column of Table 3 presents the implied fraction of drivers in the sample who have been drinking.¹⁵ To the extent that estimates of the relative risk of drinking drivers are downward biased in columns 1 and 2, the fraction of drinking drivers will be biased upward. The estimates of drinking drivers ranges from a high of 21.4 percent in 1984 down to a low of 15.3 percent in 1993. There is a discernible downward trend in the fraction of drinking drivers over the period. The null hypothesis of a constant share of drinking drivers across all years is rejected at the .01 level.¹⁶

Table 4 breaks down the estimates by hour of the day. Once again, the relative risk for two-car fatal crashes is stable, and equality across all hours cannot be rejected. One-car crashes are again less stable. The relative risk of drinking drivers appears to be somewhat lower between 8 and 10 p.m., perhaps as a result of a less lethal composition of drinking drivers on the roads

¹⁵ These estimates are derived in two steps. First, θ and λ are estimated, allowing the ratio of sober to drinking drivers to vary by state-year-hour-weekend. Then θ and λ are fixed at the MLE estimates from the first step and the optimal N is estimated. N is then transformed into the percentage of drinking drivers with the appropriate standard errors calculated using the delta method.

¹⁶ The estimated fraction of drinking drivers is broadly consistent with numbers obtained from other approaches. In a national roadside survey, Lund and Wolfe (1991) report 8.3 percent of those who provide BACs had alcohol levels greater than .05. If all of those who refused to provide BACs had been drinking, then the total fraction of drinkers in their sample could be as high as 16.3 percent.

during these early hours. The peak hours for drinking and driving are between 1 a.m. and 3 a.m., as reported in column 3. More than one-quarter of drivers appear to have some alcohol in their system during these hours. This is more than twice as high as before 10 p.m. and after 4 a.m.

Table 5 reports the sensitivity of the basic results on drinking and driving to the alternative assumptions discussed in Section III as well as to various forms of measurement error. The top row of the table is the baseline estimates for comparison purposes. Each of the remaining rows reflects a different violation of the assumptions. If crash avoidance matters in two-car crashes (violating A3) and drinking drivers are 25 percent less successful in getting out of the way ($\mu=1.25$), then the estimated impact of alcohol rises to ten in two-car crashes and eleven in one-car crashes. Since the identified parameters are θ/μ , λ/μ , and N/μ , the one- and two-car driver risks are simply scaled up by the relative crash avoidance factor μ .¹⁷

The likely manner in which violations of equal mixing (A2) would occur would be through an increased probability of same type interactions beyond that given by the binomial distribution. To consider such violations, we augment the interaction stage of the model. If N_D/N is the fraction of drinking drivers on the road, then in the augmented model the probability that two drivers passing are both drinking drivers is $(1+\Delta)(N_D/N)^2$ rather than $(N_D/N)^2$, i.e. drinking drivers are Δ times more likely to interact with one another than would be suggested by the binomial distribution. Sober drivers are assumed to be precisely enough more likely to interact with sober drivers to maintain the assumption that the overall fraction of interactions by driver type reflects

¹⁷ Although one-car crashes are not directly affected by violations in this assumption, the one-car crash estimates rely on the fraction of drinking drivers on the road which is identified through the two-car crashes. As a result, λ/μ is identified rather than λ , paralleling θ/μ 's identification rather than θ .

the fraction of drivers on the road. Given any value of Δ corresponding to some degree of unequal mixing, we can estimate this expanded model in order to assess the sensitivity of our estimates to violations of A2. Increased clustering of same-type drivers corresponds to lower predicted values of R in the model. To offset this effect and explain the observed R in the data, θ must increase (and N decrease). The result in the third row of Table 5 shows that a ten percent increase in drinking-drinking interactions would lead to an approximately twenty-five percent increase in the estimates of one- and two-car relative risks.

An implicit assumption of the model is that driver types are known. In fact, as discussed in section III, we use the police officers' assessment of the drinking status of involved drivers, which is a potentially imperfect measure. The remaining rows of Table 5 examine the sensitivity of our estimates to three alternative forms of driver misclassification. The fourth row of the table corresponds to a case in which five percent of the observations are misrecorded due to classical measurement error (i.e., a drinking driver is as likely to be mistakenly reported as sober as vice-versa), the estimated risk of drinking drivers will be biased towards zero by approximately forty percent. The sensitivity of our results to classical measurement error is due to the fact that the small number of drinking-drinking crashes have the greatest impact on the estimates. Given the rarity of drinking-drinking crashes relative to drinking-sober crashes, classical measurement error will exaggerate the number of drinking-drinking crashes, i.e. the true number of drinking-drinking crashes is even smaller than that observed in the data. Previous research suggests, however, that classical measurement error is not the most likely scenario. Lund and Wolfe (1991) present evidence that police officers systematically report drinking drivers to be sober, but not vice-versa. If five percent of drinking drivers are mis-reported in this manner, then the true θ , although still

biased, is not as sensitive as is the case with classical measurement error. A final misclassification story assumes that reporting errors are correlated within a particular crash. Different police officers may have divergent standards for classifying an individual as drinking due to the methods used to determine drinking status or the officer's skill in identifying drinking drivers. Referring back to equations (13) and (14), it is clear that such systematic measurement error will impart a large downward bias on θ because any misclassification of this type will shift crashes from the numerator to the denominator of R . For instance, if five percent of police officers filling out crash reports always report both drivers sober, even if they are drinking, then the true θ is almost thirty-five percent higher than our baseline estimate.

Examining Other Driver Traits

Although the comparisons presented thus far are limited to drinking versus sober drivers, the model is equally applicable to other comparisons. Table 6 presents results for a range of other dimensions. The top row of Table 6 presents the results from drinking and sober drivers as a baseline. The second row presents a comparison of drivers with BACs greater than .10 (as opposed to all drinking drivers) relative to sober drivers. As noted earlier, due to sample selection concerns we include only accidents in state-year pairs where a high fraction of drivers are tested. As would be expected the results for legally drunk drivers are even stronger than for all drinking drivers. Drivers over the legal BAC limit of .10 have a relative risk for fatal crashes that is 13 to 17 times higher than sober drivers. This number is almost twice as great as for all drinkers, including those who are not legally drunk. Approximately 7.5 percent of the drivers on the road during the hours we examine appear to be over the .10 limit, where this number is again

most appropriately interpreted as an upper bound. The estimates are relatively imprecise because many states have poor records of BAC testing, leading them to be excluded from the sample.

Rows 3-5 consider other categories into which drivers can be divided. Drivers under the age of 25 are 50-80 percent more likely to cause fatal crashes. The estimate of the fraction of young drivers on the road (34 percent) appears reasonable. Grossman et al. (1993) reports that those under the age of 25 accounted for 20 percent of drivers in 1984, but younger drivers are likely to make up a disproportionate fraction of night-time drivers. Males are 3.5 times as dangerous as female drivers for two-car crashes and are at two times higher risk for one-car fatal crashes. As will be demonstrated below, however, a large fraction of this gap is due to more frequent drinking and driving among males. Finally, those with bad driving records (either two or more minor blemishes on their driving record in the past three years, or one or more major blemish) appear to be more than twice as dangerous on the roads. Although Table 6 reports significant increases in risk associated with other driver characteristics, it should be stressed that alcohol usage is far and away the best predictor of increased crash risk.¹⁸

The results presented thus far fail to take into account possible interactions between the various risk factors such as drinking status, gender, and age. To the extent that risk factors are correlated with one another, the results of Table 6 may be misleading. Table 7 reports estimates that allow for interactions between drinking status and the other risk factors. In order to do so,

¹⁸ Consistent with our results, insurance premiums tend to be higher for young drivers, those with bad previous driving records, and male drivers. Auto insurance is highly regulated, however, so it is difficult to know if the magnitude of premium differences in an unregulated market would correspond to our estimates. A further complication is that insurance premiums take into account not only fatal crashes, but also non-fatal accidents and potential property damage.

we expand our model to allow for four driver types rather than two. All of the intuition from the two-type model presented in Section II continues to hold in the n -type case. The top four rows of Table 7 allow for differential fatal crash risk for young and old drivers who have or have not been drinking. Sober drivers over the age of 25 are the lowest risk drivers and serve as a baseline. Interestingly, young sober drivers are almost three times as likely to cause fatal crashes as older sober drivers, but age has little impact on fatal crash risk among drinking drivers. Most likely, this reflects the fact that young drinking drivers tend to have low BACs relative to older drinking drivers. In the last column, we estimate that roughly one in four young drivers has been drinking in our sample, compared to one in seven older drivers.

The middle four rows of Table 7 report the results of interacting drinking status with driver gender. Sober-female drivers are the safest. Sober-male drivers are 36 percent more likely to cause fatal two-car crashes and 10 percent more likely to cause fatal one-car crashes. Males who have been drinking are almost ten times more dangerous than sober females, and are a 60 percent greater risk than drinking females. Note how different these results are from the male vs. female results in the previous table which ignored the role of alcohol. There, males were more than three times as likely to cause a fatal crash. The results in Table 6 make clear that the gap between male and female drivers shrinks substantially when drinking status is also taken into account.

The last four rows of Table 7 show interactions between drinking status and past driving record. Sober drivers with bad past records are almost twice as likely to cause two-car fatal crashes as sober drivers with clean records. Interestingly, however, the impact of driving record shrinks substantially in percentage terms among those who have been drinking. Those with bad

driving records who have been drinking are roughly 25 percent more likely to cause fatal two-car crashes as drinking drivers with good records.

Section V: Conclusions

This paper presents a new methodology for measuring the relative risk of drinking drivers. Unlike previous approaches, our strategy requires only data on the observed number of fatal crashes. Although a priori it would seem impossible to separately identify relative crash risk from the fraction of drivers on the road that have been drinking using such limited data, we demonstrate that there is a hidden richness in the information contained in two-car crashes. As a consequence, with only the most skeletal assumptions we are able to identify the key parameters of interest. Drinking drivers (including those not legally drunk) are at least 8-9 times more likely to cause fatal crashes than are sober drivers, although sensitivity tests suggest that this lower bound could nontrivially understate the true value. For drivers with BACs greater than .10, that ratio is at least fourteen to one. In the presence of unequal mixing of drinking and sober drivers on the roads or measurement error in classifying drivers, these estimates represent relatively loose lower bounds. Peak hours for drinking and driving are between 1 a.m. and 3 a.m. The fraction of drinking drivers on the roads fell significantly over the sample period, but the relative risk of drinking drivers did not systematically change. Young drivers, men, and those with bad previous driving records are more dangerous on the roads, but none of these factors approach alcohol in terms of importance in explaining fatal crash risk.

The methodology we present provides a simple, flexible, inexpensive tool for analyzing driver risk. Unlike past survey methods which are expensive to administer and subject to

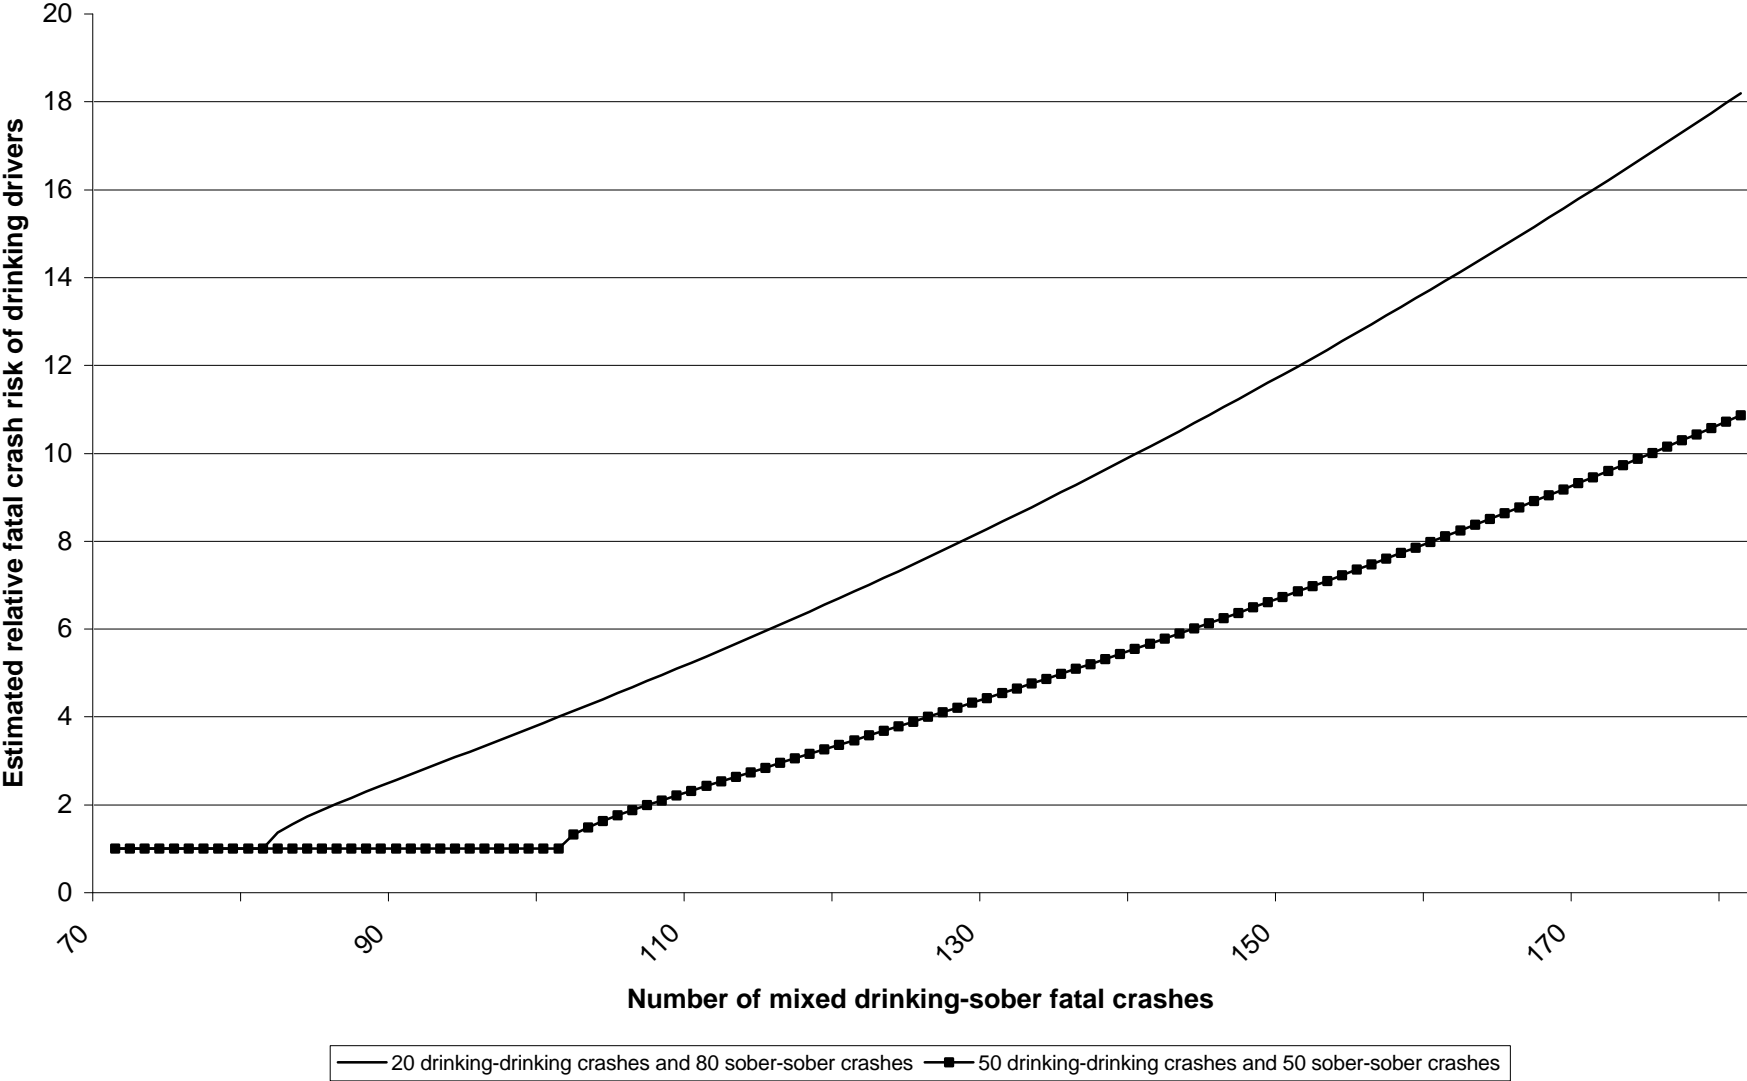
important non-participation biases, our approach is easily implemented using only fatal accident data that are readily available. Our method can be applied to any driver characteristic, any localized geographic area, and any time period. The variety of specifications presented suggest that the method is robust.

The estimates we obtain can be used to inform public policy debates on road safety. For instance, with further honing, it would be possible to determine fatal crash risk using our method for fine gradations of BAC levels. Such estimates are critical to development of the appropriate level at which to set legal BAC limits. Also, given the flexibility of our method, it is possible to generate estimates of the fraction of drinking drivers on the road and the relative fatal crash risk for a particular state and year. These estimates can then be used as left-hand-side variables in policy regressions similar to those in the literature (e.g. Cook and Tauchen 1982, Asch and Levy 1987, Saffer and Grossman 1987, Homel 1990, Chaloupka et al. 1993, Grossman et al. 1993, Ruhm 1996).¹⁹ Unlike past studies, which only capture a reduced-form relationship between policies and fatal crashes, our method makes it possible to differentiate between the different channels through which such reductions can occur (either fewer drinking drivers, or drinking drivers becoming relatively safer). For example, harsh penalties for a third drunk driving conviction may affect both channels by both deterring drinking drivers and by encouraging those who choose to drink and drive to drive more cautiously to avoid detection. On the other hand, raising alcohol taxes should only affect the number of drinking drivers, but not the care that they exert when driving.

¹⁹ Although these estimates are measured with error, as long as this error is not correlated with the right-hand-side variables in such a regression, the coefficients will not be biased (although the precision of the estimates will be reduced).

Perhaps more importantly, our approach to answering this question may provide insight to researchers on topics far afield from that which we have addressed in this paper. By exploiting the hidden richness of the data arising from interactions between participants, we are able to identify relationships in the data that may have seemed beyond reach *ex ante*. Variations on our approach may be useful in examining other issues where agents interact, e.g search models, the transmission of AIDS, oligopoly pricing, the relationship between special interest groups and politicians. The parallel between our model and these more complicated economic settings is not exact because the random matching in our traffic model is less plausible in other cases (e.g. Flinn and Heckman 1982). Nonetheless, with the adoption of an appropriate alternative matching model, extensions of this model may prove relevant.

Figure 1: Estimated Relative Risk of Drinking Drivers as a Function of Observed Crash Mix



Bibliography

- Asch, Peter, and David T. Levy, 1987, "Does the Minimum Drinking Age Affect Traffic Fatalities?" *Journal of Policy Analysis and Management* 6(2):180-192.
- Borkenstein, R. F., R. F. Crowther, R. P. Shumate, W. B., Ziel, and R. Zylman, 1974, "The Role of the Drinking Driver in Traffic Accidents," *Blutalkohol* 11(Supp. 1):1-132.
- Chaloupka, Frank J., Henry Saffer, and Michael Grossman, 1993, "Alcohol Control Policies and Motor Vehicle Fatalities," *Journal of Legal Studies* 22(January):161-186.
- Cook, Philip J., and George Tauchen, 1982, "The Effect of Liquor Taxes on Heavy Drinking," *Bell Journal of Economics* 13:379-390.
- Dunbar, James, A., Antti Penttila, and Jarmo Pikkarainen, 1987, "Drinking and Driving: Choosing the Legal Limits," *British Medical Journal* 295(December 5): 1458-1460.
- Fell, James C., and Carl E. Nash, 1989, "The Nature of the Alcohol Problem in U.S. Fatal Crashes," *Health Education Quarterly* 16(3) Fall: 335-343.
- Flinn, Christopher, and James Heckman, 1982, "New Methods for Analyzing Structural Models of Labor Force Dynamics," *Journal of Econometrics* 18:115-168.
- Grossman, Michael, Frank J. Chaloupka, Henry Saffer, Adit Laixuthai, 1993, "Effects of Alcohol Price Policy on Youth," NBER Working Paper No. 4385.
- Homel, Ross J., 1990, "Random Breath Testing the Australian Way: A Model for the United States?" *Alcohol, Health, and Research World* 14(1):70-75.
- Hurst, Paul M., David Harte, and William J. Frith, 1994, "The Grand Rapids Dip Revisited," *Accident and Analysis Prevention* 26(5):647-654.
- Lehman, R. J., Arthur C. Wolfe, and R. D. Kay, 1975, Development of a Computer Archive of Roadside Breathtesting Surveys in the United States, Ann Arbor, MI: University of Michigan Highway Safety Research Institute.
- Linnoila, M. and M. J. Mattila, 1973, "Drug Interaction on Driving Skills as Evaluated by Laboratory Tests and by a Driving Simulator," *Pharmakopsychiat* 6:127-132.
- Liu, Simin, Paul Z. Siegel, Robert D. Brewer, Ali H. Mokdad, David A. Sleet, and Mary Serdula, 1997, "Prevalence of Alcohol-Impaired Driving: Results from a National Self-reported Survey of Health Behaviors," *JAMA* 277(2): 122-125.

- Lund, Adrian K., and Arthur C. Wolfe, 1991, "Changes in the Incidence of Alcohol-Impaired Driving in the United States, 1973-1986," *Journal of Studies on Alcohol* 52(4):293-301.
- Rice, Dorothy P., Sander Kelman, and Leonard Miller, 1991, "The Economic Cost of Alcohol Abuse," *Alcohol Health and Research World*, vol. 15 (4): 307-316.
- Ruhm, Christopher, 1996, "Alcohol Policies and Highway Vehicle Fatalities," *Journal of Health Economics* 15 (4; August): 435-454.
- Saffer, Henry, and Michael Grossman, 1987, "Drinking Age Laws and Highway Mortality Rates: Cause and Effect," *Economic Inquiry* 25(July):403-417.
- Zador, Paul L., 1991, "Alcohol-Related relative Risk of Fatal Driver Injuries in Relation to Driver Age and Sex," *Journal of Studies on Alcohol* 52(4):302-310.
- Zylman, R., 1973, "Youth, Alcohol, and Collision Involvement," *Journal of Safety Research* 5:58-72.

Table 1: Summary Statistics for Fatal Crashes in the Sample
 One-and two-car crashes between 8 p.m. and 5 a.m., 1983-1993

Variable	Mean
Total number of fatal one-car crashes	101,701
Total number of fatal two-car crashes	41,068
Percent of all drivers in fatal crashes:	
Reported to be drinking by police	56.9
Male	83.0
Under age 25	44.3
Bad previous driving record	37.2
Reported to be drinking and male	49.0
Reported to be drinking and under age 25	26.8
Reported to be drinking and bad previous driving record	24.4
Percent of fatal one-car crashes with:	
One drinking driver	66.2
One sober driver	33.8
Percent of fatal two-car crashes with:	
Two drinking drivers	13.7
One drinking, one sober driver	53.4
Two sober drivers	32.9
Percent of fatal one-car crashes in restricted sample with high BAC reporting with:	
One legally drunk driver	56.0
One sober driver	44.0
Percent of fatal two-car crashes in restricted sample with high BAC reporting with:	
Two legally drunk drivers	5.4
One legally drunk, one sober driver	45.6
Two sober drivers	49.0

Notes to Table 1: Means are based on one- and two-car fatal crashes in Fatal Accident Reporting System (FARS) data for the years 1983-93 between the hours of 8 p.m. and 5 a.m. Drinking status is based on the police classification of drivers as drinking and includes drivers who were not legally drunk, except in the bottom portion of the table where the sample is restricted to crashes in state-year pairs where at least 95 percent of the drivers that police report to have been drinking were given blood-alcohol (BAC) tests. A bad driving record is defined as two or more minor blemishes (a moving violation or previous accident), or one or more major blemishes (previous DWI conviction, license suspension, or license revocation) in the last 5 years.

Table 2: The Relative Likelihood of Causing a Fatal Crash: Drinking vs. Sober Drivers
 Allowing for differing restrictions on the unit of observation
 over which “equal mixing” is imposed

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Relative two-car fatal crash risk for drinking drivers (θ)	3.92 (0.14)	4.91 (0.16)	5.11 (0.16)	5.35 (0.17)	5.57 (0.17)	6.01 (0.18)	6.84 (0.20)	8.04 (0.23)
Relative one-car fatal crash risk for drinking drivers (λ)	5.81 (0.13)	6.16 (0.13)	6.50 (0.14)	6.72 (0.14)	6.91 (0.15)	7.31 (0.15)	8.09 (0.17)	9.11 (0.19)
Unit of observation over which “equal mixing” of drivers is imposed	All data	hour	hour x year	hour x year x weekend	hour x region x year	hour x region x year x weekend	hour x state x year	hour x state x year x weekend
Degrees of freedom used in estimation	3	11	101	200	893	1778	3430	6652
log likelihood	-99584	-96912	-96349	-95758	-95262	-94258	-92408	-89991

Notes: Estimates in the table are the estimated relative one- and two-car crash risk for drinking versus sober drivers. The reported values are maximum likelihood estimates of equation (11). Drinking status is based on police assessments of the presence of alcohol; those who have been drinking, but are not legally drunk, are included in the drinking category. The unit of observation is state x year x day x hour. Data is for the years 1983-93 between the hours of 8 pm and 5 am, with some state-year observations omitted due to data problems. For a number of reasons presented in the paper, these estimates are likely to be lower bounds on the increased risk of drinking drivers. Moving from left to right in the table, the unit of observation over which the assumption of “equal mixing” of drinking and sober drivers imposed is relaxed. As this restriction is relaxed, theory predicts that the estimated coefficients should increase, but are still likely to reflect a lower bound on the true values. Standard errors in parentheses.

Table 3: Estimates of the Relative Crash Risk of Drinking vs. Sober Drivers:
 Allowing for Risk to Vary across Years

Year	Relative two-car fatal crash risk for drinking drivers (θ)	Relative one-car fatal crash risk for drinking drivers (λ)	Implied fraction of drivers that have been drinking (8 pm to 5 am)
1983	8.10 (0.93)	12.11 (1.02)	.186 (.016)
1984	6.95 (0.68)	11.36 (0.80)	.214 (.013)
1985	9.40 (0.87)	13.88 (0.98)	.173 (.012)
1986	7.42 (0.68)	12.29 (0.82)	.200 (.011)
1987	8.29 (0.71)	7.26 (0.44)	.198 (.011)
1988	7.16 (0.66)	7.57 (0.49)	.188 (.012)
1989	8.01 (0.74)	7.76 (0.51)	.177 (.011)
1990	7.69 (0.73)	7.72 (0.52)	.182 (.012)
1991	8.89 (0.88)	8.60 (0.63)	.161 (.012)
1992	7.34 (0.80)	8.08 (0.64)	.164 (.013)
1993	7.58 (0.85)	7.98 (0.66)	.153 (.014)
F-test: Equality of coefficients across years	8.80	93.92	20.49

Notes: Values are maximum likelihood estimates of equation (11), allowing the relative crash risk of drinking and sober drivers to vary by year. The implied fraction of drinking drivers is calculated based on the estimates in the first two columns and the means of the data in our sample (between the hours of 8 pm and 5 am, 1983-1993). A driver does not have to be legally drunk to be categorized as drinking. Note that in the estimation, the fraction of drinking drivers is allowed to vary by state x year x weekend x hour. Thus, the estimates in this table are comparable to Column 8 of Table 2. Coefficients for each year are estimated separately. For the reasons stated in the paper, the values in the first two columns are likely to be lower bounds, and the value in the final column is likely to be an upper bound. Standard errors in parentheses. The reported F-test is asymptotically distributed χ^2 with 10 degrees of freedom. The .05 critical value for this statistic is 18.3.

Table 4: Estimates of the Relative Crash Risk of Drinking vs. Sober Drivers:
Allowing for Risk to Vary by Time of Day

Time of day	Relative two-car fatal crash risk for drinking drivers (θ)	Relative one-car fatal crash risk for drinking drivers (λ)	Implied fraction of drivers that have been drinking (9 pm to 5 am)
8 pm -9 pm	7.04 (0.60)	7.45 (0.49)	.125 (.008)
9pm - 10pm	7.66 (0.63)	7.85 (0.49)	.133 (.008)
10 pm -11 pm	7.92 (0.63)	9.03 (0.55)	.147 (.008)
11 pm -midnight	8.32 (0.64)	9.82 (0.56)	.171 (.009)
midnight - 1am	9.21 (0.72)	10.84 (0.60)	.198 (.009)
1 am - 2 am	7.95 (0.69)	8.85 (0.49)	.266 (.011)
2 am - 3 am	6.96 (0.65)	8.30 (0.48)	.288 (.012)
3am - 4 am	7.96 (0.97)	10.23 (0.84)	.211 (.016)
4 am - 5 am	8.50 (1.29)	11.58 (1.33)	.127 (.021)
F-test: Equality of coefficients across hours	8.20	34.30	224.12

Notes: Values are maximum likelihood estimates of equation (11), allowing the relative crash risk of drinking and sober drivers to vary by hours of the day. The implied fraction of drinking drivers is calculated based on the estimates in the first two columns and the means of the data in our sample (fatal crashes between the hours of 8 pm and 5 am, 1983-93). A driver does not have to be legally drunk to be categorized as drinking. Note that in the actual estimation, the fraction of drinking drivers is allowed to vary by state x year x weekend x hour. Thus, the estimates in this table are comparable to Column 8 of Table 2. Coefficients for each hour of the day are estimated separately. For the reasons stated in the paper, the values in the first two columns are likely to be lower bounds, and the value in the final column is likely to be an upper bound. Standard errors in parentheses. The reported F-test is asymptotically distributed χ^2 with 8 degrees of freedom. The .05critical value for this statistic is 15.5.

Table 5: Sensitivity of the Estimates to Violations of the Modeling Assumptions
and the Presence of Measurement Error

	Relative two-car fatal crash risk for drinking drivers (θ)	Relative one-car fatal crash risk for drinking drivers (λ)	Implied fraction of drivers that have been drinking (9 pm to 5 am)
Baseline	8.04 (0.23)	9.11 (0.19)	.178 (.004)
Drinking drivers 25 percent less efficient at avoiding crashes initiated by other drivers	10.06 (0.29)	11.39 (0.24)	.147 (.003)
Drinking drivers 10 percent more likely to interact with drinking drivers than sober drivers	10.18 (0.28)	10.93 (0.22)	.153 (.003)
5 percent of drivers misclassified due to classical measurement error	13.22 (0.21)	15.09 (0.08)	.125 (.005)
5 percent of drinking drivers misclassified as sober	10.46 (0.15)	11.59 (0.06)	.169 (.006)
In 5 percent of <i>crashes</i> all drivers are reported as sober, regardless of true drinking status	12.15 (0.37)	12.66 (0.29)	.157 (.004)

Notes: The baseline specification in the top row corresponds to column 8 of Table 2. The other values reported in the table are estimates of the true parameters implied by the coefficients in the baseline specification if that baseline specification is contaminated in the named manner. Crash avoidance represents a violation of assumption A3. Non-equal mixing violates assumption A2. The last three rows of the table report three different types of measurement error.

Table 6: Estimating Relative Fatal Crash Risks Based on Other Driver Characteristics

Comparison Groups	Relative two-car fatal crash risk for first category named	Relative one-car fatal crash risk for first category named	Implied fraction of drivers in the first category named (8 pm to 5 am)
All drinking vs. sober	8.04 (0.23)	9.11 (0.19)	.178 (0.004)
Legally drunk vs. sober	14.07 (1.25)	17.76 (1.32)	.075 (.011)
Under age 25 vs. all others	1.51 (0.16)	1.79 (0.10)	.342 (.008)
Male vs. female	3.51 (0.22)	2.03 (0.05)	.715 (.005)
Bad driving record vs clean driving record	2.46 (0.13)	1.90 (0.06)	.252 (.006)

Notes: Values in the first two columns are maximum likelihood estimates of the fatal crash risk of the first group named relative to the second group named. All specifications assume equal mixing by hour x state x year and therefore are comparable to column 8 of Table 2. Estimates are based on the same data sample used in Table 2, except in the comparison of legally drunk to sober drivers, where only those state- year pairs that report blood alcohol for 95 percent of those drivers identified as drinking by the police are included in the sample. Standard errors in parentheses.

Table 7: Relative Fatal Crash Risk: Allowing for Interactions
between Drinking and other Driver Characteristics

Driver classification	Two-car fatal crash risk relative to baseline category	One-car fatal crash risk relative to baseline category	Implied fraction of drivers in the first category named (9 pm to 5 am)
Under 25 x drinking	12.10 (0.63)	14.24 (0.54)	.077 (.012)
Over 25 x drinking	10.95 (0.49)	10.79 (0.34)	.111 (.007)
Under 25 x not drinking	2.90 (0.14)	2.51 (0.07)	.230 (.003)
Over 25 x not drinking	1.00	1.00	.581 (.006)
Male x drinking	9.39 (1.06)	8.99 (0.50)	.198 (.011)
Female x drinking	5.71 (0.86)	5.70 (0.51)	.041 (.005)
Male x sober	1.36 (0.17)	1.10 (0.06)	.603 (.006)
Female x sober	1.00	1.00	.159 (.010)
Bad driving record x drinking	10.01 (0.58)	9.83 (0.43)	.077 (.009)
Clean driving record x drinking	8.12 (0.41)	8.50 (0.30)	.112 (.008)
Bad driving record x sober	1.93 (0.11)	1.37 (0.04)	.215 (.004)
Clean driving record x sober	1.00	1.00	.595 (.005)

Notes: Values in the first two columns are maximum likelihood estimates of the fatal crash risk of the named group relative to the named baseline group (i.e the group with a relative risk defined to be equal to one). All specifications assume equal mixing by state x year x day x hour and therefore are comparable to column 8 of Table 2. Estimates are based on the same data sample used in Table 2. Standard errors in parentheses.