SAMPLE SELECTION IN THE ESTIMATION OF AIR BAG AND SEAT BELT EFFECTIVENESS*

Steven D. Levitt and Jack Porter*

Abstract—Because data are collected for only fatal crashes, it is difficult to accurately measure seat belt and air bag effectiveness. The use of safety devices influences survival rates which in turn determine whether a crash is included in the sample, leading to sample selection bias. We propose a simple solution to the selection problem: limiting the sample to crashes in which someone in a different vehicle dies. Empirically, we find seat belts more effective and air bags to be less effective than previously found. The cost per life saved through seat belts is approximately $30,000, compared to $1.8 million for air bags.

I. Introduction

ROAD fatalities remain the leading cause of death among those aged six to 34 in the United States, in spite of the fact that the number of traffic deaths has declined dramatically over the last two decades. Figure 1 shows that traffic fatalities peaked at more than 54,000 in 1972 and have since fallen to roughly 40,000 annually. This decline is even more remarkable in light of the large increases in the total volume of travel. Fatalities per vehicle mile traveled are less than one-third as high as the level in the late 1960s.

Increased seat belt usage (for example, Graham et al. (1997), NHTSA (1984), and Orsay et al. (1988)) and the proliferation of airbags (for example, Graham et al. (1997), Lund and Ferguson (1995), and NHTSA (1996)) are two factors that have been identified as contributing to the declining death toll on the roads.[1] NHTSA (1996) estimates seat belt usage rates of 58% to 68% in recent years, up from 11% in 1980. Air bags, first available on passenger vehicles in 1987, were installed in more than half of all new cars sold by 1992 (NHTSA, 1996). Dual air bags became mandatory under federal law in all 1998 model-year passenger cars and all new 1999 model-year light trucks. Even in advance of federal requirements, however, air bags have been standard equipment on almost all vehicles in recent years. Consumer demand for air bags appears strong, despite their cost ($410 for dual airbags according to Graham et al. (1997)).

Past estimates of seat belt effectiveness span a wide range. Robertson (1976), for instance, cites nineteen studies on belt effectiveness with estimates ranging from an 8% to an 86% reduction in death or injury relative to those not wearing seat belts. NHTSA (1984) concludes that seat belts are 45% to 55% effective in reducing fatalities. Although different methodologies are sometimes used, the standard approach to measuring seat belt effectiveness is to identify a sample of crashes and compare outcomes among those with and without seat belts (for example, Kaplan and Cowley (1991), Marine et al. (1994), Orsay et al. (1988), and Robertson (1976)). One important shortcoming of this approach, which we later discuss at length, is the possibility of sample selection. Typically, only particular kinds of crashes are included in a sample (for example, those crashes with fatalities or crashes involving injuries that require an ambulance). If seat belt usage reduces injury severity, then such sample selection will tend to bias downward the measured benefits of seat belts.

Although it is recognized that air bags pose risks of both injury (Hollands et al., 1996; Morris & Borja, 1998) and death (NHTSA, 1996), recent research and government evaluations continue to support the effectiveness of air bags. NHTSA estimates that air bags saved more than 3,000 lives between their introduction and September 1998. Graham et al. (1997) concludes that the cost effectiveness of air bags is on par with other medical and public-health interventions.

Studies of air bags have generally used one of three empirical approaches. The first approach analyzes fatality rates per registered vehicle for vehicles with and without airbags (Lund & Ferguson, 1995; Ferguson, Lund, & Greene, 1995). An important potential weakness of this approach is the endogeneity of vehicle choice. Faced with an option, those consumers who value safety most highly may disproportionately purchase vehicles with air bags.[2] This form of selection emerges clearly in the data. For example, in our data set, comparing drivers with and without air bags who are involved in fatal two-vehicle crashes, those with air bags are 19% less likely to be reported to be drinking by the police, 11% less likely to have a recent speeding ticket or accident, 14% less likely to have been convicted of driving while intoxicated or to have had a license suspension, 12% less likely to be male, and 13% less likely to be under the age of 25.[3] All of those characteristics

Received for publication August 5, 1999. Revision accepted for publication November 7, 2000.

* University of Chicago and American Bar Foundation, and Harvard University, respectively.

We would like to thank Gary Chamberlain, Austan Goolsbee, John Graham, Lawrence Katz, and two anonymous referees for helpful comments. All remaining errors are our own. Financial support of the National Science Foundation and the John Lee Fund are gratefully acknowledged. Previn Mankodi provided outstanding research assistance.

[1] Other factors that have been identified as contributing to reduced fatalities are reductions in drunk driving due to increased enforcement, changing attitudes, and changes in the minimum legal drinking age (Grossman et al., 1993; Ruhm, 1996; Saffer & Grossman, 1987; Zobeck et al., 1990), lower speed limits (NHTSA, 1989), although Lave and Elias, 1997 disagree), child safety seats (NHTSA, 1998), and mandatory helmet laws for motorcycle riders (NHTSA, 1996).

[2] In theory, the opposite type of selection could also occur. The most dangerous drivers might disproportionately choose cars with air bags because they will obtain the greatest reduction in death rates. Empirically, however, this does not appear to be the case.

[3] In addition to selection of drivers based on observable characteristics such as sex, there is likely to be selection on unobservables as well. For instance, among men of a given driving age and similar previous driving records, those valuing safety most highly may be more likely to choose vehicles with air bags.

The Review of Economics and Statistics, November 2001, 83(4): 603–615

FIGURE 1.—MOTOR VEHICLE FATALITIES 1966–1994

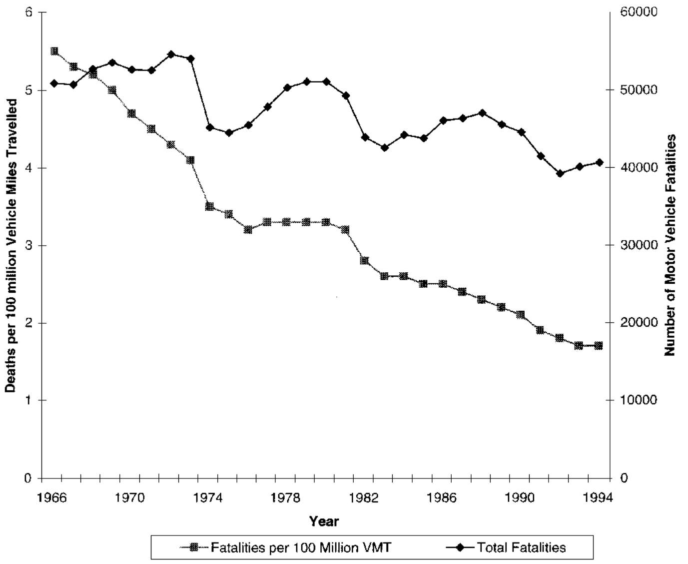

(driving while intoxicated, bad past driving record, being young or male) are associated with greater likelihoods of fatal crash involvement (Fell & Nash, 1989; Levitt & Porter, 2001; NHTSA, 1998). Thus, this approach is likely to exaggerate any causal impact of air bags on fatality reductions.

The second approach that is frequently used to measure air bag effectiveness involves comparing ratios of the number of deaths in frontal versus nonfrontal collisions for vehicles with and without air bags. This approach is known as a "frontal/nonfrontal comparison" and is a form of difference-in-difference estimator. Air bags are designed to protect occupants only in frontal crashes. If the ratio of deaths in frontal versus nonfrontal crashes is lower for vehicles with air bags relative to those without air bags, this difference is interpreted as lives saved due to air bags. Using variations on this approach, researchers have found air bags to be 20% to 35% effective in reducing death rates in frontal crashes (Braver et al., 1997; Kahane, 1996; NHTSA, 1996; Zador & Ciccone, 1993).

Because this analysis is based on ratios of numbers of deaths rather than crash survival rates, it relies critically on the assumption that the proportion of frontal crashes is constant for drivers with and without air bags. This assumption, however, appears to be false. As noted above, "safe" drivers (that is, sober, female, over 25, and having a good past driving record) are disproportionately represented among those with air bags. Safe drivers, however, are more likely to be hit than to hit another vehicle, and thus are under-represented in frontal crashes. For instance, in two-car crashes in our data set, being sober is associated with a 72% reduction in the proportion of frontal to nonfrontal crashes, wearing a seat belt, having a previous drunk driving conviction or license suspension, and being female are associated with 25%, 33%, and 28% reductions, respectively. These observable driver characteristics—which are related not to differential survival rates in frontal and nonfrontal crashes, but rather to the frequency of such crashes—yield estimates similar to or larger than those for air bags.[4] Thus, it appears that the frontal/nonfrontal approach confounds differences in crash survival rates with the frequency with which frontal crashes occur, leading to exaggeratedly large estimates of the benefits of air bags.[5]

A third approach to measuring the impact of safety devices is known as the "double pairs" comparison (Evans, 1986). The ratio of driver fatalities to passenger fatalities in crashes in which neither occupant has an airbag is compared to the ratio of driver and passenger fatalities in crashes in which neither have airbags. As long as the set of crashes involving vehicles with exactly one airbag is the same as that involving vehicles with zero or two airbags, this approach provides a consistent estimate of air bag effectiveness.[6] The primary drawback of this approach is that, with almost all vehicles equipped with dual air bags in recent years, the method can be applied to only older model cars.

Previous researchers have employed these indirect approaches to measuring the impact of seat belts and air bags due to an important data limitation: the most comprehensive data set is the Fatality Analysis Reporting System (FARS), which provides extensive information on all passengers in virtually all U.S. crashes, but only in accidents in which a death occurs. The obvious problem in using these data to analyze the benefits of air bags and seat belts is sample selection. Seat belts and air bags influence the probability of death, which in turn determines whether or not a crash is included in the data set.[7] As economists have long understood in other contexts (Angrist & Krueger, 1999; Heckman, 1979; Heckman et al., 1996; Heckman et al., 1999), sample selection leads to biased estimation. Failing to account for sample selection leads to estimates that systematically understate the benefits of effective life-saving devices. Empirically, the extent of sample selection in the fatal crash data is enormous. Crashes with a fatality account for only 0.5% of all reported crashes and less than 2% of reported crashes with injuries (NHTSA, 1998). Moreover, in almost 90% of fatal crashes, there is a single fatality. If

[4] These lower instances of death in frontal crashes for safe drivers is not attributable to the fact that more of the safe drivers have air bags. Similar results are obtained when controlling for air bag status.

[5] On the other hand, if air bags provide benefits in nonfrontal crashes (or there is classification error so that some frontal crashes are reported as nonfrontal), this approach will underestimate the true effectiveness of air bags.

[6] For instance, the composition of crashes among people who drive older cars that have no air bags may be different if they are also less likely to be safe drivers.

[7] Other data sets are available, but the problems posed by these data sets are even more severe. The crashes included in these data are a representative sample of crashes reported to the police, not of all crashes. Crashes in which injuries occur are more likely to be reported, leading to sample selection if the probability and extent of injury is a function of seat belts and airbags. In addition, much less information is available for these accidents.

TABLE 1.—PROBABILITY OF DEATH AS A FUNCTION OF NUMBER OF OCCUPANTS IN CRASH

	Total Number of Occupants in All Vehicles Involves in the Crash:				
Sample	One Occupant	Two Occupants	Three Occupants	Four Occupants	Five Occupants
All crashes	1.000	0.538	0.393	0.318	0.269
Anyone else dies in crash	—	0.142	0.157	0.150	0.138
Anyone dies in a different vehicle	—	0.068	0.074	0.070	0.070

Values in table are computed using all passengers in all vehicles for crashes in which at least one vehicle occupant dies. Crashes in which the only fatality is a pedestrian, motorcyclist, or other nonmotorist are excluded. Unlike later tables, the sample used in this table is not limited to frontseat passengers in two-vehicle crashes. Data are for crashes over the period 1994–1997 as reported in FARS.

that individual had not died, the crash would be excluded from the data set.

In this paper, we use a simple identification strategy that allows us to directly estimate the impact of seat belts and air bags on crash survival rates, despite sample selection in the data. Sample selection arises because a given individual's seat belt usage affects his or her probability of death, which in turn influences whether the crash is included in the data. The key insight is that as long as anyone else dies in the crash, it is included in the fatal accident data regardless of what happens to others in the crash. We propose to restrict the sample to cases in which anyone dies in another vehicle in the crash. By focusing on this subset of crashes, sample selection is eliminated under the conditions derived in section II.[8] Perhaps counterintuitively, the sample selection problem, which arises because observations are excluded from the data set, is overcome by further restricting the data that are used.

Table 1 demonstrates how sample selection distorts the data. The top row of the table presents the probabilities of death in the raw data as a function of the total number of individuals occupying all vehicles in the crash.[9] By definition, when there is only one vehicle occupant, the probability of death is equal to one; if the individual did not die, the crash would be excluded from the data set. As the number of occupants increases, sample selection becomes decreasingly severe as evidenced by the steady decline in the probability of death. When five motorists are involved in a crash, each individual's probability of death is 26.9%. The falling death rates in the top row are not a consequence of decreasing crash severity as demonstrated by the bottom two rows of the table. Restricting the sample to cases in which anyone else dies in a crash results in a probability of death that is nearly constant at 15% as the number of occupants increases. Conditional on anyone dying in a different vehicle, individual death rates hover around 7% regardless of the number of total motorists.

[8] One could include any individual who is involved in a crash in which someone else dies, but the required identifying assumptions are more difficult to justify for occupants of the same vehicle relative to occupants of different vehicles.

[9] Unlike later results that are restricted to two-vehicle crashes, these results are for all crashes regardless of the number of vehicles. Only crashes in which some vehicle occupant dies are included. (There are some crashes in the data in which the only fatality is a pedestrian, bicyclist, or motorcyclist.)

Empirically, using data on fatal crashes in the United States from 1994 to 1997, we find that correcting for sample selection bias dramatically increases the measured effectiveness of seat belts. We find seat belts to be substantially more effective than has generally been the case in past research. Wearing a seat belt reduces the risk of death in our sample by 60% to 70%. Air bags are found to be approximately 15% effective in lowering death rates in direct frontal crashes, but appear to yield little or no benefit in partial-frontal or nonfrontal crashes. Our air bag estimates are on the low end of the spectrum of previous estimates in the literature.

The outline of this paper is as follows. Section II develops the theoretical model and identification strategy. Section III uses data on fatal crashes to estimate the impact of seat belts and airbags on probabilities of death and injury, correcting for sample selection. Section IV offers interpretation of the coefficients in terms of lives saved and cost-benefit ratios, and section V concludes.

II. Theoretical Model and Identification Strategy

Our goal is to estimate the effectiveness of seat belts and air bags in preventing death in motor vehicle crashes. In this section, we adopt a potential outcomes framework (Heckman, 1990; Rubin, 1990) to formally consider the causal effects of safety devices on accident fatalities. Let S and A denote seat belt and air bag usage indicators. For example, if individual i wore a seat belt but did not have an air bag in a crash, then $S_i = 1$ and $A_i = 0$. Also, let $Y_i^{S,A}$ denote the potential outcome indicator for whether individual i would have died in the crash as a function of seatbelt and air bag usage. For instance, if $Y_i^{1,0} = 0$ and $Y_i^{0,1} = 1$, then person i would have survived the crash if he or she had been wearing a seatbelt but did not have an air bag, but not vice versa. This notation is used for defining causal effects because it encompasses both the actual outcome and the counterfactual outcomes.

In practice, we are never able to observe causal effects for any single individual because, in any given crash, only the actual outcome is observed, and not the potential outcomes had safety device status been changed. As a result, we can only hope to identify a mean causal effect over some population, for example, $E(Y_i^{1,0} - Y_i^{0,0})$. In theory, there are many potentially interesting causal effects to measure.

In the next section, we present estimates of mean safety device causal effectiveness, but first we establish conditions for validity of the proposed estimation method.

A. *Measuring Causal Effects in a Hypothetical Data Set without Sample Selection Problems*

Before discussing the actual estimation problem we face because of data limitations, it is useful to briefly consider an idealized data set that is not subject to the sample selection problems found in FARS. This data set, like FARS, would include all crashes in which someone dies. Unlike FARS, however, it would also include any crash in which there exists an alternative safety device configuration (for example, no seat belt, no airbag) in which someone would have died. In other words, the hypothetical data set would include all possibly fatal accidents. Stated formally, a crash is included in the set of possibly fatal accidents if and only if k individuals $i_l, \ldots, i_k$ are involved in a crash and

$$\max\{Y_{i_1}^{0,0}, Y_{i_1}^{1,0}, Y_{i_1}^{0,1}, Y_{i_1}^{1,1}, Y_{i_2}^{0,0}, \ldots, Y_{i_k}^{1,1}\} = 1. \qquad (1)$$

The only crashes not included in the set of possibly fatal accidents are those in which everyone would have survived, regardless of safety device use. For any individual i, we adopt the notation $P_i = 1$ if equation (1) is satisfied, meaning that the individual was in a potentially fatal crash.[10]

If this idealized data set were available, then the natural condition for identification of a causal effect would be that, conditional on the set of observed individual and crash characteristics, safety device usage is as good as randomly assigned. Stated formally, where X_i is a vector of observed individual and crash characteristics (other than seatbelt and airbag use):

Assumption A1: $(Y_i^{0,0}, Y_i^{1,0}, Y_i^{0,1}, Y_i^{1,1})$ *and* (S_i, A_i) *are jointly independent conditional on* X_i *and* $P_i = 1$.

If assumption A1 holds, the relevant causal effects can be identified by simply conditioning on the other observed characteristics X and comparing differences in mean death rates by safety device status. For instance, the causal effect of wearing a seat belt when no air bag is present is simply

$$\begin{aligned} E(Y_i^{1,0} - Y_i^{0,0}|X_i, P_i = 1) \\ = E(Y_i|S_i = 1, A_i = 0, X_i, P_i = 1) \\ - E(Y_i|S_i = 0, A_i = 0, X_i, P_i = 1). \end{aligned}$$

A number of factors need to be considered when assessing whether assumption A1 is reasonable. This assumption is violated if safety device usage is correlated with unobservable individual or crash characteristics that affect an individual's potential outcomes (conditional on the observed factors). Correlation with unobserved crash characteristics could violate the assumption in one of three ways. First, if drivers feel safer when wearing seat belts, this may induce them to drive more recklessly (Peltzman, 1975) and consequently have more-severe crashes on average. If this is the case, then the estimation of the causal impact of seat belts conditional on a crash will be biased downward.[11] A second way in which assumption A1 might be violated is if safety conscious people both wear a seat belt and drive more cautiously, resulting in a negative correlation between seat belt status and crash severity. However, note that it is conditional independence that is assumed in A1. For example, if we directly observed safety consciousness, then we could condition on it and it would not pose an estimation problem. Of course, safety consciousness is not directly observable, but age, gender, and past driving record are among the variables observed, as are vehicle type, time of day, and urban versus rural. These characteristics might reasonably proxy for safety consciousness. Furthermore, various crash characteristics are also observed, such as the type of impact, relative vehicle weights, and the speed limit, which might directly control for crash severity. The more such characteristics are observed, the more believable the conditional independence assumption becomes, motivating the "kitchen sink"-type approach we employ. Moreover, conditioning on $P = 1$, the subset of crashes severe enough that someone could have died, further mitigates this concern. For example, suppose that seat belt wearers are likely to be in less-severe accidents. That correlation would be problematic for A1 only if within the subset of possibly fatal crashes seat belt wearers are more likely to be in the less-severe ones. That is, more-frequent involvement in fender-bender accidents does not create a problem for A1, because those crashes are not of the possibly fatal type.

The final way in which assumption A1 could be violated is through a correlation between safety device use and other unobserved individual or vehicle characteristics. For instance, suppose seat belt wearers are more likely to have health problems or ride in lighter cars, and these factors in turn influence their potential outcomes in crashes. If health status and car weight were observed and included in the conditioning set, then this channel of correlation would be eliminated. As before, every health-related factor and the precise car weight are not directly observable, but we condition on as many observable individual and car characteristics (such as age, gender, and car model) as possible to control for such channels of correlation.

[10] Note that, if seat belts and air bags are effective, then $Y_i^{0,0} \geq \max\{Y_i^{1,0}, Y_i^{0,1}, Y_i^{1,1}\}$ (if an individual would die in a crash using a seat belt and/or an air bag, then he or she would also die without any safety device), in which case the criterion for being in the PFA data set could be simplified to $\max\{Y_{i_1}^{0,0}, \ldots, Y_{i_k}^{0,0}\} = 1$.

[11] Note that the typical Peltzman (1975) effect—more crashes due to seat belt use—does not pose any problems for our estimation.

B. Sample Selection in FARS

The actual data set available, FARS, differs from the hypothetical data set previously described in that it includes only those accidents in which someone actually dies. Crashes in which someone did not die, but would have died had they been using an alternative safety device configuration, are excluded. The root of the sample selection problem in FARS is the direct link between safety device use and inclusion in the sample. For instance, if seat belts are effective, then an individual wearing a seat belt is more likely to survive a crash and thus is less likely to be included in FARS. All of the cases in which a safety device fails to prevent a death are included in the data set, but only a subset of the cases in which a life is saved are recorded. By excluding these saved lives from the sample, safety devices will tend to appear less effective than they actually are.

Let $Q_i = 1$ be an indicator variable denoting whether someone actually died in the crash in which i is involved.[12] This condition is the FARS equivalent of $P_i = 1$ in the hypothetical data set. Following the approach described earlier, one would like to identify a causal effect such as $E(Y_i^{1,0} - Y_i^{0,0}|X_i, Q_i = 1)$. When all potentially fatal crashes are included in the data set, identification is made possible by assumption A1. The corresponding assumption for FARS requires independence between potential outcomes and safety device use conditional on observables and selection into FARS as follows.

Assumption A1′: $(Y_i^{0,0}, Y_i^{1,0}, Y_i^{0,1}, Y_i^{1,1})$ *and* (S_i, A_i) *are jointly independent conditional on* X_i *and* $Q_i = 1$.

If assumption A1′ were true, then identification of causal effects in the FARS data set follows in a straightforward manner from the earlier discussion. Note, however, that assumption A1′ is unlikely to hold. In particular,

Theorem 1: If A1 holds and safety devices are effective, then A1′ is violated.

PROOF: See appendix A.

This theorem is a direct result of the sample selection problem created in the formation of the FARS data set. FARS is a subset of the possibly fatal accidents because FARS includes only those accidents in which a fatality actually occurs. Relative to the set of possibly fatal accidents, FARS excludes some crashes in which a safety device saves a life. If the random assignment assumption holds in A1, then the exclusion of such crashes from FARS will naturally induce a correlation between safety devices and potential outcomes (conditional on observables) in the remaining crashes in FARS. By this theorem, if safety devices are as good as randomly assigned in the set of potentially fatal crashes and if safety devices have some effectiveness, then the usual method of identifying causal effects cannot be followed in FARS.

C. Sample Selection Correction

Because assumption A1′ is unlikely to hold, we propose an alternative approach to identifying causal effects: we propose using a subset of FARS for which the random assignment assumption is reasonable. We focus on individuals in two-car crashes. An individual is included in the subset if any individual in the other car died in the crash.[13]

If individual i is involved in a two-car accident, let W_i denote an indicator function equal to one when someone in the other car in the crash dies. Then define assumption A2 as follows.

Assumption A2: W_i *and* (S_i, A_i) *are jointly independent conditional on* X_i *and* $(Y_i^{0,0}, Y_i^{1,0}, Y_i^{0,1}, Y_i^{1,1})$.

This assumption requires that—conditional on other factors—safety device usage in one car is independent of the fatality outcome of individuals in the other car involved in a crash.

Theorem 2: If assumptions A1 and A2 hold, then causal effects can be consistently estimated using the subset of the FARS data in which someone dies in the other vehicle in the crash.[14]

PROOF: The result is a straightforward application of Bayes' theorem.

By limiting our attention to crashes in which someone dies in the other car, the link between one's own safety device usage and inclusion in the sample is eliminated. Regardless of my personal outcome, the crash will be included in the data set.

Assumption A2 seems quite reasonable. Certainly, safety device use in one car should not directly affect survival in a different car, so possible violations would have to work through indirect channels. For instance, seat belt use might be correlated with driving a safer car, and safer cars might be more likely to crash into other safer cars. Here, however, it is important to note that only conditional independence is required. So, by conditioning on car model, the correlation through driving safer cars would no longer constitute a violation. By conditioning on enough individual and vehicle characteristics, it is reasonable that safety device use in one car is independent of the unobserved individual and vehicle

[12] Formally, if individuals $i_1, \ldots, i_k$ are involved in a crash, $Q_i = 1$ denotes max $\{Y_{i_1}, \ldots, Y_{i_k}\} = 1$, where Y_{i_j} is the fatality outcome observed for individual i_j.

[13] Because FARS contains all fatal accidents, such a set is truly a subset of FARS. Also, because the criterion for inclusion in the subset depends only on observed outcomes (and not potential outcomes), the subset is well identified and observable within FARS.

[14] Formally, the conclusion of the theorem is $(Y_i^{0,0}, Y_i^{0,1}, Y_i^{1,0}, Y_i^{1,1})$, and (S_i, A_i) are jointly independent conditional on X_i and $W_i = 1$.

characteristics in the other car that might influence crash survival after conditioning on the observed characteristics.[15]

Before actually proceeding to the estimation, it must be noted that the approach we adopt has an important limitation. The best we can do is to estimate the mean seat belt effectiveness conditional on a given vector of characteristics, X, and average over crashes severe enough that $W = 1$. Note, however, that the subset of crashes in which someone in the other vehicle dies may have a distribution of crash severity that is very different from the universe of crashes. How well safety devices work is likely to be a function of crash severity. For example, in the most gentle crashes, even unbelted occupants will not die, so seat belts are completely ineffective (with respect to preventing fatalities); in the most severe crashes (such as a convertible plunging off a cliff), seat belts may also be useless. Thus, the estimate we obtain on safety device effectiveness may not be readily generalizable outside our sample.

III. Estimating the Impact of Air Bags and Seat Belts

We estimate the model of the previous section using FARS data for fatal crashes occurring in the years 1994–1997. These data contain detailed information on virtually every crash in the United States in which a fatality occurs among either vehicle occupants or pedestrians. The sample of vehicle occupants actually used in the estimation that follows is limited in a number of ways. First, we exclude backseat passengers, both because of low death rates for such passengers and the absence of air bags. Second, we exclude children in car seats. The risks posed by air bag deployment when car seats are placed in the front seat are well established (NHTSA, 1996). As this fact became public knowledge, the prevalence of car seats placed in the front seat declined dramatically. In terms of predicting future effectiveness of air bags, excluding car seats is likely to provide a more accurate measure. Third, the sample is limited to two-vehicle crashes. One-vehicle crashes are not included because the sample selection correction requires a fatality in another vehicle. Crashes with three or more vehicles (7% of all fatal crashes) are not included because of concerns that crash severity may vary dramatically across vehicles in a manner that is difficult to control for. Fourth, crashes in which the only fatalities are motorcyclists, bicyclists, or pedestrians are also dropped because such accidents are likely to pose little risk to vehicle occupants. Fifth, we exclude occupants of large trucks from the sample as well as vehicles of model years prior to 1991. In both cases,

[15] Alternatively, safety device use in one car could be correlated with crash-level factors that influence survival in the other car. Here it is important to note that we are conditioning on both observable crash factors and potential outcomes. Having potential outcomes in the conditioning set makes violations of this kind particularly difficult to fathom. The potential outcomes would seem to summarize far more about *i*'s crash severity than safety device use. Even Peltzman-type effects or adverse selection stories would not create a violation due to the presence of the potential outcomes in the conditioning set.

TABLE 2.—SUMMARY STATISTICS

	Full Sample		With Sample Selection Correction (Anyone Dies in Other Vehicle)	
Variable	Mean	Standard Deviation	Mean	Standard Deviation
Individual-Level characteristics				
Die in crash	0.393	0.488	0.093	0.291
No seat belt; no air bag	0.198	0.399	0.162	0.368
Seat belt only	0.410	0.492	0.444	0.497
Air bag only	0.099	0.299	0.072	0.259
Both seat belt and air bag	0.293	0.455	0.322	0.467
Driver	0.708	0.455	0.739	0.439
Male	0.578	0.494	0.629	0.483
Age	41.3	20.7	36.5	16.8
Vehicle-Level Characteristics				
Model year	1,992.8	2.0	1,993.0	2.0
Vehicle weight	4,058	2,039	4,735	2,291
Difference in vehicle weights	2,634	2,238	2,596	2,343
Direct frontal impact	0.538	0.499	0.729	0.444
Partial frontal impact	0.206	0.404	0.182	0.386
Non-frontal impact	0.256	0.391	0.089	0.302
Vehicle type: automobile	0.607	0.488	0.453	0.498
Driver previous minor violations	0.359	0.480	0.383	0.486
Driver previous major violations	0.071	0.256	0.073	0.260
Crash-Level Characteristics				
Year	1,995.7	1.1	1,995.7	1.1
Speed limit 55 mph or greater	0.577	0.494	0.572	0.495
Time of day: 6 AM to 8 PM	0.717	0.451	0.706	0.455
Time of day: 8 PM to 1 AM	0.217	0.412	0.231	0.422
Rural	0.610	0.488	0.611	0.487

Summary statistics are for frontseat occupants involved in two-vehicle crashes in which a motorist is fatally injured between 1994 and 1997. Occupants of large trucks and vehicles built prior to 1990 are excluded from the sample. Refer to the text for other sample restrictions. A data appendix with the precise manner in which all variables were constructed is available on request from the authors.

air bag installation is extremely low. Note, however, that if a newer vehicle crashes into a large truck or an older vehicle, the occupants of the newer vehicle are included in the sample. None of the basic results are sensitive to these exclusions.

There are serious quality concerns in the FARS data, especially for the variables measuring seat belt and air bag status: 8% of the individual-level observations are missing information as to whether a seat belt was worn. Those individuals are excluded from the sample. As we later report, the results obtained are not particularly sensitive to their inclusion under different sets of assumptions. The variable measuring air bag status is constructed using the VINDICATOR software developed by NHTSA, which extracts air bag information from the vehicle identification numbers (VINs) included in FARS. Approximately 10% of vehicles in the FARS data set fail to yield a valid match using VINDICATOR and are dropped from the sample.

Our baseline sample after these exclusions includes approximately 42,000 individuals. Means and standard deviations for this sample are reported in columns 1 and 2 of table 2. A more complete description of variables as well as information about the way they were constructed is available from the authors upon request.

As the model in the preceding section demonstrates, it is necessary to correct for sample selection. We do so by further restricting the data set to occupants of vehicles in which anyone in the other vehicle dies in the crash. Summary statistics for the selection-corrected samples are presented in the remaining columns of table 2. This restricted sample has roughly 21,000 observations. There are a few important differences between the full sample and the selection-corrected sample. First, death rates in the full sample are four times higher. This is primarily due to the artificial inflation of death rates due to sample selection. Second, passengers in vehicles with frontal impact and in larger vehicles are over-represented in the restricted sample. Vehicles that strike other vehicles head-on are more likely to inflict fatalities, as are larger vehicles.

The basic estimating equation used in the analysis is

$$Y_{jvc} = \alpha + \beta_1 Seat\_belt_{jvc} + \beta_2 Air\_bag_{jvc} + X_{jvc}\Gamma + V_{vc}\Theta + Z_c\Lambda + \epsilon_{jvc}, \tag{4}$$

where j indexes individual vehicle occupants, v corresponds to a given vehicle, and c reflects a particular crash. The dependent variable Y is an indicator variable equal to one if the occupant is killed, and zero otherwise. In addition to the seat belt and air bag measures, we include a vector of individual-level characteristics, X (age, sex, seat position); vehicle-level controls, V, for the driver's vehicle and/or the other vehicle in the crash (the type of vehicle that is crashed into, vehicle weight, measures of the weight differential between the vehicles and the squared weight differential, model year, driver's past driving record, and other driver's past driving record); and crash-level factors, Z (year of crash, a speed limit indicator, urban versus rural, time of day indicators).[16] All of the results presented are based on linear probability models with robust standard error corrections for heteroskedasticity and within-vehicle correlation. Probit and logit estimation yield similar results.

The specification in equation (4) imposes two important functional form restrictions. First, it assumes that there is no interaction between seat belt and air bag effectiveness. We focus primarily on the average causal effect of wearing a seat belt or having an air bag. We do not report separate estimates for air bag effectiveness when a seat belt is or is not worn. Empirically, we find no evidence that the use of a seat belt influences the effectiveness of air bags or vice versa. Thus, looking at the average causal effect provides the simplest summary of our data without compromising the interpretation of the results.[17] Second, we model the covariates as entering in a linear-additive form. We explore the sensitivity of our results to different sets of covariates and restricted subsamples in our analysis.

Because the effectiveness of seat belts and air bags may be a function both of vehicle size/type and the manner of impact (for example, air bags are designed to protect occupants only in frontal crashes), we present separate estimates for automobiles and other vehicles, and divide crash impacts into direct-frontal, partial-frontal, and nonfrontal collisions. A direct-frontal impact is one in which the principal vehicle impact is at twelve on the clock face. A partial-frontal impact is one in which either the initial or principal vehicle impact is between ten and two on the clock face, excluding those crashes that qualify as directly frontal.

To simplify interpretation of the empirical results, we report the fraction of lives saved through the use of seat belts or air bags, rather than the raw regression coefficients. For instance, the value reported in the table for seat belts is (β_1/probability of death with no seat belt and no airbag). The magnitude of the raw regression coefficients alone, without a scaling adjustment for average crash severity, is not particularly informative, although we always provide the normalizing denominator probability so one could calculate the coefficient estimates. Standard errors for the fraction of lives saved are calculated using the delta method.

Tables 3, 4, and 5 present regression results for direct-frontal, partial-frontal, and nonfrontal crashes, respectively. In each table, the first three columns are estimates based on the full sample, and the last three columns correct for sample selection by restricting the sample to cases in which someone dies in the other vehicle. We present results with no covariates, a limited subset of covariates, and the full set of controls. The top panel of the table corresponds to automobiles, and the bottom panel shows results for larger vehicles. Only the seat belt and air bag coefficients are reported in the table. (The other variables in the regression are discussed later.)

We begin with automobiles involved in direct-frontal crashes (the top panel of table 3). Ignoring sample selection (the first three columns), seat belts reduce death rates slightly less than 40% (off a baseline probability of death for those with neither seat belts nor air bags of 0.61). Air bags are estimated to reduce death in the full set of direct-frontal crashes by 15% to 22%. Controlling for sample selection (columns 4 through 6) increases the estimated effectiveness of seat belts, but has almost no effect on the air

[16] We have also experimented with including an interaction between seat belts and air bags that would allow a differential impact of air bags depending on whether a seat belt is worn. No systematic patterns emerged, although the interaction term was frequently statistically significant. The estimated overall impact of seat belts and air bags is not affected by the inclusion of the interaction; therefore, we have omitted the interaction term for the sake of simplicity of exposition.

[17] Later in the paper, we present a total of eighteen different specifications with the sample selection correction. The coefficient on a seat belt-air bag interaction is positive and insignificant in eight of those specifications, negative and insignificant in seven instances, and negative and statistically significant at the 0.05 level in three cases. The three cases in which the coefficient is negative and statistically significant are in nonfrontal crashes, a setting in which there is little theoretical reason to expect any interaction to be present because air bags are not designed to provide benefits in such crashes.

TABLE 3.—FRACTION OF LIVES SAVED BY SEAT BELTS AND AIR BAGS RELATIVE TO DEATH RATES WITH NO SAFETY DEVICE: DIRECT-FRONTAL CRASHES

	Full Sample			With Sample Selection Correction (Anyone Dies in Other Vehicle)		
	(1)	(2)	(3)	(4)	(5)	(6)
Automobiles						
Seat belt only	0.37	0.39	0.36	0.50	0.52	0.51
	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Air bag only	0.22	0.16	0.15	0.19	0.16	0.14
	(0.01)	(0.01)	(0.02)	(0.03)	(0.03)	(0.04)
Probability of death with no seat belt, no air bag		0.610			0.259	
Number of observations		12,106			6,752	
Utility vehicles, vans, and small trucks						
Seat belt only	0.64	0.64	0.58	0.68	0.69	0.65
	(0.02)	(0.01)	(0.02)	(0.03)	(0.03)	(0.03)
Air bag only	0.07	0.10	0.14	0.09	0.12	0.17
	(0.02)	(0.02)	(0.03)	(0.04)	(0.04)	(0.05)
Probability of death with no seat belt, no air bag		0.354			0.121	
Number of observations		10,641			8,493	
Set of covariates included?	None	Limited	Full	None	Limited	Full

Values in table are the fraction of lives saved relative to a vehicle occupant with no seat belt or air bag. The sample is limited to frontseat occupants of vehicles involved in two-vehicle crashes in which a motorist dies. The first three columns correspond to the whole sample; the last three provide sample selection-corrected estimates. The values in the table are computed from linear probability regressions of an indicator variable for death on seat belt and air bag dummies. The raw probability of death for individuals with neither a seat belt nor an air bag are also reported. Columns 1 and 4 are coefficients from simple regressions with no covariates. Columns 2 and 5 add a partial set of covariates: age dummies, gender, past driving record controls, vehicle weight, and whether the person was the driver. Columns 3 and 6 include a full set of covariates: seat position, sex, age, the type of vehicle crashed into, speed limit, vehicle weight, weight differential between vehicles, squared weight differential, past driving record of both drivers, time of day, rural, model year of vehicle, and year of crash. Full results of the specifications in column 6 are reported in table 6. Direct frontal crashes are those in which the principal point of impact is completely frontal (twelve on the clock face). White standard errors accounting for within-vehicle correlations are shown in parentheses. Standard errors are computed using the delta method.

bag coefficients. Adding covariates to the specification has little impact on the seat belt estimates, but reduces the impact of air bags.

For larger vehicles (the second panel of table 3), seat belts appear more effective, but air bags are generally less effective. Correcting for sample selection somewhat increases the measured impact of both seat belts and air bags.

Table 4 presents the regression results for partial-frontal crashes. The structure of the table mirrors that of table 3. Controlling for sample selection has a greater impact on the seat belt coefficients than was the case in the previous table. In automobiles, seat belts appear to be only 30% effective in the full sample, but are twice as effective with the sample selection correction. In utility vehicles, vans, and trucks,

TABLE 4.—FRACTION OF LIVES SAVED BY SEAT BELTS AND AIR BAGS RELATIVE TO DEATH RATES WITH NO SAFETY DEVICE: PARTIAL-FRONTAL CRASHES

	Full Sample			With Sample Selection Correction (Anyone Dies in Other Vehicle)		
	(1)	(2)	(3)	(4)	(5)	(6)
Automobiles						
Seat belt only	0.29	0.32	0.31	0.59	0.62	0.62
	(0.02)	(0.02)	(0.02)	(0.05)	(0.05)	(0.05)
Air bag only	0.08	0.03	0.06	0.07	0.06	0.03
	(0.02)	(0.02)	(0.02)	(0.06)	(0.06)	(0.07)
Probability of death with no seat belt, no air bag		0.634			0.283	
Number of observations		5,282			1,776	
Utility vehicles, vans, and small trucks						
Seat belt only	0.60	0.60	0.56	0.75	0.77	0.72
	(0.02)	(0.02)	(0.02)	(0.05)	(0.04)	(0.05)
Air bag only	−0.03	0.02	0.10	−0.09	−0.05	−0.00
	(0.03)	(0.03)	(0.04)	(0.09)	(0.08)	(0.10)
Probability of death with no seat belt, no air bag		0.490			0.147	
Number of observations		3,419			2,028	
Set of covariates included?	None	Limited	Full	None	Limited	Full

Values in table are the fraction of lives saved relative to a vehicle occupant with no seat belt or air bag. The sample is limited to frontseat occupants of vehicles involved in two-vehicle crashes in which a motorist dies. The first three columns correspond to the whole sample; the last three provide sample selection-corrected estimates. The values in the table are computed from linear probability regressions of an indicator variable for death on seat belt and air bag dummies and an interaction between those two variables. The raw probability of death for individuals with neither a seat belt nor an air bag are also reported. Columns 1 and 4 are coefficients from simple regressions with no covariates. Columns 2 and 5 add a partial set of covariates: age dummies, gender, past driving record controls, vehicle weight, and whether the person was the driver. Columns 3 and 6 include a full set of covariates: seat position, sex, age, the type of vehicle crashed into, speed limit, vehicle weight, weight differential between vehicles, squared weight differential, past driving record of both drivers, time of day, rural, model year of vehicle, and year of crash. Partial-frontal crashes are those in which either the initial or principal point of impact is between ten and two on the clock face (excluding crashes qualifying as direct-frontal). White standard errors accounting for within-vehicle correlations are shown in parentheses. Standard errors are computed using the delta method.

TABLE 5.—FRACTION OF LIVES SAVED BY SEAT BELTS AND AIR BAGS RELATIVE TO DEATH RATES WITH NO SAFETY DEVICE: NONFRONTAL CRASHES

	Full Sample			With Sample Selection Correction (Anyone Dies in Other Vehicle)		
	(1)	(2)	(3)	(4)	(5)	(6)
Automobiles						
Seat belt only	0.21	0.24	0.23	0.65	0.67	0.64
	(0.01)	(0.01)	(0.01)	(0.06)	(0.06)	(0.06)
Air bag only	−0.03	−0.05	−0.03	−0.02	−0.02	0.02
	(0.02)	(0.02)	(0.02)	(0.07)	(0.07)	(0.09)
Probability of death with no seat belt, no air bag		0.672			0.317	
Number of observations		8,320			950	
Utility vehicles, vans, and small trucks						
Seat belt only	0.54	0.53	0.51	0.78	0.78	0.77
	(0.02)	(0.02)	(0.02)	(0.08)	(0.08)	(0.09)
Air bag only	−0.01	0.04	0.12	−0.15	−0.07	−0.10
	(0.04)	(0.03)	(0.04)	(0.14)	(0.13)	(0.18)
Probability of death with no seat belt, no air bag		0.604			0.096	
Number of observations		2,551			902	
Set of covariates included?	None	Limited	Full	None	Limited	Full

Values in table are the fraction of lives saved relative to a vehicle occupant with no seat belt or air bag. The sample is limited to frontseat occupants of vehicles involved in two-vehicle crashes in which a motorist dies. The first three columns correspond to the whole sample; the last three provide sample selection-corrected estimates. The values in the table are computed from linear probability regressions of an indicator variable for death on seat belt and air bag dummies and an interaction between those two variables. The raw probability of death for individuals with neither a seat belt nor an air bag are also reported. Columns 1 and 4 are coefficients from simple regressions with no covariates. Columns 2 and 5 add a partial set of covariates: age dummies, gender, past driving record controls, vehicle weight, and whether the person was the driver. Columns 3 and 6 include a full set of covariates: seat position, sex, age, the type of vehicle crashed into, speed limit, vehicle weight, weight differential between vehicles, squared weight differential, past driving record of both drivers, time of day, rural, model year of vehicle, and year of crash. Nonfrontal crashes are all crashes that are not direct-frontal or partial-frontal. White standard errors accounting for within-vehicle correlations are shown in parentheses. Standard errors are computed using the delta method.

seat belt effectiveness rises from approximately 60% to 75% when taking into account sample selection. Air bags are much less effective in partial-frontal crashes than in direct-frontal crashes, as would be predicted. Indeed, after the sample selection correction, none of the air bag coefficients are statistically significant, and the point estimates are actually negative for non-automobiles.

Table 5 presents the results for nonfrontal crashes. The structure of the table is identical to the preceding two tables. The results for seat belts parallel the findings on direct-frontal and partial-frontal crashes: large increases in seat belt effectiveness when correcting for sample selection and overall success rates of seat belts of 65% to 80%. There is no evidence that air bags provide any protection to occupants in nonfrontal crashes.

The other covariates in the regressions underlying tables 3 through 5 are plausibly estimated. Full regression results corresponding to column 6 of table 3 are reported in table A1 in appendix A. (Full results for all specifications are available on request from the authors.[18]) Drivers are two to three percentage points more likely to die than right-frontseat passengers. Gender is not an important predictor, but the probability of death is an increasing function of age. Even controlling for the weight differential of the vehicles involved in the crash (occupants of heavier vehicles fare better), survival probabilities are greater when crashing into an automobile than a utility vehicle, van, or small truck. The chance of death is higher during the night, in rural areas, and on roads with posted speed limits of 55 miles per hour or more. Bad previous driving records (both for your own driver and for the driver of the other vehicle) are associated with slightly higher death rates. There does not appear to be a systematic trend in death rates for later model-year vehicles once seat belts and air bags are controlled for, suggesting that other vehicle safety design innovations between 1990 and 1997 have not had a dramatic impact on crash survival.

[18] The signs and magnitudes of the covariates are generally similar across different types of crash impacts. The air bag coefficient is the only one that is highly sensitive to crash type.

A. *Comparison to Previous Estimates*

It is useful to compare the magnitude of our estimates for seat belts and air bags to previous values in the literature. The commonly accepted range of seat belt effectiveness in reducing fatalities is from 45% to 50% (Evans, 1986; Graham et al., 1997; Kahane, 1996; NHTSA, 1996). Our estimates correcting for sample selection range from 50% to 78%, with a median of 67%. Thus, we find seat belts to be substantially more effective than previous estimates. This result does not appear to be an artifact of the sample we use: ignoring sample selection, we actually obtain lower benefits of seat belts than previous studies.

Our estimates for air bags, however, are lower than previous values. We find air bags to be roughly 15% effective in direct frontal crashes. In comparison, NHTSA (1996) reports a value of 31%, and Zador and Ciccone (1993) finds a 28% reduction. Braver et al. (1997) find a 20% effectiveness for right-front passengers when excluding those under the age of ten (which makes their sample

TABLE 6.—SENSITIVITY ANALYSIS OF THE REDUCTION IN DEATH ESTIMATES FOR SEAT BELTS AND AIR BAGS IN DIRECT-FRONTAL CRASHES

	Automobiles		Utility Vehicles, Vans, and Light Trucks	
Sub-Group Analyzed	Seat Belt Coefficient	Air Bag Coefficient	Seat Belt Coefficient	Air Bag Coefficient
Baseline	0.51	0.14	0.65	0.17
	(0.03)	(0.04)	(0.03)	(0.05)
Add make and model dummies	0.55	0.10	0.70	0.09
	(0.03)	(0.06)	(0.03)	(0.08)
Add vehicle-fixed effects	0.37	0.15	0.40	0.24
	(0.09)	(0.08)	(0.13)	(0.15)
Driver	0.51	0.10	0.66	0.13
	(0.04)	(0.04)	(0.03)	(0.06)
Front-seat passenger	0.50	0.30	0.61	0.45
	(0.06)	(0.08)	(0.07)	(0.18)
Males	0.53	0.12	0.66	0.16
	(0.04)	(0.06)	(0.03)	(0.06)
Females	0.49	0.17	0.63	0.14
	(0.05)	(0.05)	(0.06)	(0.11)
Exclude age <16	0.50	0.14	0.64	0.17
	(0.03)	(0.04)	(0.03)	(0.05)
Top third of crashes by severity	0.42	0.07	0.50	0.28
	(0.05)	(0.06)	(0.05)	(0.08)
Middle third of crashes by severity	0.57	0.20	0.78	−0.08
	(0.05)	(0.06)	(0.05)	(0.12)
Bottom third of crashes by severity	0.64	0.10	0.85	0.25
	(0.08)	(0.11)	(0.04)	(0.08)
Limit sample to model years 1990–1994	0.52	0.13	0.66	0.22
	(0.03)	(0.04)	(0.03)	(0.07)
Missing data coded as wearing seat belt	0.48	0.15	0.60	0.19
	(0.03)	(0.04)	(0.03)	(0.06)
Missing data coded as not wearing seat belt	0.46	0.17	0.60	0.21
	(0.03)	(0.04)	(0.03)	(0.06)

Specifications in this table are variations on the results reported in column 6 of table 3. The values in the table are the fraction of lives saved by safety devices relative to occupants with no seat belt or air bag. For fuller description, see notes to table 3. Direct-frontal crashes are those in which the principal point of impact is completely frontal (twelve on the clock face). White standard errors accounting for within-vehicle correlations are shown in parentheses. Standard errors are calculated using the delta method.

comparable to ours). As with seat belts, the difference between our results and previous research is not due to a difference in samples. When we apply the standard "double-pair comparison" methodology to our sample of crashes, we obtain estimates of air bag effectiveness of 42% in direct-frontal crashes, higher than any of the studies just cited. Moreover, we find little or no benefit of air bags for crashes other than direct-frontal impacts, whereas small positive estimates in such crashes characterize the previous literature. Extrapolating our coefficients to all fatal crashes and taking into account the relative frequencies of frontal and nonfrontal crashes, we obtain an overall air bag effectiveness of 8%. Graham et al. (1997) surveys the literature and concludes that 13% is the best available estimate for adults based on the previous literature, an estimate that is 60% higher than ours.

B. *Sensitivity Analysis*

Table 6 presents sensitivity analysis from a range of alternative specifications. Only direct-frontal, selection-corrected specifications with the full set of covariates are included in table 6. Results are once again broken down by automobiles versus other vehicles. The top row of the table shows the baseline estimates taken from column 6 of table 3. The next two rows add indicator variables to soak up potential unobserved heterogeneity. Adding make and model dummies to the regressions has little impact on the results. Including vehicle-fixed effects reduces the measured effectiveness of seat belts, but increases the coefficient on air bags. The precision of the estimates decreases substantially because more than 70% of the vehicles in the sample have only a single passenger and thus provide no information when vehicle-fixed effects are included.

The next nine rows of table 6 focus on subsets of the data to isolate potential differences in safety device effectiveness across population subgroups or types of crashes. Comparing the results for drivers and frontseat passengers, seat belts look equally effective, but air bags appear substantially more beneficial to passengers than drivers. This result appears plausible in light of the fact that the air bag deploys from the dashboard rather than the steering wheel for frontseat passengers. The greater distance between the air bag and the occupant both increases the likelihood that the air bag will have inflated in time to be of use, and lessens the chance that the occupant will be close to the explosion that triggers the air bag with potential negative consequences. There is little evidence of differential impacts across men and women. Excluding those under the age of sixteen does not substantially change the results, but it is important to bear in mind that children in car seats are excluded from our baseline sample.

To determine whether the effectiveness of safety devices varies with crash severity, we categorize individuals in crashes by the predicted severity of the impact. Using information on crash characteristics (such as relative vehicle weights, the posted speed limit, time of day, urban versus rural), we run a probit to determine the predicted likelihood of death for each individual. We then divide the sample into thirds according to crash severity. The results estimated on each of these three subgroups are reported in the table. Seat belt effectiveness monotonically declines as crash severity increases for both automobiles and other vehicles. The greater overall effectiveness of seat belts in larger vehicles is consistent with the fact that, on average, occupants of larger vehicles do not suffer as severe of impacts. The coefficient on air bags does not exhibit any systematic pattern.

Air bags went from being almost nonexistent to standard equipment over a short period of time. As a consequence, most of the variance in the presence of air bags comes across rather than within model years. To the extent that there are other changes in vehicles over time that we are not adequately capturing with covariates, the air bag estimates may be biased.[19] To test this possibility, we restrict the

[19] Stated more formally, for those with and without air bags, there is little overlap in the distribution of the propensity score (Rosenbaum & Rubin, 1983) because model year is a good predictor of air bag status. In contrast, the covariates are not good predictors of seat belt status, so there is less

sample to vehicles built between 1992 and 1994, the period over which air bag prevalence rose from 18% to 68% for automobiles in our sample. As shown in table 6, for this limited subset of vehicles, both seat belt and air bag coefficients remain similar to baseline estimates.

As noted earlier, 8% of the individuals have missing values for seat belt usage. In the results presented thus far, these observations were dropped. The bottom two rows of the table present results if the missing values are included under the assumption that all of these individuals were wearing seat belts, or that none were wearing seat belts. The estimated seat belt effectiveness falls only slightly under each of these assumptions. The air bag coefficients are not greatly affected.

IV. Calculating the Number of Lives Saved by Seat Belts and Air Bags

In terms of understanding the value of seat belts and air bags, it is useful to translate the effectiveness estimates into numbers of lives saved. Calculations of lives saved do not require estimates of the fraction of drivers wearing seat belts (or driving vehicles equipped with air bags), but rather depend solely on r reported in the tables and the number of actual deaths with and without the safety device.[20] Let D_S and D_0 equal, respectively, the number of individuals who die with and without a given safety device. The number of lives saved as a result of the current level of usage of that safety device is equal to $(r/(1 - r))D_S$. The number of cases in which seat belts failed to save lives provides the key to how many lives they actually saved, for a given r. Using 0.60 as a (conservative) overall estimate of seat belt effectiveness and extrapolating our estimates out of sample to the set of all fatal crashes, we calculate that the lives of slightly more than 15,000 frontseat occupants were saved by seat belts in 1997 alone. Given that there were almost 50,000 motor vehicle-related fatalities (including pedestrians, backseat passengers, and so on), our estimates suggest that, without seat belts, total fatalities would have been at least 30% higher. As important as seat belts are in saving lives, however, increased usage can explain less than half of the observed decline in fatality rates per vehicle mile traveled between 1980 and 1997.

Lives saved due to air bags are substantially lower. Using an effectiveness of 0.15 for air bags in direct-frontal crashes and zero in other crashes yields an estimate of roughly 550 lives saved in 1997, or less than 5% of the number of lives saved by seat belts.

Given the comparatively limited effectiveness of air bags, the possibility that seat belt usage declines with the availability of an air bag becomes an important public policy concern. If one in eight seat belt wearers decided not to buckle up because of a mistaken notion that the air bag would provide complete protection, the net impact of air bags on saving lives would then be negative. There does not appear to be any evidence for this kind of behavioral response, either in survey data (Williams, Wells, & Lund, 1990) or in our sample. In the raw data, seat belt usage is higher among those with air bags. After controlling for an extensive set of observable variables, including whether other frontseat occupants in the vehicle are wearing seat belts, there does not appear to be any systematic relationship between the availability of air bags and seat belt usage.

In addition to calculating the number of lives saved, one might also be interested in the additional number of lives that would be saved if occupants who currently do not use seat belts and air bags were to use them. This value is given simply by rD_0. If all frontseat occupants had worn seat belts in 1997, almost 11,000 further deaths would have been averted. If all vehicles were equipped with dual air bags, an extra 1,700 lives would have been saved.

Given the estimates above, it is clear that seat belts are an extremely good investment from the perspective of cost-benefit analysis. The annual expenditure on equipping vehicles with seat belts is roughly $500 million, yielding a crude estimate of the cost per life saved of roughly $30,000.[21] In comparison, more than $4 billion dollars are spent annually on air bag installation and maintenance. If all vehicles had dual air bags, approximately 2,250 lives would be saved annually, for an average cost per life saved of $1.8 million. Note that these estimates understate the true benefits of both seat belts and air bags because they focus solely on deaths prevented, ignoring any impact the safety devices have in reducing injury severity. Nonetheless, at least based on these rough calculations, air bags do not appear to be a particularly effective investment in public health.[22]

V. Conclusion

This paper presents estimates of the effectiveness of seat belts and air bags in saving lives that overcome sample

concern that the results will depend critically on the way in which the covariates enter the regression.

[20] As a sidelight, we note that it is possible given r and the number of deaths with and without seat belts to back out an estimate of the fraction of frontseat passengers wearing a seat belt. That fraction is equal to $(D_S/(1 - r))/(D_S/(1 - r)) + D_0)$. Based on all frontseat occupants dying in fatal crashes in 1997 (not just our limited sample of two-vehicle crashes), we estimate a seat belt usage rate of 58%. This number is lower than estimates based on survey responses, but virtually identical to values obtained in observational studies in which researchers actually tallied seat belt usage by the roadside (NHTSA, 1996).

[21] This calculation assumes a steady state in the number of vehicles on the road, so that the flow of new vehicles offsets the number of vehicles retired annually. If this is the case, then the annual investment in seat belts is constant over time, allowing one to calculate the cost per life saved by simply dividing the annual expenditure on seat belts by annual lives saved.

[22] Graham et al. (1997) performs a much more careful cost-benefit analysis that includes reductions in injuries. Their conclusion with respect to air bags is somewhat more favorable than ours, both because injuries are considered and because they assume a higher effectiveness rate of air bags. In the sample of crashes we examine, seat belts and air bags are not nearly as successful in eliminating injuries or reducing their severity as they are in averting death. It may be the case, however, that, in less severe crashes, injury reduction is more effective.

selection bias inherent to fatal-crash data. We find that wearing a seat belt reduces the likelihood of death by roughly 60%, and air bags reduce the probability of death by approximately 16% in direct-frontal impacts and 9% in partial-frontal impacts. Based on our estimates, seat belts are more effective than is generally thought, whereas air bags are less effective. If our estimates are correct, roughly 15,000 lives were saved by seat belt usage in 1997, along with roughly 550 lives saved by air bags. The benefit-cost ratio of seat belts is more than fifty times greater than that of air bags. More generally, this paper presents an example of how nonstandard approaches can provide useful estimates even when naive estimation provides coefficients that are clearly biased.

REFERENCES

Angrist, Josh, and Alan Krueger, "Empirical Strategies in Labor Economics" (ch. 23), in Orley Ashenfelter and David Card (Eds.), *Handbook of Labor Economics,* vol. 3. (Amsterdam: Elsevier, 1999).

Braver, E. R., S. A. Ferguson, M. A. Greene, and A. K. Lund, "Reductions in Deaths in Frontal Crashes among Right Front Passengers in Vehicles Equipped with Passenger Air Bags," *JAMA* 278:17 (1997), 1437–1439.

Evans, L., "The Effectiveness of Safety Belts in Preventing Fatalities," *Accident Analysis and Prevention* 18 (1986), 229–241.

Fell, James C., and Carl E. Nash, "The Nature of the Alcohol Problem in U.S. Fatal Crashes," *Health Education Quarterly* 16:3 (1989), 335–343.

Ferguson, S. A., A. K. Lund, and M. A. Greene, *Driver Fatalities in 1985–1994 Air Bag Cars* (Arlington, VA: Insurance Institute for Highway Safety, 1995).

Graham, John D., Kimberly Thompson, Sue Goldie, Maria Segui-Gomez, and Milton Weinstein, "The Cost-Effectiveness of Air Bags by Seating Position," *JAMA* 278:17 (1997), 1418–1425.

Grossman, Michael, Frank J. Chaloupka, Henry Saffer, and Adit Laixuthai, "Effects of Alcohol Price Policy on Youth," NBER working paper no. 4385 (1993).

Heckman, James, "Sample Selection Bias as a Specification Error," *Econometrica* 47:1 (1979), 153–161.

———, "Varieties of Selection Bias," *American Economic Review, Papers and Proceedings* 80 (1990), 313–338.

Heckman, James, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd, "Sources of Selection Bias in Evaluating Social Programs: An Interpretation of Conventional Measures and Evidence on the Effectiveness of Matching as a Program Evaluation Method," *Proc. Natl. Acad. Sci. USA* 93 (November 1996), 13416–13420.

Heckman, James, Robert LaLonde, and Jeffrey Smith, "The Economics and Econometrics of Active Labor Market Programs" (ch. 31), in Orley Ashenfelter and David Card (Eds.), *Handbook of Labor Economics,* vol. 3, (Amsterdam: Elsevier, 1999).

Hollands, C. M., F. K. Winston, P. W. Stafford, and S. J. Shochat, "Severe Head Injury Caused by Airbag Deployment," *Journal of Trauma* 41:5 (1996), 920–922.

Imbens, Guido, and Joshua Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica* 62 (March 1994), 467–475.

Kahane, C., *Fatality Reduction by Air Bags: Analyses of Accident Data Through Early 1996* (Washington, DC: National Highway Traffic Safety Administration, 1996).

Kaplan, Burton, and R. Adams Cowley, "Seatbelt Effectiveness and Cost of Noncompliance among Drivers Admitted to a Trauma Center," *American Journal of Emergency Medicine* 9:1 (1991), 4–10.

Larkin, Gregory L., and James E. Weber, "Letter to the Editor," *JAMA* 279:7 (1998), 506.

Lave, Charles, and Patrick Elias, "Resource Allocation in Public Policy: The Effects of the 65-MPH Speed Limit," *Economic Inquiry* 35 (July 1997), 614–620.

Levitt, Steven, and Jack Porter, "Estimating the Effect of Alcohol on Driver Risk Using Only Fatal Accident Statistics," *Journal of Political Economy,* forthcoming (2001).

Lund, A. K., and S. A. Ferguson, "Driver Fatalities in 1985–1994 Air Bag Cars," *Journal of Trauma* 38 (1995), 49–475.

Marine, William, Edward Kerwin, Ernest Moore, Dennis Lezotte, Anna Baron, and Michael Grosso, "Mandatory Seatbelts: Epidemiologic, Financial, and Medical Rationale from the Colorado Matched Pairs Study," *Journal of Trauma* 35:1 (1994), 96–100.

Morris, Michael S., and Leticia Borja, "Letter to the Editor," *Journal of Trauma: Injury, Infection, and Critical Care* 44 (January 1998), 238–239.

National Highway Traffic Safety Administration, *Final Regulatory Analysis: Amendment to FMVSS No. 208, Passenger Car Front Seat Occupant Protection* (Washington, DC: NHTSA, 1984).

———, *The Effects of the 65 MPH Speed Limit through 1988: A Report to Congress* (Washington, DC: NHTSA, 1989).

———, *Effectiveness of Occupant Protection Systems and Their Use* (*Third Report to Congress*) (Washington, DC: NHTSA, 1989).

———, *Traffic Safety Facts 1997* (Washington, DC: NHTSA, 1998).

Orsay, Elizabeth, Timothy Turnbull, Mary Dunne, John Barrett, Patricia Langenberg, and Charles Orsay, "Prospective Study of the Effect of Safety Belts on Morbidity and Health Care Costs in Motor-Vehicle Accidents, *JAMA* 260:23/30 (1988), 3598–3603.

Peltzman, Sam, "The Effects of Automobile Safety Regulation," *Journal of Political Economy* 83 (August 1975), 677–725.

Robertson, L. S., "Estimates of Motor Vehicle Seat Belt Effectiveness and Use: Implications for Occupant Crash Protection," *American Journal of Public Health* 66 (September 1976), 859–864.

Rosenbaum, Paul, and Donald Rubin, "The Central Role of the Propensity Score in Observational Studies for Casual Effects," *Biometrick* 70:1 (1983), 41–55.

Rubin, Donald, "Comment: Neyman (1923) and Causal Inference in Experiments and Observational Studies," *Statistical Science* 5 (1990), 472–480.

Ruhm, Christopher, "Alcohol Policies and Highway Vehicle Fatalities," *Journal of Health Economics* 15 (August 1996), 435–454.

Saffer, Henry, and Michael Grossman, "Drinking Age Laws and Highway Mortality Rates: Cause and Effect," *Economic Inquiry* 25 (July 1987), 403–417.

Williams, A. F., J. K. Wells, and A. K. Lund, "Seat Belt Use in Cars with Air Bags," *American Journal of Public Health* 80:12 (1990), 1514–1516.

Zador, Paul L., and Michael A. Ciccone, "Automobile Driver Fatalities in Frontal Impacts: Air Bags Compared with Manual Belts," *American Journal of Public Health* 83:5 (1993), 661–665.

Zobeck, Terry S., Gerald D. Williams, Bridget F. Grant, and Darryl Bertolucci, "Years of Potential Life Lost and Other Trends in Alcohol-Related Fatal Traffic Crashes: 1977–1987, *Alcohol Health and Research World* 14:1 (1990), 63–68.

APPENDIX: PROOF OF THEOREM 1:

PROOF: Fix a conditioning set $X = x$. A1 implies

$$\frac{\Pr(Y^{0,0} = 1, S = 1, A = 0|X = x, P = 1)}{\Pr(Y^{0,0} = 1, S = 0, A = 0|X = x, P = 1)} = \frac{\Pr(Y^{1,0} = 1, S = 1, A = 0|X = x, P = 1)}{\Pr(Y^{1,0} = 1, S = 0, A = 0|X = x, P = 1)}$$

We will show that, when conditioning on $Q = 1$ rather than $P = 1$, the above equality becomes a strict inequality violating the independence in assumption A1′. Note, if $S = 0$, $A = 0$, $Y^{0,0} = 1$, then $Y = 1$ and $P = Q = 1$. So

$$1(Y^{0,0} = 1, S = 0, A = 0, P = 1) = 1(Y^{0,0} = 1, S = 0, A = 0, Q = 1) \text{ (almost everywhere)}.$$

Similarly, $1(Y^{1,0} = 1, S = 1, A = 0, P = 1) = 1(Y^{1,0} = 1, S = 1, A = 0, Q = 1)$. Also, safety device effectiveness implies $Y^{0,0} \geq \max\{Y^{1,0}, Y^{1,0}, Y^{1,1}\}$, so $Y^{1,0} = 1$ implies $Y^{0,0} = 1$. If $S = 0$, $A = 0$, $Y^{1,0} = 1$, then $Y = Y^{0,0} = 1$ and $Q = 1$. Thus,

$$1(Y^{1,0} = 1, S = 0, A = 0, P = 1) = 1(Y^{1,0} = 1, S = 0, A = 0) = 1(Y^{1,0} = 1, S = 0, A = 0, Q = 1),$$

where the first equality follows by $Y^{1,0} = 1$ implying $P = 1$. Finally, $Y^{0,0} = 1$ implies $P = 1$, so

$$1(Y^{0,0} = 1, S = 1, A = 0, P = 1) = 1(Y^{0,0} = 1, S = 1, A = 0) = 1(Y^{0,0} = 1, S = 1, A = 0, Q = 1) + 1(Y^{0,0} = 1, S = 1, A = 0, Q = 0)$$

and

$$\Pr(Y^{0,0} = 1, S = 1, A = 0, P = 1 | X = x) = \Pr(Y^{0,0} = 1, S = 1, A = 0, Q = 1 | X = x) + \Pr(Y^{0,0} = 1, S = 1, A = 0, Q = 0 | X = x) > \Pr(Y^{0,0} = 1, S = 1, A = 0, Q = 1 | X = x),$$

$$\frac{\Pr(Y^{0,0} = 1, S = 1, A = 0 | X = x, P = 1)}{\Pr(Y^{0,0} = 1, S = 0, A = 0 | X = x, P = 1)} = \frac{\Pr(Y^{0,0} = 1, S = 1, A = 0, P = 1 | X = x)}{\Pr(Y^{0,0} = 1, S = 0, A = 0, P = 1 | X = x)} = \frac{\Pr(Y^{0,0} = 1, S = 1, A = 0, P = 1 | X = x)}{\Pr(Y^{0,0} = 1, S = 0, A = 0, Q = 1 | X = x)} > \frac{\Pr(Y^{0,0} = 1, S = 1, A = 0, Q = 1 | X = x)}{\Pr(Y^{0,0} = 1, S = 0, A = 0, Q = 1 | X = x)} = \frac{\Pr(Y^{0,0} = 1, S = 1, A = 0 | X = x, Q = 1)}{\Pr(Y^{0,0} = 1, S = 0, A = 0 | X = x, Q = 1)},$$

and so

$$\frac{\Pr(Y^{1,0} = 1, S = 1, A = 0 | X = x, Q = 1)}{\Pr(Y^{1,0} = 1, S = 0, A = 0 | X = x, Q = 1)} = \frac{\Pr(Y^{1,0} = 1, S = 1, A = 0 | X = x, P = 1)}{\Pr(Y^{1,0} = 1, S = 0, A = 0 | X = x, P = 1)} = \frac{\Pr(Y^{0,0} = 1, S = 1, A = 0 | X = x, P = 1)}{\Pr(Y^{0,0} = 1, S = 0, A = 0 | X = x, P = 1)} > \frac{\Pr(Y^{0,0} = 1, S = 1, A = 0 | X = x, Q = 1)}{\Pr(Y^{0,0} = 1, S = 0, A = 0 | X = x, Q = 1)},$$

and assumption A1′ does not hold.[23] ☒

[23] It is assumed that there exists x such that $\Pr(Y^{0,0} = 1, S = 1, A = 0, Q = 0 | X = x) > 0$, $\Pr(P = 1 | X = x) > 0$ and $\Pr(Q = 1 | X = x) > 0$, which is clearly true in the real crash environment so we do not include this as an explicit assumption in the claim.

TABLE A1.—FULL ESTIMATES OF THE RAW REGRESSION COEFFICIENTS FOR DIRECT FRONTAL CRASHES

	Automobiles		Utility Vehicles, Vans, and Light Trucks	
Variable	Coefficient	Standard Error	Coefficient	Standard Error
Seat belt	−0.128	0.013	−0.076	0.007
Air bag	−0.036	0.010	−0.020	0.007
Driver	0.028	0.009	0.020	0.006
Male	−0.010	0.008	0.001	0.005
Age 16–40	0.015	0.016	0.014	0.008
Age 41–65	0.075	0.017	0.047	0.010
Age 65+	0.226	0.024	0.163	0.021
Crash into utility, van or small truck	0.058	0.018	0.048	0.007
Speed limit less than 55 mph	−0.091	0.009	−0.035	0.005
Vehicle weight ($\times 10^3$)	−0.046	0.009	0.000	0.002
Weight differential * lighter ($\times 10^3$)	0.019	0.017	0.054	0.018
(Weight differential)^2 * lighter ($\times 10^8$)	−0.098	0.295	−0.599	0.263
Weight differential * heavier ($\times 10^3$)	−0.040	0.010	−0.015	0.004
(Weight differential)^2 * heavier ($\times 10^8$)	0.714	0.139	0.078	0.030
Previous minor incidents (own driver)	0.002	0.009	0.008	0.005
Previous major incidents (own driver)	0.012	0.015	0.006	0.011
Previous minor incidents (other driver)	0.010	0.109	−0.001	0.006
Previous major incidents (other driver)	0.025	0.017	0.021	0.012
Time of day: 6 AM to 8 PM	−0.088	0.019	−0.049	0.015
Time of day: 8 PM to 1 AM	−0.065	0.020	−0.023	0.016
Rural	0.043	0.009	0.018	0.005
Model year: 1996	−0.033	0.025	−0.016	0.014
Model year: 1995	−0.000	0.024	−0.005	0.014
Model year: 1994	0.002	0.025	−0.011	0.014
Model year: 1993	−0.000	0.025	−0.005	0.015
Model year: 1992	−0.018	0.025	−0.007	0.016
Model year: 1991	−0.010	0.026	−0.010	0.016
Model year: 1990	0.010	0.026	−0.013	0.016
Crash year: 1997	−0.015	0.013	0.010	0.008
Crash year: 1996	−0.005	0.013	−0.001	0.008
Crash year: 1995	−0.009	0.013	0.003	0.008
Constant	0.417	0.043	0.136	0.027
Number of observations	6,752	—	8,493	—
R^2	0.125	—	0.092	—

Values in the table are the raw regression coefficients underlying the specifications reported in column 4 of table 3 (direct-frontal crashes with the sample selection correction). See notes to table 3 for a fuller description of the sample. White standard errors accounting for within-vehicle correlations are shown in parentheses.