Lecture Note 15: Risk, Safety Regulation and the Value of a Statistical Life

David Autor, Massachusetts Institute of Technology

14.03/14.003, Microeconomic Theory and Public Policy, Fall 2010

1 Risk and Safety Regulation

[Most of this material in this subsection is for self-study and will not be covered or tested in class.]

- We have so far studied risk as an individual-level consumer problem. But it is also a societal problem.
- How much risk should we subject ourselves to? Alternatively, how much safety should we demand? And how much are we willing to pay for it?
- Let's be clear: safety is a 'good' and we buy it by giving up other things (time, adventure, money, convenience). The notion that "the more safety the better" (or "safety first") is not a sound economic concept. Ultimately, there is an optimal level of safety (for an individual or a society) and we could consume either too little or too much.
- We will take three angles of attack on this question.

  1. We'll first consider product safety and liability law. Some very simple models establish how the market might optimally make choices about safe products, and how the legal regime (specifically, liability law) affects the level of safety provided by the market and who bears the cost.
  2. We'll next consider reasons why the market might not provide the optimal level of safety, and consider possible public policy responses to these 'market failures.' These include:

     (a) Private mechanisms (these are actually a market response) including information provision, warranties and reputation

(b) Licensing requirements
(c) Facilitating provision of information
(d) Requiring provision of information
(e) Establishing legal liability standards
(f) Limiting or banning products

3. We'll finally consider an interesting empirical evaluation of the willingness of society to pay for safety (at the margin). This is the paper by Ashenfelter and Greenstone, "Using Mandated Speed Limits to Measure the Value of a Statistical Life."

2 Legal liability

- It's useful to start with a few simple legal concepts.

- Let's say that there is a product that gives utility $U > 0$ and has probability $p > 0$ of harming the buyer with monetized damages of $d > 0$. Assume that $U > MC + pd$, where MC is the cost of production. Consumers would purchase this product, despite the risk of harm.

- Question: Although consumers would be willing to buy the product 'as is,' *should* (normatively) the manufacturer make the product safer? The answer depends on the manufacturer's cost.

- Assume the manufacturer's cost of making the product completely safe is $b > 0$. Clearly, if $b < pd$, economic efficiency demands that the manufacturer *should* do it.

- Next question: Will the manufacturer make the product safer if it is economically efficient for her to do so? The answer to this question may depend on the legal liability regime.

2.1 Liability Regime 1: Strict Liability

- Consider first a legal regime of **strict liability**. Here, the manufacturer is monetarily responsible for *any* harm the consumer suffers from the product, regardless of why it occurs.

- Under this regime, one of two things will occur:

  1. If $b \leq pd$, the manufacturer will make the safety modification.

2. If $b > pd$, the manufacturer will not modify the product, but it will reimburse consumers for harms they suffer.

- Assuming the product is sold competitively, the price will be equal to

$$P = MC + \min[b, pd],$$

where MC is the marginal cost of production (prior to the safety modification).

2.2 Liability Regime 2: Caveat Emptor

- An alternate regime is 'caveat emptor' – buyer beware. Here, the buyer faces the full risk of any damage incurred.
- You might assume that this regime provides insufficient incentives for manufacturers to make products safe. But consider carefully.
- If the buyer is fully informed about the attributes of the product, his willingness to pay is:

$$U - pd \text{ if the product is unsafe,}$$
$$U \text{ if the product is safe,}$$

where U is the monetized utility value of consumption.

- The profit maximizing firm will therefore find it optimal to spend b to make the product safe if $b \leq pd$. It will not spend b to make the product safe if $b > pd$.
- This outcome is identical to the strict liability case above. The efficient action occurs under either liability regime. If the product is safe, its equilibrium price cannot exceed U. If it is unsafe, it's equilibrium price cannot exceed $U - pd$. The price again incorporates the riskiness of the product.

2.3 Liability Regime 3: Comparative negligence

- If either extreme of liability law produces the same outcome, and this outcome is efficient, why worry?
- A problem arises when both parties can take actions to avert the harm. In many cases, both victim and 'victimizer' have access to precautions that reduce the probability and magnitude of harm. For example, a manufacturer can put a blade guard on a lawn mower

so that consumers are unable to stick their fingers in the blade while it is turning. On the other hand, consumers could spend a moment to verify that the mower is off before reaching under it. Which party should do what?

- Consider a case where either party—the manufacturer or the consumer—can avert the loss. Their costs of abating the loss are $b_m, b_c > 0$ where m, c denote manufacturer and consumer.
- Assume we are in a *strict liability* regime.
  - Now, the manufacturer will make the product safe if $b_m < pd$ and not otherwise.
  - Let's say, however, that $b_m > pd > b_c$. It's not efficient for the manufacturer to avert the risk, but it is efficient for the consumer to do so.
  - Will the consumer avert the risk? No. Assuming that the harm, d, is fully reimbursed to the consumer, he is indifferent as to whether or not the harm occurs.
  - Therefore, although the consumer could efficiently avoid the harm, he will not (necessarily) do so.[1]
- Assume we are in a *caveat emptor* regime.
  - Now, the consumer will avert the risk if $b_c < pd$ and not otherwise.
  - Let's say, however, that $b_m < b_c < pd$. It's efficient for either party to avert the risk, but it is *more* efficient for the manufacturer to do so.
  - Will the manufacturer avert the risk? Possibly.
  - The product price should still incorporate the risk associated with its use. Hence, the manufacturer may find it efficient to spend b_m and raise the price by b_c (which is what the consumer would otherwise spend to avert the risk).
  - However, it's conceivable that in this more complex case, the manufacturer would not spend b_m and the consumer would spend b_c instead.
- The general problem with liability regimes that place all of the responsibility for damages with one party or the other is that they don't necessarily provide incentives for the risk to be abated at lowest cost. The risk level could either be inefficiently high (as in the first example above) or the expenditures on risk abatement could be inefficiently high (as in the second example above).

[1] One could imagine a case where the manufacturer pays the consumer b_c to avert the harm, thereby saving the manufacturer $p \cdot d$. It is hard to imagine this contract being enforceable in practice.

- This motivates a regime of **Comparative Negligence** (sometimes called 'bilateral precaution'), where courts attempt to determine which party is 'more at fault' and allocate the loss to him or her.

- This idea sounds rather nebulous, but Judge Learned Hand proposed a formal standard of comparative negligence in 1947 that has stood the test of time. (And as economist Robert Cooter observed, Judge Hand framed his argument using a mathematical formula, thus endearing himself to economists in perpetuity.)

- Judge Hand found that a party is at fault for a loss *iff* its cost of avoiding the accident was less than the resulting harm multiplied by its probability. This is exactly equivalent to our efficiency condition above. The manufacturer should avert if $b_m < pd$ and the consumer if $b_c < pd$.

- Hence, if $b_c < b_m < pd$ and the accident occurs, the consumer would be held responsible.

- Similarly, if $b_m < b_c < pd$ and the accident occurs, the manufacturer would be held responsible.

- If $b_m, b_c > pd$, then no one is at fault.

- This rule provides incentives for the lowest cost party to avert the risk, which is exactly what we want.

- You can think of the Comparative Negligence regime as equivalent to a strict liability regime when the manufacturer has the lower cost of averting the harm and as equivalent to a *caveat emptor* regime when the consumer has the lower cost of averting the harm. Note that that neither party is found negligent when $b_m, b_c > pd$.

- Question: What if there is less than perfect enforcement, for example, the probability of legal punishment conditional on an accident occurring is $0 < \gamma < 1$?

2.4 The Limits of Liability

- The above examples might suggest that we can always achieve the efficient level of safety simply by using an efficient liability regime. But this is a bit far-fetched.

- There are a number of reasons why an efficient liability regime won't solve all safety problems—that is, the market will fail. The sources of market failure emanate from several areas:

1. Transaction costs: It may be costly to exercise the liability rights given by the law:
   - Courts may make random errors. Even if they usually get it right, they will be imperfect enforcers. In this case the liable party doesn't face the full expected cost of a loss because there is some chance of court error – meaning that the expected loss to the manufacturer is strictly less than pd.
   - It may be expensive to litigate, in which case, a consumer might not bring suit even if $b_m < pd$. This creates two problems. One, the consumer is not compensated for the harm. Two, the manufacturer does not face full incentives to avert the harm. [This observation is probably the genesis of class action suits.]
   - Courts may have trouble determining pd, b_c, b_m. This is an information gathering problem. (See below.)
2. Information failures:
   - Often purchasers are ill-informed about the quality of products. This makes it hard for them to make optimal decisions about the costs and benefits of risk.
   - In the examples above, pd, b_c, b_m were understood by consumers and manufacturers. But often consumers do not know either p, d, let alone their cost of abating that risk (b_c).
   - Manufacturers may also not know the risks (and they may even face incentives not to know this information if it could be shown in court that they knew that, for example, $b_m < pd$).
   - It may also be costly to obtain accurate information about risk (think of the case of prescription drugs). This is also a transaction cost. (See above.)
3. Non-economic considerations

- A liability regime may also want to distinguish between punishing a forbidden act and pricing a permitted act.
- A famous example: Ford Pintos sold during the early 1970s had a defect where they spilled fuel and caught fire when rear-ended at moderate speed. This design issue led to numerous burn deaths and horrendous injuries to Pinto drivers and passengers. During litigation, it came to light that:

  Internal Ford Motor Company documents... proved that Ford knew of the weakness in the fuel tank before the vehicle was placed on the market but that a cost/benefit study was done which suggested that it would be "cheaper" for Ford to pay liability for burn deaths and injuries rather than modify the fuel tank to prevent the fires in the first place. Dowie [a reporter for *Mother Jones* magazine] showed that Ford owned a patent on a better designed gas tank at that time, but that cost and styling considerations ruled out any changes in the gas tank design of the Pinto. [Center for Auto Safety, http:///www.autosafety.org/]

- According to our criteria above, Ford may have acting efficiently. That is: $b_m > pd$. Yet, many people do not feel comfortable with viewing liability as 'the price' of causing harm. For example, it is fairly uncontroversial that a parking ticket is the 'price' of parking illegally, but this analogy only extends so far.
- In this case, we might still use regulation to avert risk even if liability provides efficient incentives.
- [It's also possible in the case of the Pinto that $b_m < pd$ but Ford believed it faced less than perfect odds of losing in court. Say the odds of losing were equal to $\gamma < 1$. Imagine that $\gamma \cdot pd < b_m < pd$. In this case, it was profit maximizing but socially inefficient for Ford to allow the harm to occur.]

- For better or worse, we seem to be moving towards a strict liability standard in the U.S.
- A December 15, 2003 article in *Newsweek* calculated the cost of liability insurance built into several consumer items. These include:
  - Eight foot aluminum ladder - liability cost: $23.86 (20% of total cost)
  - Heart pacemaker - liability cost: $3,000.00 (16.5% of total cost)
  - Motorized wheelchair - liability cost: $170.00 (17% of total cost)
  - Tonsillectomy (doctor's fee) - liability cost: $191.00 (33% of total cost)
  - Two-day maternity stay - liability cost: $500.00 (15% of total cost)
- It's possible that this is efficient. But, in the case of ladders, for example, it's hard to believe that holding the manufacturer responsible for most/all ladder injuries provides efficient safety incentives to consumers.
- Liability regimes generally trade off manufacturers' versus users' incentives to take care with a product. The more liability is distributed to manufacturer, the less incentive the user has to be careful.

- Some argue that liability is the most efficient way of providing health insurance, e.g. at least ladder users are paying ladder accident insurance.

- The problem with this argument is moral hazard. If people are insured by manufacturers for the harms they inflict on themselves by misusing products, they are probably more likely to suffer these harms.

- Although some would have predicted otherwise, strict liability does *not* seem to reduce litigation.

- One liability problem is proving the connection between product and injury, e.g. did smoking *cause* you to have cancer?

- If it's hard to prove liability ex post, there may be more reason to regulate the product up front.

- But the negligence standard creates an incentive for the manufacturers to not know about the dangers of their products. i.e., cigarette companies put themselves in legal peril simply by conducting research on how carcinogenic their products are (assuming they don't then disclose this information). So, if negligence is defined by knowing that your product is dangerous, then you might decide to not test it.

- Liability also restricts range of choices available because many products may not come to market due to uncertainty or cost of potential liability.

3 Policies that Ameliorate Market Failure in the Provision of Risk: Some Loosely Constructed Notes

[Most of this material in this subsection is for self-study.]

3.1 Private market responses

- One market response is information providers and specialists (e.g., *Consumer Reports,* jewelry appraisers, auto inspection services, etc.). But we suspect that there is an undersupply of this information. And we'd like them to charge marginal cost – but if you go to *www.consumerreports.org*, you will find they are not charging $P = MC$.

- Warranties provided with product (establish a liability regime for that purchase):
  - It's costly to write a warranty contract

  - It will be incomplete
  - Enforcing the warranty is expensive
  - Warranties can also give rise to moral hazard (might you abuse a product because you are not liable for harm to it or to you?).

- Brand names – the manufacturer's reputation can provide an implicit and imperfect warranty, e.g., Honda wouldn't sell an unreliable car because it would damage the brand name.

3.2 Regulation

If the private market doesn't solve the problem, it may be efficient to regulate. If so, there are a number of factors that one would consider in deciding regulations:

1. *How costly is it to get the information?*
   Buying a house: you don't want to take a risk of buying a house that has hidden flaws, so you hire a building inspector to check the house's condition. But you can't do this to 'inspect' a new pharmaceutical.

2. *Ease of verification*
   Some goods are purchased often, so you can determine after one or two uses whether or not it's a good product (an "experience good"). This is true for toothbrushes. It's not true for autos, drugs, etc. In these cases, it's costly to acquire the information (see above).

3. *Technological range of choice ('ease of variation')*
   Some products are intrinsically dangerous. You cannot make a 'safe' knife. In this cases, regulation may significantly diminish the value of the product.
   But some products can be made much safer without raising cost or reducing functionality to any significant degree (e.g., defectively designed portable cribs).

4. *Extent of moral hazard*
   If the product is bound to be misused if the manufacturer is liable for harm, this may increase the case for regulation. For example, if auto manufacturers were held responsible for all auto accidents, people would drive less safely.

5. *Extent of preference for variety*
   There are some choice restrictions that government can make where almost no one will

mind. For example, no one minds banning dangerous cribs. But there are some safety restrictions that limit desirable consumer choice. For example, children's pajamas are required to be made from flame retardant fabrics. These fabrics are uncomfortable to sleep in and nasty to touch.

6. *Paternalism*
(This is not generally viewed as an economic criterion – though this is changing.) Government may believe that consumers should not be able to make certain choices for themselves (for example, the ban on marijuana use).

Now, consider types of regulation.

3.2.1 Facilitating provision of information

One way to improve market outcomes surrounding risk and safety is to facilitate the provision of information. One way to do this is to standardize the definition of certain terms, for example:

- Gas mileage
- What you can call 'cheese'
- Tire longevity ratings

Standardizing product categories facilitates credibility of product information. You may not be required to use the categories, but if you choose to use them, you must meet the definitions.

- For example, 'organic produce' has a regulatory definition. You cannot call produce 'organic' simply because it is (mostly) made up of carbon and hydrogen molecules.
- This type of regulation (creating standard categories) does not substantially restrict the range of consumer choice.
- In many cases, it may be a relatively efficient policy

3.2.2 Requiring information to be provided

- Next Step: Require the information be provided – e.g. the mileage of a new car.
- But doesn't market provide the incentive to provide this information so that certain manufacturers can capitalize on it?
- Would there be "low tar" cigarettes sold without the FDA?

- The government does not set standards for computer processor speed, memory capacity, hard drive space, monitor quality, etc. But this information is readily provided and there is substantial agreement on its validity and usefulness.

3.2.3 Licensing: Setting a minimum quality threshold

- Standards are frequently used in many types of trades and some professions. There are licensed contractors, barbers, dermatologists, teachers, truck drivers, etc.
- The existence of a licensing standard implicitly states that government can tell who has 'adequate' skills in their occupation.
- There is still likely to be a large continuum of quality above the licensing threshold.
- It's possible that the government may also set the standard too high (or too low). What if I'd prefer to hire an unlicensed teacher, day-care provider, carpenter, carpet-installer, manicurist at a lower price? .

3.2.4 Limiting or banning certain products

- What are the arguments for banning products altogether?
- Externality arguments: you don't have adequate incentive not to buy dangerous products that might hurt me, e.g., houses that burn down easily and set neighborhoods or fire.
- Information conservation arguments:
  - Instead of regulating over the counter medications, you could simply provide complete information on the label.
  - But it would be very costly for consumers to use this information (they would need a medical degree).
  - It may be efficient for the government to study the best available information and then only permit certain products
  - This yields a loss of consumer choice but it economizes on information processing
- On the other hand: some reasonable products are banned because some people have misused them.
  - For example, diving boards are almost completely gone from most public pools (more liability than regulation).

  - During the 1970s, cars were engineered so that they could not be started if a passenger was not wearing his/her seat belt. This was incredibly irritating.

- The regulatory process is itself costly, e.g. drug certification. This results in the development of fewer drugs, and creates huge fixed costs.

- There are also large transactions costs associated with the prescription system, e.g. every time I want to get my kids a dose of penicillin for an ear infection, the insurance company (or the parent) must pay for a $125 doctor's visit.

- Also, 'safe/unsafe' is an arbitrary distinction. There is a continuum.

- The spread of AIDS and subsequent efforts to develop treatments changed the public's thinking about the costs and benefits of allowing dangerous drugs to be sold prior to full testing.

4 Interesting Example: Should Infant Seats be Required on Airplanes?[2]

[For self-study.]

The U.S. National Transportation Safety Board [NTSB] is considering a regulation wherby all children younger than two years old would be required to travel in child safety restraint systems (CSR) on airplanes. This ruling would require adults to purchase seats specifically for children younger than two years instead of allowing these children to travel on an adult's lap for free (which is the current regulation).

- Consider first: Assume that parents have full information about the risks of traveling with their babies on their laps. Is it socially efficient to allow parents to make a decision to buy a seat for their babies or should the government make it mandatory?

- The NTSB calculates that there are 6.5 million 'enplanements' (plane trips) per year by children younger than 2 years. It also calculates that the expected number of child fatalities averted per year by use of CSR is 0.4. Assuming a price of an airline ticket for a child of $200 (and ignoring the direct cost of the CSR), what would be the total cost per saved life?

[2] This section draws on Newman, Thomas B., Brian D. Johnson and David C. Grossman. 2003. "Effects and Costs of Requiring Child-Restraint Systems for Young Children Traveling on Commercial Airplanse." *Archive of Pediatric & Adolescent Medicine*, 137, October, 969–974. This article will be posted to the 14.03 web site.

- Answer:

$$(6,500,500 \times \$200)/0.4 = \$3.25 \text{ billion per life saved}$$

So, from a pure cost-benefit perspective, this does not appear a attractive. If you believe that saving one life is worth $3.25 billion, consider that you could save many more lives by spending $3.25 billion elsewhere, for example, by improving drinking water quality in the developing world or by building fences around backyard swimming pools (many multiples more children die each year from accidental backyard drownings than from plane crashes.)

- Returning to airplane safety seats... Consider that some families will switch from air to car travel or vice versa depending on the relative costs of the two. Some relevant estimates (from the article cited above):

  - Approximately 6% of families with infants will choose to drive rather than pay $200 for the child seat on the plane. (This seems a very conservative estimate.)
  - The average net increase in car travel per enplanement for families switching from planes to cars is 300 miles.
  - The average vehicle occupancy for extra trips is 2.4.
  - Each person in the car has a risk of car death of 30 percent of the national average of one fatality per 100 million vehicle miles traveled.

- From these numbers, the annual effect of the seat mandate on motor vehicle deaths is:

$$\frac{6,500,000 \times 0.06 \times 2.4 \times 300}{(0.3)^{-1} \times 100,000,000} = +0.84 \text{ deaths}$$

The number of motor vehicle deaths caused by the policy is twice as large as the number of airplane deaths averted. If we assume that more than 6% of families switch from plane to cars, the policy only looks worse.

- For this reason, the Federal Aviation Administration [FAA] has so far resisted the NTSB's regulatory recommendation.

- Not all supporters of the NTSB policy find this type of argument compelling. For example, Ralph Nader and Wesley Clark in their 1994 book, *Collision Course: The Truth about Airline Safety*, write

"The argument in support of the FAA's resistance to the NTSB [National Transportation Safety Board] recommended rule mandating child safety seats is unreasonable on its face, and ridiculous in its justification. *It protects theoretical children driving in cars at the expense of real flesh-and-blood infants whose safety is unquestionably compromised when flown as a lap-baby* [italics by Autor not by authors].”

- A high ranking regulatory official in the Food and Drug Administration said of the child seat policy and the statistics above that,

  “It identifies a classic regulator’s dilemma of which risks to protect against. While the NTSB may well recognize that there could be more auto fatalities if they require car seats, those fatalities will not be blamed on them. Assuming the study above is accurate, if the NTSB does the right thing from the point of view of mitigating “total risk,” they face the prospect of getting all of the blame for allowing child fatalities on aircraft and none of the credit for preventing child fatalities on the road. Of course, if the airlines wanted to provide seats for kids under age two at a nominal cost, they might at least break even financially because parents would fly more and kids could fly safely.”

- The attorney above recommends that airlines should subsidize infant tickets rather than charge for them. For example, imagine that the FAA paid the $200 per seat so that parents incurred no additional cost to travel with infants in *CSR’s*. Would this make the policy any more attractive from a cost-benefit perspective? Clearly not. It would still cost $3.25 billion per life saved.

- Now, consider a policy where the FAA instead paid parents $200 to travel with their children *on their laps* **instead of** driving in their cars? This policy appears perverse: it would surely *increase* the number of infants killed in aircraft accidents! (What would Ralph Nader say?)

- But, if this policy were feasible, it would be more efficient from a lives-saved-per-dollar perspective.

  – As we know, the chance of an airplane fatality for an infant is $0.4/6,500,000$ per enplanement

  – The chance of an auto-fatality is $0.3/100,000,000$ per mile.

  – The crossover point where an airplane trips is safer than a car-trip is

$$\frac{M \cdot 0.3}{100,000,000} > \frac{0.4}{6,500,500} \Rightarrow M > 20.5$$

- This means (very roughly) that any trip over 6.2 miles is safer on an airplane than in a car! So, if we want to minimize travel fatalities, we should pay people to fly rather than drive. (And requiring *CSRs* on airplanes would be a waste of money under such a policy – better to use the money to subsidize more people to fly rather than drive.)
- Note that we are ignoring the direct risk of fatality from airplane crashes. Over 1987 to 1996, this was estimated at $1/5,000,000$ (that is, your chance of death on a randomly chosen flight on an international jet from the U.S. over 1987 to 1996 was one in five million). So, adding this in:

$$\frac{M \cdot 0.3}{100,000,000} > \left(\frac{0.4}{6,500,500} + \frac{1}{5,000,000}\right) \Rightarrow M > 87.2$$

- So, for trips over 88 miles, it's safer to fly. And we are assuming that drivers who choose to drive instead of fly have only 30 percent of the normal accident risk. If we assume they have the *average* accident risk, the crossover point is 26.2 miles.

5 Estimating the Statistical Value of a Human Life

- This is a topic that makes non-economists uncomfortable. But for policy analysis, there is no way around it. How much should society spend, at the margin, to save a 'statistical life?'

- A statistical life is a probabilistic concept. When we save a statistical life, we reduce the number of deaths by one *in expectation.* The value of a statistical life (VSL) is clearly very different from what we would spend to save a specific individual who was in grave danger of death. [And it is emphatically *not* the answer to the question, "What would I have to pay you to kill you?"]

- It is critical to have some knowledge of the Value of a Statistical Life because we must make regularly societal decisions about how much risk we should tolerate and how much we should spend in tax revenue or how much we should curtail freedom of choice to abate risk.

- In general, economic reasoning says that society should undertake projects that cost less than the VSL per life saved and should not undertake projects that cost more than the VSL per life saved.

- There is no 'correct' answer to the value of a statistical life. This value is something that arises out of people's preferences (and their wealth). The only thing that is certain is that the value of a statistical life is not infinite.

- How do you get a credible estimate of the VSL? Not easy. Asking people will not be very informative (they'll be horrified). But the Weak Axiom of Revealed Preference says that we can observe the VSL from the trade-offs that people (or governments) make between cost and safety.

- Speed limits are one place where that choice is very apparent. The faster people drive, the less time they spend getting from place to place. Since time has value, going more slowly is costly in foregone opportunities. However, going faster increases the probability of death.

5.1 Context

- Prior to 1973, speed limits in the U.S. were set by states. There was no national speed limit.

- With the oil crisis in 1973, the federal government imposed a national speed limit of 55 MPH.

- Although this was probably not the intention, highway fatalities fell 15 percent the following year (a reduction of nearly $10,000$ fatalities!).

- [Fatalities were also trending downward before and after 1973. This may in large part reflect advances in auto safety.]

- In 1987, with oil prices low, the federal government allowed states to raise their speed limits to 65 MPH if they wished to.

- 37 states raised their speed limits in 1987 and 3 more did so in 1988.

- It's critical for the research design that *not all* states raised their speed limits. If they had, the revealed preference argument would not be relevant. The 'treatment' here is to *offer* state legislatures the option to relax speed limits in their states. If no state *or* every state took up this offer, we wouldn't be able to bracket the implied *VSL*. If no state took up, then we could only infer that the expected loss in life was not worth the time savings. If all states took up, we would be inclined to infer that they are probably also constrained at 65 miles per hour. Hence, we would not learn about their unconstrained preference for safety versus time-savings.

5.2 Research design

- Though there is considerable technical material in this paper, the research design is straightforward.
- The plan:
  1. Contrast the change in fatalities in adopting versus non-adopting states.
  2. Contrast the change in actual speeds traveled in adopting versus non-adopting states.
  3. Use these two contrasts to develop an estimate of the hours saved in driving time per statistical life lost.
- Now, multiply time saved by some monetary value per hour to obtain an estimate of the VSL. Ashenfelter and Greenstone use the state mean wage as the value of an hour saved. We can discuss whether this is appropriate.
- A&G refer to their approach as an 'instrumental variables' estimation, and this is one valid way to interpret it. The adoption of the higher speed limit raises speed (the endogenous variable) and raises fatalities (the outcome variable) by raising speed but probably does not affect fatalities through any channel other than speed.
- What's unusual about the setup is that the decision to 'take-up' the higher speed-limit is *chosen* by states – it is not randomly assigned (this is unlike a conventional IV).
- The choice aspect is crucial for interpreting the results through the lens of Revealed Preference. Revealed Preference allows us to say that any state that *chose* to take up the higher speed limit *must* have valued the time savings at greater than or equal to the lives lost (otherwise, by Revealed Preference, it would not have made this choice). If this time savings was $1 million per life lost, then the VSL could be no higher than $1 million.
- There is also an important discussion in the paper of whether political decision making about speed limits is efficient. It's crucial to the interpretation to know whether:
  - Legislators roughly understood the trade-offs between time-savings and safety when deciding on the speed limit.

- Legislators' choices roughly represent the preferences of 'average' citizens ("the median voter") rather than of some interest group that has very different preferences about the VSL.

- Why isn't it enough to assume that individuals optimally choose their speed as a function of time savings and safety?

5.3 Theoretical framework

- See figures drawn in class.
- States face a Production Possibility Frontier in Time Saved-Lives Saved space.
- They want to choose their most preferred point on this frontier
- If the speed limit is capped at 55 MPH, states may not be able to select their optimal point on the PPF.
- The 1987 law expands the feasible choice set.
- For states that choose to move to the new location on the PPF, we can say that this point is Revealed Preferred to the old location.
- We can observe the gains they make in time savings and the loss of life as they make this movement. That forms the basis for our calculations

6 Results

See:

- Figure 1
- Figure 3
- Figure 4
- Table 1
- Table 2
- Table 3 (bottom row)

- Table 4 (bottom row)
- Table 6 (panel A)
- Table 7