Lectures 1 and 2 - Introduction and a First Application: The Minimum Wage Debate and Causal Inference in Economics

David Autor

14.03 Fall 2004

1 Introduction to the class

This is an intermediate course in microeconomic theory. I assume that you've all had 14.01. And it's helpful if you've also taken some statistics or econometrics.

The class is organized around three themes:

1. Economic theory – What does it say? What is it good for?

2. Causality – What do we mean by this? And how do we know it when we see it?

3. Empirical tests – Economic theory is a way of organizing facts and predicting and interpreting patterns in the world. It doesn't mean much without data. Data allow us to test the relevance of theory and to calibrate the magnitudes of theoretical predictions (theoretical predictions are often unambiguously signed (positive/negative) but rarely come with magnitudes). We'll be analyzing numerous experiments and quasi-experiments in the light of the theory we discuss.

Definition: Quasi-experiment – Events that unintentionally creates conditions similar to a randomized experiment. Examples: identical twins are separated at birth: a quasi-experiment that might be used to analyze nature-nuture questions. 200 million people buy 1 lottery ticket each and 100 of them win the lottery. Could use this event to evaluate the effect of wealth on happiness, health, marital dissolution, obesity.

Why use quasi-experiments? Because many of the key economic questions center around major life choices and outcomes: health, wealth, education, risk. These things are not usually open to experiment, so we look for chance events in the real world that approximate the experiment we would run if it were ethically feasible

1.1 Brief discussion: Methodology of Economics – Or 'Why Economic Theory'

Positive Economics

- The study of "what is". A descriptive endeavor free of value judgements.
- Build models to make sense of and generalize the phenomena we observe.
- Make predictions based on those models.

Normative Economics

- Assessing "what ought to be done". Making economic policy prescriptions.
- Sometimes positive economics gives us all the tools we need to say that one policy is preferable to another.

  (Q: When is this the case? A: When one policy is Pareto superior to another. Not too many of these)
- Definition: Pareto Improvement - A choice/policy/outcome that can make at least one person better off without making anyone else worse off. This is a pretty timid moral criterion, though it's not entirely uncontroversial.
- In any case, Pareto improving policy options are very rare (why – theory says that people should already have made those types of improvements). Most policy choices involves value judgements, ethical preferences, trade-offs among competing goals (e.g., employment and inflation; equity and efficiency).
- Though economic theory rarely tells you what policies to choose, it often makes the trade-offs clear.

1.1.1 Strength of economic approach to social science

- *Rigorous*: assumptions are stated, methods are formal, conclusions are internally consistent.
- *Cohesive:* built on a foundation of first principles and theory.
- *Refutable*: makes strong, testable (refutable) predictions, many of which appear correct.
- *Practical*: will help you to better understand how the world works.

1.1.2 Weaknesses of the economic approach

- "Economics is marked by a startling crudeness in the way it thinks about individuals and their motivations..." - Paul Krugman
- E.g., strong, simplifying assumptions that are often unpalatable and cannot be completely right (e.g., people act rationally to pursue self-interested - distinct from selfish - objectives...)

1.1.3 But there is some strength in this weakness

- We have a model of the world. It's called the world - and it's too complicated to be useful.
- Economic theory typically presents a very simplified, highly stylized world. But this can be quite helpful.
- *Friedman:* "The test of the validity of a model is the accuracy of its predictions about real economic phenomena, *not* the realism of its assumptions".
- *Friedman:* "A hypothesis is important if it 'explains' much by little".
- Our approach: simple models, significant insights.

1.1.4 Three significant insights of economic approach

- "People doing the best with what they have". Gets you a long way in understanding human action - both positively and normatively. It's surprising just how much of human interaction you can understand simply by assuming that people are trying to make the best choices for themselves. Also note: Alternative assumptions appear much less attractive...
- Equilibrium - Market 'aggregates' individual choices to produce collective outcomes that are sometimes *spectacularly different* from individual decisions.
- Properties of equilibrium can be evaluated using the criterion of efficiency:
  - Individuals are trying to make best choices for themselves given options.
  - Does market equilibrium produce the best possible outcome over all people (i.e., is it Pareto efficient)? There is no obvious reason to assume that it would – i.e., that we couldn't do better with central planning than the haphazard result of everyone making choices independently for selfish reasons.

> Yet, one of the stunning insights of economics is that under some conditions, the market often does produce the 'best possible' outcome.
> And, where it does not, theory provides an explanation for why and may give guidance on how to get to a better outcome.

Good outcomes: fundamental welfare theorems: self-interest produces efficiency
Bad outcomes: externalities, market failures.

2 Course requirements and expectations

1. Readings – At least one-third of the class will be given to discussing recent published papers in economics in some detail. You must read these articles before class to be prepared for class discussion. I don't expect you to understand the entire article, but I want you to understand: the question, the basic method, and the key results.
2. Class participation - Matters. (See syllabus)
3. 6 problem sets - One is dropped automatically. The other 5 each count for 5 percent of your grade. They are all due on Thursdays at 5pm. No late problem sets are accepted – that's what the automatic drop rule is for. So, don't waste your ammunition early in the semester. Problem sets are a mixture of standard, formal problems and questions based on the readings and lectures.
4. 3 exams, each 25 percent. Two in class, one in exam period. Exams will focus on material since prior exam.
5. On attending class. A large chunk of the class material is not in the textbook. This is esp. true of the readings, but also some of the theoretical material we will cover. If you don't regularly attend class, you will have difficulty on problem sets and exams. If you were planning to only show up for exams, I strongly recommend against taking the class.
6. Class is not graded on a strict curve. Everyone can do well (or badly). Class participation credit is given after grades are assigned. If you do only the minimal amount, you'll probably get a C. If I think you are heading for a D, I will recommend that you drop the class. But I can't help you if you after the drop date.
7. Support outside of class:

- Recitations are held every Friday, one at 9am and one at 10am. Here, your TAs will clarify class material, help to prep for exams, and review the problem sets. The first two recitations – including the 1^{st} Friday – will cover math tools for 14.03. I'll be using these tools in class but I won't be covering them in detail. So, please don't miss these.
- TA office hours. Held every Wednesday from 2pm to 4pm in ???. All problem sets are due on Thursdays.
- Autor office hours. Tuesdays from 3pm to 4pm in my office.
- Questions on class topics and problem sets. Use the class web site, which we'll monitor every day. Do not email us with substantive, class-related questions (personal issues can be handled by email).

3 A first example: The minimum wage and employment debate

Agenda

1. Textbook model of competitive labor market
2. Impact of minimum wage on textbook model
3. Assumptions behind this model
4. What happens when we relax an assumption - price taking?
5. Impact of minimum wage when employers have market power.
6. Testing the textbook model and alternatives
7. Natural experiments in economics
8. Card and Krueger article.

Minimum wages: A venerable topic in economics and area of ongoing controversy.

4 Textbook model of wages and employment

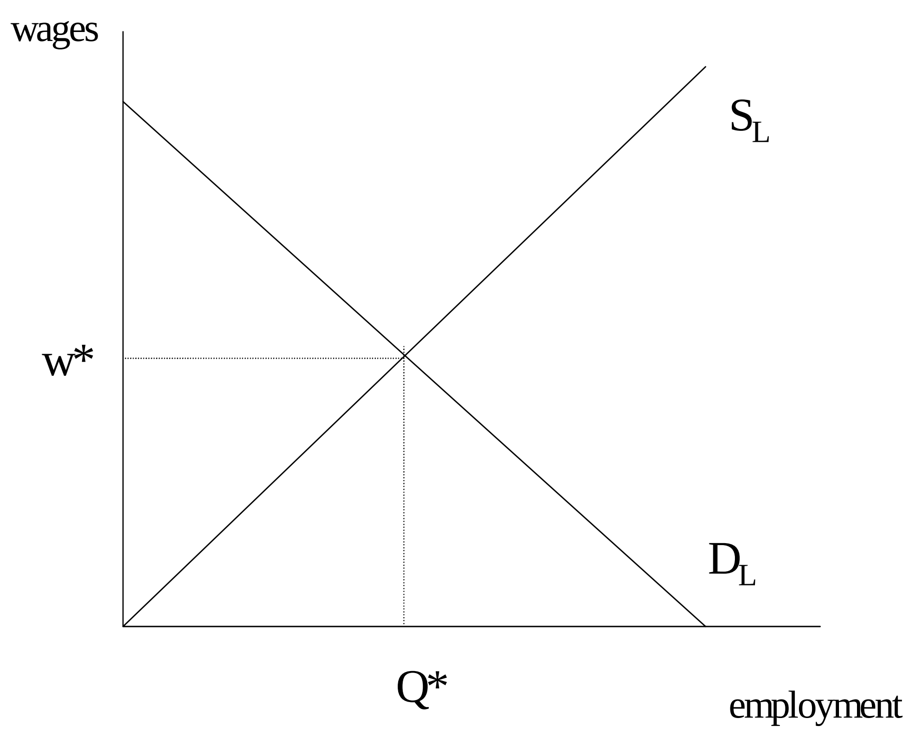

Labor Supply curve: all of the potential workers in the labor market, arranged according to their "reservation wage," which is the lowest wage they will accept to take a job (from low to high)

Labor Demand curve: all potential employers in labor market, arranged according to their willingness to pay for a worker (from high to low)

Q: What is the key variable in this model? The wage or the number of employed workers?

A: Neither. They are both outcomes, endogenous variables.

Definition: ENDOGENOUS: internally determined. An outcome as opposed to a cause.

Definition: EXOGENOUS: externally determined. A causing or forcing variable.

What happens when we impose a minimum wage in this labor market?

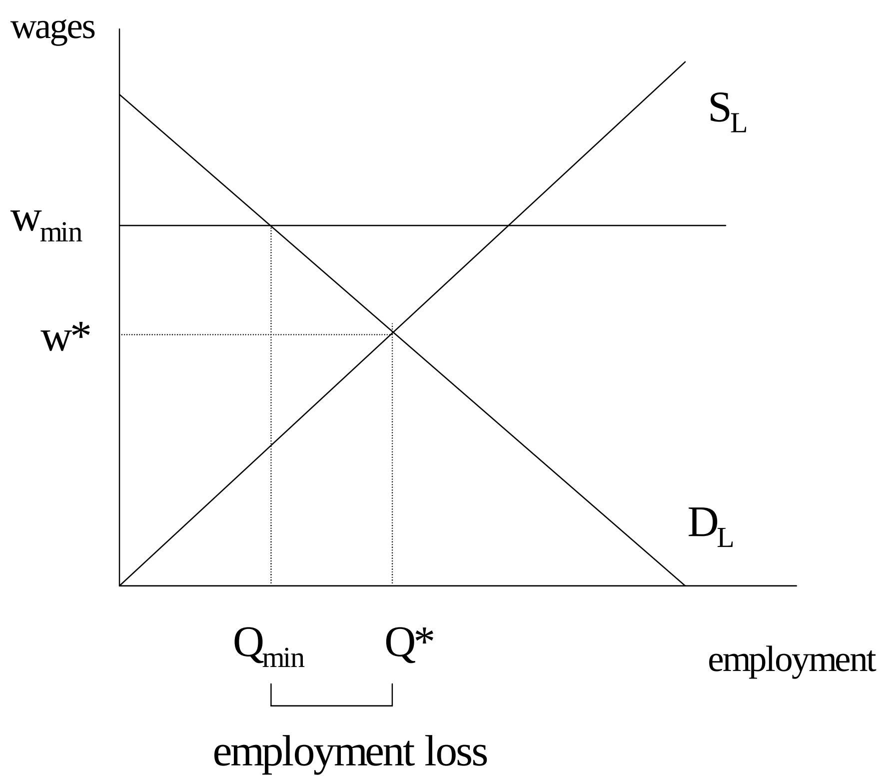

Wages:

$$w_{\min} > w^*$$

Employment:

$$Q_{\min} < Q^*$$

Q: if this model is right why would you ever want to impose a minimum wage?

<u>One answer: Total earnings</u>

$$w_{\min} Q_{\min} \gtrless w^* Q^*$$

Total worker earnings may increase even if employment falls.

Q: What does this depend on?
A: Elasticity of demand:

$$\eta = \frac{\partial Q}{Q} \frac{w}{\partial w} \gtrless -1$$

If proportional increase in wages larger than (induced) proportional decline in employment $\Longrightarrow$ wagebill increases

What is the primary assumption behind the textbook model?

Individual "price-taking firm"

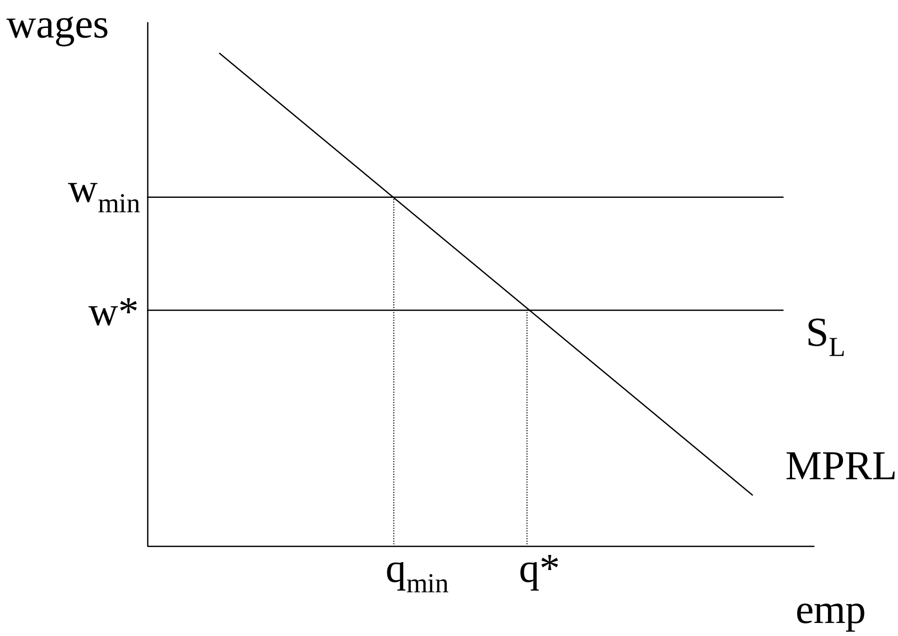

MRPL = Marginal Revenue Product of Labor $\Longrightarrow$ "what the marginal worker produces". It is decreasing in employment due to decreasing returns in the production function.

How did we conclude that the firm sets:

$$MRPL = w^*$$

Recall firm's optimization problem:

$$\max\ \pi = pf(L) - w(L)L$$

where p is the product price, $w(L)$ is the wage necessary to "call forth" L workers.

$$\begin{aligned}\frac{\partial \pi}{\partial L} &= p\frac{\partial f(L)}{\partial L} - w(L) - \frac{\partial w(L)}{\partial L}L = \\ &= pf\prime(L) - w(L) - w\prime(L)L = 0\end{aligned}$$

The FOC can be rewritten as:

$$pf\prime(L) = w(L) + w\prime(L)L$$

where:

- $pf\prime(L)$ = the marginal revenue product of labor ($MRPL$)
- $w(L)$ = wage
- $w\prime(L)L$ = variation in the wage per additional worker × *total* workforce

Q: What is the key assumption of a competitive model?

$$w'(L) = 0 \Longleftrightarrow price\ taking\ firms$$

How does firm choose employment when it is not price taker?

It chooses employment according to the above FOC:

$$pf\prime(L) = w(L) + w\prime(L)L$$

If $w\prime(L)L \neq 0$ then firms must pay all workers more for each additional worker it hires.

Q: How do we know this means lower employment?

Intuitively, a higher price must mean lower quantity demanded.

Optional:

$$\begin{aligned} pf\prime(L) &= w(L) + w\prime(L)L \\ w &= MRPL - \frac{\partial w}{\partial L}L \\ 1 &= \frac{MRPL}{w} - \frac{\partial w}{\partial L}\frac{L}{w} \\ 1 &= \frac{MRPL}{w} - \frac{1}{\eta} \end{aligned}$$

where η is the elasticity of labor supply, $\frac{1}{\eta}$ is the elasticity of wages WRT employment perceived by the single firm. In the case of a competitive labor market the elasticity of wages wRT. employment is perceived by the single firm as zero.

$$w = \frac{MRPL}{1 + \frac{1}{\eta}}$$

The more elastic is labor supply the lower the wage relative to marginal product of labor.

4.1 Conventional case: Individual Price Taking Firm

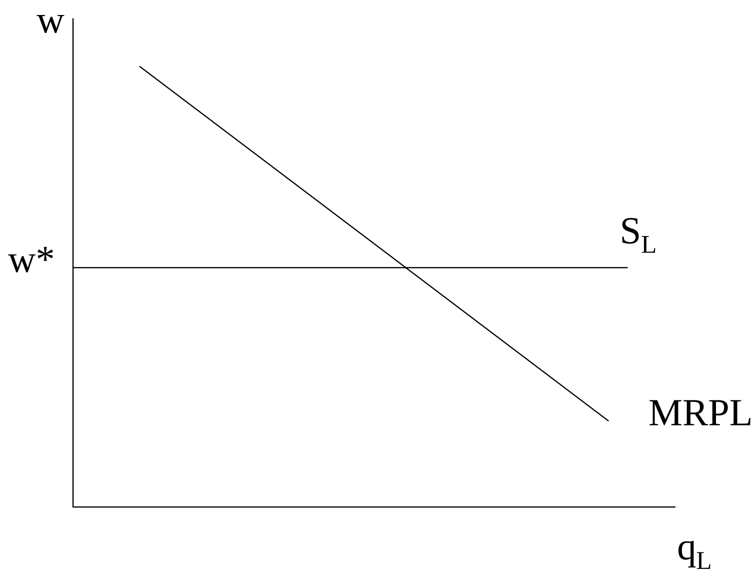

Notice that the labor supply curve is upward-sloping at the market level, but it is flat as perceived by the single firm.

4.2 Monopsonistic employer

The labor supply curve for a monopolist is upward sloping. To get one more worker, the monopolist must raise the wage by a small amount. Assuming that all workers receive the same pay (i.e., the late-comers don't get paid more), then the marginal cost of the next worker is not simply her wage but the wage increase given to all of the other ('infra-marginal') workers.

Hence, the marginal labor cost curve for this firm is even more upward sloping. The additional cost for each worker is given by the higher wage of that worker and by the increase in wage given to the entire pool of workers.

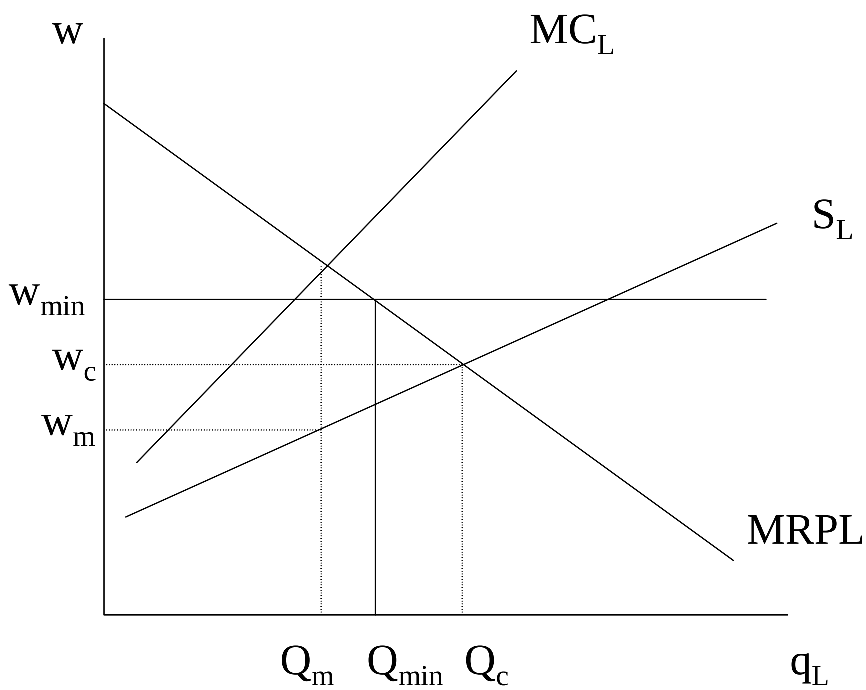

- In this example, implementation of a binding minimum wage raises wages and employment.

$$w_{\min} >> w_m$$
$$Q_{\min} >> Q_m$$

- How does this work?

  Firm is now a price-taker for labor at $w_{\min}$, hence labor supply is "perfectly elastic". That is, there are an 'unlimited' number of workers available (are far as any one firm is concerned) at the going wage. Therefore the monopsonist sets:

$$w = MRPL$$

  since the choice of the quantity of labor has no impact on the level of wages.

Q: Does raising minimum wage to monopsonists always increase wages and employment?

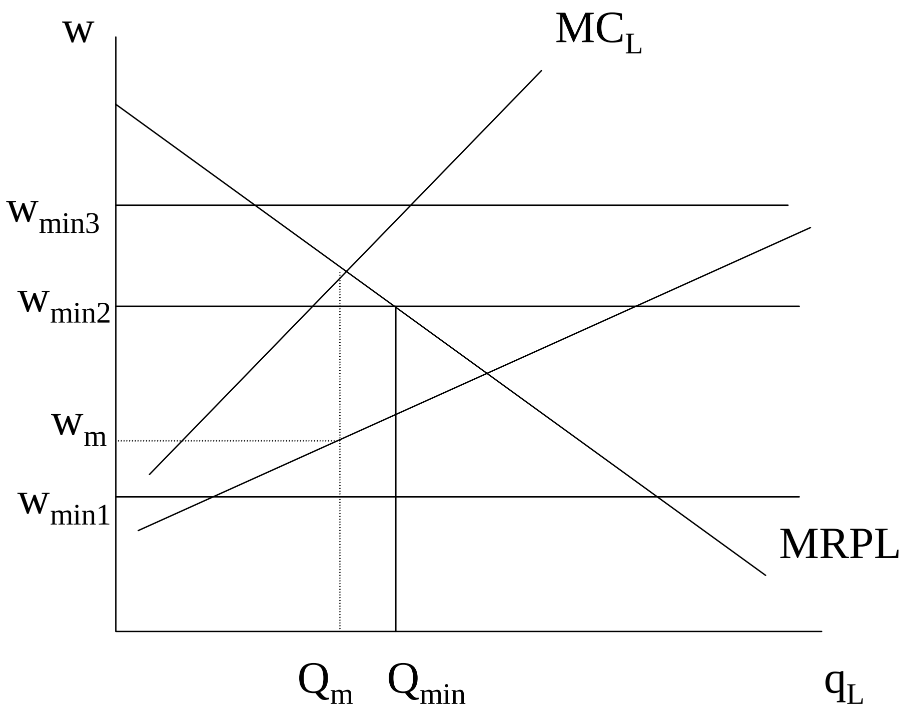

$w_{\min 1}$ - Introduction of a minimum wage at this level has no effect because the minimum wage is below w_m and hence doesn't bind

$w_{\min 2}$ - Introduction of a minimum wage at this level raises wages and employment

$w_{\min 3}$ - Introduction of a minimum wage at this level raises wages but reduces employment

5 Monopsony

If monopsony were present in the labor market, where would you expect to find it? (Remember the criterion: the firm's own labor demand changes the market wage.)

We would expect to find a monopsony in the following markets:

- Company towns
- If skills are very specific, e.g. IBM mainframe repair technicians
- 'Captive' labor markets, spouses of soldiers based away from home.
- Fast food restaurants located in nearby towns in NJ and PA?

5.1 Testing for monopsony in the labor market

- How do we go about testing the monopsony vs competitive model of the labor market?
- Q: What's the key empirical implication that distinguishes these models?

  A: In the competitive model: w $\uparrow$ and employment $\downarrow$

  In the monopsonistic model: w $\uparrow$ and employment $\uparrow$
- So how do you test this implication?
  - We can look across different states and ask ourselves the following question: is employment higher where wages are higher?

    Let's suppose you find the following pattern:

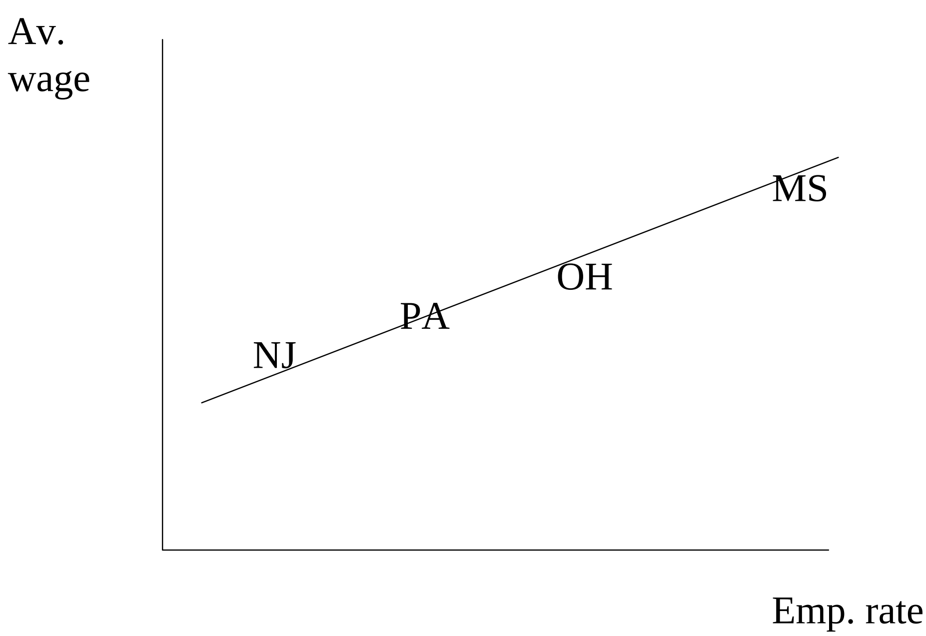

Q: Would this convince you? What's the problem with the wage here?

A: We don't know why it is different across states. There could be different demand and supply schedules.

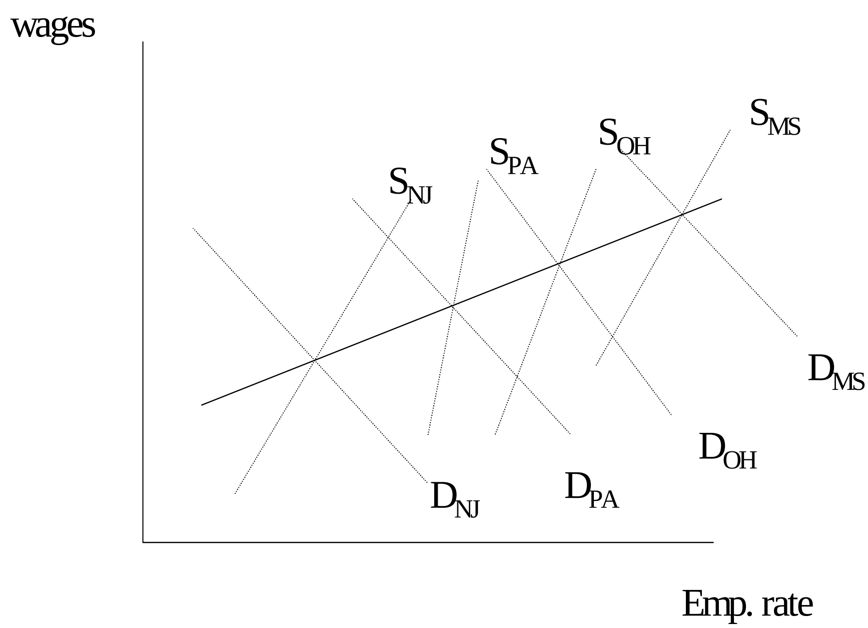

Since both employment and wages are endogenous outcomes - determined by both supply and demand - this picture tells us nothing about the impact of minimum wages on employment.

– How do we overcome this problem? We need an experiment. But what type of experiment? One in which wages go up. What about something that shifts supply out? No.

  By exogenously manipulating wages we can study impact on employment to infer the slope of the relationship between wages and employment (downward sloping → competitive market, upward sloping → monopsony)

– An idea: New Jersey Minimum Wage Law.

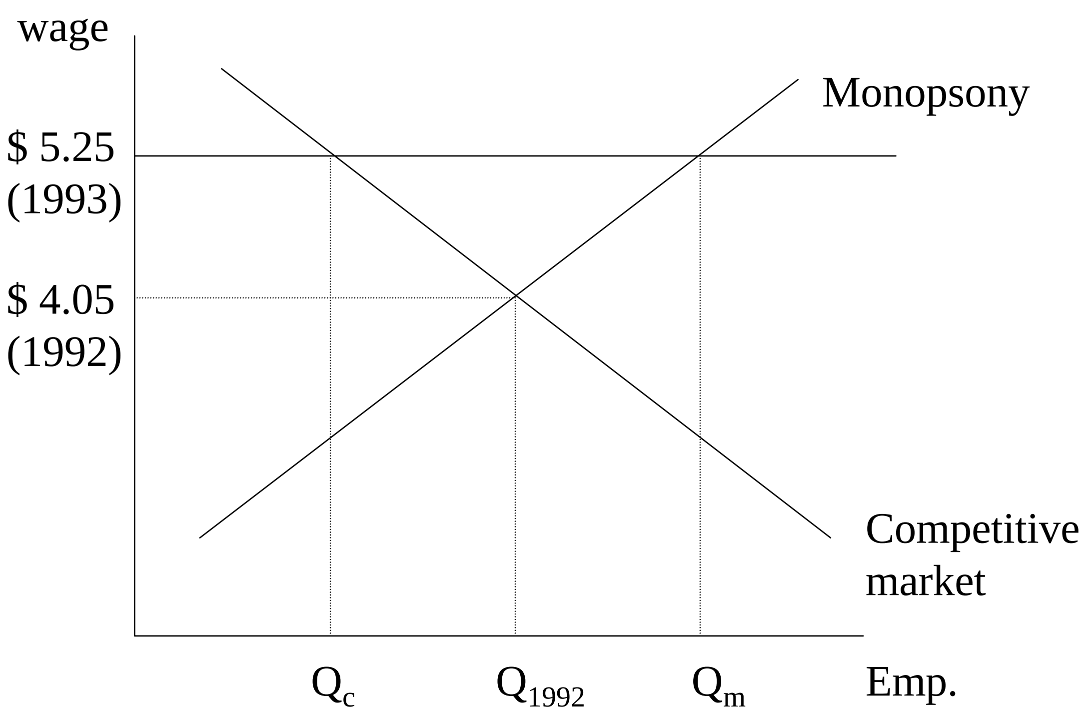

Notice that the conditions under which the introduction of a minimum wage raises employment in a monopsonistic market are only locally satisfied – that is, raising the minimum wage by "too wage" will reduce employment even in the monopsonistic setting.

By looking at the change in employment after the adoption of the minimum wage , we can see if employment has gone up (monopsony) or down (competitive market).

Notice that we don't ever observe the supply or the demand curve. These are pieces of conceptual apparatus that we use to guide thinking. All we observe is the wage level and the quantity employed.

Before discussing the Card and Krueger paper, we need to step back and talk about causal inference in social science. This is a topic that is more subtle than it initially appears.

5.2 Causal inference in social science

Much of social science (psychology, anthropology, sociology, political science, epidemiology and large parts of economics) concerns analyzing correlations among variables – i.e., the correlation between education and income, the correlation between obesity and heart disease, the correlation between happiness and longevity.

Correlation describes the statistical relationship between two observed variables. Correlation has no necessary relation to cause and effect. You can measure the correlation between happiness and longevity with great precision and yet know nothing about how making someone happier affects their longevity (maybe the same gene that causes longevity causes happiness, so making someone happier would not increase their longevity.)

In 14.03, we are not generally interested in these correlational questions. Our goal is to analyze causal questions: (Why: Because science advances through analyzing cause and effect relationships, not (primarily) by documenting correlations.)

Causal questions:

- What is the effect of education on income?
- What is the effect of obesity on heart disease?
- What is the effect of happiness on longevity?

These questions are much harder to answer than the correlational questions. Correlations are readily measured from observational data. We can never directly observe causal effects (to be explained).

The effect of a cause involves two ingredients that are not present in studying correlations:

1. First, a causal effect intrinsically depends on a time dimension. When we say, what is the causal effect of X on Y, we mean, what value would Y have taken if X were some other value? It's easiest to phrase this for a binary cause, $X \in \{0, 1\}$. So, if we know that $X = 1$ and $Y = k$, the causal question is, what would Y have been if $X = 0$ instead. That's a question that involves rolling back time to a different counterfactual reality.

2. Second, a causal effect is always intrinsically measured relative to some alternative cause. In this case, we are comparing $X = 0$ to $X = 1$. It's not meaningful to ask what is the causal effect of X on Y without (at some level, perhaps implicitly) specifying what alternative X we are comparing it to; the counterfactual Y depends on the counterfactual X.

Some notation:

Let Y_i be the outcome of interest for unit i, where i could be a person, a cell, a drop of water, a sovereign country.

We want to consider two possible outcomes for i. Let Y_{0i} be the outcome if $X = 0$ and Y_{1i} be the outcome for $X = 1$.

Thus, for every unit i, we can imagine two potential outcomes $\{Y_{0i}, Y_{1i}\}$ that we woudl observe if the unit were treated ($X_i = 1$) or untreated ($X_i = 0$).

The causal effect of X on Y for unit i is therefore $T_i = Y_{1i} - Y_{0i}$ (where T stands for Treatment Effect).

The problem that this immediately reveals is that we never observe $Y_{1i} - Y_{0i}$.

Instead, we observe $Y_i = Y_{1i} \cdot X_i + Y_{0i} \cdot (1 - X_i)$.

This is the Fundamental Problem of Causal Inference:

It is impossible to observe the value Y_{1i} and Y_{0i} on the same unit i and so it is impossible to directly observe the causal effect of X on Y for unit i.

Solutions to the Fundamental Problem of Causal Inference?

1. Assume temporal stability and causal transience. If the causal effect of X on Y is the same at any point in time (now, the future) and the causal effect of X on Y is reversible (so having once exposed Y to X doesn't permanently change the effect of X on Y), then we can observe $Y_{1i} - Y_{0i}$ simply by repeatedly changing X from 0 to 1 (e.g., flipping a light switch on and off).

   Notice that temporal stability and causal transience are postulates; they cannot be directly tested.

2. Assume unit homogeneity. If the Y_{1i} and Y_{0i} are identical for all i, then we can measure the causal effect simply by calculating $Y_{1i} - Y_{0j}$ for $i \neq j$. Again, unit homogeneity isn't normally a verifiable assumption. But for laboratory conditions, it might be reasonable (e.g., experimenting on two molecules of water).

3. Statistical solution. This is different from the other two.

   - Although we may never be able to observe $T_i = Y_{1i} - Y_{0i}$, we could potentially estimate some type of population average for

     $$T = E\left(Y_{1i} - Y_{0i}\right).$$

   - But since we can never observe $Y_{1i} - Y_{0i}$, the quantity $E\left(Y_{1i} - Y_{0i}\right)$ is also not directly estimable.

   - More subtle: There is no reason to assume that $E\left(Y_{1i}|X=1\right) = E\left(Y_{1i}|X=0\right)$ and $E\left(Y_{0i}|X=1\right) = E\left(Y_{0i}|X=0\right)$. So, simply calculating $E\left(Y_{1i}|X=1\right) - E\left(Y_{0i}|X=0\right)$ will not generally give us $E\left(Y_{1i} - Y_{0i}\right) = T$. This is because whether $X_i = \{0, 1\}$ may be correlated with Y_{i1}, Y_{i0}.

     More specifically:

     $$E\left(Y_{1i}|X=1\right) - E\left(Y_{0i}|X=0\right) = \underbrace{E\left(Y_{1i}|X=1\right) - E\left(Y_{0i}|X=1\right)}_{T_i} + \underbrace{\{E\left(Y_{0i}|X=1\right) - E\left(Y_{0i}|X=0\right)\}}_{Bias},$$

     Example: let's say Y is the number of mathematical expressions you can differentiate in an hour after 4 years of college and X is an indicator variable for whether or not you attended MIT. If we administered math tests at random, we would certainly find that $\hat{T} = E(Y_1|X = MIT) - E(Y_0|X = NOT\ MIT) > 0$, i.e., MIT students do more calculus in an hour than non-MIT students.

Is $\hat{T}$ a valid estimate of the causal effect of attending MIT on calculus skills? Definitely not.
Students who are skilled in calculus choose to come to MIT, and they would be more skilled than the average student in calculus, regardless of whether they attended MIT.
So, $E(Y_{0i}|X = MIT) > E(Y_{0i}|X = NOT\ MIT)$. That is, students who attended MIT would have done better at math than students who didn't MIT even if the attendees had not attended MIT. So $\hat{T} > T$. [See expression above.] So, we do not get a valid causal estimate of the effect of attending MIT on math skills of the students who attended MIT by comparing them to students who didn't attend MIT.

The substantive problem is that attendance at MIT is endogenous. Students come to MIT in part because they are good at math. So, it's not reasonable to assume that non-MIT students are a valid comparison group for MIT students.

- Let's say instead that we picked a large number of $i's$ at random and randomly assigned half to $X_i = 1$ and half to $X_i = 0$. Then this pretty much guarantee (unless we are very unlucky) that $E(Y_{1i}|X = 1) = E(Y_{1i}|X = 0)$ and $E(Y_{0i}|X = 1) = E(Y_{01i}|X = 0)$. Consequently:

$$T = E(Y_{1i} - Y_{0i}) = \underbrace{E(Y_{1i}|X = 1) - E(Y_{0i}|X = 1)}_{T_i} + \underbrace{\{E(Y_{0i}|X = 1) - E(Y_{0i}|X = 0)\}}_{bias\ =0}.$$

In this case, randomization removes the bias term. Random assignment makes it's reasonable to believe that $E(Y_{0i}|X = 1) - E(Y_{0i}|X = 0) = 0$.

In summary, randomization potentially solves the causal inference problem by making the treatment status $X_i = \{0, 1\}$ independent of potential outcomes: $E(Y_{1i}), E(Y_{0i}) \perp X_i$ meaning $E(Y_{1i}|X = 1) - E(Y_{1i}|X = 0) = 0$, $E(Y_{0i}|X = 1) - E(Y_{0i}|X = 0)$. This is the idea of a 'control group' – the group not receiving the treatment provides an estimate of the counterfactual outcome for the treated group.

To solve the Fundamental Problem of Causal Inference in Economics, we almost always use the statistical solution. That's because, for human behavior, it is rarely plausible that either of the other two solutions applies: temporal stability + causal transience or unit homogeneity.

By contrast, the statistical solution is almost certain to work. However, for ethical and/or practical reasons, we can rarely use it ... So, we look for quasi-experiments.

<u>Difference-in-difference estimates</u>

Often, we don't simply measure the level of Y but it's change as a function of X.

So, if we have a treatment and control group, we form

	Before	After	Δ
Treat	$Y_{0i,0}$	$Y_{1i,1}$	ΔY_i
Control	$Y_{0j,0}$	$Y_{0j,1}$	ΔY_j

Why do we want to make a pre-post comparison?

We do not need to do this if we have a very large population of treatment and control units to work with. We could simply calculate $T = E\left(Y_{1i,1}|T_i = 1\right) - E\left(Y_{0j,1}|T_i = 0\right)$.

However, we often don't have very large samples of treatment and control individuals to work with.

Let's say we are assessing the effect of a new drug treatment on cholesterol levels. We could pick 20 people each for the treatment and control groups, give the treatment group the drug treatment and the control group the placebo, and then compare the average cholesterol level between these two groups.

There is nothing wrong with this approach. But we might be concerned that, just by chance, these two groups started out with slightly different cholesterol levels.

In this case, we could also take baseline data – prior to drug treatment – just to be sure that they were comparable.

Let's say the baseline averages were comparable but not identical. The treatment group, for chance reasons, had a slightly lower cholesterol level than the treatment group. We'd be concerned that our experiment would be biased in favor of the finding that the treatment lowered cholesterol (since the treatment group started with a better outcome).

It's that concern that motivates us to compare the <u>change</u> in cholesterol in the treatment group to the change in cholesterol in the control group.

This approach subtracts off initial differences that could potentially prove confounding in small samples. It allows us to focus on the improvement in the treatment group relative to the control group.

So, more formally, let's say that prior to treatment:

$$\begin{aligned} Y_{0i,0} &= \alpha_i \\ Y_{0j,0} &= \alpha_j. \end{aligned}$$

We would hope that $\alpha_i = \alpha_j$, but this doesn't have to be the case.

Now, imagine that after treatment, we observe

$$Y_{1i,1} = \alpha_i + \delta + T,$$

where T is the causal effect and δ is any effect of time. For example, cholesterol levels may tend to worsen over time as people age.

So, if we take the first difference for Y_{1i}, we get:

$$\Delta Y_i = Y_{1i,1} - Y_{0i,0} = (\alpha_i - \alpha_i) + \delta_i + T$$

This obviously does not recover T. But it does remove the 'level effect' α_i.

Similarly, $\Delta Y_j = (\alpha_j - \alpha_j) + \delta_j$. Differencing removes the level effect for group j.

If we are willing to believe that the time effect operates identically on the treatment and control groups, $\delta_i = \delta_j = \delta$, then we have

$$\Delta Y_j = Y_{1i,1} - Y_{0i,0} = \delta,$$

and

$$\Delta Y_i - \Delta Y_j = T + \delta - \delta = T.$$

So, the difference-in-difference estimator allows us to potentially recover the causal effect when: the treatment and control groups are not entirely identical and when there is a potentially confounding effect of time.

5.3 Now let's go back to New Jersey

Let $Y_{nx,(b,a)}$ be employment in New Jersey <u>b</u>efore or <u>a</u>fter introduction of the minimum wage.

Let X be the minimum wage treatment. $X = 0$ is untreated (Federal minimum wage) and $X = 1$ is treated (New Jersey minimum wage).

So, if we want to estimate the causal effect of the minimum wage hike on New Jersey employment, we could calculate:

$$\hat{T} = Y_{n1,a} - Y_{n0,b}\,,$$

which is simply the before/after change in New Jersey employment.

What do we think of this estimate of the causal effect?

Requires assumption of temporal stability: were it not for the minimum wage hike, New Jersey employment would have remained unchanged.

Is this plausible? Probably not. In our previous example, $\hat{T} = T + \delta$

So, what do we need to improve on this experiment?

We could select a group of states at random and assign the minimum wage increase to half of them and not to the other half. Then, we could compare employment in each group of states.

A problem here is that this experiment is not available to us. But it's a good idea.

Another idea is to select a single state that we think is closely comparable and use it as our 'control group' – in this case, that state is Pennsylvania.

We take baseline data in both states, then compare the change in NJ to the change in PA. That's our difference-in-difference estimator.

5.3.1 Card & Krueger experiment

This paper is a widely cited study of the impact of the minimum wage on employment levels.

This study created huge controversy among economists and arguably caused millions of workers to get a raise from the Clinton administration in 1995.

April 1, 1992: in New Jersey the minimum wage rose from \$4.25 to \$5.05 per hour (this is a sizable increase)

Eastern Pennsylvania (bordering NJ) didn't raise the minimum wage. Maintained the Federal minimum wage of \$4.25 per hour.

Card & Krueger surveyed 410 fast food restaurants.

In this experiment the timing of the experiment is the following:

Before: Feb-Mar 1992

After: Nov-Dec 1992

The setup:

	Before	After	Δ
NJ	$Y_{n0,b}$	$Y_{n1,a}$	ΔY_n
PA	$Y_{p0,b}$	$Y_{p0,a}$	ΔY_p

$$\hat{T} = \Delta Y_n - \Delta Y_p$$

Table 3 in the paper shows "Per store employment"

	Before	After	Δ
NJ	20.44	21.03	$\Delta Y_n = +0.59$
PA	23.33	21.37	$\Delta Y_p = -2.16$

$\hat{T} = 0.59 - (-2.16) = 2.76$ with a standard error of 1.36 (so, it is statistically significant at the 5 percent since the t-ratio is ≈ 2.0).

The paper contains many more tests, but this is the basic result: $2.76 \approx 13.5\%$ increase in employment in NJ relative to PA.

Interpretations:

1. Monopsony

   Other interpretations:

2. Hungry teens

3. Motivational effects

4. Confounding variables (shocks to PA that are not accounted for in the test)

5. Wrong venue (why did they study fast food?)