# Lecture 1 — Introduction and a First Application: The Minimum Wage Debate

David Autor, MIT Department of Economics

14.03/14.003 Microeconomic Theory and Public Policy, Fall 2016

# 1 Introduction to 14.03/14.003

This is an intermediate course in microeconomic theory and its application to real world policy problems. The class assumes proficiency with economic theory at the 14.01 level as well as multivariate calculus. It is also quite helpful if you have taken some statistics or econometrics. For those who have not, the Teaching Assistants, will provide a primer on some key statistical and mathematical concepts during the initial recitation. In addition, the handout Math Tools for 14.03/003 (on the class website), offers a review of some mathematical concepts that you'll be using in lectures, problem sets, exams. I also recommend the brief but helpful Mathematical Appendix of the required text by Samiran Banerjee, *Intermediate Microeconomics: A Tool-Building Approach*.

This class is organized around three themes:

1. **Economic theory:** Where does it come from, what does it predict, and in what ways is it useful?

2. **Causality:** What is it, and how do we measure or estimate causal effects?

3. **Empirical applications:** Economic theory is a way of organizing facts and interpreting and patterns in the world. This class will use data to test theory and use theory to interpret data. We will analyze numerous randomized experiments and quasi-experiments in the light of economic theory.

**Definition 1.** Randomized experiment (also frequently called a Randomized Controlled Trial or RCT). In the statistical theory of design of experiments, randomization involves randomly allocating the experimental units across the treatment groups. (Source: Wikipedia, http://en.wikipedia.org/wiki/Randomized_experiment).

Example: If an experiment compares a new drug against a proven existing drug, patients would be allocated to either the new drug or to the existing drug using randomization. A comparison of outcomes among patients allocated the new drug and those allocated the existing drug would provide an estimate of the causal effect of the new drug relative to the existing drug. (Note, it would not provide an estimate of the causal effect of the new drug relative to no treatment unless there was also a placebo group in the experiment.)

**Definition 2.** Quasi-experiment. An event that unintentionally creates conditions similar to a randomized experiment.

Example: One million people buy one lottery ticket each and one hundred of them win. This quasi-experiment could be used to evaluate the effect of unanticipated wealth increases on happiness, health, marital dissolution, obesity. (Indeed a recent high profile paper did just this!)

There is an impressive diversity of experiments and quasi-experiments that economists have applied to analyze important questions in social science. You may ask, why would a researcher

use a quasi-experiment instead of an RCT? In reality, RCTs are the gold standard of evidence and economists use them frequently. Indeed, we will study multiple RCTs during the semester. It's also the case that many key economic questions center either around major life choices and outcomes such as health, wealth, education, and risk, or 'macro-scale' treatments such as international trade, civil war, or epidemic disease. For both ethical and practical reasons, these outcomes are often not suitable for randomized experimentation. In such cases, we look for chance events in the real world that approximate the experiment we would conduct if it were ethically or practically feasible.

## 2 A first application: The minimum wage and employment

Rather than start the class with a discussion of economic methodology, we'll start with an application and return to the big picture when that is complete (sometime in the second lecture). The application we'll consider is the impact of the minimum wage on wages and employment. The costs and benefits of legislated minimum wages are a central policy topic in economics—an area of ongoing controversy and active policymaking (e.g., President Obama would like to raise the national minimum wage from \$7.25 to \$10.10 per hour, a substantial increase, and to increase it further to \$12.00/hr by 2020. The Republican party vehemently opposes this change). Here's a mini-outline of what we'll cover on this topic in the first two lectures:

1. Textbook model of competitive labor market
    (a) Impact of minimum wage on employment in the textbook model
    (b) Assumptions behind this model
2. What happens when we relax a key assumption: price-taking by firms
    (a) Impact of minimum wage on employment when employers have market power
3. Testing the textbook model and alternatives
4. Economic experiments
    (a) The Fundamental Problem of Causal Inference (FPCI)
    (b) Overcoming the Fundamental Problem
    (c) Notation for causal inference
    (d) Estimating causal effects using "Differences-in-Differences" (DD)
5. The Card and Krueger minimum wage study
    (a) Interpretation and discussion

## 3 Textbook model of wages and employment

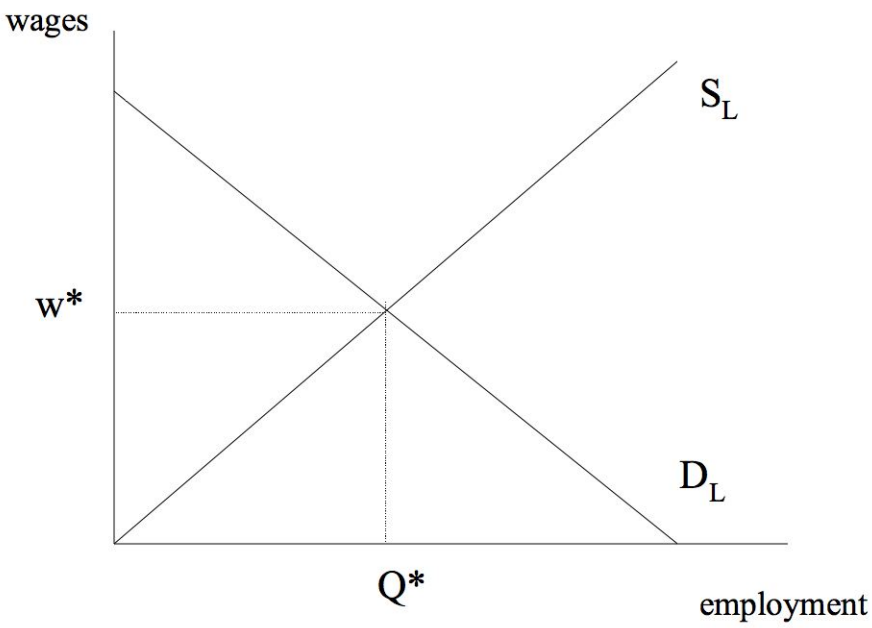


**Definition 3. Labor supply curve.** All potential workers in the labor market, arranged from low top high according to their "reservation wage," (the lowest hourly wage at which they are willing to work)

**Definition 4. Labor demand curve**. All potential employers in labor market, arranged according to their willingness to pay (hourly) for a worker

- Q: What is the key outcome variable in this model: the wage or the number of employed workers? A: *Neither*. They are simultaneously determined. Another way to say this is that these outcomes are *endogenous*.
- In the example above, the demand and supply curves are *exogenous*. The equilibrium wage and employment levels are *endogenous*.

**Definition 5. Endogenous**. Internally determined. An effect rather than a cause.

**Definition 6. Exogenous**. Externally determined. A causing or forcing variable.

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

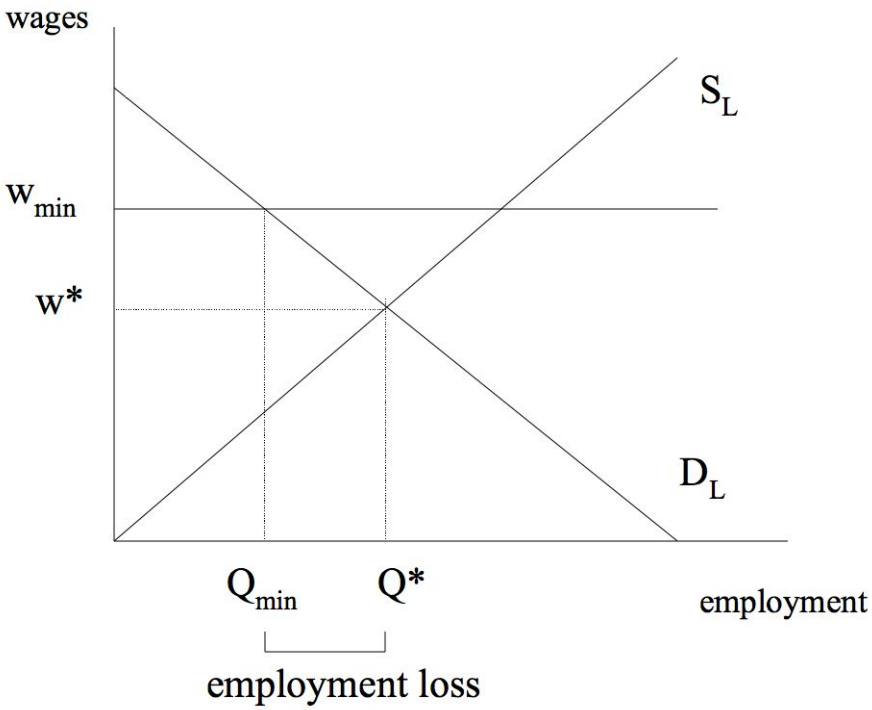


- Wages:

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

  Employment:

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

- If this model is right why would a policymaker ever want to impose a minimum wage?
- One possible answer: A binding minimum wage could raise total worker earnings even if it reduced employment

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

  Total worker earnings may increase even if employment falls.

- What does this depend on? The *elasticity* of demand:

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

  If the proportional increase in wages is larger than the (induced) proportional decline in employment $\Longrightarrow$ wage-bill increases. Specifically, if $\eta > -1$ (i.e., $|\eta| < 1$), then a 1% rise in wages reduces employment by less than 1%, so total wages paid (wages $\times$ workers) rises.

**Definition 7. Elasticity**. The ratio of the proportional change in a variable $Y$ caused by a given proportional change in a variable $X$. So for example, if the elasticity of $Y$ with respect to $X$ is 2, then a 1% increase in $X$ causes a 2% increase in $Y$.

### 3.1 Why do minimum wages reduce employment? Revisiting the theory

- What is the primary assumption in the textbook model that yields the prediction that (binding) minimum wages always and everywhere reduce employment?
- The answer is *price-taking* behavior, both in labor and product markets. That is, the price of the good the firm is producing does not fall if the firm makes a few more, and the prevailing wage the firm faces does not rise if it hires a few more workers. Formally, product demand and labor supply are both *perfectly elastic* as the far as the firm is concerned.
- We say "as far as the firm is concerned" because each firm is infinitesimal relative to the size of the market, meaning that its own labor demand cannot affect the wage level, though of course, the aggregate demand of all firms in the market does affect the wage level.

**Individual "price-taking" firm**

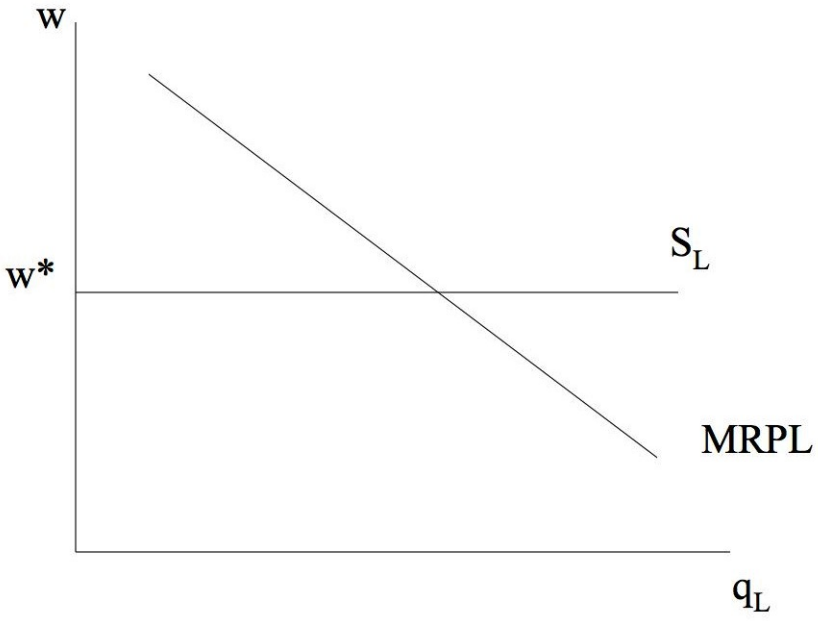


- MRPL = Marginal Revenue Product of Labor $\Longrightarrow$ "What the marginal worker produces."
- We normally assume that at any given firm, MRPL is decreasing in employment due to decreasing returns in the production function. All else equal, the next worker produces marginally less than the prior hire. This could be because the most important tasks are always done first, so adding more workers means that some less important tasks are also accomplished.
- You learned in 14.01 that the firm equates the Marginal Revenue Product with the wage:

$$MRPL = w^*.$$

Where did that come from?

- Recall the firm's profit maximization problem, which is to maximize the difference between revenues and costs (i.e, profits). Assume that the firm's only input is labor. Denote the first derivative of a function $f(\cdot)$ by $f'(\cdot)$ and the second by $f''(\cdot)$. The firm's problem is:

$$\max \ \pi = p \cdot f(L) - w(L) \cdot L,$$

where $p$ is the product price, $w(L)$ is the wage necessary to "call forth" $L$ workers, and $f(L)$ is the amount of output produced.

- We assume that $f'(\cdot) > 0$ and $f''(\cdot) < 0$, so an additional worker always raises output, but marginal productivity declines as we add workers. Note that $p$ is not a function of $L$, meaning we assume that the price of output is taken as given for this problem (it's exogenous).

- Differentiate this expression with respect to $L$ and set it equal to zero. (Why zero? At the optimum, this derivative must equal zero. If not, the firm would want to adjust $L$ further. If the marginal profit were positive, the firm would want to hire more labor. If the marginal profit were negative, the firm would want to hire less labor.)

$$\frac{\partial \pi}{\partial L} = p \cdot \frac{\partial f(L)}{\partial L} - w(L) - \frac{\partial w(L)}{\partial L} \cdot L = 0$$

Rearranging:

$$pf'(L) = w(L) + w'(L)L$$

where:

  - $pf'(L)$ is the marginal revenue product of labor ($MRPL$)
  - $w(L)$ is the equilibrium wage, also equal to the hourly cost of the last worker hired
  - $w'(L)L$ is the additional increment to *total labor costs* incurred by hiring one more worker (excluding the payments to that one last worker). This increment is equal to the product of the the marginal wage increase and the size of the firm's *entire* work force.

- This third term is potentially important. It says that each additional worker hired (each "marginal" worker) could potentially raise the cost of all of the previous workers hired ("infra-marginal" workers). Why? If all workers are paid a single wage ($w(L)$), and calling forth an additional worker raises that wage, then the cost of the additional worker is not simply $w$ but $w + w'(L) L$.

- Contrast to the key assumption of the benchmark *competitive* model:

$$w'(L) = 0 \Longleftrightarrow \text{Price taking firms}$$

No firm is large enough to raise the market wage simply by hiring a few more workers.

- If the firm is a price taker in the labor market, it chooses employment so that:

$$pf'(L) = w^*,$$

where $w^*$ is the market wage, which the firm takes as given. In other words $w = MRPL$

- How does firm choose employment when it is *not* price taker? According to the FOC above:

$$pf'(L) = w(L) + w'(L)L$$

If $w'(L) \neq 0$ then a firm must pay *all* of its workers a higher wage as it hires each additional worker.

- Here's one convenient way to express this result

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

where $\eta$ is the elasticity of labor supply (the percent change in labor supply for a 1 percent change in the wage) as experienced by the single firm. For a price-taking firm, $\eta \to \infty$, meaning that $1/\eta \to 0$. So, if a firm is a price taker, the wage is exactly equal to MRPL (since the denominator of the above expression is equal to one). If the firm is not a price taker in the labor market, then the wage it pays is *strictly less* than MRPL.

- Why *less* than the MRPL? As we'll show below, the conditions above imply that the marginal cost of labor exceeds the wage because each marginal hire raises the cost of all prior (AKA inframarginal) hires: $MC_L > w$. The profit maximizing firm that has market power will choose $MC_L = MRPL$. But since $MC_L > w$, this implies that $w < MRPL$. Details below.

### 3.2 Conventional case: Individual Price Taking Firm

Let's return to the conventional case of individual price taking firms.

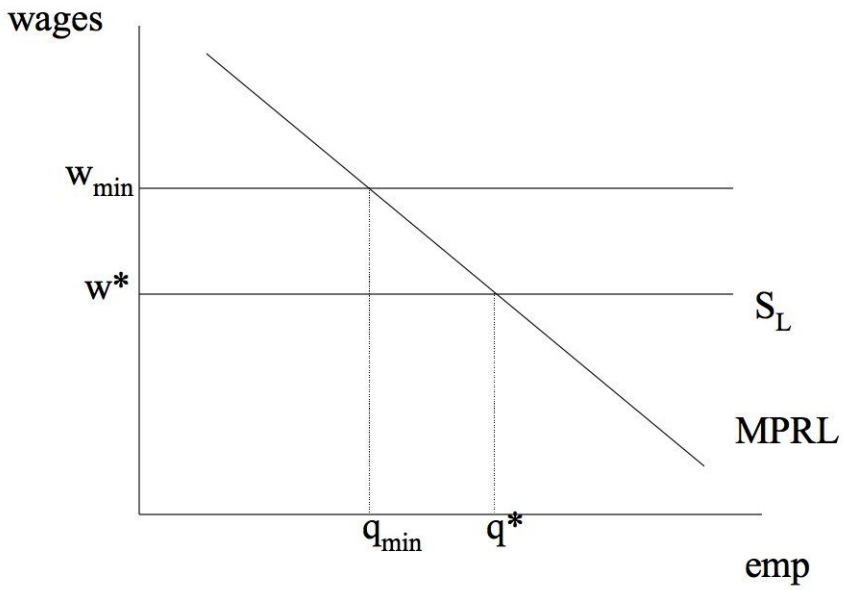


- Notice that the labor supply curve is upward-sloping at the market level, but it is flat as perceived by the single firm.
- If we imposed a binding minimum wage in this market ($w_{\min} > w^*$), each firm in this market would reduce its quantity of workers demanded.

### 3.3 Monopsonistic employer

**Definition 8.** Monopsony. "One buyer, many sellers." More generally, monopsony is a case where an agent (firm or consumer) is not a price taker in a market in which it is a buyer. Its own demand affects the price it pays. (Conversely, monopoly is a case where a firm is not a price taker in a market in which it is a seller. Its own supply affects the price it commands in the market.)

- The labor supply curve for a monopsonist is upward sloping. To obtain one more worker, the monopsonist 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), the marginal cost of the next worker is not simply her wage but also the wage increase given to all of the other inframarginal workers.
- Thus, the marginal cost of labor curve for a monopsonistic firm is *even more* upward sloping than its labor supply curve. 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.

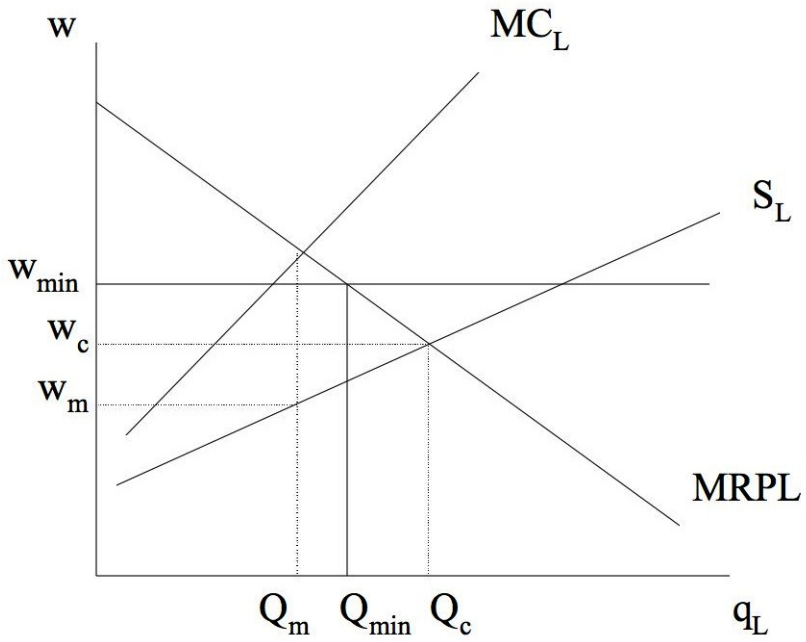


- What happens if we impose a binding minimum wage on a monopsonistic employer?
- One case is illustrated above. In this example, implementation of a binding minimum wage raises wages *and* employment:

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

- Why does that happen? The firm is now a *price-taker* for labor at $w_{\min}$. That is, there are an *unlimited* number of workers available (as far as any one firm is concerned) at the going wage—so, labor supply to the firm is perfectly elastic at the minimum wage, just as in the standard competitive case (without a minimum wage). The firm chooses

$$w = MRPL$$

  because its choice of the quantity of labor has no impact on the wage level. Thus, paradoxically, raising the minimum wage can raise both wages and employment in a monopsonistic labor market.

- Does raising minimum wage to monopsonists *always* increase wages and employment? Definitely not.

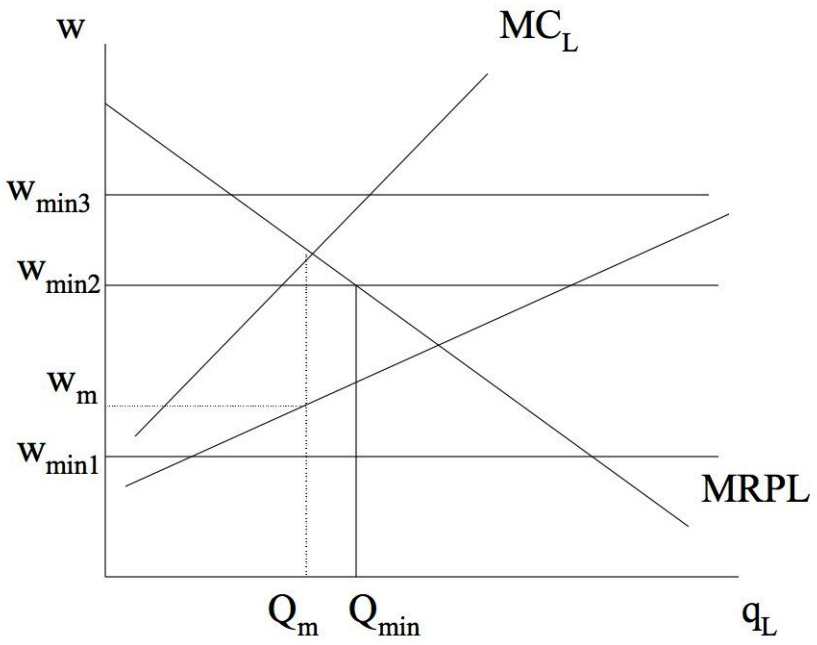


1. $w_{\min 1}$ - Introduction of minimum wage $w_{\min 1}$ has no effect because the minimum wage is below $w_m$ and hence doesn't bind
2. $w_{\min 2}$ - Introduction of minimum wage $w_{\min 2}$ raises wages and employment
3. $w_{\min 3}$ - Introduction of minimum wage $w_{\min 3}$ raises wages but reduces employment

- So, this bit of simple theory presents an interesting possibility. It is conceivable—though not necessarily likely—that a mandated minimum wage could raise both earnings and employment. If so, this is a policy that many policymakers would support (though not most businesses; this is a redistribution of wealth from firms and their customers to workers).

## 4 Testing for monopsony in the labor market

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.)

- "Company towns" such as old mining towns, where the mining company was the only employer
- Cases where skills are very specific, e.g. Tesla automobile technicians
- "Captive" labor markets, e.g., spouses of soldiers at rural military bases or in remote island locations.
- Fast food restaurants located in nearby towns in New Jersey and Pennsylvania? [Clearly, this setting is the least likely place to find monopsony, which makes it an interesting place to look.]

### 4.1 Testing for monopsony in the labor market

- How do we go about testing the monopsony vs competitive model of the labor market?
- *Focus on the key empirical implication that distinguishes these models:*
  - In the competitive model, an increase in the minimum wage always reduces employment: $W \uparrow \rightarrow L \downarrow$
  - In the monopsonistic model, an increase in the minimum wage (may) raise employment: $W \uparrow \rightarrow L \uparrow$
- How do you test this implication?
- We could look across different states and ask ourselves the following question: Is employment higher in states where wages are higher?
- Let's suppose you find the following pattern:

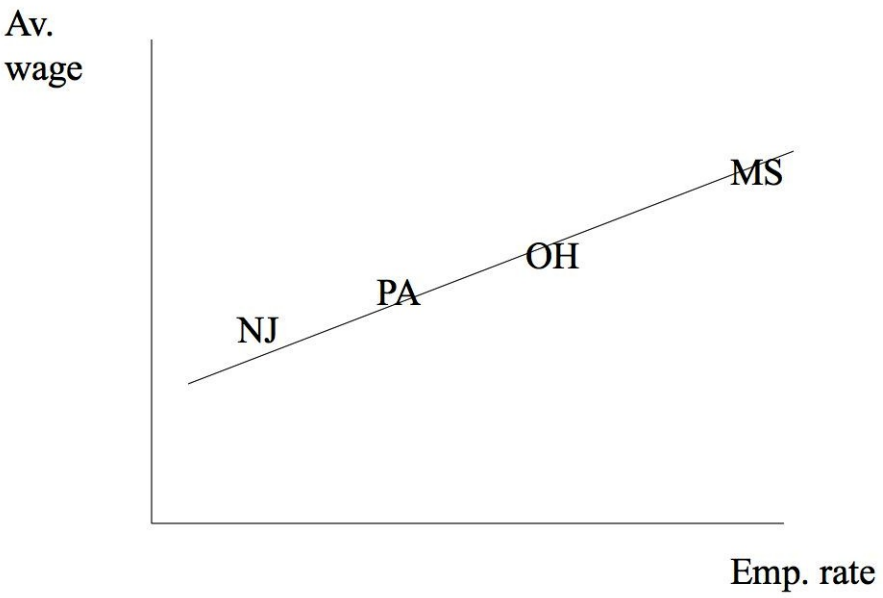


  Would this convince you that higher wage levels *caused* higher employment? I hope not!
- What's the problem with the wage here? We don't know why the wage differs across states. For example, there might be different demand and supply schedules in each state (after all, we don't think the labor market for fast food workers is a nationwide one; presumably, people won't move cross-country for a job at McDonalds).

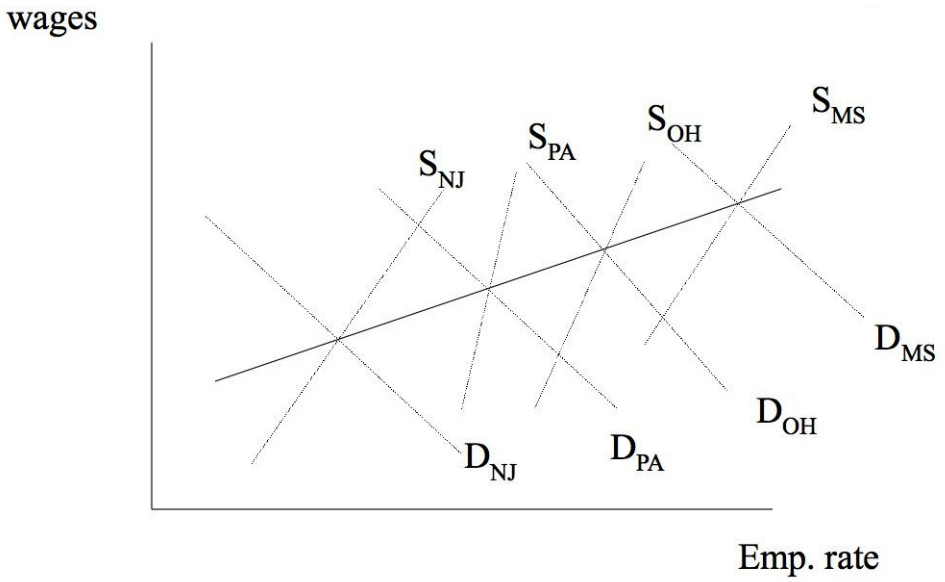


Since both employment and wages are endogenous outcomes—determined by both supply and demand—this picture tells us essentially nothing about the impact of minimum wages on employment.

- A further problem: While we may believe in the existence of supply and demand curves as an outcome of market processes, we do not ever see these curves. What we *observe* is the equilibrium wage level and the quantity employed. Thus, we cannot directly see whether or not individual firms face upward sloping labor supply (as would occur if they had monopsony power).

- How do we overcome this problem? We need an experiment, specifically, one in which wages are raised exogenously. (Why not an experiment that shifts supply inward so that fewer workers are willing to work at any given wage? Consider the empirical implications in both the monopsonistic and competitive cases.)

- If we could exogenously manipulate the minimum wage, we could study its impact on employment to infer the slope of the relationship between wages and employment (downward sloping $\rightarrow$ competitive market, upward sloping $\rightarrow$ monopsony)

### 4.2 An idea: New Jersey's 1993 minimum wage change

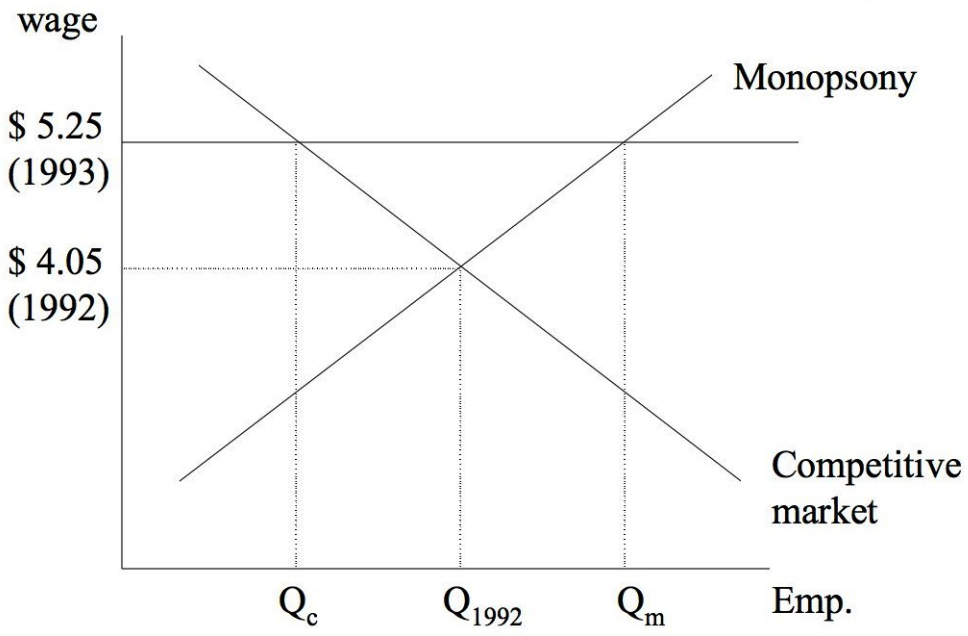


- Notice that the conditions under which the introduction of a minimum wage raises employment in a monopsonistic market are only locally satisfied—that is, raising the minimum wage by "too much" will reduce employment even in the monopsonistic setting. Thus, if we find that an increase in the minimum wage raises employment, this is *sufficient* but not *necessary* to establish the existence of monopsony.
- By studying the pre/post change in employment following the adoption of the minimum wage, we can explore whether employment rises (monopsony) or falls (competitive market).
- Before we explore this relationship empirically, we need to take a moment (okay, a half hour) to discuss causal inference.

## 5 A word on 14.03/003 class requirements and expectations

1. Readings—About one-third of lecture time will be devoted to discussing important published (and sometimes unpublished) papers in economics in some detail. If a reading is marked with a †, you are responsible for preparing the paper prior to class. I do not expect mean that you will memorize and digest the entire article. I do expect you to read the Abstract, Introduction and Conclusion so that you can answer the following questions: what is the research question; how did the authors go about answering the question (i.e., what was the research design); what conclusions did they reach?
2. Class participation and cumulative in-class quizzes each count for 5% of your grade.

3. There will be 6 problem sets, with your lowest score dropped automatically. The other 5 p-sets each count for 6 percent of your grade. Most are due at 5pm on the designated date, though the second p-set is due at **9am at 10/3/2016** (there's a scheduling reason for this). No late problem sets are accepted—that's what the automatic drop rule is for. (Recommendation: don't waste your freebie early in the semester.) Problem sets are a mixture of formal problems and questions based on the readings and lectures. *All assignments must be electronically submitted via the class website.* (This ensures that everything is submitted on time and nothing gets lost.)

4. There will be about a half-dozen in class quizzes (again, one is dropped). These quizzes will be brief, non-technical checkups on your knowledge of what's going on in the class (readings, lectures, problem sets). I also cold call in class to help overcome your natural shyness (and equally natural sleepiness).

5. There will be two in-class exams and one final exam, each 20 percent of your grade. I do not re-test the same material on subsequent exams. Thus, each exam roughly covers one third of the material in the class. Of course, much of this material is cumulative, so I would not advise you to willfully forget each third of the class after the corresponding exam.

6. On attending class. This class is meant to be participatory. The two excellent recommended textbooks (Banerjee and Angrist and Pischke) are meant to serve as references and do not cover most of the in class material. The lecture notes fill in more of the picture, but they are also not a complete guide to what happens in class. If you don't regularly attend class, you will likely have difficulty on problem sets and exams (and you will be marked down for missing quizzes and for not participating in class). If you were planning to only show up in 14.03/14.003 for exams, I recommend against taking the class.

7. The class is not graded on a strict curve. Everyone can do well (or badly). In general, students receive A's, B's, and C's in the class.

8. If you do only the minimal amount of work in 14.03/14.003, you'll probably get a C. If I think you are heading for a D, I will contact you to recommend that you drop the class. I cannot help you if your grades plummet after the MIT drop date, however.

9. Support outside of class:

    - Two recitations are held every Friday, one hour apart. During the recitation, 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 review statistical concepts used in 14.03/14.003. Unless you have taken a class on probability and multivariate statistics (e.g., regression) such as 14.32, it's likely that some of this material will be unfamiliar to you. I'll be using these tools in class but I won't be covering them in detail.

- There are two TA office hours, and one instructor office hour per week (each 90 minutes). TAs will try to meet with you by appointment if you are unable to make office hours. But the office hours should be your first recourse.
- Questions on class topics and problem sets. Use the class web site/Wiki, which we'll monitor.
- *Please do not email us with substantive, class material related questions; these are the for the web site. Of course, personal issues should be handled by email or in person (not on the Wiki).*

14.03 / 14.003 Microeconomic Theory and Public Policy
Fall 2016