Externalities, the Coase Theorem and Market Remedies

David Autor

14.03 Spring 2004

1 Introduction

- You were introduced to the topic of externalities in 14.01. An externality arises when an economic actor does not face the 'correct price' for taking a specific action. The correct price of an action is the marginal social cost of that action. As we discussed during the section on General Equilibrium, when markets work properly, they align private costs and benefits with social costs and benefits. When private benefits differ from social benefits (either higher or lower), externalities result.
- Some classic externalities include:
  - Traffic: When I take the highway, I increase congestion for other drivers, a negative externality. Since the toll on the Mass Turnpike does not vary with traffic conditions, I probably face the wrong price of driving on the highway (too high at non-peak hours, too low at peak hours). As a result, I use the Pike 'too much' during peak hours and not enough during non-peak hours.
  - Disease transmission: When I decide whether to have my children receive flu shots, I consider whether the cost of the inoculation in time, money, discomfort is worth the reduced risk. I probably do not consider that by protecting my children from the flu, I also protect the children at their school. Because my private benefit of the shot does not incorporate the external social benefit of the shot, I am less motivated than I "should be" to get my children inoculated. It is therefore likely that too few children receive small pox vaccines.

    Ironically, there are other parents who recognize that, because most parents *do* get their children inoculated, other kids may be reasonably protected even without receiving an inoculation. Hence, these parents free-ride on the positive externality, and are even less motivated to get a shot than the parents who do not consider the positive externality they are generating. This exacerbates the problem.

- Pollution: Because clean air is not priced, I pay essentially no cost to pollute the air. When I decide whether to drive to work or take the train, my marginal cost of driving (fuel plus wear and tear on my car) probably does not incorporate the social cost of additional pollution. Because my private cost is lower than the social cost, I will likely drive 'too much' relative to a case where I faced the full marginal social cost.

Are these externalities never 'internalized' by the market?

2 The Coase Theorem

- Until the publication of Ronald Coase' 1960 paper, "The Problem of Social Cost," most economists would have answered yes. Coase made them reconsider that view.
- Coase gave the example of a doctor and a baker who share an office building. The problem: the baker's loud machinery disturbed the doctor's medical practice. The doctor could not treat patients while the baker's machinery was running.
- The standard economic reasoning (at the time) was that the baker should have to compensate the doctor for the harm he was causing since he was 'causing' the externality. Having the baker provide compensation would correct the externality imposed on the doctor.
- But is this reasoning complete? Coase pointed out that one could re-frame this problem as follows: a doctor sets up his office in a new building and after moving in notices that the baker's machinery is too loud for him to conduct his practice. He demands that the baker shut down his operation due to the disturbance.
- Who is responsible for the externality in this case? One can legitimately argue that the doctor is creating an externality by requiring the baker to bake in silence. The baker's noise can be viewed as an 'input' into his production of baked goods. Without it, the baker could not perform his work. So perhaps the doctor should accommodate the baker, either by moving his practice or by installing soundproofing.
- If this reasoning is valid, then who should compensate whom? From a legal point of view, the answer may be clear. From an economic point of view, the answer is indeterminate based only on the information provided.

- Consider the following additional information. The baker could buy quieter machinery for $50. The Doctor could soundproof his walls for $100. Economic efficiency demands that the lowest (marginal) cost solution that achieves the objective is the right solution. The baker should buy quieter machinery.
- So, does this mean that the baker should have to pay to abate? It does not.
- Consider the following scenarios:
    1. The town council assigns the doctor the right to control the noise level in the building. He notifies the baker that he needs quiet to work. The baker spends $50 for quieter machinery.
    2. The town council assigns the baker the right to make as much noise as needed to do his work. The doctor complains about the noise and the baker points out that he has the right to make as much noise as he likes. Will the doctor now spend $100 to sound proof his office? If the doctor and baker can negotiate readily, they should arrange for the doctor to pay the baker $50 to buy quieter machinery.
- As this example demonstrates, the efficient economic outcome should occur regardless of which party is 'responsible' for the externality. In either case, quieter baking machinery is purchased.
- However, the legal framework does matter. If the 'sound rights' are assigned to the doctor, the baker spends loses $50. If the sound rights are assigned to the baker, the doctor spends $50.
- So the Coase theorem says the following. If (1) property rights are complete (so, in our example, one party clearly owns the 'sound rights') and (2) parties can negotiate costlessly (so the doctor and baker don't come to blows), then the parties will always negotiate an efficient solution to the externality. The law determines who pays this cost, but the outcome is the same. (Note the parallelism with the Welfare theorems: efficiency and distribution are separable problems.)
- The Coase theorem implies that the market will solve externalities all by itself unless: (1) property rights are incomplete (for example, no one owns the air) or (2) negotiating is costly (for example, the entire population owns the air, but all citizens cannot simultaneously negotiate about pollution levels).
- The Coase theorem is often misinterpreted to suggest that the market will solve all externalities. This is not true, and Coase will probably go to his grave railing against the 'Coaseans' who make this claim.

- Rather, the Coase theorem suggests that the market can *potentially* solve externalities *if* property rights are clearly assigned and negotiation is feasible.

- In some cases, this is clearly infeasible.

  - Airlines cannot realistically negotiate with individual homeowners for overflight rights to their houses, even though these overflights do create externalities.
  - I cannot negotiate with all handicapped drivers for the use of an empty handicap space in an emergency, even though I'd be glad to pay these drivers handsomely to rent the parking space.

- A key inference that follows from the Coase Theorem is that the best solution to resolving an externality may not be to regulate the externality out of existence but rather to assign property rights or facilitate bargaining so that the relevant parties can find an economically efficient solution.

3 Remedying pollution: Three approaches

- Consider two oil refineries that both produce fuel, which has a market price of \$3 per gallon (assume demand is infinitely elastic so that this price is fixed regardless of the quantity produced).

- Assume that each refinery uses \$2 in raw inputs (crude oil, electricity, labor) to produce 1 gallon of fuel.

- In addition, each plant produces smog, which creates \$0.01 of environmental damage per cubic foot.

- The amount of smog *per gallon of fuel* produced differs at the two plants:

$$\begin{aligned} s_1 &= y_1^2, \\ s_2 &= \frac{1}{2}y_2^2, \end{aligned}$$

  where y_1, y_2 denote the number of gallons of fuel produced at each plant. So, plant 2 pollutes only $\frac{1}{2}$ as much for given production.

- Assuming initially that there are no pollution laws. In this case, each plant will produce as many gallons as it can until it runs out of capacity (since it makes \$1 profit per gallon). Assume each plant can produce 200 gallons.

3.1 Competitive outcome

- What will firms optimally do in the absence of any legal framework for resolving the externality? The problems for the respective firms are:

$$\begin{aligned}\max_{y_1} \pi_1 &= y_1 \cdot (3-2) \; s.t. \; y_1 \leq 200, \\ \max_{y_2} \pi_2 &= y_2 \cdot (3-2) \; s.t. \; y_2 \leq 200, \\ y_1^* &= y_2^* = 200.\end{aligned}$$

- Each firm ignores the social damage from its smog production (notice that s_1, s_2 do not enter into the firms' profit maximization problems). Hence, pollution is $s_1 = 40,000$, $s_2 = 20,000$. The negative pollution externality is \$400 and 200 from plants 1 and 2 respectively.

3.2 Welfare maximizing case

- Before analyzing how to correct this externality, we need to determine the 'optimal' level of pollution. In this case, optimal pollution is non-zero. More generally, not all activities that generate externalities should be stopped. But if these activities generate negative (positive) externalities, then social efficiency generally suggests that we want to do less (more) of them than would occur in the free market equilibrium.

- To get the socially efficient level of fuel production, we want to equate the marginal social benefit of the last gallon of fuel to the marginal social cost.

- What is the social benefit? It is \$3. This comes from the infinitely elastic demand curve.

- The marginal social cost of production is \$2 in raw inputs plus whatever pollution is produced.

- The efficiency condition is $MB_s = MC_s$, marginal social benefit equals marginal social cost.

- We therefore want it to be the case that at the margin, that no more than \$1 of environmental damage is done per gallon of fuel produced. Consequently, no plant should produce more than 100 cubic feet of smog per gallon of fuel.

- (Note: no plant should produce *less* than this amount either since the pollution is indirectly beneficial: it is an 'input' into the production process; less pollution means less fuel production).

- Imagine that each plant faced the private *plus* social costs of production. If so, we could rewrite the previous profit maximizing conditions as:

$$
\begin{aligned}
\max_{y_1} \ \pi_1 &= y_1 \cdot (3-2) - 0.01 y_1^2 \quad s.t.\ y_1 \leq 200, \\
\max_{y_2} \ \pi_2 &= y_2 \cdot (3-2) \ - 0.005 y_2^2 \ \ s.t.\ y_2 \leq 200, \\
y_1^{**} &= 50,\ y_2^* = 100.
\end{aligned}
$$

- When Plant 1 is producing 50 gallons, the marginal gallon produces 100 cubic feet of smog, which causes $1.00 in environmental damage. More pollution than this would be socially inefficient since fuel sells for $3 and uses $2 in raw inputs to produce. For Plant 2, the corresponding production is 100 gallons because this plant produces less smog per gallon.

We now have an efficient benchmark for welfare maximization.

How do we get plants to produce the socially efficient level of pollution? Three classes of regulatory solution are possible. Each has different properties.

3.3 Command and control ('quantity') regulation

- 'Command and control' regulation is the traditional approach to limiting externalities. This approach sets numerical quantity limits on activities that have external effects. It is often called 'quantity' regulation.

- The most common command and control regulation is simply banning an activity – 'though shalt not litter.' But as we know, just because an activity has external effect does not mean it should be banned outright–only that too much or too little relative may be done relative to the social optimum.

- Much command and control regulation recognizes this point, and so permits some positive amount of an activity, but less than a private actor would otherwise undertake.

- How does this apply to the example above? We know the optimal quantity of production for each plant from our calculations above. We could therefore pass a law that says "Plant 1 may produce 50 gallons of fuel and Plant 2 may produce 100 gallons of fuel." This will achieve exactly the desired result.

- But this kind of regulation is clumsy.

  - It's difficult to write laws that regulate the behavior of each plant individually.

- Once passed, such laws are difficult to modify as technology or pollution costs change.
- If the law cannot be written to regulate each plant's output differentially, further inefficiencies will result.
- [For example, calculate as an exercise the optimal amount of fuel production to permit these two plants to produce if the regulator must apply the same numerical production cap (in fuel) for each plant. (Hint: the answer is not 75 gallons.) This is actually a commonplace case: regulators can set average output at the industry level but cannot further regulate the behavior of individual plants. As your calculation shows, this leads to inefficiencies where some regulated plants pollute 'too much' and other regulated plants pollute 'too little' relative to the efficiency condition that $MB_s = MC_s$.]

- Despite these weaknesses, command and control regulation is the most common approach taken for regulating externalities.

3.4 Pigouvian tax ('price regulation')

- An alternative approach is to use the price system to 'internalize' the externality.
- We know from above that the marginal social cost of pollution is \$0.01 per cubic foot of smog. If we charged firms for polluting, the social cost would be incorporated in the private cost. Done correctly, firms will make optimal choices.
- This type of tax is known as a Pigouvian tax after the economist Pigou who suggested it.
- Specifically, it we set the pollution tax at $t =$ \$0.01 per cubic foot of smog, then each plant would choose the optimal quantities due to its profit maximization. In other words

$$\begin{aligned}\max_{y_1} \pi &= y_1(3-2) - t \cdot y_1^2, \text{ where } t = 0.01 \rightarrow y_1^p = 50 \\ \max_{y_2} \pi &= y_2(3-2) - t \cdot \frac{1}{2} y_1^2, \text{ where } t = 0.01 \rightarrow y_2^p = 100\end{aligned}$$

- This solution achieves the desired result with arguably less complexity. Facing this tax, plants will choose the efficient amount of production. *We do not have to write a separate law for each plant.*
- Note that this problem is made especially simple by the assumption that the marginal damage of pollution is always \$0.01 per cubic foot. If the marginal damage varied with the amount of pollution (plausible), then setting the right tax schedule would be much harder. For example,

if pollution above a certain threshold caused mass extinction but pollution below this level did little harm, this Pigouvian taxation scheme would be *quite* risky. Setting the tax slightly too low would result in calamity.

3.5 Assigning property rights: The Coasean approach

- The Pigouvian tax idea does not really use the Coase theorem. It aligns private and social costs by pricing the social costs, thereby causing firms to internalize these costs. The tax arguably does use property rights – the state is now selling firms the right to pollution at price \$0.01 per cubic foot. But the Pigouvian solution does not create conditions for negotiation among firms. The state sets the price and collects the tax.

- The Coase theorem suggests that we may be able to do even better. If property rights are fully assigned, then the regulatory body should not, in theory, have to be involved. Instead, parties will negotiate among themselves to find the lowest cost solution to correcting the externality.

- How can this insight be applied? Because firms do not own pollution rights, there is not an efficient negotiation over the how much pollution is generated. This motivates the idea of *selling pollution rights.*

- Using the algebra above, we can calculate that the 'optimal amount of pollution' is $50^2 + \frac{1}{2}\left(100^2\right) = 7,500$ cubic feet of smog. This is the socially efficient quantity of pollution that that results from producing the optimal quantity of fuel.

- In this example, the government could issue $7,500$ "permits to pollute" 1 cubic foot of smog. These permits could be used by the permit holder to pollute, or could be sold by the permit holder to another refinery so it could pollute instead.

- How would this work? Consider two cases.

- First, the permits are all given to Plant 2, the more efficient refinery. It could do the following:

    1. Produce 122.4 gallons of fuel (pollution is $\frac{1}{2} \cdot 122.4^2 \simeq 7,500$). Its profits would be \$122.40.
    2. Produce 100 gallons of fuel (pollution is $\frac{1}{2} \cdot 100^2 = 5000$) and sell its $2,500$ remaining permits to Plant 1 (assuming that this transaction is next to cost-less). With these $2,500$ permits, Plant 1 could produce 50 gallons of fuel (pollution will be $50^2 = 2,500$). Since its profits are \$1 per gallon, it would pay up to \$50 for these permits. Plant 2 would therefore make \$150 in profits by using $5,000$ permits and selling $2,500$ others.

3. Plant 2 could also implement any mixture of these two options, including selling all of its permits to Plant 1. But you should be able to demonstrate to yourself that Plant 2 cannot do better than the 2nd option above.

- Now, instead, assume the permits are given to Plant 1, the less efficiency refinery. It could do the following:

  1. Produce 86.6 gallons of fuel (pollution is $86.6^2 \simeq 7,500$). Its profits are \$86.60.
  2. Sell all of the permits to Plant 2, the more efficient plant. Plant 2 will pay up to \$1 per gallon produced. Hence, Plant 1's profits would be \$122.4 dollars.
  3. It could keep 2,500 permits and sell 5,000 to Plant 2. Here profits would be \$150 because Plant 1 would produce 50 gallons and Plant 2 would produce 100 gallons and pay up to \$100 for the privilege.

- [Optional:You could demonstrate this result to yourself more rigorously by calculating each plant's marginal willingness to pay for permits as a function of its output. For the first gallon of production, plant 1 is willing to pay up to \$1 per permit: it produces only 1 cubic foot of pollution and makes \$1 of profit. Similarly, for the first gallon of fuel, plant 2 is willing to pay \$2 per permit: it needs only 1/2 of a permit to make the first \$1 in profit. But willingness to pay falls off rapidly with each additional gallon of fuel produced since the number of permits required per gallon of fuel rises with output. The equilibrium market price, p^*, for permits is where the marginal willingness to pay for permits is equated between the two plants at fuel output Y_1^{**}, Y_2^{**}, and total permits consumed is 7,500. At this price, each plant is indifferent between selling/buying additional permits and producing/not-producing the next gallon of fuel. At equilibrium p^*, the quantity of permits consumed will *not* be identical at the two plants. That's because at any given non-zero quantity, plant 2 is always willing to pay more for the next permit than plant 1. The marginal willingness to pay of the two plants can only be equated when $Y_1 > Y_2$. Hence, there will be a market clearing price p^* where $Q(p^*) = 7,500$, $y_1 = 50$ and $y_2 = 100$]

- The key result: regardless of which plant receives the permits, the key economic outcome is identical: fuel produced, pollution produced, *and* (surprisingly) the allocation of production of pollution (and fuel) across plants.]

- Why does this equivalence hold? Once pollution property rights are assigned, the plants negotiate to achieve the most efficient solution to the externality.

- What differs between the two allocations is: which plant makes the profits (a transfer among plant owners).
- This cap and trade example demonstrates the power of the Coase theorem. By assigning property rights to pollution, the government allows the market to correct the externality.
- And as the Coase theorem indicates, the exact distribution of property rights to interested parties (Plant 1 or Plant 2) has no effect on economic efficiency. But it does greatly affect the distribution of profits across the two plants (or their owners). This allocation problem is a major political stumbling block to implementing cap and trade regulations generally: how do lawmakers assign the initial ownership rights to pollution (or other negative externalities)?

4 Comparison of the three methods of abating an externality

- These three methods – command and control, Pigouvian taxation, and cap and trade – have identical economic consequences. But they are not identical from a regulatory perspective.
- Command and control regulation requires intimate knowledge of the production structure of each plant. It is cumbersome to implement and to get right. Sometimes it is not feasible or legal to regulate firm's behavior at the plant level.
- The Pigouvian taxation has the advantage that plants will optimal choose the level of pollution that maximizes their profits, including the cost of the Pigouvian tax. But Pigouvian taxes are risky when the marginal social cost of pollution varies with the quantity – for example, above a certain threshold everyone dies. In these cases, it is difficult and possibly risky to attempt to set the tax exactly right.
- Cap and trade regulation has several special virtues:
  1. Like command and control, it allows the regulator to set the amount of pollution to whatever level is desirable (the Pigouvian tax will not do this unless the regulator knows the cost structure exactly).
  2. Like the Pigouvian tax, cap and trade is comparatively simple to implement since the regulator does not need to write a separate law for each plant.
  3. Unlike the other means of regulation, it causes firms to optimally reallocate pollution among themselves through trade (as the Coase theorem predicts). Even if the regulator does not

know firms' cost structures, the cap and trade system will cause the least polluting firms to do the majority of the production since its social cost of production is lowest.

4. If the regulator cares specifically about *which* plant does the polluting, however, cap and trade will not generally achieve the desired result. This case might be relevant if introduction of a cap and trade rule caused substantial geographic concentration of pollution (let's say all the low-cost polluters in the U.S. just happen to be located in Harvard Square).

- There is an article by Schmalensee et al. on your syllabus that discusses the creation of the market for Sulfur Dioxide emissions in the United States. This policy experiment is a triumph of economic reasoning applied to environmental regulation. You may be interested.

5 Summary

- Externalities are a source of economic inefficiency. But they are also potentially correctable through the market.
- The Coase theorem identifies the two conditions needed for an efficient market solution: complete property rights and zero (or low) transaction costs.
- Sometimes these conditions can be approximated by assigning property rights, thereby creating a market for the externality.
- Understanding why externalities persist in equilibrium comes down to identifying why the Coase theorem does *not* hold in a specific circumstance.
- Rectifying the externality often means finding a way to restore market conditions so the Coase theorem will hold. When that isn't feasible, external quantity regulation (like command and control regulation) may be needed.

6 An example: The market for real estate brokers

- As we've discussed this semester, the price system solves an informational problem: determining how much of a good should be produced and how much should be consumed. Production should occur until the marginal willingness to pay is equated with the marginal cost of production. When prices rise, more should be produced and/or less consumed. When prices fall, less, more should be consumed and/or less produced. Prices provide signals to consumers and producers about how to adjust production and consumption. These signals continually push the market back towards equilibrium.

- What happens to supply and demand when prices are not set (in a logical manner) by market forces? The Hsieh and Moretti paper, in the *Journal of Political Economy*, 2003, provides one great example.

- The market for real estate brokers appears to be cartelized by the real estate brokerage industry. Brokerage commissions are fixed across time and space at six percent of the selling price of the property, regardless of the price of the property, the state of the market (active, slow), the experience of the broker (old, young), the number of competing brokers available (a glut, a shortage), the brokerage services the seller's desire, etc. It is hard to explain this fixity by any mechanism other than collusion. Collusion appears coordinated and enforced, as Hsieh and Moretti discuss, by use of a national sales database (MLS) that publishes the brokerage commissions charged on every sale. Brokers may enforce the cartel by penalizing one another for price discounting and by shunning sellers who attempt to sell their homes without a broker (even though self-sellers often advertise their willingness to pay the 'selling broker's' 3% of the deal).

- This fixed commission structure creates a bizarre market pricing scheme for a real estate sales. The 'price' of a sale is always higher on more expensive properties, even though these properties may not take more work to sell. Much stranger: when real estate appreciates, realtors' fees rise. This means that whenever rents or sale prices increase, there is an automatic spike in the transaction price that is generated by the fixed commission structure.

- Is this efficient? It seems unlikely. When the price of housing appreciates, this does not make it intrinsically harder for brokers to sell – in fact, it may be easier since rising prices signal a 'hot market.' If so, rising commissions may 'over-compensate' brokers relative to their opportunity costs. In economic jargon, brokers may earn '**rents**' on sales.

- A well-known problem with rents is that they create incentives for 'rent-seeking.' If someone is handing out free money, people will expend real resources to get some of it. For example, they may stand in line. And if there is a lot of free money to be had, the line will be very long – so much so that the last person in line may be just indifferent between getting the free money and going home. It is possible for no one to earn rents in equilibrium because the rents are entirely dissipated by rent-seeking. This is much worse than a simple case where rents are earned: substantial resources are consumed to reach an equilibrium where no one gains from rents. (Can it be worse than that? Yes. There is no reason why the sum of resources expended on rent-seeking is bounded by the amount of rents available. If a thousand people each waste a dollar (for example, by waiting in line) seeking \$100 in rents, then \$900 is lost on rent-seeking.)

- Could this case be relevant in the real estate market? Consider the following simple conceptual model:

  1. There is a supply of houses on the market that depends on the transaction price of a real estate deal. Denote $H(P)$ as the number of houses for sale at price P, with $H'(P) < 0$.
  2. There is a supply of real estate agents, $R(P)$, willing to provide services at price P, with $R'(P) > 0$.
  3. In free market equilibrium, P^* solves $H(P^*) = R(P^*)$: the supply of brokers is equated with the demand from sellers to sell their houses.

- Define the following quantities:

  1. The sum of sales commissions is the number of houses times the transaction price. Notice that this is also the sum of wages paid to brokers.

  $$\ln(C) = \ln(P \cdot H(P)) = \ln P + \ln H(P).$$

  2. The average realtor wage is the sum of commissions divided by the number of brokers.

  $$\ln(w) = \ln[P \cdot H(P)/R(P)] = \ln P + \ln H(P) - \ln R(P)$$

- In free market equilibrium, the realtor average wage is:

$$\ln(w^*) = \ln\left(\frac{P \cdot H(P^*)}{R(P^*)}\right) = \ln P^* + \ln H(P^*) - \ln R(P^*).$$

- But we won't reach this free market equilibrium because P is *not* set by supply and demand. It is set by the price of housing itself.

- Index the price of housing by HPI, for Housing Price Index. HPI measures changes in the real cost of housing; a 1 percent rise in the HPI means a 1 percent rise in the cost of housing: $\partial \ln P/\partial \ln HPI = 1$.

- How does an exogenous rise in HPI affect the wages of real estate brokers? Your first instinct might be that

$$\partial \ln w/\partial \ln HPI = 1.$$

  But this is not quite right. The reason is that an exogenous increase in the transaction price will reduce the number of houses for sale *and* increase the supply of realtors.

- In particular:

$$\frac{\partial \ln C}{\partial \ln P} = 1 + \frac{\partial \ln H(P)}{\partial \ln P} < 1,$$

so commissions will rise less than one for one with prices.

- In addition, new brokers will enter the market.

$$\frac{\partial \ln R}{\partial \ln P} > 0.$$

- Putting these together:

$$\frac{\partial \ln w}{\partial \ln P} = 1 + \frac{\partial \ln H(P)}{\partial \ln P} - \frac{\partial \ln R}{\partial \ln P} < 1 \text{ and } \gtrless 0.$$

- In words, broker wages must rise less than one for one with house prices. Moreover, it's possible for broker wages to *stay the same or fall* when prices are artificially increased by a change in housing prices. If a sufficient number of new brokers crowds into the market when HPI rises, broker income per capita may remain unchanged (or fall).

- If this occurs, it has a second direct implication. Define Realtor productivity as

$$\gamma = \frac{H(P)}{R(P)},$$

equal to sales per realtor. What is $\partial\gamma/\partial(\ln P)$?

$$\frac{\partial \gamma}{\partial \ln P} = \frac{\ln H(P)}{\partial \ln P} - \frac{\ln H(R)}{\partial \ln P} < 0.$$

Sales per broker fall as housing prices rise. More brokers chase a fixed amount of business, which reduces productivity per broker.

- In what sense does this set of outcomes reflect an externality? The problem is pure business stealing. Increased Broker supply in response to an HPI-induced price rise does *not* increase the number of houses sold nor (to a first approximation) raise the well-being of sellers. It simply sends more realtors into the market to chase a fixed quantity of deals. Output in the realty sector is unchanged by the rise in realtors. But the opportunity cost of realtor's labor in other sectors is lost.

- This scenario is a purely *dissipative* externality. Society is made worse off by the entry of new real estate brokers into the sector (due to their opportunity costs). But neither incumbent *nor* entrant brokers *nor* homeowners benefit from rising prices (at least in the case where $\frac{\partial \ln w}{\partial \ln P} \leq 0$).

6.1 Putting this hypothesis to the test

- This simple conceptual framework is surprisingly straightforward to test.
- It has three main empirical implications. As housing prices rise:

  1. The number of real estate brokers increases
  2. Productivity per broker falls
  3. Wages of brokers rise by much less than the price of houses

- The figures from Hsieh-Moretti make the case very clearly.
- Are there any reasons to think that the entry of new brokers in response to rising housing rise is *not* a pure social waste?
- How could this market structure be altered to produce a more socially efficient outcome?
- Assume that you were constrained to keep the current fixed commission structure in place. Is there any regulatory action that could be taken to make the market operate more efficiently given this constraint?