Lecture Note 6: Applied Competitive Analysis Part I: Taxation and Market Equilibrium

David Autor, MIT and NBER

14.03/003 Microeconomic Theory and Public Policy, Fall 2025

(compiled October 5, 2025)

1 Consumer demand and market demand

- To develop consumer theory, we've analyzed the choices of **individuals** who take prices and income (or utility) as **given** (exogenous) when making optimizing decisions
- But where do prices come from?
- Where does income come from?
- We're going to answer those questions in two steps over the course of four lectures
  1. How do prices emerge in a market for a single good, taking incomes as given. This is *partial equilibrium applied competitive analysis*
  2. How are prices and incomes determined simultaneously. This is *general equilibrium competitive analysis*. This is *not* the subject of the current lecture note, nor the next one. It will be covered in Lecture Note 8.

1.1 Where do prices come from?

- Why does a gallon of water cost less than a gallon of gas–even though water is surely more essential?
- Why does a gallon of Coke (usually) cost more than a gallon of gas, even though Coke is mostly water?
- A potential explanation
  - *Prices are set at the margin:* Prices reflect *marginal* willingness to pay (AKA, marginal value to consumer) not *average* willingness to pay
  - The price for something that is essential to life can be very low if that thing is abundant, so marginal utility value is low
  - The price for something that is inessential to life can be quite high if that thing is scarce, so marginal utility value is high
- All of these points are visible in the aggregate labor market diagram that we drew in the first lecture, i.e., the canonical Marshallian cross, and that is repated below.

Textbook model of wages and employment

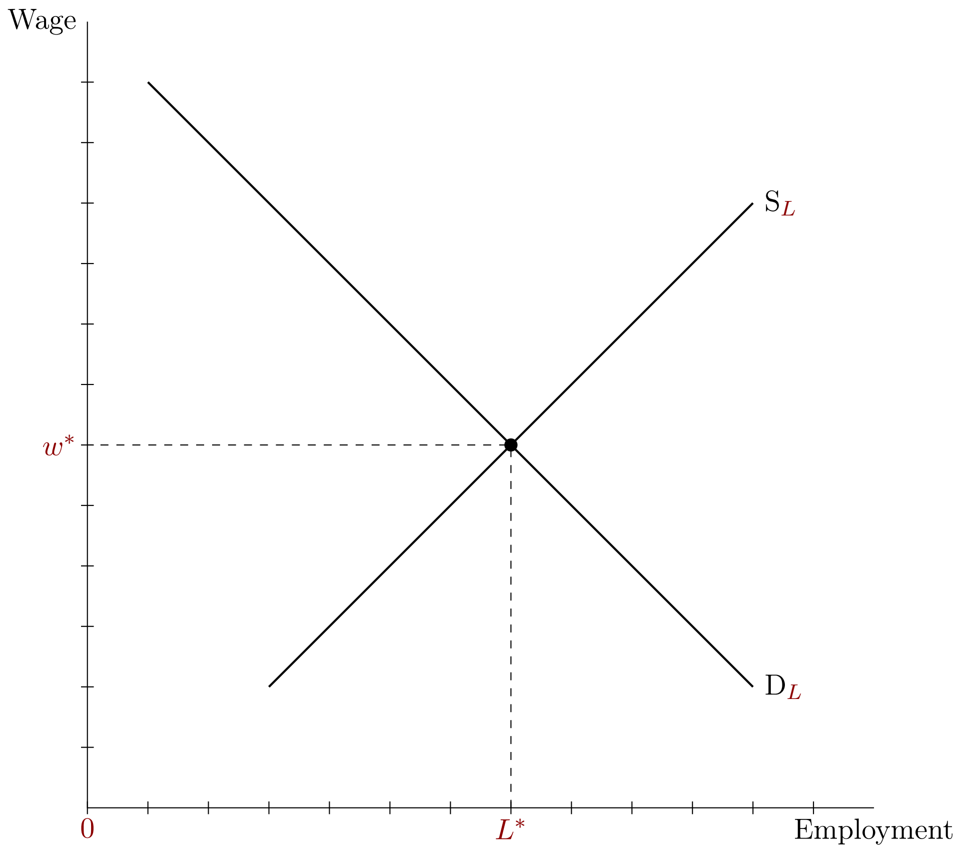

Notice the following

- The wage is the market clearing *price* of labor
- w^* is not determined solely by the 'cost' of supply labor (the supply curve) nor by the 'value' of labor (the demand curve) but by intersection of the two
- At this point of intersection, w^* every worker who wants to work at wage w^* is employed, and every employer who wishes to hire a worker at wage w^* is able to hire one. w^* is the only point in this figure that has this property. Any other wage w' will generate either more workers seeking work than are desired by employers (wage too high) or more employers seeking workers than are available (wage too low)
- The sum of producer and worker surplus is maximized: the area between the supply and demand curves (to the left of l^*) is as large as it could be. That means that the solution is Pareto efficient. There is no deviation from this point that could make someone better off without making another party worse off. If we were not at this point, there would be ways to grow the pie, potentially making someone better off without making anyone worse off.

- w^* is *not* the average product of the workers who are employed. (That average is the midpoint of the points labeled w^* and Wages.)
- w^* is also *not* the average reservation wage of the workers who are employed. (That average is the midpoint between the origin and w^*.)
- Instead, w^* is the equal to the reservation wage of the marginal worker and the willingness to pay of the marginal firm. Restated, it is the point where the reservation wage of the marginal (last employed) worker is equal to the marginal revenue product of the marginal (lowest willingness to pay) employer (among those with positive employment).
- The wage in this figure could *potentially* be as high as the y-intercept of the demand curve. It could *potentially* be as low as the y-intercept of the supply curve.
- Where the equilibrium wage falls between these two bounds depends on the position of the supply curve relative to the demand curve. In short, it's set at the margin—where the marginal worker is indifferent between working and not working and the marginal employer is indifferent between hiring and not hiring.
- Thus, the **price** (wage) in this market is the *solution* to a constrained *maximization* problem: maximize the wage subject to every worker who wants to work being able to find a job
- The price (wage) in this market is also the *solution* to a constrained *minimization* problem (the dual): minimize the wage subject to the constraint that every firm that wants to find a worker can find one.
- Finally, notice that *this problem solves itself.* Atomistic workers and atomistic firms that are interacting in a fully competitive labor market will attain the solution w^* (and $l^* = L(w^*)$) without any centralized planner required to coordinate the solution.
- Restated: the competitive, decentralized, self-attaining equilibrium in this (idealized) labor market is *Pareto efficient.*

Now that we have established the brilliance and the beauty of the properties of this simple market equilibrium, we are going to consider some variations. The first variation is tax incidence.

2 Who really pays that tax? The question of incidence

Benjamin Franklin famously said that "In this world, nothing is certain except death and taxes." And with advances in human longevity, perhaps we cannot even count on death.

But taxes will always be with us. Think of all of the taxes you pay (or will pay): income tax, sales tax, property tax, excise tax, taxes on utility services (including landlines, mobile lines, cable), gas tax, liquor tax, gas guzzler tax (yes, this exists), and others I'm surely not thinking of.

But does it strike you as unfair that some taxes are placed on consumers and others on producers? Shouldn't oil companies pay taxes on gas, not consumers? Doesn't taxing workers for their labor income penalize work? Why not tax employers instead?

To answer this question analytically, we need to ask who actually bears the *burden* of a tax. Restated in economic parlance, on whom does the *incidence* of a tax fall, and why?

Let's take a simple example. Write the labor supply curve in a labor market as $L_s(w)$, with $L_s'(\cdot) > 0$. Write the labor demand curve as $L_d(w)$, with $L_d'(\cdot) < 0$. Equilibrium in this market is determined by the 'market-clearing' wage that equates supply and demand

$$w^* : L_s(w^*) = L_d(w^*)$$

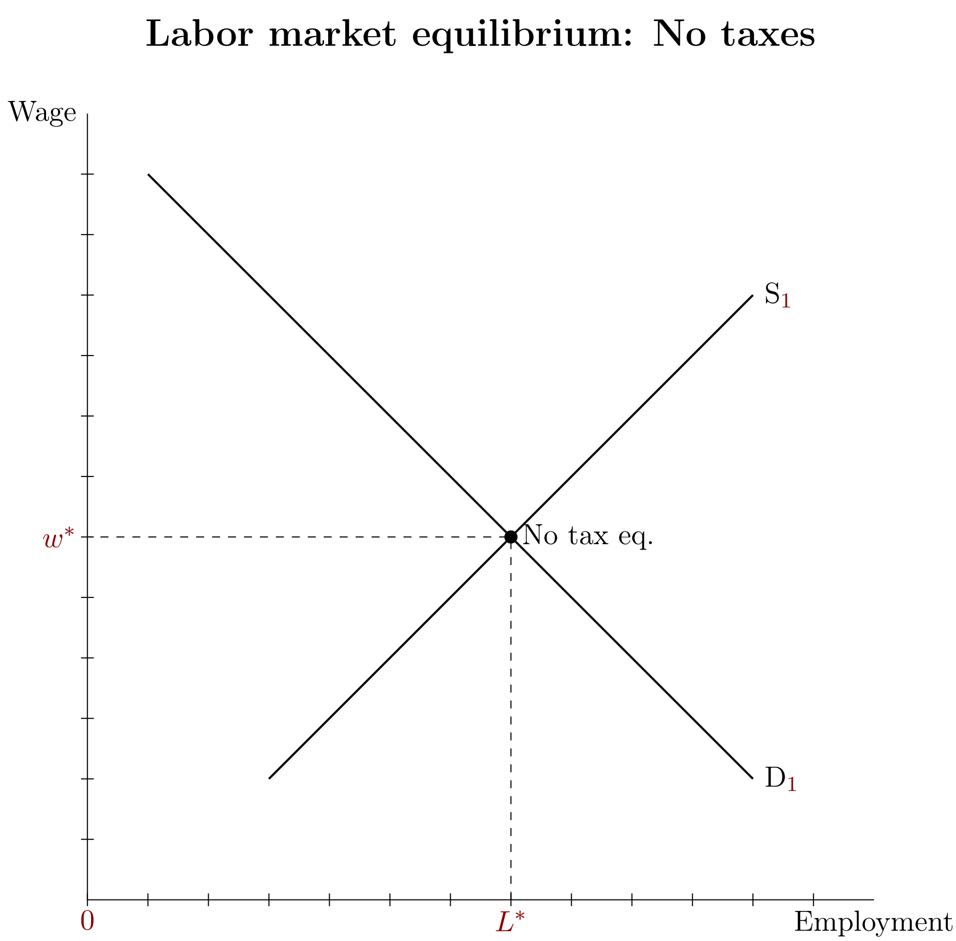

2.1 Effects of taxation on wages/prices

Imagine that we impose an income tax: workers must pay a tax of $\$\tau$ if they work. (It is far more realistic, of course, to assume that workers (or firms) pay a proportional tax, e.g., $\tau\%$ of wages. But this setup makes the figures and algebra messy—unless we transform everything into logs, which most people find confusing. So please go with this simplification.)

We can solve for the new equilibrium with this income tax in place:

$$w_\tau^W : L_s(w_\tau^W - \tau) = L_d(w_\tau^W)$$

The new equilibrium wage is w_τ^W is the one where labor supply and labor demand are equated after accounting for the fact that workers keep only $w_\tau^W - \tau$ although w_τ^W is paid by the firm. (Note that the W superscript on w_τ^W indicates that the tax is paid by the worker.)

There is no reason to expect that w_τ^W is equal to w^* in the competitive equilibrium. From the perspective of the firm, the labor supply curve just shifted upward, so that to get the identical number of workers as in the pre-tax regime, the firm would have to pay a higher wage—specifically, $w^* + \tau$.

Similarly, from the perspective of the worker, to receive the identical wage as in the pre-tax regime, the firm would have to raise the wage to $w^* + \tau$. But of course, this will not happen. Since the firm has a downward sloping labor demand curve, it employs fewer workers at the new higher wage. It will not raise wages by the full amount of the tax. But neither will it reduce the wage by the full amount of the tax. Why? If a firm hires fewer workers, the marginal product of labor at the firm is higher, so it's willing to pay more for the marginal worker.

Thus, in equilibrium: $w_\tau^W > w^* > (w^* - \tau)$ and $L_s(w_\tau^W - \tau) < L_s(w^*)$.

Labor market equilibrium: Workers pay income tax of τ

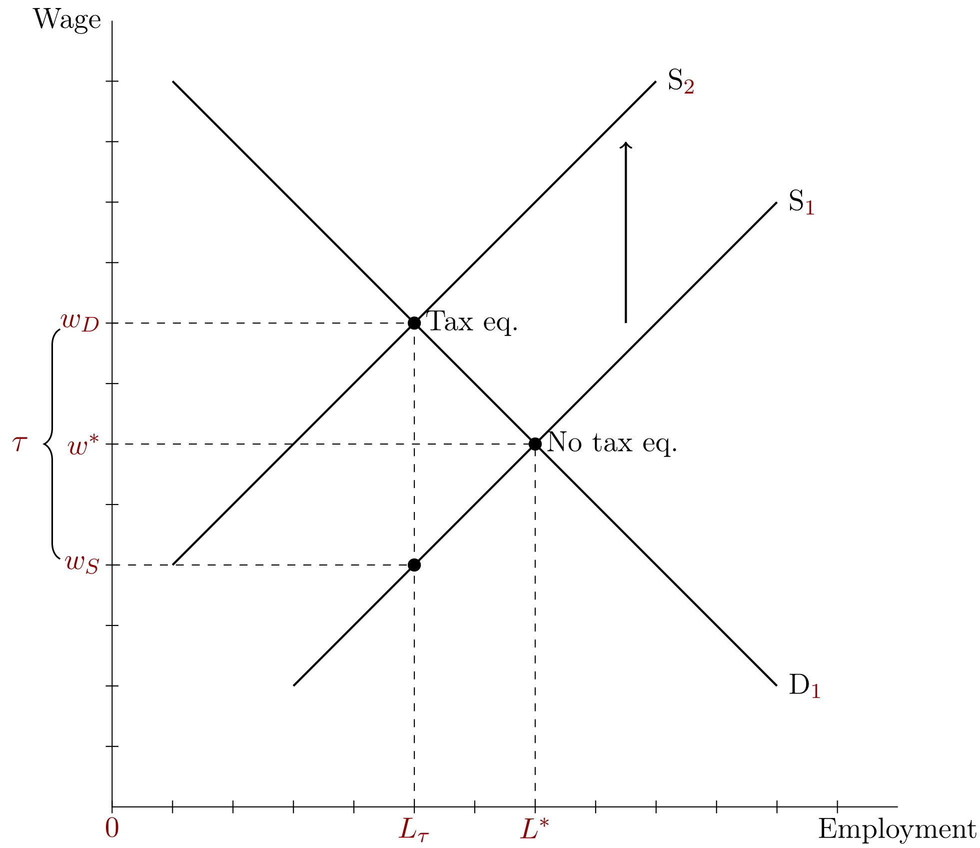

What if, instead of charging the tax to the worker, we charged the tax to the firm? If the firm wants the worker to receive w, it must be spend $w + \tau$, where τ is paid to the government. We can solve for this equilibrium as well:

$$w_\tau^F : L_s(w_\tau^F) = L_d(w_\tau^F + \tau)$$

After a moment's reflection, you can see that it must be the case that $w_\tau^F + \tau = w_\tau^W$, and similarly $w_\tau^F = w_\tau^W - \tau$. That is, whether the tax is placed on workers or firms, the equilibrium wage is the same. The gap between what the firm pays what the work receives is τ, no matter who nominally pays the tax worker or firm. The nominal incidence of the tax (who pays) is uninformative about its economic incidence. Who ultimately pays the tax—the incidence of τ on workers versus firms—depends on the interaction between the supply and demand curves.

Labor market equilibrium: Firms pay payroll tax of τ

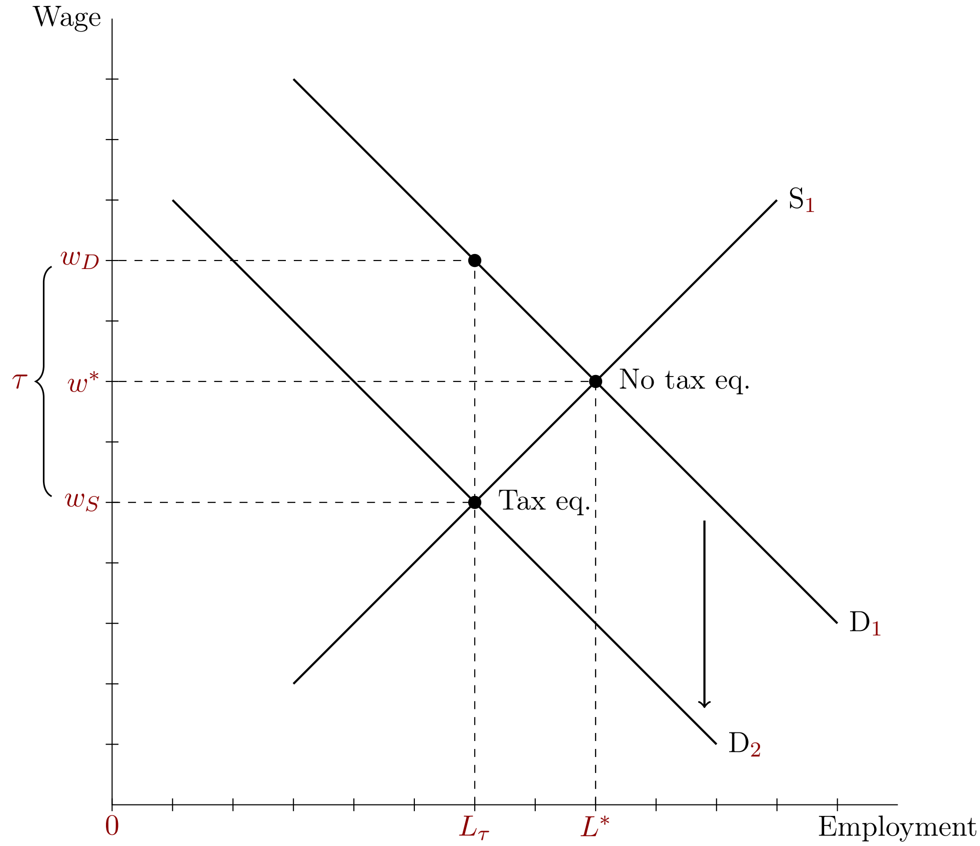

So, if the nominal incidence of the tax *doesn't* determine its economic incidence (i.e., what share is borne by firms and what share is borne workers), what *does* determine it? The answer is logical and intuitive. You will derive it on your problem set.

2.2 Effect of taxation on quantities

Notice that the tax does not merely create a gap between the wage paid by the firm and received by the worker, it also affects employment. Specifically, $L^*_\tau \leq L^*$, with strict inequality (except in edge cases). That is, the tax almost necessarily reduces employment.

Is this reduction in employment undesirable? In general, the answer is *yes*. Unless the tax is being used to correct an externality (e.g., perhaps employment creates un-priced pollution), the reduction in employment constitutes a deadweight loss (DWL). That DWL is visible in the figure above as the triangle from the right of L^*_τ, encompassing the area that was previously above the supply and below the demand curve. That area but is now empty because employment in this region has been curtailed by the tax.

Thus, this figure highlights a fundamental property of taxation: by changing prices, taxes also change behavior. If the goal of income (or payroll) taxation is to raise revenue *not* to reduce employment, then the reduction in employment is an undesirable side-effect of taxation—a distortion. Restated: if the initial market equilibrium was Pareto efficient, the taxed equilibrium is not.

Some taxation is necessary, so the distortionary effect of taxation is not a case against taxes *per se*. But the fact that taxes create DWLs relative to an otherwise Pareto efficient benchmark strongly suggests that we should think carefully about *what* to tax and *how* to tax it. As it turns out, some taxes are more distortionary than others. We will discuss this in class.

3 Taxation, rebates, and behavioral change

In many cases, we *intentionally* imposes taxes to change behavior. For example, we might want people to burn less fossil fuel, drink fewer sugar-sweetened beverages (SSBs), not drive during rush hour, or run their energy-intensive appliances at non-peak hours. We can accomplish these goals by taxing these activities, i.e., imposing gas taxes, SSB taxes, time-varying road tolls, time-of-day energy pricing.

But this invariably creates a conundrum. Lower income individuals and households spend a larger share of their budgets on these items (gas, soft drinks, tolls, utilities), so these taxes are regressive—they take a larger share of income of the poor than the wealthy.[1] A natural solution to this conundrum is to simultaneously impose a tax, and make a transfer to low-income households to hold income constant. Of course, if the policy simply *rebates* the tax to low-income households, then it's not a tax at all. So, the transfer has to be independent of the tax paid.

Let's consider the economics of this policy. Specifically, we want to ask what is the effect on consumer welfare of pairing a tax with a rebate. To make this problem as crisp possible, we will consider a consumer whose tax rebate is exactly equal to her tax payments. This could happen purely by coincidence, though we'll solve for it explicitly in this example.

- Consider a tax τ on each unit of good X that is *fully rebated* to the consumer:

$$\tau \times d_x(P_x + \tau, P_y, I + R) = R. \tag{1}$$

- This tax is revenue neutral for consumer; rebated exactly the amount paid in taxes ($R = \tau d_x(\cdot)$)

[1] Wealthy people spend a much smaller share of income on necessities and a larger share on luxuries. In addition, they save more, both in absolute and relative terms.

- Again, we are assuming that the consumer is not choosing X with the expectation that all tax paid will be refunded. Perhaps the rebate is equal to the average tax paid, and this consumer happens to purchase the average amount of X.

- Notice that this tax alters the *price ratio* faced by the consumer but does not change her budget set

- You can check that the consumer spends the original budget I by writing:

$$(P_x + \tau) \times d_x(P_x + t, P_y, I + R) + P_y \cdot d_y(P_x + \tau, P_y, I + R) = I + R.$$

- Subtracting (1) from both sides, we get

$$P_x \cdot d_x(P_x + t, P_y, I + R) + P_y \cdot d_y(P_x + t, P_y, I + R) = I.$$

  Hence, the consumer is on the original budget set.

- But as long as the consumer changes the consumption bundle in response to the tax-ratio (i.e., as would occur for any utility function satisfying the standard 5 axioms), then the consumption bundle is *altered* by the tax:

$$\begin{aligned} d_x(P_x + \tau, P_y, I + R) &\neq d_x(P_x, P_y, I), \\ d_y(P_x + \tau, P_y, I + R) &\neq d_y(P_x, P_y, I). \end{aligned}$$

- In words, the consumer will be consuming on a different point on the original budget set I under the 'taxed' price ratio.

- This is illustrated in the figure below
  - The original budget set is the ray connecting points I/p_y, I/p_x
  - Point A is the consumer's original chosen point on the non-taxed budget set
  - With the tax in place, the consumer's budget side rotates to become the ray connecting points I/p_y, $I/(p_x + \tau)$. Point A is no longer attainable
  - Given the increase in the price of X, the consumer will very clearly substitute away from X (unless X is Giffen)
  - Our goal is to ensure that the consumer gets his or her tax money back: no matter what new value of X she chooses, we will ensure that she's back on her original budget set

– This is a slightly tricky problem since we must adjust the rebate in response to the choice of X. Yet, as we adjust the rebate, the chosen quantity of X will change (due to the income effect)

– Nevertheless, this target point will exist because there *must* be an indifference curve that is on the original budget set but tangent to the new price ratio. This follows from completeness. (It the consumer stops purchasing X altogether, we are then at a corner solution—but also back on the original budget without any rebate. The choice is still well defined.)

– Let's say that this new chosen point is B: At point B, the consumer is consuming less X than she would without the tax (i.e., at A). But her expenditure is unaffected. *She is on the original budget set.*

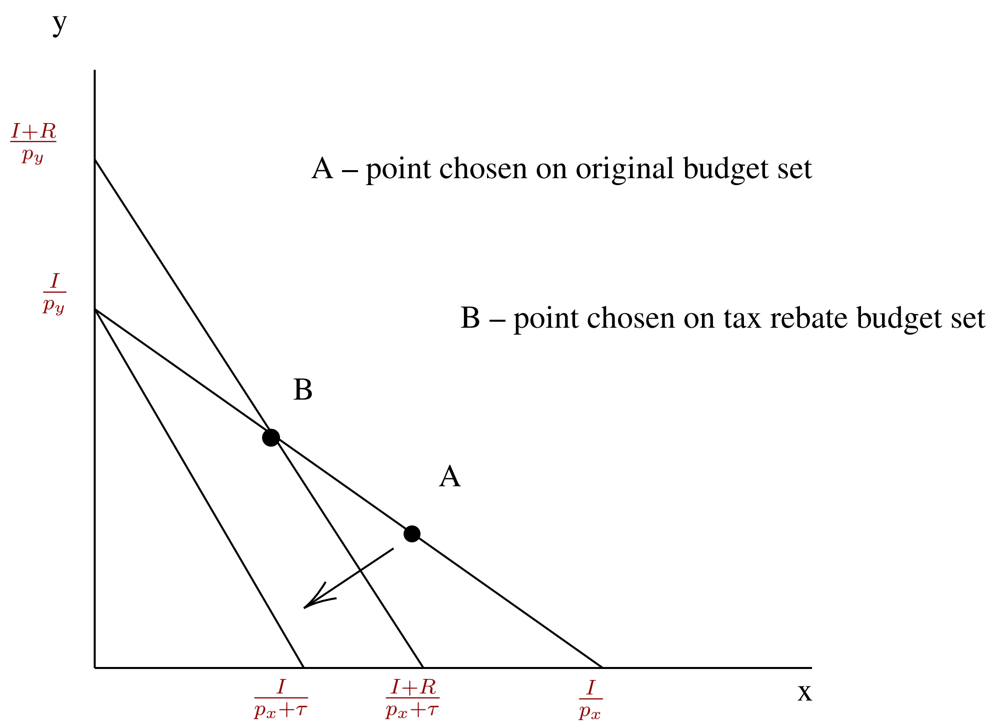

Question for you to contemplate: *Is the consumer worse off, better off, or indifferent as a result of this policy tax + rebate policy, or is the answer indeterminate?*