



Math and Economics with *AI at Your Side*

AI at Your Side · Lecture 3
Sebastian Galiani · Raul A. Sosa


Kyiv School of Economics



Book website

Today: One Model, Taken Apart

Three blocks, and the middle one is most of the lecture



🕒 The lecture's claim: **a model you cannot break is a model you have not understood.**

💡 Every answer from Claude you will see today is a real chat, **condensed to fit**, stamped with the model and the date. Block 2 builds Becker in front of you and then pushes until something gives.

One Rule, Before Any Model

Never ask AI to solve a model you have not tried to understand.

Solve it for me

You get a derivation that reads well and no way to tell whether it is right. Every line looks equally reasonable, which is exactly what fluency does.

Solve it with me

You arrive with an attempt and a specific place where it broke. Now you can tell whether the answer is any good.

 Today we do the second thing twice: once on the statistics behind a regression, and once on a whole economic model, which takes most of the lecture.

Block 1 · The Tutor

Four ways to find out what you do not know

Mode 1: The Diagnosis

Before studying, find the gap



Claude Sonnet 5 · clean session · Sep 23, 2026

condensed · my answers are ours, written for the demo

I want to really understand the statistics behind a regression. Ask me five questions, from the most basic to the sharpest, and tell me where my gap is.



Claude

[...] **1. The mechanics.** [...] **2. Where the randomness comes from.** You have one dataset and one $\hat{\beta}$. What does the standard error of $\hat{\beta}$ describe? [...] **3. Separating the assumptions.** [...] **4. Multiple regression.** [...] **5. The sharp one.** Suppose the true relationship $E[y|x]$ is nonlinear, and you run a linear regression anyway. [...]

[...] 2. It's the spread of the data around the regression line [...]



Claude

[...] you describe the regression as a summary of the dataset in front of you, not as an estimator with a sampling distribution aimed at a population target. Questions 2, 3, and 5 all trace back to that. [...] **2. This is the core gap.** You described the residual standard deviation [...]

Clean session: a new chat, with no tools, no web search and no memory of other chats.

Mode 2: Three Minutes on p-values

What AI is unusually good at: catching a misconception you did not know you had



Claude Sonnet 5 · clean session · Sep 23, 2026

condensed · my definition is ours

Ask me to define a p-value in my own words. Then list the most common wrong readings and tell me which one I made.



Claude [...]

The p-value is the probability that my result is just due to chance. So if $p = 0.03$, there's only a 3% chance that the effect isn't real.



Claude

[...] The p-value is the probability of seeing data at least as extreme as yours, **assuming the null hypothesis is true** (and the model's other assumptions hold). [...]

- | | |
|--|---|
| 1. "The probability the result is due to chance." [...] | 4. "A small p-value means a large or important effect." [...] |
| 2. "The probability the null hypothesis is true." [...] | 5. "A large p-value proves no effect." [...] |
| 3. "1 - p is the probability the effect is real or the alternative is true." [...] | 6. " $p < 0.05$ means the result will replicate." [...] |
| | 7. " $p = 0.03$ means a 3% chance I'm making an error." [...] |

You made **#1 and #2**. [...] (and #3 by implication). **A corrected version of your statement** "If there were truly no effect, I'd see data this extreme or more only about 3% of the time." [...]

Mode 3: Auditing a Sentence

🔍 It finds the mistake, you fix it

A sentence from a draft, and every word that claims more than the regression can show

Claude Sonnet 5 · clean session · Sep 23, 2026

condensed · the sentence is ours

Here is a sentence from my draft: "The regression proves that education causes 8% higher wages, and it is significant, so it matters." Which words claim more than the regression can show? Tell me what is wrong with each one, but do not rewrite the sentence for me.

🗨️ Claude

"proves": Regression doesn't prove anything. [...]

"causes": This is the biggest overreach. [...]

"8%": [...] a point estimate with a confidence interval. [...]

"education": [...] Years of schooling, credentials, and school quality are different treatments. [...]

"higher": [...] The regression compares different people [...], not the same person before and after.

"significant": [...] a significant coefficient can be significantly wrong. [...]

"so": This is a logical leap. [...]

"it matters": [...] That's a judgment your argument has to support [...]



It flagged eight words and did not rewrite the sentence, as asked. You can do this to your own prose in twenty seconds.

Mode 3: Auditing a Derivation

🔍 It finds the mistake, you fix it

An interval 0.0036 wide, with a standard error of 0.031

Claude Sonnet 5 · clean session · Sep 23, 2026

condensed · the first of two runs

I computed this interval by hand. Check it step by step and tell me where I went wrong, without giving me the corrected result.

$$\hat{\beta} = 0.082, \text{SE} = 0.031, n = 1,200$$

$$95\% \text{ CI} = 0.082 \pm 1.96 \times 0.031 / \sqrt{1,200} = [0.0802, 0.0838]$$

 Claude

[...] **Step 3: Standard error.** This is where the mistake is. You divided 0.031 by $\sqrt{1,200}$, but 0.031 is already a standard error. [...] You applied the $\sqrt{n}$ adjustment twice.

[...] **Sanity check:** [...] It implies a t-statistic of roughly 90, but the reported SE gives a t-statistic of only about 2.6. [...]

✅ Both runs found it, and neither gave the corrected interval. **Redone by hand:**
 $0.082 \pm 0.061 = [0.021, 0.143]$.

Mode 4: Let the AI Be the Student

Learning by teaching: explaining is the cheapest way to find your own gaps



Claude Sonnet 5 · clean session · Sep 23, 2026

condensed · my lines are ours

[...] Be a sharp student: stop me at the first thing I get wrong, and ask me a question instead of correcting me. [...] What that means is that there is a 95% probability that the true return to education is between 0.021 and 0.143. [...]



Claude

Hold on [...]. In your setup, what is random and what is fixed? Once you've computed those two endpoints from your sample, is the true return still something that could land inside or outside the interval? If not, what is the 95% describing?

The procedure. What is random is the interval.



Claude

That makes sense to me: the interval is the random thing, and the 95% describes the procedure over repeated samples. [...] Could you restate that sentence so it fits what you just told me?



It asked a question instead of correcting me, twice, so the fix was mine: 0.021 to 0.143 either contains the true return or it does not, and 95% is how often intervals built this way do.

Block 2 · A Model From the Inside

Becker, built from the inside and then pushed until it gives

What Does Sleeping Cost?

Three questions to get into the model



What does
sleeping 7 hours cost?



What does
streaming 2 hours cost?



What does
the **“free” scroll** cost?

💡 **Becker (1965):** nothing is free even when nothing is paid. Labor economics already priced leisure at the wage. Becker put that price inside every activity: the goods it uses up plus what you give up earning while you do it. The paper is part of the program behind his 1992 Nobel.

The Model *Harness*

The harness (Lecture 1's five fields) aimed at an economic model, and the first thing it answers



Claude Sonnet 5 · clean session · Sep 23, 2026

condensed

Context: I'm a second-year economics student. I can solve a consumer's problem with a budget constraint. I suspect "not having time" is an economic problem because nobody can buy more hours, but I don't see how time enters the math.

Task: Derive Becker's 1965 model of the allocation of time with me, step by step.

Constraints: Do not give me the result. One step per turn, and wait for my answer before the next one.

Output: In each turn, one step, the intuition behind it, and one question for me.

Verification: At each step, give me one check I can do myself, such as a special case or the units. Write UNSURE next to anything you are not sure Becker's paper actually says.



Claude

Step 1: Utility comes from "commodities" that households produce, not from market goods [...]

$Z_i = f_i(x_i, T_i)$ [...] UNSURE whether that is exactly his notation on every page. [...]



Saying what you know and where you are stuck makes Claude start from your own frontier.

The Model, Properly

Becker 1965: you do not consume what you buy, you produce it

💡 Nobody enjoys groceries. What you enjoy is *dinner*, and dinner is something you **produce** out of groceries and an hour of your own time.

⚙️ The production function $Z_i = f_i(x_i, t_i)$

Activity Z_i	Market goods x_i	Your time t_i
Dinner at home	groceries	an hour of cooking
A night's sleep	the bed, the room	seven hours

✅ This is the move that puts time inside the price of everything. Once every activity eats hours, hours are an input you can run out of.

From Two Constraints to One

The algebra, one line at a time



Money. Spending cannot exceed income

$$p \cdot x \leq w \cdot T_w + V$$



Time. The day does not stretch

$$T_w + T_c = T$$

T_w hours worked, T_c hours kept, V income that does not come from work.

1. Solve the time one for hours worked

$$T_w = T - T_c$$

2. Substitute it into the money one

$$p \cdot x \leq w(T - T_c) + V$$

3. Distribute and move the negative term

$$p \cdot x + w \cdot T_c \leq w \cdot T + V$$



Money is bought with time, which is why the money constraint alone was never enough. On the right is **full income**, all T hours at w plus V . It is a yardstick, not a paycheck. On the left, every hour you keep, charged at w . **Check it:** $T_c = T$ gives $p \cdot x \leq V$, the no-work constraint.

The Full Price, and What the Condition Says

With the algebra done, the hard part is saying what it means

Write the inputs *per unit* of the activity: b_i of goods and τ_i hours. Then every activity carries a **full price**, the goods plus the hours at your wage:

$$\pi_i = p_i b_i + w \tau_i$$

Since $x_i = b_i Z_i$ and $T_c = \sum_i \tau_i Z_i$, the left side of step 3 on slide 14 is $\sum_i \pi_i Z_i$. Maximize $U(Z_1, \dots, Z_n)$ subject to $\sum_i \pi_i Z_i = wT + V$, and the condition is the one you know from first year:


$$\frac{MU_i}{\pi_i} = \lambda \quad \text{with } \lambda \text{ what one more dollar of full income is worth to you.}$$

 **In words:** at your best choice, the last dollar of full price buys the same satisfaction whatever you spend it on.

On the notation: Becker writes T_i for our t_i and t_i for our τ_i , and per-unit inputs are constant only under constant returns.

At \$20 an Hour, the Expensive Dinner Is the Cheap One

What you pay, against what it costs you



Illustrative figure, with groceries at \$8 and an hour of cooking, delivery at \$22 and five minutes, and a wage of \$20 an hour.

Where it flips

Cooking costs \$8 plus an hour, delivery \$22 plus five minutes. Below about \$15 an hour, cooking is the cheap one again.

Same rule, different value of an hour

It is why lawyers hire cleaners and retirees shop around: older households really do pay less for the same groceries (Aguiar and Hurst, 2007).

Now Raise the Wage

Ask it to draw the change, then read what it says moved



Claude Sonnet 5 · clean session with code · Sep 23, 2026

condensed

In Becker's 1965 model, a home-cooked dinner takes \$8 of groceries and one hour of my time. A delivered dinner takes \$22 and five minutes. I earn \$20 an hour. Plot the full-income line, cooking against delivery. Now raise the wage from \$20 to \$30 and label what moved.

 Claude

[...] **Assumption I added:** you didn't give a full-income figure, so I assumed 60 discretionary hours per period and no non-labor income. That makes full income $F = wT$, or \$1,200 at \$20 and \$1,800 at \$30. [...]

What moved when the wage went from \$20 to \$30


Slope: the line is steeper [...]. **Cooking intercept (x-axis):** it moves from 42.9 to 47.4. [...] **Delivery intercept (y-axis):** it moves from 50.7 to 73.5. [...] **Dotted gray line:** this shows the wage rise with full income held at \$1,200. It is a pure pivot, with the cooking intercept falling to 31.6 and the delivery intercept to 49.0. [...]

The first of two runs. With code: it could write and run Python. Its chart is on the next slide.

What It Drew



The chart from the same answer, with its legend read out on the right



The chat of slide 17, Sep 23, 2026. **Decide**, the tag in the corner, is step five of the method on slide 22: do you believe what came back?

What it drew

Home-cooked dinners across, delivered ones up. Blue, the \$20 line. Red, the \$30 line: full income rises with the wage, so it rotates and moves out. Both runs drew this right, unprompted.

Now check the label

Its legend calls the dotted gray line a “*pure relative-price pivot*.” Both intercepts fall: it is the old full income, \$1,200 (60 hours at \$20), at the new prices. **A picture can be right and its label wrong.**

More or Less?

“Less” is only the substitution effect: the sign depends on a number about your tastes



Claude Sonnet 5 · clean session with code · Sep 23, 2026

condensed · the question is ours

After the raise, do I cook more or less?

 Claude

It depends on numbers the problem doesn't give, so theory can't settle it. [...]

Substitution effect (cook less): [...] A home dinner's full price rises 36%, from \$28 to \$38. A delivered dinner's rises only 3.5%, from \$23.67 to \$24.50.

Income effect (cook more): your full income rises. If every hour of your time is valued at the wage, as in my 60-hour assumption, full income rises 50%. If cooking is a normal good, you want more of it.

Which one wins With the simplest preferences (Cobb-Douglas, constant budget shares), [...] Cooking goes up by about 10% [...] I used a constant-elasticity-of-substitution (CES) version with equal budget shares on the two goods. Under it, cooking falls once the elasticity of substitution is above about 1.7. That is plausible for two ways of getting the same meal. So a dinner-only view leans toward less cooking. [...] I assumed all of your 60 hours are valued at the wage and that you have no other income. If you have non-labor income, or can't actually work extra hours at \$30, the income effect is weaker and the substitution effect dominates. [...] In practice, higher-wage households do tend to cook less. [...]

The same chat as slides 17 and 18.

The Model, Applied to an Ordinary Monday

Pricing your own schedule in wage-hours. But what if you cannot sell another hour at w ?

Claude Sonnet 5 · clean session · Sep 23, 2026

condensed

This is my Monday schedule: 8 h of work, 7 of sleep, 2 of streaming, 1 at the gym, 1 of scrolling, 45 min of commuting and 4 h 15 across meals, hygiene and errands. Price it at Becker prices: which is my most expensive activity?

 Claude

Short answer: at Becker prices, **sleep** is your most expensive activity, at 7 hours of forgone earnings. This assumes your hourly wage is the only price, since you gave no wage or goods costs. [...]

Activity	Hours	Time cost ($\times w$)			
Work	8	Not a cost, since it earns w	Streaming	2	$2w + \text{subscription}$
Sleep	7	$7w$	Gym	1	$1w + \text{membership}$
Meals, hygiene, errands	4.25	$4.25w + \text{food and other goods}$	Scrolling	1	$1w$
			Commuting	0.75	$0.75w + \text{fares or fuel}$

[...] **Work isn't priced as a cost.** In Becker's model it's the activity that generates the wage [...]. **Becker would also ask if sleep hours can be traded for work hours.** If you can't earn w at the margin, for example on a salaried job, the true opportunity cost of sleep is lower than $7w$. [...]

Break It: You Work 40 Hours or You Work Zero



What breaks is the link between the value of your hour and the wage

Claude Sonnet 5 · clean session · Sep 23, 2026

condensed · two turns · the follow-up is ours

[...] Why does time still have a shadow value if I cannot sell the ninth hour?

 Claude

The shadow value of time comes from scarcity in use, not from being able to sell it. [...] If you are **forced to work more** than you would choose, an extra hour off is worth more than w . [...] If you are **forced to work less** than you would choose, an extra hour is worth less than w . [...]

OK. So with my hours fixed, my wage goes up. [...] Does anything flip?


 Claude

The main thing that changes is that the wage stops acting as a price of time. It works only as an income shifter [...]. The wage change is equivalent to a lump-sum increase in money income [...]. So the sign flip that can occur is for time-intensive normal goods. [...]

✅ **For anyone who stays at forty hours, a raise works like a gift of the same amount.** Computed by us (Cobb-Douglas, \$100 a week not from work): two people at \$20 an hour, **with the same job, forty hours and paycheck**, value an hour at \$22.50 and \$10.00. A raise to \$30 lifts weekly income 44% (\$900 to \$1,300) and both values 44%. So does a \$400 gift.

How to Break a Model on Purpose

Five steps, and everyone skips the first one



Why predicting comes first. If you ask before you commit, whatever comes back sounds reasonable and you learn nothing. Write the guess down and you buy the chance to be wrong on the record. **Try it before the next slide:** if every hour past 40 pays 1.5w, do you expect a pile of people at exactly 40 hours?

And step five is a real step. A derivation that checks out can still describe a world you do not live in. Deciding whether you believe it is the point of the exercise.

What If Overtime Pays 1.5w?

You predicted. First it sets up the model

Claude Sonnet 5 · clean session with code · Sep 23, 2026

condensed

Take Becker's model of time, but now the first 40 hours of work in a week pay w and every hour after that pays $1.5w$. People differ only in how much they value goods relative to their own time. What does the overtime premium do to the distribution of weekly hours across people? Derive it, don't just assert it.

 Claude

[...] A person has T hours (I use $T = 168$), works h hours, and has $\ell = T - h$ hours of non-work time. Goods C are bought with earnings at a price of 1.

Preferences are $u = \theta \ln C + \ln \ell$. The parameter $\theta > 0$ is the value of goods relative to own time, and it is the only source of heterogeneity. [...] **Benchmark: no premium** [...] $h_0(\theta) = \theta T / (1 + \theta)$ Hours are strictly increasing in θ [...]

Step 1: the budget set is non-convex. The earnings schedule is convex in h , because the marginal wage jumps up at 40. [...]

 **Read θ as a taste.** It is the only thing that differs across people: the higher it is, the more someone values goods against their own time, and the more they work.

Who Stops at 40?

The same answer, steps 2 to 5: stay under 40 at w , or go past it at $1.5w$, and who picks each

Claude Sonnet 5 · clean session with code · Sep 23, 2026

condensed

 Claude

[...] **Step 2: candidate A, straight-time interior.** [...] It is valid only if $h_0 \leq 40$, which requires $\theta \leq \theta_A = 40/(T - 40) = 0.3125$.

Step 3: candidate B, overtime interior. [...] It is valid only if $h_1 \geq 40$, which requires $\theta \geq \theta_B = 80/(3(T - 40)) \approx 0.208$.

Step 4: the kink is never chosen. [...] For θ between 0.208 and 0.3125, utility falls just left of 40 and rises just right of it, so 40 is a local minimum. Unlike a rising-tax kink, this convex kink produces no bunching at 40.

Step 5: who picks what. [...] numerically gives the switching type $\theta^* \approx 0.256$. At θ^* , the person is indifferent between working **34.3 hours** and **44.9 hours**. [...] **No one works between about 34 and 45 hours.** [...]

 **Why a hole and not a pile.** At 40, cutting an hour gives up only w , and adding one earns $1.5w$. For anyone who values an hour between the two, moving either way helps, so nobody stops there. A tax rate that rises past a threshold does the opposite and piles people up at it: that pile is what “bunching” means.

Five Thousand People, Simulated

The same chat: we asked it to show the hole, with and without the premium

Claude Sonnet 5 · clean session with code · Sep 23, 2026


condensed · the request is ours

Can you show it? [...]



Claude

I simulated 5,000 people [...] Each one picks hours by brute-force search over a 0.05-hour grid. [...]



Left panel, no premium: one smooth, right-skewed hump. [...]

Middle panel, with the premium: the distribution splits in two. [...] **Nobody works between about 34 and 45 hours.** [...]

Right panel, hours by type: the two curves coincide up to $\theta \approx 0.256$. There the premium curve jumps up by about 10 hours [...]

You Predicted the Employer's Answer



The data show a pile at 40, and the model a hole: which question did you ask?

Claude Sonnet 5 · clean session with code condensed

Claude
[...] Real-world bunching at exactly 40 comes from institutional constraints (fixed schedules and employer-set hours), not from this model. [...]

Can you show it? [...]

Claude
[...] **Caveat:** the sharp hole depends on frictionless hours choice and on the log form. With real-world scheduling constraints, firms setting hours, or heterogeneity in wages or fixed costs, the gap would be smoothed out. The direction of the effect would remain.

The chat of slides 23 to 25, Sep 23, 2026: its first answer, and the end of its second.

✓ Our check

Our own solver, with 100 hours to split and \$100 of other income a week, also finds a hole: 34.9 to 44.2 hours.

🏠 And yet, in the data

Forty is by far the most common US workweek (Bick, Blandin and Rogerson, 2022), and Claude had already warned why.

🎯 If you predicted a pile at 40, the data are on your side: there, at least in part, employers pick the hours. In the model, workers do.

Four More Ways to Break It

Same five steps. Each takes one sentence out of the model

A schedule with brackets

Takes out: one wage for every hour. Past a threshold, each extra hour pays *less*, as under a progressive tax.

A fixed cost of working

Takes out: that a job costs only its hours. Commute and childcare are paid the moment you take any job.

Doing two things at once

Takes out: that each hour goes to one activity. An audiobook on the commute counts the same hour twice.

Getting better with practice

Takes out: constant returns. The more you cook, the fewer hours a dinner takes.

© **Part 7 of your problem set: pick one of these four.** Write down what you expect *before* you ask anything. We already know one thing about all four: none of them kills the price of time.

The Ownership Test

Is the model yours, or does it still belong to the textbook?

- ✓ Tonight, blank sheet, screen closed: set up the two constraints.
- ✓ Merge them into full income and write the full price of one activity.
- ✓ Say what a raise does in the model, and which assumption would change that.

🔍 Cannot do one of them? That is where mode 1 comes back: ask for five questions on Becker and see where the gap is, just as at the start of the lecture.

🎯 **Block takeaway:** the derivation you do yourself is the one you keep. Part 7 of the problem set is where you do one.

Your Turn: Coase

The same method, with less help: Problem Set 2

 **This one you take apart yourself.** As in our Substack series, *AI at Your Side: Reading Classic Papers in Economics*, do not ask the machine for the summary.

“The Problem of Social Cost” (1960), his case in our numbers. Cattle stray onto a neighbor’s field and trample the crop. A steer sells for \$300 and ruins \$100 of crop. **How many are worth raising, looking at the whole?**

Steer	1	2	3	4	5
Cost of raising it	\$100	\$150	\$200	\$250	\$300
Plus the \$100 of damage	\$200	\$250	\$300	\$350	\$400

 By a 2012 count (Shapiro and Pearse), it is the most-cited law-review article of all time.

Two Laws, One Herd

Three steers under either law, as long as bargaining costs nothing

✓ **The rule:** raise a steer while the \$300 covers its cost *plus* the \$100 of damage. That is **three**. Alone, the cattle-raiser looks at the cost only, and raises five. The unpaid damage makes the gap, and up to here this is Pigou (1920): the state steps in.

He is liable

He pays the damage, so the fourth steer costs him \$350 and brings in \$300. **He raises three.**

He is not liable

The fourth steer leaves him \$50 and costs the farmer \$100, and the fifth leaves him nothing. The farmer pays him to stop at **three**.

⚠ **The herd is the same and the payer is not**, so long as bargaining is costless. When striking the deal costs more than what is at stake, the legal rule *does* decide. Coase never called it a theorem: Stigler named it, and the costly case is the one Coase cared about.

The Harm Is Reciprocal

The idea under both laws, and what AI does to the cost of a deal

↺↻ Stopping the cattle to protect the crop imposes a cost on the cattle-raiser. The question is which use is worth more, not who is at fault.

↓ What AI cuts

Finding a counterparty, drafting, comparing offers, monitoring.

🔒 What it does not

Holdouts, and the cost of enforcing a right. Those are still institutional.

📋 Your problem set uses different numbers, so do the arithmetic yourself.

What You Take With You Today



The tutor

Find the gap, catch the wrong definition, audit what you wrote, and teach it back so it can question you.



The model

Becker from the production function up: full price, the condition, and what a raise does.



The five steps

Predict, ask, derive, simulate, decide: the method of slide 22, on any model you are learning.



Problem Set 2: Coase by hand, then with AI, and one break of Becker. Parts 1 and 2 **without** AI. Then 3, 6, 7 and one of 4 or 5 with AI, graded on your request and objection (where you pushed back or checked), and in 7 on your prediction and your decision. Due before Lecture 4.

>_ Next lecture runs code on screen. Problem Set 3 needs Python and an agent (Lecture 1: an AI that pursues a goal with tools such as code). The installation guide is in the course materials.

A model you cannot break is a model you have not understood.

Today we changed a model in our heads. **Next lecture we make one run.**

Problem Set 2, due before Lecture 4



Book website