



UNIVERSITY
OF WARSAW



FACULTY OF
ECONOMIC SCIENCES

WORKING PAPERS

No. 41/2026 (535)

ONE SPELL TO MODEL THEM ALL: FROM HOMO OECONOMICUS TO THE GENERAL EQUILIBRIUM UNIVERSE

TOMASZ KOPCZEWSKI
JAN LISICKI
TOMASZ POTOCKI

WARSAW 2026

ISSN 2957-0506



One Spell to Model Them All: From Homo Oeconomicus to the General Equilibrium Universe

Tomasz Kopczewski^{1*}, Jan Lisicki¹, Tomasz Potocki²

¹ *University of Warsaw, Faculty of Economic Sciences*

² *University of Rzeszów, Faculty of Economics and Finance*

* *Corresponding author: tkopczewski@wne.uw.edu.pl*

Abstract: Physicists have spent a century chasing a Theory of Everything and have not caught it. Economists got there first. Starting from one stylized atom, the optimizing individual, they built a full cosmology of a competitive economy with a single tool: one constrained-optimization template, applied to consumers, then to firms, then lifted whole into an Arrow–Debreu equilibrium. This paper reconstructs that move as a spell: one optimization problem to model them all. We trace what mathematical and conceptual engineering each step demands, from the axioms of homo oeconomicus to the existence theorems that hold the edifice up. Rather than prove that the equilibrium exists, the paper computes it, and a companion dashboard puts every lever of the model in the students’ hands. The aim is to show students who already have the mathematics what the monument cost to build, so they can use it, and later drill holes in it.

Keywords: general equilibrium, constrained optimization, homo oeconomicus, teaching microeconomics, methodology of economics

JEL codes: A23, B21, B41, D50, D58

1. Introduction: The Price of a Platonic Cosmology

Physicists have wanted a Theory of Everything for a century and still do not have one. Economists built theirs decades ago and mostly stopped noticing it was there. Starting from a single, radically stylized atom, the optimizing homo oeconomicus, they assembled a complete, internally consistent universe: the Arrow–Debreu General Equilibrium. The skeleton comes from Léon Walras (1874), the first to write the whole economy as one system of simultaneous equations in which every market clears at once; with him, economics began to speak the language of mathematics. Arrow and Debreu (1954) gave that skeleton a rigorous proof.

One mathematical template powers the whole edifice: One Spell to Model Them All, the constrained optimization problem, cast in its primal and dual forms. This article is the foundational overview for a microeconomics course built for an unusual audience: graduate students in Data Science, fluent in mathematics and largely innocent of economic intuition. They do not need lessons in solving an optimization problem; they can do that already. What they need is the why: why we bother, what the solution means, and what conceptual universe we had to engineer before the problem would agree to be solved at all.

Their mathematical fluency is the asset we mean to spend, fast, to build that intuition. We show how a single spell gets applied recursively to conjure an entire economic cosmology. First it is cast on the consumer: a whole system of axioms exists for the sole purpose of guaranteeing a well-behaved utility function, so that the primal problem, utility maximization, yields demand functions, and its dual, expenditure minimization, quietly builds the foundations of welfare measurement. The same spell is then cloned onto the firm: cost minimization for a given output, and profit maximization, generate the supply side. Finally, the spell is lifted to its highest register, the cosmology of general equilibrium. The Arrow–Debreu model assembles every atom, lets them interact through nothing but a vector of prices, and proves that a state of perfect, simultaneous balance exists: a point where nothing pushes any more. We then do what the textbook does not: instead of proving that this point exists, we compute it, and an accompanying dashboard lets students move every lever and watch the whole world find its new point of rest.

This is the Platonic world of microeconomics: mathematically perfect, internally flawless and, exactly as its architects intended, beautiful. A Platonic world need not be the real one, and it does not pretend to be. In Mäki's (2009) terms, a model can serve as a surrogate, a small world we study in order to learn about the large one, or it can quietly become a substitute,

studied for its own sake. Arrow–Debreu was built as the first and remains a live object of research; but part of the profession, enchanted by its beauty, has treated it as the second, and lost sight of the drab, unruly world that keeps knocking at the door.


The beauty is not free. Perfect internal validity was bought with assumptions, paid in a currency called external validity, at a brutal exchange rate. The homo oeconomicus who populates this world is not a human; it is a construct engineered to be solvable, nothing more. Perfect rationality, complete preferences, god-like computation, zero transaction costs, perfect information: the invoice, itemized. Real atoms look nothing like this. The classic defence, Friedman’s (1953), is that unrealistic assumptions do not matter as long as the model predicts well. By that test the verdict came in 2008, when the equilibrium machines built on this world did not see the crisis coming (Colander et al. 2009).

None of which is an argument for tearing the monument down. Its value is precisely its perfection: the Arrow–Debreu model was never meant to describe reality; it is the benchmark reality gets measured against. Building this ideal economy on day one gives students a fixed point. Everything that follows in modern economics (behavioural economics, institutional economics, game theory) can then be read as a disciplined campaign to drill holes in this monument, loosening one rigid assumption at a time, to let the actual, messy world back in.

2. Casting the Spell — The Grammar of the Universe

The standard starting point for any economics course is the circular flow diagram (Figure 1): households buy goods from firms, firms buy labour and capital from households, money runs one way and “stuff” runs the other. It is one of the oldest and sturdiest metaphors in the field, and it was once built out of actual plumbing: Irving Fisher’s (1892) hydraulic machine, a contraption of pipes, cisterns, and levers engineered to show prices and quantities finding their own level, the way water does. Fisher’s machine and the circular-flow diagram are the blueprints of the Platonic world: self-correcting, homeostatic, sealed. But a diagram is only a road map. It says nothing about what makes the cars want to drive.

Microeconomics lays a theory of motivation over this map, treating the flow not as plumbing but as the output of a great many decentralized, independent decisions. Here is where the spell gets cast. Its central trick: household and firm are, mathematically, the same animal. Both are atoms whose entire purpose is to optimize. One optimization problem, worn in two different costumes, drives both sides of the circular flow.

Figure 1. The circular flow of the economy

Note: Own elaboration of the standard textbook diagram (e.g. Mankiw 2021, ch. 2).

Take the household. Cast the spell in primal form, utility maximization, and the atom-consumer solves for the greatest happiness a limited budget can buy. Solve that once and you have the household's whole economic function: its demand for goods and services.

Take the firm. Cast the same spell in dual form, cost minimization, and the atom-producer solves for the cheapest way to hit a fixed level of output, given the technology on hand. Solve that and you have the firm's demand for capital and labour (the very inputs households are supplying), and one further optimization, profit maximization, tells you how much the firm will sell.

Two atoms, two puzzles, solved in isolation. Lifting the spell from atom to cosmos takes one more move: wiring them together. The wiring needs no new theory, only a handful of near-tautological market-clearing equations. Total demand for goods equals total supply of goods. Total demand for labour and capital equals total supply of labour and capital. And the income households spend must equal the income they earned selling their labour and capital in the first place. Link the consumer's output to the firm's output through these bookkeeping identities and General Equilibrium is born. The atom joins the universe, and one mathematical logic now holds the whole thing together.

Look closer and a stranger abstraction appears, something close to conceptual schizophrenia. The circular flow needs its agents to play two parts at once: consumer

and producer. Our homo oeconomicus atom is drafted into both. As a household it maximizes preferences. But who runs the firm on the other side of the flow, minimizing costs? The model's answer, a long-running sore point in methodology (Mirowski 1989; McCloskey 1998), is that the atom was never a person to begin with. It is a decision-making unit, a pure locus of optimization, deliberately built "thin": androgynous, calculating, stripped of history, background, culture, with no identity beyond its preference function. Levitt and List (2008) describe it as a mythical species, a relative of Homo sapiens that is perfectly rational, entirely selfish and able to solve any optimization problem without effort. Because the atom is a function rather than a person embedded in a society, it can run the household in the morning and the firm in the afternoon without contradiction. Robinson Crusoe himself had to be split this way: Crusoe the Consumer trading with Crusoe the Producer to find an equilibrium wage for his own labour. That split is the source of the model's power and its deepest vulnerability. Whether the "atom" is an individual, a family, or the notoriously shaky representative agent (Kirman 1992) remains the live front in economics' methodological wars, and the reason this paper exists at all.

One assumption still hides beneath the equations, more powerful than any we have named out loud: the Market itself. It appears on no diagram, yet it is the invisible, omniscient force running the whole cosmology. Writing those tautological balancing equations quietly assumes a perfectly competitive market, which is less a simplification than a claim about the world. That mechanism is assumed to clear every demand and supply exactly, which carries a controversial rider: the zero-profit condition. In the Platonic world the firm's optimization leaves nothing on the table: no supernormal profit. Capital and labour are paid precisely what they contribute, and not a cent more. Economics built an almost biblical story on top of this: each gets exactly what they contribute. Critics like Mirowski and Nik-Khah (2017) have pointed out where that leads: the Market becomes the sole, unquestionable judge of value and of truth, and any distribution of wealth, however lopsided, can be waved through as "efficient." For most of the twentieth century this perfect market was the profession's unicorn: admired, believed in, rarely interrogated. It took experimental economists, Vernon Smith (1962) chief among them, to show that efficiency, where it appears, is not conjured from nowhere. It depends on how the market's rules are built. The spell only works if someone wrote the institutional spellbook correctly.

This minimalist 2×2 skeleton can still take on passengers. Add a Government, usually in two roles: a redistributive one (taxing here, subsidizing there, rerouting flows that already

exist) and a market-correcting one, supplying public goods (national defence, clean air) that private optimization, on its own terms, cannot produce. Add a Financial Sector, cast as plumbing that perfectly channels household savings into firm investment. Neither actor has any agency of its own in this telling; both are pipes.

That completes the map. The cosmology is a graph: households, firms, government as nodes, each carrying its defining property (optimize, or redistribute), and edges of money and goods governed by the hidden institution of the Market. The last hole worth drilling, and the deepest, is whether the graph itself is stable. We have been assuming the circular flow is eternal and homeostatic. As modern macro-finance keeps discovering, this may be the single most heroic assumption in the whole cosmology. That fight belongs to the macroeconomists, not us, not today. It is their job to ask what happens when the direction of the flows reverses: when firms find it more profitable to park capital in financial markets than to build anything real, or when capital leaks out of the system entirely into something like Bitcoin, wholly detached from production and labour.

Those leaks are exactly what the Platonic model is built to ignore. The critique is not new. Before Keynes, Michał Kalecki, writing in a rather more Marxist key, had already found these holes (Kalecki 1971), arguing that income distribution and firms' financing decisions, not a clean optimization, drive the economy. Post-Keynesian economists have returned to Kalecki with System Dynamics and stock-flow-consistent models built to catch exactly the feedback loops and structural breaks a static circular flow cannot see. Understanding why those leaks are so destabilizing, though, means understanding, with full respect, the beautiful, flawed, indispensable machine they are leaking from.

3. The “Spell” Itself — Taming the Tesseract of Optimization

At the centre of this machine sits the spell of constrained optimization, and in its full mathematical dress it is formidable. The general problem: maximize or minimize an objective $f(x)$ subject to constraints $g(x) = 0$ or $g(x) \leq 0$, over a vector x of n decision variables. Its complete statement lives in the Karush–Kuhn–Tucker (KKT) conditions, a genuinely beautiful theorem: they give necessary conditions for a solution (under a mild regularity condition), sufficient ones under convexity, and handle “corner solutions” (binding inequality constraints) with total elegance.

And yet, building this universe from scratch, we barely invoke KKT at all. For students already fluent in high-dimensional mathematics this looks perverse: why keep the cannon

in the closet? Because for building intuition, the full KKT apparatus is overkill: it characterizes a solution, if one exists, and describes its properties, but gives you no geometric story to hold onto. So we perform the first great engineering act: simplify the problem to fit a simpler tool. Assume the atoms are “well-behaved” (Sections 4 and 5 spell out what that means), constraints stay binding, solutions stay interior (the consumer buys at least a little of everything), and KKT collapses down to its tamer cousin, the Lagrange multiplier method. (The full KKT machinery returns in the last tab of the dashboard, for anyone who wants the cannon.)

That is not, in fact, the biggest simplification we make. Even in its Lagrangian form the spell still applies to an n -dimensional choice space: a consumer picking among n goods, a firm combining n inputs. For a data scientist, optimizing over an n -dimensional vector is trivial: run the algorithm, get x^* , done. But x^* is just a list of numbers. It tells you nothing about how the consumer thinks, how the firm trades off inputs, what the logic of the choice actually is. This is the crisis at the centre of n -dimensionality in economics: computationally solvable, interpretively bankrupt. What does “tangency” even mean in five dimensions? What is the “slope” of an indifference hypersurface? How would anyone ever see the trade-off?

Here the tesseract earns its keep. A tesseract is a four-dimensional cube. We know, mathematically, that it exists, and we, being three-dimensional creatures, will never see it. All we get is its shadow: its projection into a world we can walk around in. The n -dimensional optimization problem is the real object, the tesseract; we know it is there, but no intuition survives contact with it. So we project it down deliberately, onto its shadow, a two-good world, x_1 and x_2 . We do so not because two goods fit neatly on a textbook page, though they do, but because the two-dimensional shadow is the only version of this problem our minds can actually hold.

That reduction is the single most important pedagogical move in the whole course, because it lets us see the unseen. The n -dimensional problem hands us a vector. The two-dimensional problem hands us geometry: indifference curves, budget lines, isoquants, isocosts, and one clean rule the whole apparatus reduces to: equilibrium sits at the point of tangency. That tangency condition, the marginal rate of substitution equal to the price ratio (or its twin for the firm, the marginal rate of technical substitution equal to the factor-price ratio), is the intuitive grammar of microeconomics, born entirely of the 2D projection. We spend the rest of the course inside this shadow, trusting that what we learn there survives the trip back up to the dimensions we cannot see. That trust is itself an act of faith. Read as physics, tangency

is a point of rest. As long as the internal trade-off differs from the market's, a small step in one direction pays; the gap acts like a force, and the force vanishes only where the two slopes meet.

The geometric rule is tangency; the algebraic tool that finds it is the Lagrangian, and it is our workhorse for the rest of the course. Engineer the problem (interior solutions, binding equality constraints), and a constrained optimization turns into an unconstrained one.

To maximize $f(x_1, x_2)$ subject to $g(x_1, x_2) = c$, write the Lagrangian $\mathcal{L}$:

$$\mathcal{L}(x_1, x_2, \lambda) = f(x_1, x_2) - \lambda(g(x_1, x_2) - c)$$

Solve it as an unconstrained problem: take the first-order partial derivatives with respect to x_1 , x_2 , and the new variable λ , and set them all to zero:

$$\frac{\partial \mathcal{L}}{\partial x_1} = 0 \quad ; \quad \frac{\partial \mathcal{L}}{\partial x_2} = 0 \quad ; \quad \frac{\partial \mathcal{L}}{\partial \lambda} = 0$$

Out comes everything worth having: the optimal choices x_1^* and x_2^* , and the optimal multiplier λ^* . For a data scientist x_1^* and x_2^* are the whole answer. For an economist they are half the story: the real payoff is not in the pedestrian demands x_1^* and x_2^* , it is in the mysterious λ^* .

Mathematically, λ^* is the shadow price of the constraint, the derivative of the optimized objective $f(x_1^*, x_2^*)$ with respect to c . In plain language it answers a question that shows up everywhere in economics: how much more of my objective could I get if my constraint relaxed by one unit? It measures exactly how much the constraint hurts. That one number is the key to the dual problem, and to all of welfare economics. Apply it to a consumer and it gets a specific name, the marginal utility of income, but underneath the label, “shadow price” is what it always was.

4. Engineering the Consumer

Our first job is to build the atom that occupies the Households node. It plays two roles. As a supplier, the modeller hands it an exogenous endowment of capital (K) and labour (L). As a consumer, it sells that K and L into the factor market, collects a nominal income M , and spends it on goods x_1 and x_2 in the product market.

Which raises the first reduction of complexity: what are x_1 and x_2 , exactly? Introductory textbooks usually turn childish here (“ice cream” and “chocolate”, or Robinson Crusoe’s coconuts) and get mocked for it, missing the point entirely. The real n -dimensional world,

where a consumer chooses among millions of distinct goods, is computationally tractable and interpretively empty. Projecting it onto a 2D shadow (Section 3) is not a claim about reality; it is an intuition engine. x_1 and x_2 stand in for everything we are willing to pay for.

The atom wants these goods (“want” is what drives it), but “want” is a messy, unobservable, psychological state, and we cannot model it directly. We engineer a usable proxy instead, in three moves that are the archetype of every microeconomic model built since.

Move one: do not model want, model preference. Assert, by construction, that preferences are rational, not in the everyday sense of “sensible” but in the strict sense defined by the axioms of rational choice. Completeness: the atom can always compare A and B. Transitivity: if A beats B and B beats C, A beats C. Add Continuity (no sudden jumps), which is what allows a utility function to exist at all, and, for mathematical convenience, Non-Satiation (more is always better). That is the price of admission: an atom that is internally consistent and perfectly predictable, at the cost of resembling no actual human, who is regularly indecisive and often irrational.

Move two: collect the reward. Because preferences are now well-behaved, we can pull off a genuine trick: representing the entire consistent ranking with a mathematical object, the utility function $U(x_1, x_2)$. It is purely ordinal: a ranking device, not a measuring cup. $U(A) = 20$ and $U(B) = 10$ says nothing about liking A “twice as much”, only that A is preferred to B. The function itself can be drawn as a 3D surface, a mountain of happiness on which higher elevations mark more preferred bundles; its contour lines are the indifference map.

Move three, the most pragmatic and the most disappointing: feed the Lagrangian spell an actual objective function. What specific form of $U(x_1, x_2)$ represents “rational preferences” as we just defined them? One might hope for something psychologically deep. What we get, in a pattern that will recur, is a plain disappointment: we do not discover the function, we pick one, and we almost always pick the simplest algebra that satisfies the axioms: the Cobb–Douglas function. $U = x_1^a x_2^b$ is not a marvel of psychology; it is a workhorse of algebra, beautiful for its simplicity rather than its realism. The weights a and b (usually normalized to sum to one) are not deep psychological parameters but bookkeeping devices that make the math behave: the consumer always spends a constant share $a/(a + b)$ of income on good 1, whatever the prices, whatever the income. The algebra is clean, which is why it is the default choice for every CGE model in captivity.

That same simplicity is its fatal flaw. Cobb–Douglas hard-codes an elasticity of substitution between the two goods of exactly one, a very specific and rather implausible claim about how ice cream relates to chocolate. Try to model near-perfect substitutes (Coke and Pepsi) or near-perfect complements (a left shoe, a right shoe) and Cobb–Douglas simply fails. The fix is the Constant Elasticity of Substitution (CES) function, a student’s nightmare to manipulate algebraically but conceptually gorgeous. One new parameter, ρ , dials the elasticity of substitution to whatever is needed: as $\rho \rightarrow 1$, perfect substitutes; as $\rho \rightarrow -\infty$, perfect complements (Leontief); and, in a small moment of mathematical grace, as $\rho \rightarrow 0$, the whole unwieldy CES function collapses back into Cobb–Douglas. Cobb–Douglas is not a rival model but one special case of a larger truth. Fancier utility functions exist, certainly, but nine times out of ten a working microeconomic model reaches for Cobb–Douglas or CES, and pays for the privilege in simplicity.

Project the 3D function, Cobb–Douglas or CES, down into its 2D shadow, the indifference map. Every line on it connects bundles (x_1, x_2) that deliver identical utility. Its slope is the marginal rate of substitution (MRS): the atom’s purely internal, psychological trade-off, answering how much x_2 it will surrender for one more unit of x_1 while staying exactly as happy. One more engineering assumption keeps the math honest: convexity, or diminishing MRS: the more x_1 you already have, the less x_2 you will give up for another unit. It is the assumption of convex preferences, drawn as a shape. It should not be confused with diminishing marginal utility, which depends on the particular numbers a utility function assigns and so has no place in an ordinal world.

The atom is now fully built: a want (a convex indifference map) and an internal trade-off (the MRS). Confront it with the external world, the budget constraint, $p_1x_1 + p_2x_2 = M$, and its slope, $-p_1/p_2$, becomes the market’s own trade-off: how much x_2 must actually be given up, out in the world, for one more unit of x_1 .

The spell resolves cleanly: the Lagrangian finds the highest indifference curve that still touches the budget line, and that single point of contact is the tangency, the grand and intuitive solution of all microeconomics. The atom has its optimal bundle, (x_1^*, x_2^*) .

How does the algebra deliver this geometry? Take $U = x_1^a x_2^b$ as the objective and $p_1x_1 + p_2x_2 = M$ as the constraint. The Lagrangian:

$$\mathcal{L} = x_1^a x_2^b - \lambda(p_1x_1 + p_2x_2 - M)$$

Turn the crank and take three first-order conditions:

$$\frac{\partial \mathcal{L}}{\partial x_1} = ax_1^{a-1}x_2^b - \lambda p_1 = 0 \quad \Rightarrow \quad ax_1^{a-1}x_2^b = \lambda p_1$$

$$\frac{\partial \mathcal{L}}{\partial x_2} = bx_1^ax_2^{b-1} - \lambda p_2 = 0 \quad \Rightarrow \quad bx_1^ax_2^{b-1} = \lambda p_2$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = M - p_1x_1 - p_2x_2 = 0 \quad \Rightarrow \quad p_1x_1 + p_2x_2 = M$$

The third condition just returns the budget constraint. The magic sits in the first two: divide one by the other, and λ , its algebraic job finished, cancels, leaving pure economics:

$$\frac{ax_1^{a-1}x_2^b}{bx_1^ax_2^{b-1}} = \frac{\lambda p_1}{\lambda p_2} \quad \Rightarrow \quad \frac{ax_2}{bx_1} = \frac{p_1}{p_2}$$

The left side, $\frac{ax_2}{bx_1}$, is the ratio of marginal utilities, the MRS by definition. The right side is the price ratio. The algebra has forced:

$$MRS = \frac{p_1}{p_2}$$

Which is the proof: the algebraic solution is the geometric tangency. The slope of the indifference curve equals the slope of the budget line. Then comes the slightly chilling interpretive payoff: the atom's internal, psychological world (how much it wants to trade) has been forced into exact alignment with the external, objective world (how much the market demands it trade). The atom's mind is calibrated to the market. The market, in effect, is inside its head. The same point can be reached without calculus, one step at a time. Suppose the atom would give up five units of x_2 for one more unit of x_1 ($MRS = 5$), while the market asks for only one ($p_1/p_2 = 1$). Every such swap pays. Because the indifference curve is convex, the gain shrinks with each step, and the swapping stops where $MRS = p_1/p_2$: the point where nothing pushes any more.

We took an unobservable want for chocolate and ice cream, ran it through a factory of axioms, and came out the other end with a precise, calculable demand function, $x^*(p_1, p_2, M)$, the first edge of our graph, ready to be wired into the general equilibrium.

What did all that engineering buy us? The Marshallian demand functions, and for Cobb–Douglas, beautifully (almost suspiciously) simple ones:

$$x_1^*(p_1, p_2, M) = \left(\frac{a}{a+b}\right) \frac{M}{p_1}$$

$$x_2^*(p_1, p_2, M) = \left(\frac{b}{a+b} \right) \frac{M}{p_2}$$

This is the prize, and now the holes want drilling. Notice what is missing: demand for x_1 does not depend on p_2 , demand for x_2 does not depend on p_1 : the cross-price elasticity is zero. In this world the price of ice cream has no bearing whatsoever on demand for chocolate. This is not a discovery; it is a hard-coded consequence of the function we chose.

Income behaves the same way: double M , and consumption of x_1 and x_2 both double. Income elasticity is exactly one, everywhere. No necessity goods (bread, bought proportionally less of as you get richer), no luxury goods (yachts, bought proportionally more of). Every good is a perfectly unit-elastic “normal” good, and the income share spent on good 1, $p_1 x_1 / M$, is a fixed constant, $a / (a + b)$. A billionaire and a pauper spend identical shares of income on food. That is the price tag on the simplicity we bought.

And yet, holes and all, the model gives us exactly what the cosmology needed, a stable and tractable edge for the graph: a demand function $x^*(p_1, p_2, M)$ connecting the nodes, half the circular flow now built. The spell also runs backwards. Its dual asks the opposite question: what is the least money that buys a given level of utility at given prices? The answer, the expenditure function $e(p, U)$, looks like mere bookkeeping here; in Section 8 it becomes the scoring system of welfare.

One link is still missing. The demand function depends on M , but what is M ? Until now it arrived as manna, an exogenous number, no questions asked. The household is not only a consumer; it also owns things. It owns capital (K), which it lends at interest r , and so saves. It also owns time, sold on the labour market at wage w , working L hours. So $M = rK + wL$: the income the household earns selling what it owns. Its income is no longer given; it is set by the market prices of capital and labour, the very things the atom is holding.

One open question remains: where do r and w come from? Answering it means taking off the consumer’s hat and putting on the producer’s. Time to cast the spell a second time, on the other great atom in the universe: the Firm.

5. Engineering the Producer (The Spell in Dual Form)

We have built the Households node and the demand edges flowing out of it, but those edges, $x^*(p_1, p_2, M)$, hang on a dependency we have not paid off. Household income M is earned selling endowments of capital and labour at prices r and w , and those prices come

from the demand side we have not built yet: the Firms node. Time to put homo oeconomicus to work in its second, schizophrenic role: the producer.

The producer is a consumer run in reverse. It does not consume goods for utility; it consumes factors, K and L , to produce output, Y . The symmetry is close enough to cast the identical spell in dual form: instead of maximizing utility for a fixed budget, minimize cost for a fixed output. The objective is no longer $U(x_1, x_2)$ but $C(K, L) = rK + wL$; the constraint is no longer the budget $p \cdot x = M$ but the technology, $\bar{Y} = f(K, L)$.

The engineering runs in perfect parallel. Where the consumer needed the axioms of rational choice to guarantee a well-behaved utility function, the producer needs an equivalent set of properties of technology to guarantee a well-behaved production function $f(K, L)$, and they turn out to be nearly the same axioms wearing new names. Monotonicity: more inputs, at least as much output, the firm's version of non-satiation. Convexity ("diminishing MRTS"): the isoquant, the producer's indifference curve, bows towards the origin, meaning a firm using ever more labour has to give up ever less capital to hold output constant. It is the firm's version of convex preferences and, once again, the assumption that, made strict, buys a unique tangency.

Which functions, by sheer coincidence, satisfy exactly these properties? Cobb–Douglas and CES, cloned wholesale. $Y = AK^\delta L^{1-\delta}$ is the same equation as before, re-dressed: U becomes Y , A is total factor productivity, δ and $1 - \delta$ are output elasticities rather than preference weights. One clean, algebraically obedient function, playing both leading roles in the economy.

One difference, though, changes everything downstream. Utility was ordinal: we never cared what U equalled, only its ranking, so any increasing transformation (U^2 , $\ln U$) left the optimal choice untouched. Production is cardinal. $Y = 20$ is exactly twice $Y = 10$; take the log of the production function and revenue, profit and returns to scale lose their meaning.

Ordinal versus cardinal forces a question the consumer model never had to answer: returns to scale. Double every input, K and L , and what happens to Y ? The textbook story is a U-shaped average cost curve: increasing returns first (specialization), diminishing returns later (bureaucracy). Messy, realistic, and something our Platonic cosmology cannot survive.

The hidden institution from Section 2, the perfect market, needs the zero-profit condition to hold, and increasing returns undermine it: the first firm to grow slightly bigger gets more efficient, undercuts everyone, and the industry drifts towards a monopoly. So we perform

the most brutal engineering act in the paper: we outlaw the U-shape. Production functions are permitted diminishing returns, or, at the outside limit, constant returns to scale ($\delta + (1 - \delta) = 1$ in Cobb–Douglas), and nothing more. Diminishing returns leave a profit, a rent that must be paid to the owner of some fixed factor; only constant returns make zero profit automatic, and the model of Section 7 is built on them. An entire dimension of external validity, the real dynamics of firm size and industrial organization, is sacrificed on the altar of one condition: the zero-profit clause that keeps the cosmology from tearing itself apart.

With that assumption bolted down, the rest of the derivation runs on the same rails as Section 4. Project the production function down into its 2D shadow, the isoquant map, the firm's version of the indifference map, each curve tracing every (K, L) combination that yields the same $\bar{Y}$. Its slope is the marginal rate of technical substitution (MRTS): how much K can be traded for one more unit of L while output stands still. Then confront the firm with the market: the isocost line, $C = rK + wL$, the firm's version of the budget line, its slope $-w/r$ the external, market-set trade-off.

Cast the spell and minimize $C = rK + wL$ subject to $AK^\delta L^{1-\delta} = \bar{Y}$:

$$\mathcal{L} = rK + wL - \lambda(AK^\delta L^{1-\delta} - \bar{Y})$$

Three FOCs:

$$\frac{\partial \mathcal{L}}{\partial K} = r - \lambda \frac{\partial Y}{\partial K} = 0 \quad \Rightarrow \quad r = \lambda \frac{\partial Y}{\partial K}$$

$$\frac{\partial \mathcal{L}}{\partial L} = w - \lambda \frac{\partial Y}{\partial L} = 0 \quad \Rightarrow \quad w = \lambda \frac{\partial Y}{\partial L}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda} = \bar{Y} - AK^\delta L^{1-\delta} = 0 \quad \Rightarrow \quad AK^\delta L^{1-\delta} = \bar{Y}$$

Solving this system directly for K and L is real work. Skip it, the way we skipped it for the consumer: divide the first FOC by the second, and λ , here the marginal cost of production, cancels, leaving the pure economics:

$$\frac{r}{w} = \frac{\lambda(\partial Y / \partial K)}{\lambda(\partial Y / \partial L)} \quad \Rightarrow \quad \frac{r}{w} = \frac{\partial Y / \partial L}{\partial Y / \partial K}$$

The right side is the ratio of marginal products, the MRTS by definition. The left side is the market's price ratio:

$$MRTS = \frac{w}{r}$$

The algebra forces the same tangency as before. The firm settles on its cost-minimizing mix, (K^*, L^*) , exactly where its internal, technological trade-off (the isoquant's slope) matches the market's external one (the isocost's slope). The firm's "mind," like the consumer's, is calibrated to the market. Step by step, the story is the same: with $MRTS = 3$ and $w/r = 1$, replacing capital with labour raises output at unchanged cost, and the firm keeps swapping until the gap closes.

Out comes the final prize, the conditional factor demand functions, the edges that finish the graph:

$$K^* = K(\bar{Y}, w, r)$$

$$L^* = L(\bar{Y}, w, r)$$

These tell us exactly how much K and L a firm demands to hit a given output $\bar{Y}$ at market prices w and r : a black box that turns inputs into output, fully wired into the price system. The circular flow closes on all four edges: households supply K and L to firms; firms pay back $wL^* + rK^*$, which is precisely the income M from Section 4; firms supply output Y (which becomes x_1 and x_2); households demand it back. And with two firms in our $2 \times 2 \times 2$ model (one making x_1 , one making x_2), the total endowments of K and L must split between them:

$$K_1^*(Y_1, w, r) + K_2^*(Y_2, w, r) = K_{\text{total}}$$

$$L_1^*(Y_1, w, r) + L_2^*(Y_2, w, r) = L_{\text{total}}$$

Constant returns add one convenience that Section 7 will need. Every unit of output costs the same, so the cost side of a firm fits into one number, the unit cost $c(w, r) = rK^*(1, w, r) + wL^*(1, w, r)$: the cheapest way to pay for one unit of output at factor prices w and r .

6. Drilling Holes in the Producer

The second atom is built. Now it earns the same scrutiny we gave the first. Homo oeconomicus takes the celebrity beating: behavioural economics has made a whole, headline-grabbing career out of showing how irrational real consumers are, which in our framing is simply drilling holes in the axioms from Section 4. What gets missed is the twin: the Firm atom is engineered just as far from reality, its consequences probably larger, and almost nobody bothers to critique it, because production does not sell magazine covers the way behavioural quirks do. A few holes are worth drilling here too.

First: what is Y , actually? We have defined it as a single, measurable output, but the Y of a modern economy is tons of steel and terabytes of data and financial advice and a streaming subscription, all at once. The aggregation problem is every bit as bad as it was for consumer goods.

Second, and worse: what are the factors of production, really? We have reduced the tesseract of production to K and L alone, a view almost Marxist in its bluntness, crediting the wealth of nations to “capital” and “labour” and nothing else. Where does human capital go, the skill inside a worker’s head? Or social capital, the value of a trusted network? New inputs, K_2 , K_3 ? And here is the real crack: our K and L get used up in production (labour is spent, capital depreciates), while human and social capital accumulate with use. The more knowledge gets used, the more of it there is. A static, reversible optimization has no grammar for a resource that grows the more you spend it.

Which leads to the most damning hole of all: this model cannot represent time, or process, of any kind. A timeless, perfectly reversible function has no room for business’s most basic fact of life: the sunk cost. Take Cobb–Douglas again: $Y = AK^\delta L^{1-\delta}$. Mathematically, if $K > 0$ and $L > 0$, then Y must be positive, with no way around it. That “harmless” algebraic fact is actually an absurd claim: it makes a start-up phase mathematically illegal. No R&D, no construction, no setup: no stretch of time in which a firm spends on capital and labour ($K > 0$, $L > 0$) and produces nothing ($Y = 0$). In the Platonic world, production is instantaneous and automatic. A cost cannot be sunk into a project whose model has no concept of “before” and “after”; the model is perfectly reversible, and running it forwards or backwards makes no difference. Plug in any $K > 0$ and $L > 0$ and Cobb–Douglas hands you positive output before the ink on the invoice is dry: an eighteen-month biotech burn on lab equipment and PhDs, the years a factory spends under construction, the whole grinding gap between spending and shipping: all vanished, mathematically outlawed. In this universe Rome goes up before lunch, and not one bricklayer breaks a sweat, because nothing is ever sunk when nothing ever precedes anything else.

One more hole, and it explains why large, complicated firms exist at all. $Y = f(K, L)$ assumes one homogeneous output, so it has no way to express economies of scope, the plain fact that Ford makes cars and trucks more cheaply together, sharing plants and suppliers, than two separate firms ever could apart. And the constant-returns assumption from Section 5, $\delta + (1 - \delta) = 1$, was never a discovery. It was chosen backwards, from the answer the cosmology needed, because it is the only way to keep the zero-profit condition alive: allow

increasing returns and the efficient outcome may well be one firm producing everything, a natural monopoly, which demolishes the “institution” of perfect competition the whole model was built to protect.

The unholy secret of the producer atom, in the end, is the same as the consumer’s: its properties were never discovered. They were engineered backwards from the cosmological answer the model needed to reach.

7. Assembling the Cosmology — The General Equilibrium

We have built both atoms, the consumer (Section 4) and the producer (Section 5), and drilled the holes worth knowing about (Section 6). The One Spell has been cast twice and yielded every edge of the graph: the consumer’s demand functions x^* , the producer’s conditional factor demands K^* and L^* . All that is left is to wire them together, lift the spell to its final register, and watch the cosmology click into place.

This is the Platonic world in its finished form: the $2 \times 2 \times 2$ General Equilibrium model, built on the pedagogical groundwork of Mixon and Hopkins (2008) and Peng (2009), who showed how to build small-scale CGE models on nothing more exotic than a spreadsheet. To make the model worth playing with, and to set up any policy analysis worth doing, we add one last piece of heterogeneity: the Households node stops being monolithic and splits into two atoms.

- The “Rich” Households: endowed mostly with capital (K_{rich}), a little labour (L_{rich}).
- The “Poor” Households: endowed mostly with labour (L_{poor}), a little capital (K_{poor}).

One assumption, and suddenly income M stops being uniform. The rich earn $M_{\text{rich}} = rK_{\text{rich}} + wL_{\text{rich}}$; the poor earn $M_{\text{poor}} = rK_{\text{poor}} + wL_{\text{poor}}$: different sources, different values. The model can finally talk about distribution, not just efficiency.

With both atoms in place, assemble the machine with the tautological market-clearing conditions from Section 2 (everything must balance), and out come six core equations:

S1 (Labour Market Equilibrium): total labour demand from Firm 1 and Firm 2 equals total labour supply from Rich and Poor households.
 $(L_1^* + L_2^* = L_{\text{rich}} + L_{\text{poor}} = L_{\text{total}})$

S2 (Capital Market Equilibrium): total capital demand from Firm 1 and Firm 2 equals total capital supply from Rich and Poor households.

$$(K_1^* + K_2^* = K_{\text{rich}} + K_{\text{poor}} = K_{\text{total}})$$

S3 (Good x_1 Market Equilibrium): total demand for x_1 from Rich and Poor equals Firm 1's supply. ($x_{1,\text{rich}}^* + x_{1,\text{poor}}^* = Y_1$)

S4 (Good x_2 Market Equilibrium): total demand for x_2 equals Firm 2's supply.

$$(x_{2,\text{rich}}^* + x_{2,\text{poor}}^* = Y_2)$$

S5a and S5b (Zero-Profit / “Biblical” Conditions): in each firm the price equals the unit cost, $p_1 = c_1(w, r)$ and $p_2 = c_2(w, r)$ (Section 5). This is the hidden institution of the Market, leaving no supernormal profit anywhere. Summed over both firms, with the factor markets clear, the two conditions give the biblical identity itself, $p_1 Y_1 + p_2 Y_2 = wL_{\text{total}} + rK_{\text{total}}$: the value of everything produced equals every payment made to every factor.

Six equations, and a closed, homeostatic, hydraulic machine. How do we solve it? Fisher's water found its own level; we have to find, by hand, the prices p_1, p_2, w, r and the outputs Y_1, Y_2 that clear this system. Walras's Law does most of the lifting. Only relative prices matter, so we set one price as numeraire, $w = 1$ (a choice with a classical ring, making labour the measure of all value), and everything else is relative to it. Households spend everything they earn, so if $n - 1$ markets clear, the n th clears automatically. The count therefore closes exactly: five unknowns (p_1, p_2, r, Y_1, Y_2) and six equations, one of which follows from the other five. The atoms choose quantities from preferences and technology; the Market supplies the prices. Prices do all the remaining work; they are, once again, the sole arbiters of value for goods, labour, and capital alike.

Here the course splits from the standard textbook. The textbook goes High Church: it does not ask students to solve the system, only to believe it is solvable, and presents the fixed-point theorems of Arrow and Debreu (Brouwer's, Kakutani's) as scripture: a genuine crowning achievement of twentieth-century mathematics, proving that a price vector clearing every market must exist, under our engineered assumptions.

For Data Science students that approach is a pedagogical disaster: abstract, non-constructive, and an answer to a question they never asked (“does a solution exist?”) instead of the one they always ask (“how do I find it?”).

So we take the engineering path instead. Existence, per Arrow–Debreu, is a given; our job is to compute it, not prove it. The whole system collapses into one optimization problem students already understand: minimize the sum of squared residuals.

$$SSR = S_1^2 + S_2^2 + S_3^2 + S_4^2 + S_{5a}^2 + S_{5b}^2$$

Equilibrium, by definition, is where every $S_i = 0$, which is exactly where SSR hits its minimum of zero. A cosmological problem becomes a standard numerical optimization a computer clears in milliseconds. Each residual is a force: on a market, it is the excess demand that pushes the price up or down; in a firm, it is the gap between price and unit cost, a profit that draws output in or a loss that drives it out. SSR is therefore literally the sum of squared forces, and equilibrium is the place where all of them vanish at once.

This is the Computable General Equilibrium (CGE) model, taught for decades through painstaking, error-prone spreadsheet exercises. For a student fluent in R it is trivial: `optim()` finds the minimum of SSR , and the equilibrium prices fall out. A numerical minimum counts as an equilibrium only when every residual is zero within a stated tolerance, so the dashboard also solves a reduced form of the same system by root-finding, as an independent check. The solver, though, finds the point as a mathematician would; a market has no such shortcut. That the point exists is guaranteed: fixed-point theorems give the Platonic world solid mathematical foundations. The strongest template is Banach's (1922) contraction principle, with exactly one equilibrium reached by repeated correction from any starting point. Its assumptions proved too restrictive for a competitive economy, so Arrow and Debreu (1954) turned to the line of fixed-point theorems running from Brouwer (1912) to Kakutani (1941), which guarantee that the point exists but say nothing about how anyone gets there. Walras's own story of how markets find it, an auctioneer raising prices wherever demand exceeds supply (the *tâtonnement*), need not converge: Scarf (1960) built a three-good economy whose prices, adjusting continuously, circle the equilibrium for ever. The deeper weakness of the spell was shown in the 1970s: optimizing atoms give aggregate excess demand almost no structure of its own, so neither uniqueness nor stability of the equilibrium follows from the grammar of the atom (Sonnenschein 1972; Mantel 1974; Debreu 1974). The profession, as Mirowski and Nik-Khah (2017) tell it, largely looked away and carried on, ignoring the malcontents, as if nothing so terrible could happen to the world it had built. On the role of these theorems, and on how the meaning of equilibrium has drifted since, from a point to cycles, regions and chaos, see Kopczewski and Lisicki (2026).

That is the real payoff. Trade the abstract existence proof for a concrete computational method, and the static Platonic model stops being a museum piece and becomes a laboratory. Wrap the R solver in a Shiny app, hand students sliders for tax rates, technology parameters, household endowments, and let them watch the entire cosmology, every price and every quantity, snap to a new equilibrium in real time. That, finally, is how the intuition gets built. The One Spell dashboard that accompanies this article does exactly that (see the Appendix).

8. Where the Fun Begins — Unleashing Leviathan

The Platonic model is complete, coherent, mathematically gorgeous and, honestly, dull. So far the only fun on offer (*ceteris paribus*, always) is sliding a technology parameter δ or a preference weight a and watching the world snap to a new perfect equilibrium. Instructive. Not policy.

The fun starts when students take on a role the model was never built to anticipate: an active Leviathan. Enter the Government, the third actor from the circular flow.

Set aside, for now, its role supplying public goods; that hole gets drilled elsewhere. What matters here is its most basic, hydraulic function: redistribution. Our Leviathan, for the moment, is a system of pipes drawing money from one part of the machine and pouring it into another.

Give it levers: tax handles on every flow the $2 \times 2 \times 2$ model defines:

- Consumption taxes (product markets): a sales tax on x_1 (t_{x1}) and on x_2 (t_{x2}).
- Income taxes (factor markets): a tax on capital income (t_k), a tax on labour income (t_l).
- Production taxes (Firms): a tax on firm revenue (t_p).

Every tax lands in one bag, T , then gets redistributed to households in whatever proportions the policy chooses. Subsidize the poor:

- $\text{Subsidy}_{\text{poor}} = s_{\text{poor}} \cdot T$, and $\text{Subsidy}_{\text{rich}} = (1 - s_{\text{poor}}) \cdot T$.

Those levers blow apart the tidy six-equation system. S1 and S2 survive unchanged (factor demand still depends on w and r), but income for both groups now reads:

$$M_{\text{poor}} = r(1 - t_k)K_{\text{poor}} + w(1 - t_l)L_{\text{poor}} + s_{\text{poor}} \cdot T$$

$$M_{\text{rich}} = r(1 - t_k)K_{\text{rich}} + w(1 - t_l)L_{\text{rich}} + (1 - s_{\text{poor}}) \cdot T$$

Consumers no longer see the market price p_1 ; they pay $p_{1,\text{consumer}} = p_1(1 + t_{x1})$, which rewrites S3 and S4: the demand functions from Section 4 now run on the new M and the new

consumer prices. Firms no longer keep p_1 either; they receive $p_{1,\text{producer}} = p_1(1 - t_p)$, which rewrites the zero-profit conditions S5a and S5b: $p_j(1 - t_p) = c_j(w, r)$.

And the new variable T demands one more equation, S6, the government budget constraint: Leviathan cannot print money, only spend what it collects:

S6: $t_{x1}p_1Y_1 + t_{x2}p_2Y_2 + t_k r K_{\text{total}} + t_l w L_{\text{total}} + t_p(p_1Y_1 + p_2Y_2) - T = 0$: everything the five taxes collect is paid out as transfers.

Seven equations (S1–S4, S5a, S5b, S6), six unknowns ($p_1, p_2, r; Y_1, Y_2$; the redistribution budget T), one equation again redundant by Walras's Law, and a new objective:

$$\text{SSR} = S_1^2 + S_2^2 + S_3^2 + S_4^2 + S_{5a}^2 + S_{5b}^2 + S_6^2$$

A complete digital hydraulic machine, levers and valves fully specified. Launch the Shiny app, grab the tax sliders, run the comparative statics. Move one lever, hold the rest still, compare the two equilibria: *ceteris paribus* turns the model into a statement of cause and effect, $x \rightarrow y$. The real world, where other things are never equal, rarely grants such statements (Hausman 1992). But “what happens” needs a definition.

This is the last, most consequential trick in the cosmology: economists want to price everything. The simulation shows real changes readily enough (do the poor buy more x_1 , does r fall?), but that does not answer the actual policy question: is the new state better? Answering that means measuring well-being, and the only measure on hand, the utility function U , is purely ordinal. Move the poor from $U_{\text{poor}} = 10$ to 12 and the rich from $U_{\text{rich}} = 100$ to 110, and there is no honest way to say who gained more: the poor's “2 utils” and the rich's “10 utils” are not the same currency, and might not be a currency at all.

So economists invented a trick for monetizing an unmeasurable feeling: the equivalent variation (EV). Suppose a policy moves a household from an original state (U_0 , prices P_0) to a new one (U_1 , prices P_1). EV asks one counterfactual question: how much money, handed over (or taken away) back in the old world at the old consumer prices P_0 , would have made this household exactly as happy as the new policy does, at U_1 ?

That sum of money is the EV, the monetary translation of a utility change. It is the payoff of the dual problem from Section 4 (the expenditure function $e(p, U)$), which finally cashes out as a scoring system:

$$EV = e(P_0, U_1) - e(P_0, U_0)$$

Because EV is hard, cardinal, and denominated in money, results become comparable. A hypothetical policy that produces $EV_{\text{poor}} = +150$ and $EV_{\text{rich}} = -400$ lets us say something precise: this policy helped the poor by the equivalent of a 150-unit cash transfer, and cost the rich the equivalent of 400. Money makes the two numbers comparable in size, not in moral weight: whether a loss of 400 to the rich counts for more than a gain of 150 to the poor is a judgement that no utility function can make. That is the final, comparable score our Shiny app produces, and the moment the Platonic model stops merely describing the world and starts evaluating policy in it.

9. The Two Towers — The Gospel of the Platonic World

If students keep exactly one thing from this whole overview, it will probably be the One Spell, the same optimization logic driving consumer and producer alike. That is a mechanical insight, though, not yet a philosophy.

The real gospel of the Platonic world, the ideological story that has shaped homo oeconomicus for a century, does not live in the optimization problem itself. It lives in the Two Fundamental Theorems of Welfare Economics.

These are the Two Towers of the cosmology. On one level they are simply the abstract mathematical conditions under which the spell successfully builds a stable universe. On another, deeper level, they are narratives: the story economists tell themselves about how the world should work, the source of the profession's deepest convictions, and, as we argue here, a beautiful, seductive story that a great many economists mistake for reality.

9.1. *The First Theorem (The Tower of Efficiency)*

The First Theorem, the White Tower, is the formal proof of the Invisible Hand: any general (Walrasian) equilibrium produced by the spell, every atom optimizing, every market clearing, is Pareto efficient. It is the founding laissez-faire story: let the hidden institution of the perfect market run, and the outcome is efficient, full stop. It proves, mathematically, that a decentralized system of purely selfish atoms produces a collective outcome where nothing is wasted; nobody can be made better off without making someone else worse off. It says nothing about fairness, only about efficiency. In our machine the claim can be watched rather than believed: at the equilibrium of Section 7 the rich and the poor read their MRS off the same price ratio, and both firms read their MRTS off the same factor-price ratio,

so no swap is left between anyone that would make one better off without hurting the other. Every force is zero, and efficiency is simply the name for that silence.

This theorem sets the economist's default setting: the market is efficient, and the burden of proof shifts onto anyone who wants to touch it. Government intervention (Leviathan, Section 8) is, by default, a distortion creating inefficiency; its only legitimate job is repairing a proven market failure: externalities, public goods, monopoly, the very holes we have been drilling, which stopped the theorem from holding in the first place. The Market is the baseline of perfection. Government is a costly patch. The machine of Section 8 shows why: a sales tax drives a wedge between what the consumer pays, $p_i(1 + t_{x1})$, and what the firm receives, so the two trade-offs no longer meet, a gap stays open, and the force that would close it is blocked by the tax.

9.2. The Second Theorem (The Tower of Distribution)

The Second Theorem is the secret, more radical twin, the Dark Tower. It states that any Pareto-efficient allocation (any point on the contract curve) can be reached as a competitive equilibrium, provided preferences and technologies are convex and the planner can perform one costless, lump-sum redistribution of the initial endowments. This is the socialist theorem hiding inside a capitalist model, and its promise is breathtaking: efficiency and fairness, cleanly separable, both attainable. Which allocation is fair, the theorem does not say; that has to come from outside it.

Look at the First Theorem's efficient outcome and decide the Rich hold everything and the Poor hold nothing, and the Second Theorem answers: do not blame the Market. The mechanism is perfect; it always delivers efficiency. The problem was never the process, only the starting point: the Rich simply started with all the capital. The correct fix is not to distort the market with ongoing taxes and subsidies, as Leviathan did in Section 8, since that creates inefficiency. The correct fix is a single act, at the very beginning, like a god resetting the board: redistribute the endowments once, lump-sum, from Rich to Poor, then step back and let the market run. In the machine this is a single lever: move a share of the rich household's capital to the poor before the market opens, and solve again (the dashboard's Two towers page does exactly that). The First Theorem then guarantees the new, fairer starting point lands on an equally efficient outcome.

Together, the Two Towers preach a beautiful gospel: the Market is the perfect engine of efficiency, the Government can be the perfect engine of fairness, so long as it confines itself

to resetting initial conditions once and never touches the mechanism again. It is the Platonic fantasy where *laissez-faire*'s efficiency and social democracy's fairness were never in tension at all: just two separable jobs, tidily divided.

The trouble, of course, is that economists keep mistaking the story for reality. The conditions these theorems need, the spells cast across Sections 2 to 6 (no externalities, perfect information, convex preferences, complete markets and, most heroically, a government able to redistribute property costlessly without anyone changing their behaviour in response), are exactly the assumptions defining the Platonic world, and exactly what is missing from our own. Worse, the two theorems were always a package deal: the First Theorem's defence of the market only holds together if the Second Theorem's promise of redistribution actually gets delivered. In practice the package gets torn in half. Free-market advocates cite the First Theorem as an iron law and wave off the Second Theorem's radical redistribution as impractical. Interventionists lean on the Second Theorem's promise and skip past the god-like conditions needed to deliver it without breaking the First. What is left is not a gospel but a broken covenant: the top layer of the whole cosmology, and the source of an argument that never ends, because each side is quoting only the half of the scripture it likes.

10. Conclusion: The Map of the Platonic World (and How to Use It)

We started with a provocation: economists have a cosmology that physicists still lack. We have shown the whole universe runs on a single spell: constrained optimization, cast recursively. That one grammar built the consumer atom (Section 4), cloned itself into the producer atom (Section 5), and had its unglamorous flaws exposed in both (Sections 4 and 6). We assembled the parts into a complete, self-consistent Platonic universe of general equilibrium (Section 7), let Leviathan loose on it (Section 8), and gave it a moral philosophy (Section 9), turning a static monument into an interactive policy simulator, and revealing its real purpose all along: a CGE game. The spell builds this world; it does not guarantee that the world holds together. Nothing in the grammar of the atom promises that the equilibrium is unique, or that anyone ever reaches it (Section 7). We believe it anyway, as Plato's readers believed in perfect Forms dwelling in the place beyond the heavens (Phaedrus 247c), where Justice, Beauty and true Love exist without a flaw while their copies down here limp along. A world without gaps is easier to believe in than to find.

That is the edge of the Platonic world. It has perfect internal validity (every piece follows logically from the last, beautiful exactly as its mathematicians intended), and the price of that

beauty is a foundation of deep and often unrealistic simplification. Consider this article a map of that world, not an encyclopedia of it: we built the whole cosmology fast, to expose its architecture and its hidden costs, and we have pointed at where the holes are without climbing into all of them. What follows is not so much a bibliography as a treasure map for the student who wants to know where to dig next.

Our toy CGE model has one fatal weakness, and it is the elephant in the room. It teleports from equilibrium A to equilibrium B as if time simply did not happen: Tuesday's economy dematerializes at midnight and Wednesday's rematerializes fully formed a moment later, nobody having filed so much as a change-of-address form. Real economies do not teleport. They transition, with friction, through a fog of risk. And nothing guarantees that the transition ends at a point at all (Section 7).

For Kenneth Arrow and Gérard Debreu, two giants of twentieth-century economics, this was originally an intellectual exercise, a feat of mathematics proving a decentralized economy could, in theory, hold itself together. The spell was seductive enough that generations of economists came to believe in it not as a fable but as a blueprint, and turned into constructors, expanding the Platonic ideal into vast, resource-hungry machines. They added time, risk, and friction, and paid for the realism in the coin the model cared about most: its own beauty and transparency. Somewhere along the way, the surrogate had become a substitute (Mäki 2009).

Add Time, and the teleporting economy becomes dynamic: agents optimize over a lifetime, choosing between today's consumption and tomorrow's, turning savings, investment, and capital accumulation into endogenous processes rather than fixed endowments. Add Risk, and the future stops being known: agents optimize expected utility over a distribution of future states, and the model becomes a tool for pricing assets and insurance. Add Friction, and markets stop clearing instantly (prices go sticky, wages go rigid, moving capital costs something), which finally lets the model see business cycles and unemployment, phenomena the static version is blind to by construction.

Dynamic CGE and Intertemporal General Equilibrium belong to industrial engineering more than to a classroom. Building one is like stacking Lego bricks into something enormous: another market, another constraint, another sector, another "realistic" friction: sticky prices, adjustment costs, even, these days, a cryptocurrency market. Each brick makes local sense. The whole construction becomes a living monster, harder to solve, harder to debug, harder to explain to anyone outside the team. Central banks and ministries run entire departments

on maintaining and calibrating such machines, dynamic CGE and DSGE models, in specialized toolchains like GAMS, GEMPACK, or Dynare. Next to those, our R scripts are toys, and happily so. But every added dose of realism chips away at the model's original elegance. Time and risk as such were in the Arrow–Debreu world from the start, as goods dated and contingent on the state of the world (Debreu 1959); what breaks the First Welfare Theorem, the first of the Two Towers, is the realism that comes with them: incomplete markets, sticky prices, frictions. The constructors keep building because there are barbarians at the gate.

Not everyone signs on to the “everything clears, everywhere, always” worldview. For plenty of questions, non-equilibrium approaches do better work. System Dynamics models the economy as stocks and flows wired together by feedback loops, and its deterministic chaos routinely reveals an instability equilibrium models are built to smooth away. Agent-Based Models throw out the top-down solver entirely: markets grow bottom-up from simple rules governing thousands of heterogeneous agents, letting bubbles, crashes, and genuinely emergent behaviour happen instead of being assumed away. Lately these approaches have started anchoring themselves in real data through digital twins and data assimilation, growing well past their “Game of Life” origins. Call that a teaser for what is waiting past the equilibrium wall.

The coming years look like a genuine contest between the Constructors, building ever more elaborate equilibrium machines, and the Barbarians, simulating emergent chaos from the ground up. Which side is playing Orc and which is playing Elf depends entirely on where you are standing, and that argument belongs to another paper.

The point of this one is different. The goal was never to hand students this model as truth. It was to build, fast, a complete if leaky benchmark for people who already have the mathematics and do not yet have the intuition, constructing the whole edifice at once precisely so it can be deconstructed afterward, consciously and in order.

This article is not a stand-alone object. It is the philosophical and methodological guidebook to the interactive tools built for this course, available in two forms: an interactive R Shiny applet (our toy, the CGE model as a game; see the Appendix) that leads students by the hand through the tax sliders and shows, in real time, how the whole hydraulic machine of prices, production, and welfare reacts to their decisions; and the complete R source code for students who want to open the hood themselves. The applet shows what happens. This text explains why it happens, and what it cost to make it happen.

As the course moves into time and risk, we keep adding elements to the class model to see how its behaviour shifts (playing constructors ourselves, in other words), and those extensions will ship with their own explanations, alongside other approaches that keep the whole enterprise honestly pluralistic.

One last, practical note: none of this overview is meant to be memorized. It is a map, something to come back to across the semester. Every hole and simplification signalled here becomes, in a few weeks, the subject of its own detailed class, drilling into the rationality axiom (behavioural economics) or the production function (transaction costs, new institutional economics), and students will be able to return to this map each time and see exactly which pillar of the Platonic world is under question that week. Understanding that structure of construction, deconstruction and conscious repair is the actual goal of a modern microeconomics education, and it is the whole point of this course's "Know Thyself" method: economic models are, first and always, stories. Knowing them properly, understanding precisely what they are built on, and recognizing the narratives they tell about the world has to come before anyone is allowed to critique or dismiss them — because those narratives, in the end, are also about us.

References

- Arrow, K.J., Debreu, G., 1954. Existence of an equilibrium for a competitive economy. *Econometrica*, 22 (3), 265–290.
- Banach, S., 1922. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3, 133–181.
- Brouwer, L.E.J., 1912. Über Abbildung von Mannigfaltigkeiten. *Mathematische Annalen*, 71(1), 97–115.
- Colander, D., Goldberg, M., Haas, A., Juselius, K., Kirman, A., Lux, T., Sloth, B., 2009. The financial crisis and the systemic failure of the economics profession. *Critical Review*, 21(2–3), 249–267.
- Debreu, G., 1959. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. Wiley, New York.
- Debreu, G., 1974. Excess demand functions. *Journal of Mathematical Economics*, 1(1), 15–21.
- Fisher, I., 1892. Mathematical investigations in the theory of value and prices. *Transactions of the Connecticut Academy of Arts and Sciences*, 9, 1–124.
- Friedman, M., 1953. The methodology of positive economics, in: Friedman, M., *Essays in Positive Economics*. University of Chicago Press, Chicago, pp. 3–43.
- Hausman, D.M., 1992. *The Inexact and Separate Science of Economics*. Cambridge University Press, Cambridge.
- Kakutani, S., 1941. A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal*, 8(3), 457–459.
- Kalecki, M., 1971. *Selected Essays on the Dynamics of the Capitalist Economy 1933–1970*. Cambridge University Press, Cambridge.
- Kirman, A.P., 1992. Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2), 117–136.
- Kopczewski, T., Lisicki, J., 2026. From the point to chaos: The evolution of the concept of equilibrium in economic theory. *WNE Working Papers 29/2026 (523)*. Faculty of Economic Sciences, University of Warsaw, Warsaw. <https://doi.org/10.33138/2957-0506.2026.29.523>
- Levitt, S.D., List, J.A., 2008. Homo economicus evolves. *Science*, 319(5865), 909–910.
- Mäki, U., 2009. MISSING the world: Models as isolations and credible surrogate systems. *Erkenntnis*, 70(1), 29–43.

- Mankiw, N.G., 2021. Principles of Economics, 9th ed. Cengage Learning, Boston, MA.
- Mantel, R.R., 1974. On the characterization of aggregate excess demand. *Journal of Economic Theory*, 7(3), 348–353.
- McCloskey, D.N., 1998. The Rhetoric of Economics, 2nd ed. University of Wisconsin Press, Madison.
- Mirowski, P., 1989. More Heat than Light: Economics as Social Physics, Physics as Nature's Economics. Cambridge University Press, Cambridge.
- Mirowski, P., Nik-Khah, E., 2017. The Knowledge We Have Lost in Information: The History of Information in Modern Economics. Oxford University Press, New York.
- Mixon, J.W., Hopkins, B.N., 2008. General equilibrium analysis using Microsoft Excel. *Journal of Economic Education*, 39(4), 403.
- Peng, A., 2009. Introducing CGE models to the classroom using Excel. Working Paper 013, Department of Economics, Ryerson University.
- Scarf, H., 1960. Some examples of global instability of the competitive equilibrium. *International Economic Review*, 1(3), 157–172.
- Smith, V.L., 1962. An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2), 111–137.
- Sonnenschein, H., 1972. Market excess demand functions. *Econometrica*, 40(3), 549–563.
- Walras, L., 1874. *Éléments d'économie politique pure*. L. Corbaz, Lausanne.

Appendix. The One Spell Dashboard

The article has an interactive companion, the One Spell dashboard: https://microeconomics.shinyapps.io/One_Spell_lab/. It is an R Shiny application built on the same model as the text. It holds no data. Every number on every page is computed from the formulas of the article at the moment the reader presses Run, and the economy of the general-equilibrium pages has an illustrative calibration, listed on the Equations page, chosen for teaching rather than estimated. The R code is part of the Shiny application itself; it is also available from the authors on request and at <https://know-thyself.blog/>.

Each page asks one question and asks the reader to predict the answer before running anything. The tabs then follow the same order on every page: the mechanism step by step; the point (comparative statics as on paper: one change, two equilibria); the region (the result as a function of a parameter); the derivation; and a closing tab on what the assumption of the page makes possible and where it stops working. Every page shows the gap that acts as a force (MRS against the price ratio, MRTS against the factor-price ratio, excess demand on each market, price against unit cost in each firm) and the place where it closes. The pages and the sections they serve are listed in Table A1.

Table A1. Pages of the One Spell dashboard

Page	Question	Section
Intro	Why does anything stop?	1, 2
1. Budget	What does the market allow?	4
2. Preferences	How do we describe a want?	4
3. Optimum	Where does the consumer stop?	4
4. Technology	What may a firm be?	5, 6
5. Cost	What is the cheapest way?	5
6. Where does M come from?	Who earns what?	4, 7
7. The first machine	Can all markets clear at once?	7
8. Does the market find it?	Does a market that corrects its own prices reach the equilibrium?	7
9. Leviathan	What happens when Leviathan pulls the levers?	8
10. Who gains?	Is the new state better, and for whom?	8
11. Two towers	Can efficiency and fairness be separated?	9

Page	Question	Section
12. The whole machine	What happens when several levers move at once?	7–9
13. Beyond the model	What may travel to a real economy?	10
Equations, Questions	The formulas of the model; questions without answers	All

The dashboard can be played in two rounds. In the first, *ceteris paribus*, the reader changes one thing on a single page, compares the two equilibria and describes how that one change moved the whole system: a statement of cause and effect, $x \rightarrow y$. In the second, complexity, the reader moves several levers at once on page 12. The joint result no longer says what caused what, but it still forecasts the state of the economy, which is how the large models of central banks built on this logic are used. Page 8 goes beyond the article: it simulates Walras's auctioneer (the *tâtonnement*) on one market, in the $2 \times 2 \times 2$ economy and in Scarf's (1960) three-good economy, where the prices circle the equilibrium for ever.



UNIVERSITY OF WARSAW
FACULTY OF ECONOMIC SCIENCES
44/50 DŁUGA ST.
00-241 WARSAW
WWW.WNE.UW.EDU.PL
ISSN 2957-0506