

Economics, Evolving

Forty influential works and four changes in what economics could see

An extensive literature paper and visualization program, 1776–present. Four epochs trace economics as its object expands from coordination to control, connection, and machine-mediated cognition.

MICHAEL SCHYMURA · SCHYM.DE · AUGUST 2026 · VERSION 1.0

schym.de/economics-evolving

Economics rarely discards an idea. More often, it discovers the boundary at which the idea stops being sufficient.

Epoch	Years	Changing object
I · Coordination	1776–1936	Workshop → aggregate economy
II · Control	1936–1980	National economy → policy rule
III · Connection	1980–2015	Idea → governed network
IV · Cognition	2015–present	Capability → consequence

READING MAP

Contents

The paper can be read as four histories, one cumulative argument, or a design specification for the accompanying scrollytelling experience.

[Abstract](#)

[The argument in one page](#)

[Why these four epochs](#)

[Influence without a league table](#)

[How to read the evidence](#)

[A tangible thread through the paper](#)

[Paper structure](#)

[Epoch I — The Market Is Invented—Then It Breaks \(1776–1936\)](#)

[Epoch II — The Economy Learns the Rule, 1936–1980](#)

[Epoch III — The Digital Turn \(1980–2015\)](#)

[Epoch IV — The Intelligence Age \(2015–present\)](#)

[Cross-epoch synthesis](#)

[Visualization program](#)

[Selection appendix — influence signals and near misses](#)

[What the sequence changes](#)

[Primary bibliography](#)

[Bibliographic and publication notes](#)

Abstract

This paper reads forty influential works in economics as a sequence of changing analytical objects. Epoch I, from 1776 to 1936, asks how decentralized exchange can coordinate a commercial society. Epoch II, from 1936 to 1980, turns the aggregate economy into a system to be measured, modeled, and steered. Epoch III, from 1980 to 2015, explains how ideas, software, networks, organizations, tasks, and intangible capital create value outside the conventional factory. Epoch IV, from 2015 to the present, confronts capital that can infer and generate, while evidence about capability, adoption, productivity, organization, employment, and macroeconomic effects remains sharply uneven.

The selection is a curated canon, not a citation ranking. Each work is included because it changed the object of explanation, supplied a durable mechanism or method, generated substantial downstream work, and occupies a non-redundant place in the four-epoch argument. Influence is not endorsement. Every entry therefore asks four questions: What did the work make visible? Why did that matter? Which later debates inherited its mechanism? Where does the explanation fail or remain incomplete?

The paper advances one central claim. The history of economic thought is not a linear replacement of wrong theories by correct ones. It is a widening of the unit of analysis and a tightening of the evidence boundary. Coordination, control, connection, and cognition each solve a problem the earlier frame could not see. Each also produces a new blind spot. The current AI debate is the latest instance: technical capability is observable before its organizational, distributive, and macroeconomic consequences are known.

The accompanying visualization program turns that argument into four different reader roles—Restorer, Decision-maker, Tracer, and Juror. Its governing rule is equally simple: interactivity may reveal mechanism or precision, but it may not carry the only version of the claim. Every proposed figure therefore includes an exact or narrative fallback, a mobile contract, an evidence boundary, and a named integrity risk.

The argument in one page

Epoch	Governing problem	Expanding unit of analysis	Characteristic evidence	Boundary exposed
I · Coordination · 1776–1936	How can specialized, decentralized activity form an economic order?	Workshop → household → class and distribution → linked markets → money → welfare → aggregate economy	Institutional argument, counterfactual reasoning, formal equilibrium, historical observation	Coordination does not guarantee fairness, stability, or full employment.

Epoch	Governing problem	Expanding unit of analysis	Characteristic evidence	Boundary exposed
II · Control · 1936–1980	Can the aggregate economy be measured and steered?	National economy → structural equation → policy rule → expectation-forming agent	Axioms, national accounts, time-series estimation, probability and identification	The model's parameters can react to the policy rule and to information the controller does not possess.
III · Connection · 1980–2015	What changes when information and ideas become productive assets?	Idea → standard → network → platform side → task → workflow → intangible capital	Formal models, firm panels, organizational measures, task data, reconstructed accounts	Cheap copying does not create cheap coordination; complements, governance, and lock-in remain scarce.
IV · Cognition · 2015–present	What follows when capital can infer, generate, and recommend?	Capability → exposure → adoption → use → task productivity → organization → labor/macro outcome	Experiments, exposure classifications, workplace rollouts, firm panels, administrative data, calibration	Capability is not consequence; no single study spans the whole causal chain.

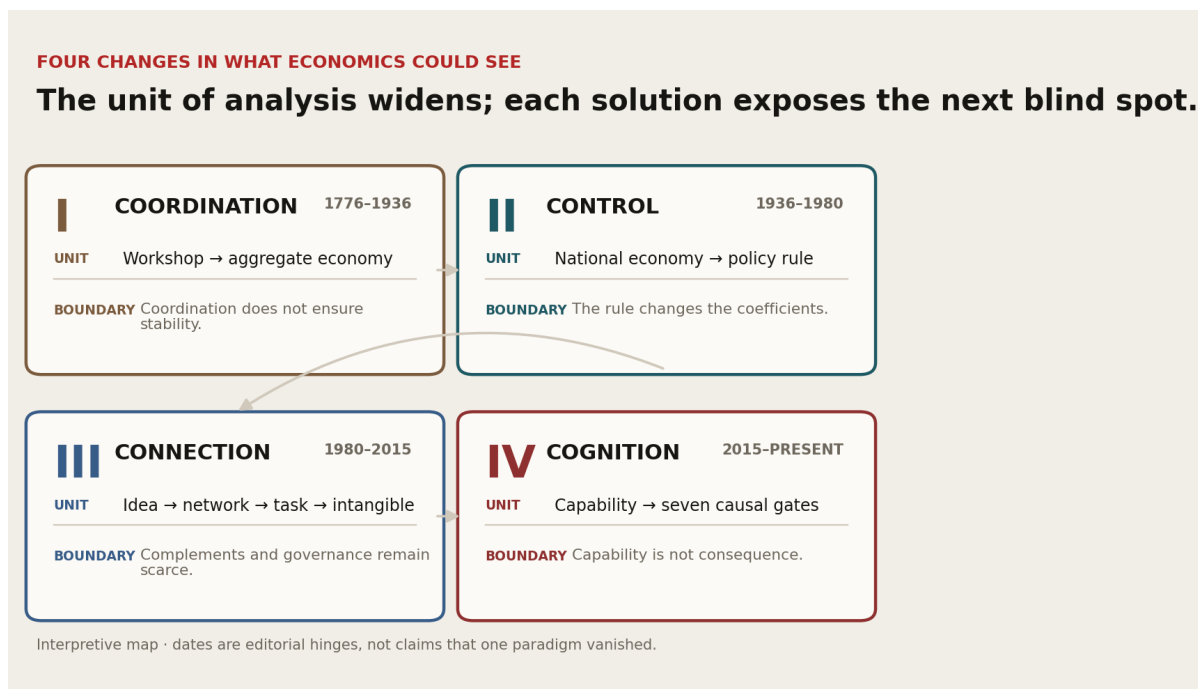


Figure 1. The four epochs shift the governing problem from coordination to control, connection, and cognition.
Source: author's synthesis.

Why these four epochs

Histories of economics usually organize themselves around schools, famous economists, methods, or policy disputes. Each approach clarifies something and conceals something else. A

school-based history makes boundaries look cleaner than they were. A biography-centered history overstates individual authorship. A method-centered history can mistake greater formalization for a simple increase in truth. A policy-centered history risks reading older work only through contemporary arguments.

The four epochs here follow a different variable: the object that economists believed they had to explain. Smith begins with specialization inside a workshop and the coordination of exchange beyond it. Keynes ends the first epoch by demonstrating that mutually consistent plans can coexist with involuntary unemployment. The second epoch makes the aggregate economy operational—first as an equilibrium panel, then as an estimated and probabilistic system—before expectations, information, and regime dependence undermine the idea of stable control. The third epoch follows value as it escapes the physical factory into knowledge, compatibility, software, platforms, tasks, organizational complements, and intangible investment. The fourth asks what changes when the productive input is not only information, but machine-mediated cognition.

This is not a claim that Smith disappears after 1936, that macroeconomics begins from nothing with Keynes, or that digital economics waited for 1980. The boundaries are hinges. A work is placed where it performs its most useful narrative function in this canon. Hayek belongs in Epoch II because dispersed knowledge is a limit on the policy machine, although the same mechanism anticipates digital coordination. Autor, Levy, and Murnane belong in Epoch III because the task lens explains the digital reorganization of labor, while that lens becomes indispensable for AI. Keynes appears once: as the work that breaks the first epoch. Hicks, not Keynes again, opens the second by making part of the break operational.

Influence without a league table

“The ten most influential papers” sounds more objective than it is. Citation counts favor recent journal articles, English-language databases, fields with dense citation practices, and works whose influence remains legible in modern metadata. They also treat endorsement, criticism, and routine invocation as the same event. Before the modern journal system, decisive interventions often appeared as books, treatises, and monographs. The paper therefore uses **works** as its accurate umbrella term, especially in Epoch I.

Selection follows five linked tests. First, paradigm effect: did the work change the problem economists tried to solve? Second, conceptual durability: does its mechanism still organize research, policy, or management reasoning? Third, downstream reach: did later work build on, formalize, test, or repair it? Fourth, cross-generational influence: does the work connect more than one period or domain? Fifth, narrative non-redundancy: does it add a function the other nine works in the epoch do not already perform?

These tests require judgment. The point is not to disguise that judgment as a score. It is to make it inspectable. Near misses matter because a canon becomes more credible when the

sacrificed alternatives are visible. The accompanying research files retain contested candidates, edition notes, selection cautions, source links, and a forty-row claim ledger. A different question would yield a different canon. A history of institutions, development, gender, public choice, finance, behavioral economics, or the Global South should not be expected to produce this list.

How to read the evidence

The paper distinguishes four kinds of statement. A **fact** is directly observed or sourced: a publication date, an experimental estimate, a sample, a series, or a documented institutional feature. An **estimate** depends on a design, model, population, and uncertainty. An **inference** is the most plausible explanation of the available facts. A **proposition** is a new, testable argument that the paper derives from the sequence.

This distinction becomes decisive in Epoch IV. A benchmark can show that a model performs a task. An exposure classification can map which tasks overlap with that capability. An experiment can identify a productivity effect on a bounded assignment. A workplace deployment can observe one organization under one implementation. Administrative data can record realized employment or earnings over a particular horizon. A calibration can translate assumptions into an aggregate range. None of these evidence forms is a defective version of another. They answer different questions.

The same discipline applies historically. A theorem demonstrates what follows under stated assumptions; it does not establish that the assumptions describe a particular economy. A stylized interactive model can reveal a mechanism; it is not an estimate or forecast. An editorial concept code can expose omission; it is not a measured intensity. The visualization program therefore labels observed, estimated, modeled, and illustrative quantities at the point of use.

A tangible thread through the paper

Each epoch receives one persistent specimen.

1. **The pin factory becomes a city.** The reader begins with specialization and adds households, land, ownership, marginal choice, linked markets, money, smoke, and an empty order book. The visual object accumulates the missing layers that eventually require macroeconomics.
2. **The policy machine reacts to its operator.** The reader moves transparent levers in a stylized economy, then changes the expectation rule and watches the mapping itself move. Control fails when behavior is treated as a fixed coefficient.
3. **The digital firm becomes a dependency structure.** The reader follows one idea through network effects, increasing returns, software, tasks, organizational redesign, platforms, and intangible capital. The number of possible links explodes faster than the capacity to govern them.
4. **One queue produces three futures.** The reader holds an observed support-task productivity result constant while changing demand, staffing, quality, and wage-sharing rules. The same technical capability can support augmentation, intensification, or contraction.

The specimens are not decorative analogies. Each is a controlled way to separate mechanism from outcome. Their common twist is that the parameter most likely to reverse the conclusion—distribution, expectations, organizational complementarity, or demand response—is not fixed by the technology itself.

Paper structure

The next four chapters follow a deliberate mechanism-led sequence and identify the transition each epoch forces; the canon tables and bibliography preserve chronology. A cross-epoch synthesis then compares recurring tensions, units of analysis, and evidence standards. The final section specifies sixteen visualization concepts, including four hero-figure consultations, data contracts, interaction behavior, accessible fallbacks, mobile reductions, and integrity boundaries. The bibliography preserves primary links and edition notes wherever the package contains them.

Epoch I — The Market Is Invented—Then It Breaks (1776–1936)




Figure 2. The title page of the 1776 *Wealth of Nations*, the archival entrance to Epoch I. Public-domain scan via Wikimedia Commons.

Canon at a glance

The table makes the selection judgment inspectable. The final column is the first, most compact influence claim carried by the underlying research record; the chapter then tests its mechanism and boundary in detail.

No.	Year	Author(s) and work	Why it enters this canon
1	1776	Adam Smith, <i>An Inquiry into the Nature and Causes of the Wealth of Nations</i>	Smith explained how specialization raises productivity and exchange coordinates dispersed plans without a central production plan.
2	1798	Thomas Robert Malthus, <i>An Essay on the Principle of Population, as It Affects the Future Improvement of Society. With Remarks on the Speculations of Mr. Godwin, M. Condorcet, and Other Writers</i>	Malthus placed population, food, wages, fertility, and mortality in one dynamic feedback mechanism.

No.	Year	Author(s) and work	Why it enters this canon
3	1817	David Ricardo, <i>On the Principles of Political Economy and Taxation</i>	Ricardo made distribution the principal problem of political economy and showed how rents, wages, and profits move within a constrained system.
4	1867	Karl Marx, <i>Das Kapital. Kritik der politischen Oekonomie. Erster Band: Buch I. Der Produktionsprozess des Kapitals</i>	Marx shifted attention from exchange alone to the ownership and organization of production.
5	1871	Carl Menger, <i>Grundsätze der Volkswirtschaftslehre</i>	Menger relocated value from labour embodied in an object to the importance of the least urgent need served by a scarce additional unit.
6	1874–1877	Marie-Esprit-Léon Walras, <i>Éléments d'économie politique pure, ou Théorie de la richesse sociale</i>	Walras transformed equilibrium from a single-market intersection into a system in which prices and quantities are jointly determined across markets.
7	1890	Alfred Marshall, <i>Principles of Economics</i>	Marshall fused classical cost reasoning with marginal demand into the practical toolkit that dominated teaching and applied price theory.
8	1898	Johan Gustaf Knut Wicksell, <i>Geldzins und Güterpreise: Eine Studie über die den Tauschwert des Geldes bestimmenden Ursachen</i>	Wicksell made money and bank credit active: a gap between the loan rate and the natural return can move prices cumulatively.
9	1920	Arthur Cecil Pigou, <i>The Economics of Welfare</i>	Pigou showed that decentralized private choices can be mutually consistent yet socially inefficient when costs or benefits spill onto others.
10	1936	John Maynard Keynes, <i>The General Theory of Employment, Interest and Money</i>	Keynes changed the unit of analysis from an isolated market to aggregate income, expenditure, output, money, and employment.

Ten works taught economics how decentralized exchange can coordinate an increasingly complex society. The tenth revealed the boundary of that achievement: a coherent market economy can still settle into mass unemployment.

A market is not yet an economy

The familiar story of early economics is too smooth. Adam Smith discovers the market, marginalism refines it, and John Maynard Keynes later adds government. The history is more

interesting because the object itself keeps changing. The workshop becomes a population. Output becomes distribution. Exchange opens into ownership and power. Individual choice expands into a simultaneous system of markets. Money enters, then social cost, and finally aggregate demand.

That widening field of vision is the central movement of this epoch. Each work makes one mechanism legible, but leaves a missing layer for the next to expose. Economics develops less like a completed machine receiving extra components than like an archive under restoration.

The ten items below are **works**, not “papers.” They are books, treatises, monographs and, in Malthus’s case, a book-length essay. The editorial canon privileges paradigm effect, durability, downstream reach and narrative non-redundancy. Its sequence is chronological, not a ranking. Influence is not endorsement.

One specimen can hold the story together. Begin with Smith’s pin workshop: a lone worker could perhaps make one pin in a day and certainly fewer than twenty, while Smith’s illustrative ten-person shop produces about 48,000. These are Smith’s numbers, not a calibrated production function. Then grow the workshop into a town: add households, land, ownership, ordered needs, linked markets, a bank, a smoking chimney and, finally, an empty order book. The physical town scarcely changes at the final step. Its economic state does.

1. Adam Smith: coordination acquires a mechanism

Smith’s *An Inquiry into the Nature and Causes of the Wealth of Nations* (1776) did not merely celebrate exchange. It explained a mechanism connecting specialization, productivity and the extent of the market. The pin workshop matters because division of labor converts a social arrangement into a production technology: tasks become narrower, dexterity rises, transitions disappear and specialized tools become worthwhile. Exchange then coordinates the resulting interdependence without a central production plan.

This displaced the mercantilist tendency to identify national wealth with treasure and trade balances. It matters now because arguments about specialization and organizational scale still turn on the relation between task decomposition and market extent. Yet Smith’s commercial society was never institution-free: justice, public works, monopoly and bargaining power belonged inside it.

The boundary is equally durable. Decentralized coordination does not prove equitable bargaining, macroeconomic stability or full employment. A town can allocate its current orders efficiently and still receive too few orders. That distinction will eventually break the epoch open.

[Primary source and edition note](#)

2. Thomas Robert Malthus: growth becomes a feedback problem

Malthus’s anonymously published first edition of *An Essay on the Principle of Population* (1798) adds households, food and time to the workshop. Resources affect wages and living conditions;

these affect fertility and mortality; population in turn changes pressure on resources. Growth is no longer an uncomplicated accumulation of output. It is a dynamic system whose behavioral response can erode the initial gain.

The contribution made political economy confront feedback rather than extrapolate a trend. Population economics and demography inherit that move. A larger harvest cannot simply be drawn as a permanent upward step; behavioral and productive responses matter.

The famous population–subsistence race is not a universal empirical law. It underweights technical change, trade, human capital, contraception, institutions and the demographic transition. The broader methodological lesson survives the failed simplicity: a constraint that appears fixed may become endogenous once people respond. Modern growth arguments still go wrong when they freeze the response parameter that carries the result. [Primary source, 1798 first edition](#)

3. David Ricardo: aggregate gain separates from distribution

Ricardo's *On the Principles of Political Economy and Taxation* (1817) changes the question from how much the town produces to who receives its product. Wages, profits and differential rent move within one constrained system. Distribution is not a ceremonial division performed after production; it affects accumulation and the trajectory of the economy.

Comparative costs provide the second move. Add a town that is absolutely more productive in cloth and pins. Specialization can still raise joint output when relative costs differ. Ricardo thus established a durable style of reasoning: construct the counterfactual and separate absolute from comparative advantage.

However, the aggregate result conceals composition. A gain from trade does not identify adjustment speed, displaced workers, strategic dependence or the groups that capture the gain. This is why Ricardo matters now: he supplies both the case for exchange and the reason to ask a second question about incidence. Efficiency and distribution are related, but they are not interchangeable conclusions. [Primary source, Sraffa–Dobb critical edition](#)

4. Karl Marx: the camera enters the factory

Volume I of Marx's *Das Kapital* (1867) moves behind the market transaction. The contract for labor power may be an exchange, but production takes place within an organization structured by ownership, authority, the working day, technology and control over surplus. Capitalism becomes a dynamic process of accumulation and recurrent disruption rather than a set of trades arriving at rest.

In the running town, the factory wall becomes transparent. Readers see who owns the machinery, who controls work, how output separates from wages and how surplus returns as investment. Market price alone cannot answer those questions.

Marx's influence across economics and the social sciences is difficult to separate from the political tradition his work generated. That reach is not the same as validation. The relation between labor values and observed prices remains contested, as do broader-corpus claims about crisis and profitability that cannot be attributed to the cited first volume alone. Later mixed economies showed that bargaining, welfare states, regulation and macroeconomic stabilization can reconfigure capitalism. They did not make ownership and power disappear. [Primary source, 1867 German first edition](#)

5. Carl Menger: value moves to the margin

Menger's *Grundsätze der Volkswirtschaftslehre* (1871) relocates value from the labor embodied in an object to the importance of the least urgent need served by an additional scarce unit. In the town, identical sacks of grain have no intrinsic rank. Their value depends on the ordered uses available to a household and on which use would be lost if one sack disappeared.

This gave choice a causal, subjective and marginal structure. Menger also placed knowledge, time and emergent institutions inside market processes. His independent route founded the Austrian tradition rather than supplying a verbal version of Walrasian equilibrium.

The modern significance is not that subjective judgment dissolves every constraint. It is that scarcity acquires meaning only in relation to purposes and available quantities. Yet a fine-grained theory of individual valuation does not by itself explain aggregate demand collapse or mass unemployment. The method illuminates why one household chooses; it does not automatically show why an entire town goes quiet. [Primary source, original German text and facsimile](#)

6. Léon Walras: the town becomes a system of equations

Walras's *Éléments d'économie politique pure* (1874–1877) changes both scale and method. Labor, pins, food, land and household budgets are no longer analyzed one market at a time. Prices and quantities are determined simultaneously across an interdependent system. General equilibrium turns the market economy into a formal coordination problem.

That program became foundational for later theories of existence, efficiency, welfare and market design. Its enduring contribution is not a claim that actual markets are perfect. It is the discipline of asking whether all plans can be mutually compatible at once, and what assumptions make such compatibility possible.

The distinction between equilibrium and adjustment is decisive. Counting equations and unknowns does not prove existence, uniqueness or stability. Walrasian *tâtonnement*—often represented by a notional auctioneer—and the convention of no trade before equilibrium remove many of the transactions through which real economies discover prices and sometimes destabilize themselves. Money, finance, institutions and unemployment remain thinly represented. The model gives the town a complete blueprint, but not yet a credible construction sequence. [Primary source, scan of the 1874 first instalment](#)

7. Alfred Marshall: economics gains a practical microscope

Marshall's *Principles of Economics* (1890) narrows the camera again, this time deliberately. Partial equilibrium holds much of the economy constant so that one market can be understood with supply, demand, elasticity, consumer surplus and different market periods. The familiar scissors metaphor matters because neither cost nor utility alone determines price; both blades do the cutting.

Marshall supplied a shared language for teaching and applied analysis, while economic time let supply response differ between the market period, short run and long run. In the pin town, demand first changes price, later production and eventually capacity.

The microscope creates its own blind spot. Income effects, increasing returns, feedback across markets, financial conditions and path dependence can invalidate the instruction to hold everything else constant. A collection of well-explained individual markets is not yet a theory of aggregate output. Marshall made price analysis usable; Keynes would ask when its local clarity becomes macroeconomic misdirection. [Primary source, 1890 first edition](#)

8. Knut Wicksell: money starts moving the real system

Wicksell's *Geldzins und Güterpreise* (1898) introduces the bank as an active institution. When the market loan rate differs from the natural return on capital, investment and prices can move cumulatively instead of returning automatically to rest. Money and credit are no longer a neutral veil draped over real exchange.

This mechanism is an ancestor of modern interest-rate stabilization, inflation targeting and the continuing debate over the natural rate, or (r^*). Its relevance lies in the gap, not in one timeless number. A central bank and private credit system can push the town along a dynamic path because financing conditions alter spending before all prices have adjusted.

The natural rate is unobservable, model-dependent, time-varying and potentially non-unique. Wicksell's original mechanism also simplifies expectations, bank balance sheets, risk and crisis. Once the stabilizing benchmark itself must be inferred, policy acquires a measurement problem. The lever may be visible. The neutral setting is not. [Primary source, original German edition](#)

9. Arthur Cecil Pigou: private coordination can produce social waste

Pigou's *The Economics of Welfare* (1920) adds a factory chimney. The pin producer and buyer can complete a mutually consistent transaction while smoke damages nearby households and the river. Private and social net product diverge. Coordination within the price system can therefore be genuine and still generate an inefficient social outcome.

This framework became the ancestor of corrective taxation and subsidies. It changes the role of government: intervention may be relevant not because exchange fails, but because it succeeds on terms that omit costs imposed on others.

The diagnosis does not manufacture an omniscient remedy. Corrective policy requires knowledge of marginal harm, incidence, enforcement and administrative capacity. Later work would add reciprocal harms, transaction costs, property rights and government failure. Pigou's town contains a social ledger, but completing it is an empirical and institutional task—not a license to write any convenient number in the missing column. [Primary source, 1920 first edition](#)

10. John Maynard Keynes: equilibrium loses its reassuring adjective

Keynes's *The General Theory of Employment, Interest and Money* (1936) is the threshold work. It closes Epoch I because it changes the unit of analysis from a collection of markets to aggregate income, expenditure, output, money and employment. The final scene in the town is not a technical breakdown. Machines work. Workers remain able. Wants have not vanished. Orders do.

The central mechanism is effective demand. Planned saving and planned investment need not be reconciled by a smoothly adjusting interest rate at full employment. They can be reconciled through falling income. Once that possibility is admitted, involuntary unemployment can persist as an equilibrium state rather than a temporary deviation awaiting a flexible wage.

This created modern macroeconomics and supplied an intellectual basis for stabilization policy. It also established a distinction that remains necessary: individual prudence can reduce aggregate income when many actors cut expenditure together. However, the later record would expose the limits of hydraulic confidence. Inflation, expectations, supply shocks, financial structure and the open economy constrain demand management. Keynes broke the presumption of automatic full employment; he did not provide a timeless setting for every policy dial. [Primary source and authorized edition record](#)

What changed across the epoch

The sequence is methodological as well as substantive. Smith uses institutional narrative; Malthus centers feedback; Ricardo sharpens counterfactual reasoning; Marx adds organization and power; Menger rebuilds value from marginal choice; Walras formalizes simultaneity; Marshall favors tractable local analysis; Wicksell adds monetary disequilibrium; Pigou separates private from social evaluation; and Keynes shows that consistency can occur below full employment.

The recurring tension is not market versus state. It is coordination versus consequence. Prices can coordinate plans without distributing gains fairly. Trades can be privately rational while imposing social costs. A mathematical equilibrium can exist without describing a stable adjustment path. Individual thrift can coexist with aggregate contraction. Once those distinctions are visible, “the market works” is no longer a complete analytical statement. One must ask: for which outcome, over which horizon, under which institutions, and at whose expense?

From diagnosis to the managed economy

Keynes belongs here only once. He is the break in the inherited frame, not the opening exhibit of a second identical gallery. Epoch II begins with the attempt to operationalize his diagnosis: national accounts, macroeconomic aggregates, formal models, statistical identification and policy instruments. The intellectual ambition changes from explaining a commercial order to steering an economy.

That ambition will generate its own contradiction. A policymaker can possess a model and still confront unstable coefficients, dispersed knowledge, expectations that react to the rule and shocks originating outside the demand system. The town has finally become visible from above. The next question is whether seeing the whole makes it controllable.

Epoch II — The Economy Learns the Rule, 1936–1980




Figure 3. Original editorial image of a hydraulic macroeconomic control room: the economy as a machine with visible levers and hidden feedback. Source: Schym editorial artwork, 2026.

Canon at a glance

The table makes the selection judgment inspectable. The final column is the first, most compact influence claim carried by the underlying research record; the chapter then tests its mechanism and boundary in detail.

No.	Year	Author(s) and work	Why it enters this canon
1	1937	John R. Hicks, <i>Mr. Keynes and the 'Classics'; A Suggested Interpretation</i>	Hicks compressed part of Keynes into simultaneous equilibrium in the goods and money markets.
2	1939	Jan Tinbergen, <i>Statistical Testing of Business-Cycle Theories. II: Business Cycles in the United States of America, 1919–1932</i>	Tinbergen turned competing business-cycle theories into an estimated model of the whole US economy.
3	1944	Trygve Haavelmo, <i>The Probability Approach in Econometrics</i>	Haavelmo gave econometrics an explicit probability foundation and separated structural relations from equilibrium correlations.

No.	Year	Author(s) and work	Why it enters this canon
4	1945	Friedrich A. Hayek, <i>The Use of Knowledge in Society</i>	Hayek recast allocation as a problem of local, tacit, and dispersed knowledge rather than optimization with all data given.
5	1954	Kenneth J. Arrow and Gérard Debreu, <i>Existence of an Equilibrium for a Competitive Economy</i>	Arrow and Debreu proved equilibrium existence for an integrated system of production, exchange, and consumption under explicit assumptions.
6	1956	Robert M. Solow, <i>A Contribution to the Theory of Economic Growth</i>	Solow replaced knife-edge fixed proportions with factor substitution and diminishing returns, making stable balanced growth possible.
7	1958	A. W. Phillips, <i>The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957</i>	Phillips documented a nonlinear inverse relation between unemployment and money-wage changes in British data.
8	1968	Milton Friedman, <i>The Role of Monetary Policy</i>	Friedman argued that expected inflation removes any permanent unemployment gain from monetary expansion, leaving only a short-run trade-off.
9	1970	George A. Akerlof, <i>The Market for 'Lemons': Quality Uncertainty and the Market Mechanism</i>	Akerlof showed that private quality information can drive good products out and shrink or destroy beneficial trade.
10	1976	Robert E. Lucas Jr., <i>Econometric Policy Evaluation: A Critique</i>	Lucas showed that parameters estimated under one policy regime may change under another because decisions incorporate expectations about the rule.

The Keynesian rupture made mass unemployment a system-level problem. The next four decades turned that diagnosis into a diagram, an estimated apparatus, and a policy ambition—then revealed the harder fact that an economy observes, interprets, and sometimes changes the rule applied to it.

A model can become useful enough to alter policy and, for that very reason, become unreliable as a description of what policy will do next. That is the central tension of this epoch. Economics learned to represent the economy as an aggregate system with measurable flows, behavioral equations, equilibrium conditions, and policy levers. It also learned, step by step, why the system was not a passive machine.

The usual summary—Keynesianism rises, monetarism answers, rational expectations wins—is too linear. It treats schools of thought as teams and misses the cumulative change in the standard of proof. Hicks made the macroeconomy drawable. Tinbergen made it estimable. Haavelmo asked what the estimates meant. Hayek questioned whether the information required by a central representation could ever be assembled. Arrow and Debreu clarified what a coherent equilibrium required. Solow stabilized the growth engine but left technology outside it. Phillips offered an empirical relation that looked like a control surface. Friedman and Phelps showed that expectations could move that surface. Akerlof demonstrated that agents may not even share the same information. Lucas then delivered the decisive inversion: a policy rule can change the coefficients on which its own evaluation rests.

This is not a story of modeling hubris followed by demolition. The machine was productive. National accounting, macroeconometric models, growth theory, and policy evaluation made questions tractable that had previously been discussed largely in prose. The deeper lesson is more demanding: tractability is an achievement, but the assumptions that create it must remain visible.

A hinge, not a duplicate

The epoch begins in 1936, although its ten-work canon begins in 1937. This is deliberate. Keynes's *General Theory* closes Epoch I because it breaks the presumption that decentralized markets reliably restore full employment. Epoch II asks what happened when that rupture became an operational research program. The boundary is therefore a hinge, not a fence: Keynes supplies the problem; Hicks and Tinbergen begin constructing the apparatus.

The ten works are ordered chronologically, not ranked. “Influential” means that a work changed the object of analysis, left a durable tool, created a downstream research or policy program, or exposed a blind spot that later economics could not ignore. Influence is not endorsement, and a ten-slot canon is necessarily selective. Samuelson, Phelps, Leontief, Stone, Coase, Kydland and Prescott, and others belong in the surrounding literature even when they do not occupy one of the ten plates on this particular machine.

From Keynes's rupture to a usable control panel

John Hicks's 1937 article, [“Mr. Keynes and the ‘Classics’; A Suggested Interpretation”](#), compressed part of Keynes's monetary theory into simultaneous equilibrium in the goods and money markets. The framework later called IS–LM gave economists a common plane on which fiscal expansion, monetary changes, interest rates, and output could be discussed. Its method was comparative statics: move a curve, find a new intersection, compare equilibria. That is an enormous gain in communicability. It is also a consequential narrowing. Time, balance sheets, changing expectations, and Keynes's more radical uncertainty become difficult to see once the economy is reduced to two schedules and one point.

The distinction matters because Hicks did not merely “explain Keynes.” He translated a difficult argument into an instrument. Instruments are powerful precisely because they suppress detail. IS–LM became the teachable panel of the postwar synthesis; it did not become a complete representation of Keynes’s economy. The policy implication is conditional: a lever appears effective inside the diagram only under the monetary regime, price behavior, expectation process, and financial structure that keep the schedules in place.

Jan Tinbergen took the next step. His 1939 League of Nations monograph, [*Statistical Testing of Business-Cycle Theories. II: Business Cycles in the United States of America, 1919–1932*](#), turned competing explanations of the cycle into a dynamic system of linked behavioral equations estimated from time-series data. The methodological shift is easy to underestimate. A diagram can organize reasoning; an estimated system can attempt to reproduce fluctuations, compare theories, and simulate intervention. The economy was no longer merely drawn as a machine. It was fitted as one.

Yet the apparent precision concealed unresolved problems. If investment, income, prices, and interest rates move together, an equation fitted to their joint movement does not automatically identify which relation is structural. Lag choices can alter results. Behavioral coefficients may not remain stable. Tinbergen’s project made macroeconomic policy analysis imaginable; it also made its identification problem unavoidable. Haavelmo would supply the probabilistic language, and Lucas would later question the stability of the object being estimated.

A machine needs probability—and never has all the information

Trygve Haavelmo’s 1944 monograph, [*The Probability Approach in Econometrics*](#), changed the relation between economic theory and observation. Economic equations were no longer deterministic propositions with an error term attached as an afterthought. They belonged inside a probability model that made estimation, testing, simultaneity, and identification parts of one system. Haavelmo’s crucial distinction is between a structural relation and an equilibrium correlation. When several variables are jointly determined, a line through observed points does not by itself tell a policymaker what would happen after an intervention.

This is the epoch’s first major move from measurement to identification. A machine with gauges is not yet a machine whose pipes are understood. Haavelmo gave economics a way to ask whether a proposed causal structure could, in principle, be recovered from the data. However, even a correctly identified structure can fail as a policy guide if its equations are not autonomous under the contemplated intervention. Probability disciplines the sensors; it does not guarantee that the wiring survives a regime change.

Friedrich Hayek’s 1945 essay, [“The Use of Knowledge in Society”](#), approached the control problem from another direction. The relevant economic knowledge is dispersed, local, and often tacit. A workshop knows that a particular input has become scarce; a household knows its immediate alternatives; a trader observes a local imbalance. Much of this information is never

recorded in a central state vector. Prices coordinate separate plans partly because they compress signals without requiring every fact to be assembled in one place.

Hayek's argument is sometimes reduced to a verdict against planning. Its more durable contribution is analytical: information architecture is an economic variable. Lower communication costs can improve a dashboard without making it complete. Aggregation can remove the very heterogeneity that explains an outcome. The policymaker's problem is therefore not only which lever to move, but which information exists, who holds it, how it is transmitted, and what is lost in translation. That question leads directly toward the information and network economics of Epoch III.

The whole system, proved and growing

Kenneth Arrow and Gérard Debreu's 1954 paper, "[Existence of an Equilibrium for a Competitive Economy](#)", gave the market system an unusually rigorous benchmark. Under explicit assumptions, an integrated economy of production, exchange, and consumption has a competitive equilibrium. Commodities can be differentiated by date, place, and state of the world, allowing the model to represent intertemporal and contingent exchange within one formal structure. The method is mathematical existence proof, not empirical estimation.

That distinction blocks several common overinterpretations. Existence is not uniqueness. It is not stability, computability, or evidence that an observed economy has reached the equilibrium. Nor does the theorem establish realism by assumption. Its value is partly diagnostic: convexity, complete contingent markets, and coherent information become a visible list of conditions. Later theories of market failure, finance, incomplete contracts, information, and computation gain precision because they can identify which gear is missing.

Robert Solow's 1956 "[A Contribution to the Theory of Economic Growth](#)" solved a different stability problem. By allowing substitution between capital and labor and imposing diminishing returns, Solow replaced the knife-edge dynamics of fixed-proportions growth with a stable balanced-growth path. Capital deepening raises output per worker, but its marginal contribution declines as capital accumulates. Sustained growth in output per person eventually requires technological progress.

The result made growth theory more stable and more honest at the same time. The model can explain convergence toward a steady state under its assumptions, but the decisive long-run fuel remains exogenous. Technology enters as something the model needs and does not yet explain. This is not a minor residual to be painted over. It is the opening through which knowledge, research and development, intangible capital, organizational complements, and endogenous innovation enter the next epoch.

The policy menu moves when it is ordered

A. W. Phillips's 1958 article, "[The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957](#)", used long-run British data to document

a nonlinear inverse association between unemployment and changes in money wages. The original dependent variable matters: this was a wage-change relation, not yet the familiar inflation–unemployment menu. Later economists and policymakers strengthened the interpretation.

The attraction is obvious. A stable historical curve appears to turn a political trade-off into a selectable coordinate: accept somewhat more inflation and obtain less unemployment. But a correlation is not automatically a structural response. Supply shocks can move it. The bargaining and expectation processes behind wage formation can change. The sample spans different institutions and regimes. Phillips created one of macroeconomics' most influential empirical objects; the controversy arose when an observed relation was treated as an invariant policy surface.

Milton Friedman's 1968 presidential address, "[The Role of Monetary Policy](#)", argued that monetary expansion cannot keep unemployment permanently below its natural rate. Workers and firms adjust as inflation becomes expected, so a short-run trade-off need not survive in the long run. Edmund Phelps independently developed the expectations and natural-rate logic and should be credited alongside Friedman. The policy mechanism is feedback: an initially surprising nominal expansion changes real decisions, but repeated use changes beliefs and therefore weakens the real effect.

Friedman also emphasized variable policy lags. A stabilizing action can arrive after the condition it addresses has changed, turning control into disturbance. The stronger conclusion is not that monetary policy is irrelevant. It is that level effects, surprise effects, and rule effects must be distinguished. The target learns, and the timing of that learning becomes part of the system.

Different information, different behavior, different coefficients

George Akerlof's 1970 "[The Market for 'Lemons': Quality Uncertainty and the Market Mechanism](#)" brought private information inside the market mechanism. When sellers know vehicle quality and buyers do not, buyers offer a price based on expected quality. Better vehicles then leave the market, average quality falls, and mutually beneficial trade can contract or disappear. The method is a compact theoretical model whose mechanism travels far beyond used cars.

The crucial contribution is not simply that information is imperfect. Information is distributed asymmetrically, and institutions respond. Warranties, certification, brands, screening, and reputation can be understood as economic devices that restore credible distinctions. A representative agent with one common information set cannot capture this process. The machine has split screens, and the differences between them affect participation, price, and welfare.

Robert Lucas completed the epoch's main inversion in 1976 with "[Econometric Policy Evaluation: A Critique](#)"⁸⁰⁰⁰³⁻⁶). Parameters estimated under one policy regime may change

under another because households and firms form decisions partly in response to the rule itself. A reduced-form relationship that predicts well inside the historical regime can therefore fail as a counterfactual when policy changes systematically.

The Lucas critique is often paraphrased as “policy does not work.” That is too crude. Its target is policy evaluation that treats behavior as invariant without explaining why. Credible counterfactuals require structures grounded in preferences, technology, constraints, and an explicit account of expectations. Even then, the empirical burden remains substantial. Microfoundations do not identify themselves, and a model can be structurally elegant while still omitting institutions, heterogeneity, or learning. Lucas did not remove the policy machine. He moved the operator inside it.

Tangible example: steer the town, then change the rule

The epoch becomes concrete in a fictional industrial town that loses 10 percent of private orders. Factories cut shifts. Household income falls, reducing consumption and supplier demand. Unemployment rises although workers, machines, and needs remain. The reader receives a control panel with a public-investment impulse, an interest-rate stance, an inflation shock, and expectation sensitivity.

In the first pass, a one-off €100 million public order moves through contractors, wages, local consumption, taxes, saving, and imports. The interface should display each round and its leakage rather than announce a mysterious multiplier. In a second pass, interest-sensitive investment and limited capacity enter. Output recovers, but inflation pressure rises as idle capacity disappears. In a third pass, the policy is repeated. Expected inflation adjusts, the employment effect weakens, and the same nominal input produces a different real path. Finally, the reader changes the policy rule itself. A model fitted to the first regime remains frozen beside the new realized path, making the forecast error visible.

The current website’s lightweight version correctly labels its outputs as dimensionless illustrative indices. It uses transparent equations:

$$y = \text{demand}(1 - 0.55e) - 0.8 \text{ rate-shock},$$

$$\pi = 2 + \text{shock} + 0.55y + 0.35e(\text{shock} + \text{demand}), \quad u = 5 - 0.45y.$$

Here, e is the reader-selected expectation sensitivity between zero and one. It dampens the real effect of a repeated demand impulse while allowing expectations to amplify inflation pressure. The phase arrows use a separate, disclosed local adjustment system around the normalized state $(y, \pi) = (0, 2)$:

$$y_{t+1} - y_t = -0.35y_t - 0.20(\pi_t - 2),$$

$$\pi_{t+1} - \pi_t = 0.28y_t - 0.22(\pi_t - 2).$$

Its negative trace and positive determinant make the displayed local system stable. These coefficients are authored teaching parameters, not empirical estimates.

These relations are teaching devices, not historical estimates. An extended version can add round-by-round town accounts and regime comparison, but it should preserve the same boundary. Every tooltip must distinguish an accounting identity, an assumed behavioral equation, a selected parameter, and an estimated relationship. Otherwise a story about model limits would quietly reproduce the problem it describes.

What the epoch settled—and what it did not

By 1980, economics had a durable macro toolkit: aggregate equilibrium, macroeconometric systems, probability-based identification, general equilibrium, steady-state growth, expectations, and information failure. It also had a stricter sequence for policy claims. Observation is not identification. Identification inside one regime is not invariance across regimes. Capability to steer is not possession of all relevant knowledge.

Several disputes remain live. Hicks's tractability must not be confused with Keynes's full argument. Phillips's wage relation must not be relabeled as a stable inflation menu. Arrow–Debreu proves existence under assumptions, not real-time convergence. Solow locates the importance of technological progress without endogenizing it. Friedman shares the expectations turn with Phelps. Lucas does not imply policy paralysis. These are not ceremonial caveats. Each one changes what can legitimately be inferred from a figure or simulation.

Transition to Epoch III: the missing fuel becomes the main object

Epoch II begins by asking how an aggregate economy can be stabilized and ends by asking what a controller can know, identify, and hold fixed. Its unresolved variables point outward. Solow's technological progress remains exogenous. Hayek's knowledge is distributed rather than centrally stored. Akerlof's information is private and institutionally mediated. Lucas's coefficients depend on the rule and on agents' interpretation of it.

The next epoch changes the unit of analysis. Value increasingly resides in knowledge, software, standards, networks, organizational complements, and intangible assets. The economy is no longer represented mainly as a hydraulic system with aggregate flows. It becomes a connected system whose returns depend on compatibility, adoption, and the structure of relationships. The surprise is that the policy machine did not simply break. It became a network—and networks are harder to steer precisely because they remember who connected to whom.

Epoch III — The Digital Turn (1980–2015)



Figure 4. Original editorial image of a physical knowledge network dissolving into connected nodes. Source: Schym editorial artwork, 2026.

Canon at a glance

The table makes the selection judgment inspectable. The final column is the first, most compact influence claim carried by the underlying research record; the chapter then tests its mechanism and boundary in detail.

No.	Year	Author(s) and work	Why it enters this canon
1	1985	Michael L. Katz and Carl Shapiro, <i>Network Externalities, Competition, and Compatibility</i>	It made compatibility and network scale central economic variables and supplied the basic model for standards wars and software ecosystems.
2	1986	Paul M. Romer, <i>Increasing Returns and Long-Run Growth</i>	It moved sustained growth inside the model and established knowledge as a productive input unlike rival physical capital.
3	1989	W. Brian Arthur, <i>Competing Technologies, Increasing Returns, and Lock-In by Historical Events</i>	It gave economics a formal account of path dependence, tipping, irreversibility, and the welfare ambiguity of standards competition.

No.	Year	Author(s) and work	Why it enters this canon
4	1990	Paul M. Romer, <i>Endogenous Technological Change</i>	It supplied the microeconomic architecture of endogenous growth and linked innovation to human capital, market incentives, and institutions.
5	1993	Erik Brynjolfsson, <i>The Productivity Paradox of Information Technology</i>	It reframed IT as a technology whose return depends on complementary investment, organizational redesign, learning, and measurement.
6	1998	Carl Shapiro and Hal R. Varian, <i>Information Rules: A Strategic Guide to the Network Economy</i>	It unified industrial-organization insights into the practical economics of Internet-era pricing, standards, rights, and lock-in.
7	2002	Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt, <i>Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence</i>	Firm-level evidence explained why identical hardware investments yield different productivity outcomes and can raise demand for skills.
8	2003	Jean-Charles Rochet and Jean Tirole, <i>Platform Competition in Two-Sided Markets</i>	It showed that price structure, not only the total price level, determines platform participation, welfare, competition, and governance.
9	2003	David H. Autor, Frank Levy, and Richard J. Murnane, <i>The Skill Content of Recent Technological Change: An Empirical Exploration</i>	It shifted labor economics from occupations to task bundles and supplied a mechanism for job redesign and changing skill demand.
10	2005	Carol Corrado, Charles Hulten, and Daniel Sichel, <i>Measuring Capital and Technology: An Expanded Framework</i>	It revealed intangible investment at roughly tangible-investment scale and changed how economists interpret saving, output, productivity, and the digital firm.

When value escapes the factory

The computer is the most visible object in the Digital Turn. It is not the most consequential one. The deeper change was economic: knowledge became a productive input, copying ceased to resemble manufacturing, compatibility made one buyer's choice depend on another's, and a growing share of the firm's capital disappeared from the physical balance sheet. The digital economy did not merely calculate faster. It altered what could be owned, reproduced, coordinated, priced, and measured.

That distinction matters because the standard technological story starts with the device and then searches for its effects. The literature of this epoch runs in the opposite direction. It identifies four changed mechanisms—non-rival ideas, increasing returns, network feedback, and organizational complementarity—and shows why a box installed in a firm is a poor measure of the value the firm can produce. The period's central finding is almost embarrassingly practical: technology is cheap to copy and expensive to absorb.

The ten works below form a mechanism-led canon, not a citation leaderboard. Each either founded a durable research program, changed the unit economists measured, or supplied a model still used to understand digital firms. Influence is not endorsement, and a famous label is not enough. The selection has to explain why the digital economy behaves differently and where that explanation begins to fail once software starts performing open-ended cognitive work.

Demand becomes connected

Michael L. Katz and Carl Shapiro's 1985 article, ["Network Externalities, Competition, and Compatibility"](#), changed the demand side of the market. In a conventional product market, one buyer's valuation can be analyzed largely in isolation. In a network market, value also depends on the number of compatible users and complements, on expectations about future participation, and on the installed base already in place. Compatibility therefore becomes an economic choice, not a technical afterthought. The mechanism is positive feedback: expected adoption raises value, higher value encourages adoption, and realized adoption validates the expectation. This framework became the basic language of standards wars, software ecosystems, and market strategies built around an early installed-base advantage.

The theory is influential precisely because it is parsimonious. That is also its boundary. User count and compatibility summarize much of the original network, while product quality is comparatively fixed. An AI network can improve or deteriorate with data, feedback, compute congestion, model updates, inference cost, and the participation of synthetic agents. More users can increase value, but they can also generate low-quality content, safety problems, or costly demand. Participation alone is no longer a sufficient statistic.

W. Brian Arthur pushed the feedback mechanism into history. In ["Competing Technologies, Increasing Returns, and Lock-In by Historical Events"](#) (1989), small stochastic differences in early adoption can move otherwise similar technologies onto different paths. Once increasing returns reinforce the leading option, the market can tip and remain locked in—even when the selected technology is not globally superior. Arthur's contribution was not the observation that history matters. It was a formal account of when history becomes persistent: early events move the system across a boundary between multiple possible equilibria. The model gave economists a disciplined language for path dependence, irreversibility, and the welfare ambiguity of market selection.

Here, too, the clean model marks the limit. Digital products can update, interoperate, and support multihoming. In AI, lock-in may sit less in the installed application than in proprietary

data, evaluation routines, compute commitments, workflow design, or human skills accumulated around it. The object that locks in has become a stack.

Knowledge refuses to diminish

Paul Romer's two contributions explain why the digital economy can exhibit increasing returns without requiring every physical input to do so. In 1986, "[Increasing Returns and Long-Run Growth](#)" treated accumulated knowledge as a productive input whose spillovers can offset diminishing returns at the level of the economy. Sustained growth no longer had to arrive as an unexplained residual from outside the model. The influential move was to show that privately accumulated inputs could create social returns beyond the firm that produced them, because knowledge can be used by one actor without being consumed by that use.

The 1986 paper leaves the production of knowledge comparatively aggregated. Romer's "[Endogenous Technological Change](#)" (1990) supplies the microeconomic architecture: profit-seeking firms invest in research, human capital helps produce new designs, ideas are non-rival but partly excludable, and temporary market power finances costly creation. The mechanism links innovation to incentives, market size, institutions, and the allocation of labor to research. It established the modern grammar of endogenous growth.

The distinction between the two papers matters. The first explains how knowledge spillovers can sustain increasing returns; the second explains why an actor would deliberately create ideas and how that activity can be financed. Neither paper is an empirical estimate of an AI production function. "Knowledge" remains a broad stock, and a design is economically usable once created. AI makes this compression difficult to maintain. Candidate generation can become cheap while verification, absorption, diffusion, energy, ownership, and implementation remain scarce. Model weights and ideas may be reusable, but compute and inference are rival inputs wrapped around them. The binding constraint may therefore move without disappearing.

The first copy is expensive; the next one is not quite free

Carl Shapiro and Hal Varian translated these mechanisms into the operating economics of the Internet in "[Information Rules](#)" (1998). Information goods combine high creation cost with cheap replication. That cost structure encourages versioning and differential pricing, while switching costs, rights, compatibility, and positive feedback make control of standards and complements strategically decisive. The book's influence came from unifying industrial-organization theory into a practical account of digital pricing and competition.

Its stylized information good is largely passive after creation. A generative system produces a new, personalized output at a non-zero inference cost, updates continuously, and may take actions rather than merely deliver content. The old "first copy versus next copy" curve has become a sequence: training, fine-tuning, retrieval, inference, verification, rights management, and continual updating. Cheap replication still matters. It no longer describes the whole cost function.

Productivity waits for the organization

Erik Brynjolfsson's "[The Productivity Paradox of Information Technology](#)" (1993) gave the era its most useful diagnostic. Computers could be visible everywhere while productivity remained difficult to find for four distinct reasons: output and input were mismeasured; learning and adjustment took time; private gains redistributed value without raising aggregate output; or implementation genuinely failed. The paper is a research synthesis, not a single causal estimate, and its influence lies in separating explanations that a slogan had bundled together. The mechanism shifted attention from installed hardware to complementary investment, organizational redesign, and lagged learning.

Timothy Bresnahan, Erik Brynjolfsson, and Lorin Hitt brought the organizational claim to firm-level evidence in "[Information Technology, Workplace Organization, and the Demand for Skilled Labor](#)" (2002). Their results describe IT, decentralized workplace organization, new products and services, and skilled labor as a complementary system. The return to one component rises when the others are present; identical hardware purchases can therefore coexist with very different outcomes. The empirical relationships are consistent with complementarity and skill-biased organizational change, but they do not turn every bundle of correlated practices into a guaranteed causal recipe. Firms choose technology and organization together, and successful adopters may differ in ways that are difficult to observe.

Carol Corrado, Charles Hulten, and Daniel Sichel then challenged the accounting boundary itself. Their 2005 chapter, "[Measuring Capital and Technology: An Expanded Framework](#)", applies a clear rule. Spending intended to raise future consumption should be treated as investment, even when the asset has no physical form. Software, databases, R&D, design, brands, firm-specific skills, and organizational capital can all generate services over time. Their estimates suggested that US business investment in intangibles was already roughly as large as tangible investment by the mid-1990s. The influence extended beyond a revised ledger. Capitalizing intangibles changes measured saving, output, productivity, and the economic description of the digital firm.

The framework also exposes a difficult empirical frontier. Intangible assets require prices, owners, service lives, deflators, and depreciation schedules that are often not directly observed. AI makes those assignments more difficult: training can resemble investment, purchased model access can resemble a service, inference can be an intermediate expense, and models or data can depreciate abruptly. A more inclusive boundary does not remove measurement judgment. It makes the judgment visible.

The firm becomes a market; the job becomes a bundle

Jean-Charles Rochet and Jean Tirole's "[Platform Competition in Two-Sided Markets](#)" (2003) showed that a platform cannot be understood by looking only at its total price level. It must bring interdependent groups on board, and participation on each side changes value on the other. The platform therefore chooses a price *structure*: one side may be subsidized while the other

pays, even when the same total revenue could be collected differently. This theoretical mechanism reshaped thinking about payments, software, media, and app ecosystems because welfare and competition depend on cross-side effects, rules, and governance as well as posted prices.

The two-sided abstraction is powerful, but AI platforms can have model, data, compute, tool, developer, user, and agent sides at once. Some participants are synthetic; ranking and defaults can be more consequential than money prices; and a vertically integrated platform may compete inside the market it governs. The pricing problem has not vanished. The number and identity of the sides have become endogenous.

David Autor, Frank Levy, and Richard Murnane changed the unit of labor-market analysis. ["The Skill Content of Recent Technological Change"](#) (2003) argues that technology acts on tasks rather than entire occupations. Computers substitute for routine cognitive and manual activities that can be expressed as rules while complementing non-routine problem solving and communication. Their theoretical model and empirical exploration supplied a mechanism for job redesign and changing skill demand that broad occupational labels conceal.

The limit is historical rather than fatal. "Routine" reflected what pre-AI computers could codify explicitly. Foundation models can now attempt language, coding, prediction, and other tasks previously classified as non-routine. Attempt is not reliable completion, however. Context, verification, tacit knowledge, liability, task recombination, and demand creation determine whether technical reach becomes substitution, augmentation, or neither. The task approach survives; its old dividing line does not.

A tangible example: build one digital firm

Imagine a fictional enterprise-software company in 1995 with a budget of 100 investment units. The visitor allocates the first tranche to product development and observes a high first-copy cost beside a reproduction cost close to zero. Three versions can be offered to different users. Shapiro and Varian explain the cost curve; Romer explains why the underlying design can be reused without being consumed.

The company must then choose an open or closed compatibility standard. Two rival products begin with identical fundamentals, but different early adopters. Replaying the same parameter settings produces different market shares when a small random event pushes one product across the tipping boundary. Katz and Shapiro explain why installed base changes value; Arthur explains why the path can persist.

Next, the firm adds independent developers and enterprise customers. The visitor can charge both sides equally, subsidize developers, or subsidize users. Equal total charges need not produce equal participation or profit because each side values the other. That is Rochet and Tirole made operational.

Finally, the remaining budget is divided among software, training, process redesign, data, and organizational capital. Buying only software produces little after the installation effect. Investing in the complements depresses conventionally measured productivity at first when those outlays are expensed, then raises economic output after a lag. Brynjolfsson supplies the diagnostic; Bresnahan, Brynjolfsson, and Hitt explain the complementary system; Corrado, Hulten, and Sichel reveal the missing asset. The last screen decomposes five jobs into tasks and lets a 2003 computer automate only explicit rules. An AI toggle expands technical reach into language and prediction, then adds error, verification cost, and accountability. Autor, Levy, and Murnane's framework works first—and then breaks in exactly the place the next epoch begins.

What the Digital Turn learned—and what it could not yet see

Taken together, the canon replaces a machine-centered account with a system-centered one. Ideas behave differently from physical capital. Network demand is interdependent. History can select an equilibrium. Cheap copying changes pricing but does not abolish fixed costs. Platforms govern several groups at once. Tasks, not job titles, are the relevant unit of technological exposure. Above all, productivity depends on organizational and intangible complements that arrive slowly and are measured badly.

This system view is the Digital Turn's durable achievement. Its unresolved assumption is more subtle: the computer remains mainly an instrument that stores, transmits, and executes encoded rules. Once it can generate candidate ideas, converse in natural language, make predictions, and take bounded actions, it enters several models as a participant rather than a passive input. Product quality becomes endogenous, the creator of an idea may be partly machine-based, the platform may host synthetic supply and demand, and a nominally non-routine task may fall inside the technical frontier.

That does not invalidate the digital canon. It changes the research question. The decisive variable is no longer whether a firm has acquired the technology, but whether capability survives verification, workflow integration, organizational redesign, accountability, and measurement. The computer finally appears in the productivity statistics only after the organization has done most of the difficult work. Intelligence will not be granted an easier entrance.

Epoch IV — The Intelligence Age (2015–present)



Figure 5. Original editorial image of a forensic workbench divided into capability evidence and realized economic outcomes. Source: Schym editorial artwork, 2026.

Canon at a glance

The table makes the selection judgment inspectable. The final column is the first, most compact influence claim carried by the underlying research record; the chapter then tests its mechanism and boundary in detail.

No.	Year	Author(s) and work	Why it enters this canon
1	2015	David H. Autor, <i>Why Are There Still So Many Jobs? The History and Future of Workplace Automation</i>	It replaced occupation-extinction forecasts with a task lens: technology substitutes for some activities while complementing workers in others.
2	2018	Daron Acemoglu and Pascual Restrepo, <i>The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment</i>	It formalized active capital as an expanding task frontier and made the direction of innovation endogenous.

No.	Year	Author(s) and work	Why it enters this canon
3	2020	Daron Acemoglu and Pascual Restrepo, <i>Robots and Jobs: Evidence from US Labor Markets</i>	It showed that automation can impose substantial local displacement even when national aggregates look calm.
4	2021	Erik Brynjolfsson, Daniel Rock, and Chad Syverson, <i>The Productivity J-Curve: How Intangibles Complement General Purpose Technologies</i>	It explained why transformative technologies can initially coincide with weak measured productivity as firms build unmeasured workflows, skills, data, and organizational capital.
5	2023	Shakked Noy and Whitney Zhang, <i>Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence</i>	Among 453 professionals, ChatGPT reduced completion time by roughly 40% and raised evaluator-rated quality by about 18%.
6	2024	Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, <i>GPTs are GPTs: Labor Market Impact Potential of LLMs</i>	It established a reproducible framework for mapping LLM capabilities onto occupational tasks and showed how software built around a model can expand the technical frontier.
7	2024	Tania Babina, Anastassia Fedyk, Alex He, and James Hodson, <i>Artificial Intelligence, Firm Growth, and Product Innovation</i>	AI-investing firms subsequently grew faster in sales, employment, and valuation, principally through product innovation.
8	2025	Daron Acemoglu, <i>The Simple Macroeconomics of AI</i>	It imposed accounting discipline on spectacular macro forecasts and estimated no more than a 0.66% TFP increase over ten years from then-available evidence.
9	2025	Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, <i>Generative AI at Work</i>	Across 5,172 support agents, the assistant increased resolved issues per hour by 15% on average, with much larger gains for less-experienced and lower-skilled workers.
10	2025	Anders Humlum and Emilie Vestergaard, <i>Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI</i>	Across 11 exposed Danish occupations, about 25,000 workers, and 7,000 workplaces, chatbot initiatives, reported time savings, and new AI tasks spread rapidly while administrative data showed precise null effects on earnings and hours.

The seven gates between capability and consequence

AI already changes bounded tasks and some production workflows. The unsettled question is how those gains travel through firms, labor markets, and national accounts. Ten works map that journey—and reveal how many links remain empirically thin.

A writing experiment can show a 40% reduction in completion time. A workplace deployment can raise resolved issues per hour by 15%. Meanwhile, a labor-market study can find no detectable change in earnings or hours. The numbers are not necessarily in conflict. They answer different questions at different scales.

That distinction defines the fourth epoch. Models can infer, generate, recommend, and increasingly execute parts of cognitive workflows. This paper proposes **active capital** as a compact label for that widening capability; it is not an established term and does not turn a model into a worker, firm, or autonomous economic subject. The harder move is to follow what happens next.

AI becomes an economic consequence only after passing through seven gates—**capability, exposure, adoption, use, task productivity, organizational response, and aggregate labor or macroeconomic effects**.




Figure 6. The seven-gate chain separates technical capability from realized labor-market and macroeconomic outcomes. Source: author's synthesis.

Each gate has its own unit, friction, and evidence. A firm may buy licenses while employees barely use them; workers may use a tool without improving verified output; task productivity may rise while the organization spends the gain on quality, new products, lower staffing, or more checking. Even large firm gains can remain small in GDP when affected tasks carry little weight or diffuse slowly.

Much public argument jumps from the first gate to the seventh. The ten works in this canon are influential partly because, taken together, they make that shortcut untenable.

From jobs to tasks—and from fate to direction

David H. Autor's 2015 essay, [*Why Are There Still So Many Jobs? The History and Future of Workplace Automation*](#), begins the epoch by changing the unit of analysis. Jobs are bundles of tasks. Technology substitutes for some activities while complementing labor in others; lower costs can also expand demand. This is why durable aggregate employment can coexist with occupational polarization and unequal gains. The essay is a synthesis, not a new causal estimate, but its conceptual influence is difficult to overstate. It made occupation-level extinction counts look analytically crude.

The task lens also places distribution at the center. Stable employment does not imply benign adjustment: an occupation can lose routine components, retain interpersonal work, and become a smaller or less attractive bundle. New tasks may emerge, yet access depends on education, geography, bargaining power, and firm-specific knowledge. What the paper cannot tell us is which tasks appear, how quickly, or who can enter them.

Daron Acemoglu and Pascual Restrepo formalized that open margin in 2018 with [*The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment*](#). Automation expands the set of tasks performed by capital and creates a **displacement effect**. Innovation may also generate new labor-intensive tasks, producing a **reinstatement effect**. Productivity, wages, employment, the labor share, and inequality can consequently move in different directions.

The decisive move is to treat technological direction as endogenous. Incentives, institutions, market power, and relative profitability help determine whether investment produces productive automation, human-complementary tools, or “so-so” substitution with modest output gains. The theory clarifies the mechanism; it does not identify which institutions produce which path.

Displacement can be local while aggregates remain calm

Theory gained a geographically concrete counterpart in Acemoglu and Restrepo's 2020 [*Robots and Jobs: Evidence from US Labor Markets*](#). The study uses differences in local industry composition and advances in industrial robots to estimate exposure across US commuting zones. In its aggregate estimates, one additional robot per thousand workers was associated with a 0.2-percentage-point reduction in the employment-to-population ratio and 0.42% lower wages.

National averages can conceal concentrated losses because workers, industries, and places are not interchangeable. Adjustment takes time, migration is costly, and new demand may arise elsewhere. The paper therefore supplies evidence at the adoption-to-labor-market end of the chain, but for industrial robots. Transferring its coefficients to generative AI would ignore differences in tasks, diffusion, complementarity, and geography.

Adoption is an investment program, not a login event

Erik Brynjolfsson, Daniel Rock, and Chad Syverson's 2021 [*The Productivity J-Curve: How Intangibles Complement General Purpose Technologies*](#) explains why even valuable technologies can arrive before their measured productivity effects. Firms must build complementary intangible capital: redesigned workflows, data, skills, governance, interfaces, incentives, and organizational knowledge. These investments consume resources before their returns appear and often sit awkwardly in conventional accounts. Measured productivity can therefore weaken before it strengthens.

For AI, the implication is practical. Purchasing access is not economically meaningful adoption. The system must enter a workflow; workers need rules for reliance, verification, and escalation; data and authority must be available; and the organization must allocate saved time. The J-curve explains delay, but cannot prove that a productivity boom waits beyond today's investment trough.

Weak productivity could reflect unmeasured investment, but also low use, unreliable output, costly checking, or narrow task relevance. The first interpretation justifies complementary investment; the second calls for selective deployment.

Capability and exposure are real—and still not outcomes

The empirical frontier shifted sharply in 2023. Shakked Noy and Whitney Zhang's [*Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*](#) randomized access to ChatGPT among 453 college-educated professionals completing recognizable writing tasks. Treatment reduced completion time by roughly 40% and increased evaluator-rated quality by about 18%. Lower initial performers gained more, compressing productivity differences within the studied task.

This is strong causal evidence for a bounded use-to-task-productivity link. It does not show that all professional writing becomes 40% cheaper, still less that employment or GDP changes by that amount. Fact-checking, interdependent workflows, repeated use, staffing, wages, and surplus allocation remain outside the design.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock widened the aperture in 2024 with [*GPTs are GPTs: Labor Market Impact Potential of LLMs*](#). Their task-exposure framework maps language-model capabilities, with and without complementary software, onto occupational activities. It created a reproducible possibility map and made the role of software complements explicit. The paper is influential precisely because it organized a fast-moving technical frontier into an economic unit—tasks.

Yet exposure is a correspondence between capability and task descriptions. It is not adoption, productive use, automation, wage change, displacement, or job loss. Reliability, cost, regulation, data, demand, and workflow redesign intervene. Calling exposure “jobs at risk” silently assumes away those variables.

Firms do not merely save labor; they choose a production model

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson's 2024 [*Artificial Intelligence, Firm Growth, and Product Innovation*](#) moves from occupational possibility to firm investment. Using AI-skilled human capital as an observable investment measure, the authors find that AI-investing firms subsequently grow faster in sales, employment, and valuation, principally through product innovation. Gains are concentrated among larger firms and coincide with greater industry concentration.

AI need not matter only through cost reduction. A firm may create products, expand demand, and hire; access to talent, data, computing, and complementary assets can simultaneously favor incumbents. The study's causal interpretation remains conditional on its instrument, and diffusion may either help smaller firms catch up or let advantaged firms pull farther ahead.

The organization enters even more directly in Erik Brynjolfsson, Danielle Li, and Lindsey Raymond's 2025 [*Generative AI at Work*](#). The staggered deployment of an AI assistant to 5,172 customer-support agents increased resolved issues per hour by 15% on average. Gains were much larger for less-experienced and lower-skilled workers, with little benefit at the top. The system appears to capture and diffuse patterns associated with expert practice.

This is more than speed. AI can alter the production of organizational knowledge by making tacit practices available to novices during work. It may also weaken incentives to train experts, independent judgment, and the supply of novel human examples. The study remains one firm in a bounded service environment; staffing, wages, expert learning, and generalization remain open.

Micro evidence reaches the macroeconomy through weights and frictions

Daron Acemoglu's 2025 [*The Simple Macroeconomics of AI*](#) imposes accounting discipline on the leap from task gains to aggregate growth. A Hulten-style framework combines the economic weight of affected tasks with task-level cost savings. Using then-available exposure and productivity inputs, the paper estimates no more than a 0.66% increase in total factor productivity over ten years.

Its value lies in showing what a macro claim must survive. Easy-to-measure tasks may not represent hard, context-dependent work, and adoption may trail capability. Conversely, scientific discovery, new tasks, products, or demand effects could be missing. The calibration may be too high or too low; either way, forecasts require explicit assumptions about affected shares, savings, diffusion, complementarity, and equilibrium response.

Anders Humlum and Emilie Vestergaard provide a different form of discipline in the 2025 working paper, revised in March 2026, [*Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*](#). NBER Working Paper 33777 was first circulated in May 2025 as *Large Language Models, Small Labor Market Effects*. Across 11 exposed Danish occupations, roughly 25,000 workers, and 7,000 workplaces, the authors document chatbot

initiatives, reported time savings, and new AI-related tasks, while administrative data show precise null effects on earnings and hours. The March 2026 revision rules out average effects larger than 2 percent two years after launch.

The null is not proof of permanent labor-market calm. The evidence covers an early horizon and selected Danish occupations; effects may emerge later, outside earnings and hours, or through margins the linked records do not capture. It still shows that adoption and task reorganization can precede measurable labor-market change. As a revised working paper, the result needs a date stamp, but the precise early null prevents task gains from becoming instant labor-market facts.

Evidence strength is not a single ladder

These works use theory, exposure classification, experiments, deployments, panels, linked records, accounting, and calibration. They should not be ranked on one scale. An experiment can identify a bounded treatment effect while saying little about employment. A model can expose a mechanism while relying on parameters. Exposure maps possibility without observing realization; labor panels capture outcomes but face selection, timing, and spillovers.

The economic unit must therefore remain visible beside every result: task, worker, firm, commuting zone, occupation, or national economy. So must the technology: industrial robot, predictive AI, language model, or complementary software system. Effect sizes across those units are not commensurable. A 15% workplace productivity estimate, a 0.42% local wage effect, and a 0.66% ten-year TFP calibration do not belong on one unqualified axis.

Recency adds another asymmetry. A 2015 synthesis has had time to shape a field. Results published or revised in 2025 and 2026 cannot be judged by raw citations and may still change. They enter this canon because they fill otherwise empty positions in the causal chain, not because their influence is already settled. The selection is consequently an editorial map as of 29 August 2026: influence is not endorsement, and “present” is a date rather than a timeless category.

Tangible example: one queue, three futures

Imagine a support operation with 100 agents and a normalized baseline of 100 resolved-issue units per hour. Turning on an assistant raises the index to 115, reflecting the average result in *Generative AI at Work*. The normalization is illustrative; the 15% treatment estimate is the evidence anchor. A cohort control reveals the reported heterogeneity: the effect is larger for novices and lower-skilled agents, and much smaller for experts.

At this point, the experiment ends and management begins. The same 15 additional units can produce three very different futures.

Future one: scale. Demand is available, so the firm holds staffing and hours constant and serves more customers. Queue length falls or throughput rises. Employment need not decline; it

may even expand if lower prices or better service increase demand. The relevant parameters are demand elasticity, capacity constraints, and whether faster resolution maintains quality.

Future two: augmentation. The firm keeps output volume broadly stable and reinvests saved time in complex cases, proactive service, model supervision, knowledge creation, or training. Measured quantity gains may look modest while quality and capability improve. New tasks partially reinstate labor, but only if the organization defines them, assigns decision rights, and rewards the work.

Future three: substitution or time pass-through. With demand fixed, the firm can reduce staffing, shorten hours, or capture the surplus as profit. These are not equivalent labor outcomes. A headcount reduction, a four-day schedule, and unchanged hours with higher margins all begin from the same task-productivity shock. Bargaining, contracts, labor scarcity, and managerial choice determine the allocation.

The interactive version should place a draggable membrane labeled **organizational choice** between the observed 115-unit state and the three downstream futures. Demand response, new-task creation, verification cost, and wage/time pass-through become sliders. Downstream numbers remain blurred until the reader crosses that membrane. Every value after the first act is labeled **illustrative scenario, not estimate**. The point is not to forecast the queue. It is to make the missing causal link tactile.

What the next literature must identify

First, research needs longitudinal evidence on **use**, not license counts or survey claims. Logs, workflow events, verification behavior, task reallocation, and output quality should be linked to worker and firm outcomes without turning workplace measurement into surveillance.

Second, economists need to measure **complementary intangible investment** directly. Data engineering, process redesign, governance, evaluation, training, and changed decision rights are part of the capital stock. Without them, the J-curve remains plausible but unquantified for AI.

Third, the field must separate **substitution, augmentation, and new-task creation** over time. Short-run assistance can become long-run automation; it can also generate demand and new specialties. The sign may reverse as models improve, prices fall, and organizations redesign work.

Fourth, distribution requires more than average productivity. Research should trace who captures rents, how wage setting responds, whether lower performers experience durable learning, and whether experts continue producing the knowledge that systems diffuse. Firm size, market power, geography, occupation, gender, age, education, and contract type may all mediate the result.

Fifth, macro estimates need representative task weights and credible diffusion paths. The field should measure hard-to-specify and verification-intensive work, not only tasks with clean digital

outputs. It also needs evidence on product innovation, scientific discovery, demand creation, prices, entry, and concentration. Cost savings are one channel, not the economy.

Finally, the relevant policy variable may be the **direction of innovation**. Tax, competition, labor, education, procurement, and research policy can influence whether firms pursue low-value displacement or create capabilities that complement people and expand useful output. Technology changes the feasible set. Institutions help choose the point inside it.

The strongest conclusion of this epoch is therefore neither a promise nor a warning. AI is already changing tasks and some workplaces. Economics has credible evidence for several gates between capability and consequence, but no paper spans the whole path. The macroeconomic settlement remains open—and, inconveniently for technological determinists, organizational choices are part of the model.

Cross-epoch synthesis

Recurring tensions and what actually changes

Across all four epochs, five tensions recur. First, decentralization versus control: Smith's coordination problem returns as Hayek's knowledge problem, platform governance, and the question of who sets an AI agent's objectives. Second, equilibrium versus process: Walrasian existence does not prove stability; Lucas shows that intervention changes the mapping; Arthur adds path dependence; AI adoption makes organizational transition itself an economic variable. Third, efficiency versus distribution: Ricardo's factor shares and Marx's power analysis reappear in skill-biased change, platform rents, local robot displacement, and the ownership of models and data. Fourth, observable surface versus latent mechanism: Wicksell's unobservable natural rate and Solow's exogenous technology are latent benchmarks or drivers, while digital intangibles and AI workflow capital are partly unmeasured assets. The distinction matters, but both cases expose what the current observational frame cannot directly see. Fifth, possibility versus realization: a market may coordinate in theory, a policy lever may work locally, a network may connect, and a model may perform a task—none of these facts alone establishes the final welfare result.

The deepest methodological change is not a straight march from prose to equations or from theory to data. It is the multiplication of valid evidence types and the need to keep their inferential boundaries explicit. A theorem proves a result under assumptions. A classification estimates technical exposure. An experiment identifies an effect on a bounded task. A workplace deployment observes an organization under a particular implementation. Administrative data records realized outcomes within a time window. A calibration aggregates assumptions. The reader should leave with a portable habit: always ask what was directly observed, at what unit, under which design, and what causal gates remain unmeasured.

Visualization program

The design system treats interaction as a form of analytical disclosure. Novelty belongs in the reader's action—not in an encoding whose authority exceeds the evidence. The sixteen concepts below are implementation-ready briefs, but they remain provisional wherever a source dataset, coding protocol, or model parameter has not yet been published.

Four hero-figure consultations

Each consultation contains exactly three candidates. “Challenger” describes a less standard form, not a presumption of superiority.

Epoch I hero — The object of economics keeps widening

Candidate	Role	Judgment
Proportional event–state timeline	Best analytical fit	Preserves irregular 1776–1936 spacing while showing the cumulative shift from workshop to aggregate economy.
Exact chronological table	Standard baseline	Highest precision and strongest static fallback, but forces the reader to reconstruct widening across rows.
Barcode timeline	Non-standard challenger	Loses. It implies event-density information the ten-work corpus does not contain and suppresses the meaning of each transition.

Epoch II hero — Policy moves the economy, and the rule moves with policy

Candidate	Role	Judgment
Phase portrait with local vector field	Best analytical fit	Shows direction and feedback around a transparent stylized policy state.
Scatterplot	Standard baseline	Familiar and exact for one state, but omits the direction of adjustment.
Connected scatterplot	Non-standard challenger	Loses. Connecting reader inputs would manufacture a time path that the illustrative model does not observe.

Epoch III hero — The digital economy is a dependency structure

Candidate	Role	Judgment
Fixed node–link diagram	Best analytical fit	Makes twelve audited conceptual dependencies traceable without assigning metric meaning to distance.
Exact edge table	Standard baseline	Accessible and unambiguous, but dependency paths must be assembled mentally.

Candidate	Role	Judgment
Force-directed network	Non-standard challenger	Loses. Emergent proximity would invite false readings of similarity or influence strength.

Epoch IV hero — No paper spans the whole causal path

Candidate	Role	Judgment
Paper-by-gate combination matrix	Best analytical fit	Keeps paper identity, evidence type, chronology, and seven noncommensurable gates visible.
Sortable evidence table	Standard baseline	Searchable and familiar, but the pattern of missing links is slower to see.
UpSet plot	Non-standard challenger	Loses for launch. Intersection counts erase paper identity and date; it becomes useful only as a secondary view if the corpus grows materially.

Visualization atlas: sixteen implementable ideas

Every brief inherits four release fields even when they are not repeated in the prose: **source readiness** (*ready*, *requires extraction*, or *requires new data*), **rights and licensing** for data and media, **update cadence** with an explicit as-of date, and a named **editorial/technical owner**. A figure cannot move from concept to production until all four fields are populated in its chart contract.

Epoch I — Archive reconstruction

I.1 Cumulative object timeline — MVP. Question: what did economics learn to see, and when? **Schema:** *work_id*, *first_publication_year*, *state*, *author*, *title*, *concept*. **Renderer:** semantic HTML + SVG proportional timeline. **Interaction:** scroll pins the axis; selecting a work adds its conceptual layer to a persistent town. **Tooltip/focus:** identical hover, focus, and tap card with year, contribution, limit, source. **Access/static:** ordered exact-value table and descriptive alt text. **Mobile:** vertical proportional rail with no horizontal scrubbing. **Integrity:** dates are observed; states are disclosed editorial coding; no causal weight is implied by spacing.

I.2 The pin workshop becomes a town — MVP. Question: which omitted layer breaks each convenient explanation? **Schema:** *step*, *work_id*, *object_added*, *mechanism*, *failure*, *repair*, *scene_asset*. **Renderer:** layered SVG/HTML cutaway. **Interaction:** a brass “theory lens” reveals labor, land, ownership, money, smoke, and idle capacity without replacing the specimen. **Tooltip/focus:** every hotspot has a labeled button and a two-sentence mechanism/limit panel. **Access/static:** six-panel annotated sequence. **Mobile:** stepper swaps

layers; no precision drag. **Integrity:** the town is an explanatory composite, not a reconstruction of historical quantities.

I.3 Two countries, two goods, four workers — Next. Question: how can comparative advantage raise total output while imposing adjustment costs? **Schema:** `country, good, labor_hours_per_unit, workers, autarky_mix, trade_price, transition_cost, group`. **Renderer:** linked production-possibility lines plus person-level allocation strip. **Interaction:** set productivity and a temporary retraining cost; compare autarky, specialization, and distributed gains. **Tooltip/focus:** show hours, output, and each group's income before/after; keyboard increments mirror sliders. **Access/static:** worked numerical table. **Mobile:** preset scenarios rather than dense dual sliders. **Integrity:** separate total gains from distribution; label the example synthetic and never imply that feasibility guarantees compensation.

I.4 Assumption X-ray — Next. Question: which analytical layers are explicit, implicit, or outside each work? **Schema:** `work_id, layer_id, status {explicit, implicit, absent, contested}, evidence_quote_ref, coder_note`. **Renderer:** accessible matrix/heatmap. **Interaction:** filter by money, power, institutions, uncertainty, externality, or aggregation; select a cell to open the coding note. **Tooltip/focus:** quote locator, edition, status definition, and uncertainty. **Access/static:** semantic table with symbols plus words. **Mobile:** one work at a time with layer chips. **Integrity:** absence means absent from the selected argument, not the author's entire corpus; dual-coder review is required.

Epoch II — Policy laboratory

II.1 Transparent macro phase portrait — MVP. Question: how do demand support, rate restraint, a supply shock, and expectations interact locally? **Schema:** `scenario_id, demand_impulse, rate_stance, supply_shock, expectation_sensitivity, output_gap_index, inflation_pressure_index, unemployment_index, response_equation_version, adjustment_equation_version`. **Renderer:** SVG vector field with HTML controls. The arrows are generated by the published local difference equations around $(y, \pi) = (0, 2)$; the highlighted state is generated by the separately published response equations. **Interaction:** levers update one highlighted state and its local arrows. **Tooltip/focus:** exact input/output readout and the equation term responsible for each movement. **Access/static:** scenario table, displayed equations, and prose interpretation. **Mobile:** chart first, full-width controls second. **Integrity:** dimensionless illustrative indices and authored teaching coefficients; never labeled estimate, forecast, or historical fit.

II.2 The coefficient that moved — Next. Question: why can an estimated policy relationship fail after a rule change? **Schema:** `regime, expectation_rule, policy_rule, coefficient_name, coefficient_value, simulated_observation`. **Renderer:** paired coefficient cards with a morphing regression line. **Interaction:** switch from surprise policy to anticipated rule and watch behavior redraw the mapping. **Tooltip/focus:** old/new parameter, governing assumption, and Lucas warning. **Access/static:** before/after equation table. **Mobile:**

two discrete regimes, no animation dependency. **Integrity:** use synthetic observations and explicit equations; animation demonstrates parameter dependence, not historical magnitudes.

II.3 One recession, four diagnoses — MVP. Question: how do Keynes, Hayek, Friedman, and Akerlof read the same symptoms differently? **Schema:** `symptom`, `school_lens`, `mechanism`, `missing_information`, `proposed_action`, `risk`. **Renderer:** four-lane causal storyboard. **Interaction:** pin a symptom—idle capacity, inflation, credit stress, quality uncertainty—and compare highlighted pathways. **Tooltip/focus:** claim, primary-work anchor, and what would falsify the diagnosis. **Access/static:** comparison table. **Mobile:** swipe-free accordion ordered by symptom. **Integrity:** these are disciplined interpretations, not claims that each author prescribed a single timeless policy.

II.4 Phillips curve time machine — Experimental. Question: how do expectations and supply shocks alter an apparent unemployment–wage/inflation relation? **Schema:** `date`, `geography`, `unemployment_rate`, `wage_or_price_change`, `expectation_proxy`, `supply_shock_flag`, `source_vintage`. **Renderer:** small-multiple scatterplots with uncertainty bands. **Interaction:** scrub data vintages and toggle original wage measure versus later inflation reinterpretation. **Tooltip/focus:** exact observation, series definition, vintage, revision status. **Access/static:** selected-vintage chart plus downloadable values. **Mobile:** preset eras and simplified small multiples. **Integrity:** source Phillips’s original UK series and modern series separately; never splice definitions or infer a stable policy menu from correlation.

Epoch III — Network room

III.1 Audited knowledge dependency network — MVP. Question: which conceptual dependencies connect the ten works? **Schema:** `node_id`, `author`, `year`, `concept`; `edge_id`, `source_id`, `target_id`, `relation`, `rationale`, `confidence`. **Renderer:** fixed-coordinate SVG. **Interaction:** selecting a node isolates incoming/outgoing paths and synchronizes the evidence card. **Tooltip/focus:** exact edge rationale, source anchors, and “distance has no meaning.” **Access/static:** edge ledger and chronological work list. **Mobile:** fixed overview then full-width node inspector. **Integrity:** only disclosed editorial edges; no line weight, centrality, or proximity claim without an auditable metric.

III.2 Standards-war tipping basin — Experimental. Question: when can small early events lock a market into one technology? **Schema:** `period`, `technology`, `installed_base`, `compatibility`, `switching_cost`, `adoption_payoff`, `random_seed`, `entrant_choice`. **Renderer:** deterministic simulation plus basin/trajectory plot. **Interaction:** replay the same seed while changing compatibility or switching cost; compare welfare after lock-in. **Tooltip/focus:** payoff components and cumulative adopters at each step. **Access/static:** three seeded scenario tables. **Mobile:** presets with step-through replay. **Integrity:** never present a random run as a forecast; disclose equations, seed, and sensitivity, and distinguish adoption dominance from welfare superiority.

III.3 The invisible balance sheet — Next. Question: which digital investments disappear when accounts treat them as current expense? **Schema:** `firm_year`, `spend_category`, `amount`, `accounting_treatment`, `persistence_assumption`, `depreciation_rate`, `source`. **Renderer:** bridge chart from reported expense to expanded intangible investment. **Interaction:** reclassify software, data, design, training, and organizational redesign under sourced depreciation assumptions. **Tooltip/focus:** amount, treatment, lifetime, and provenance. **Access/static:** reconciliation table. **Mobile:** stacked cards with one summary bridge. **Integrity:** use sourced firm or aggregate data; keep accounting reclassification separate from market valuation and disclose sensitivity to asset lives.

III.4 Platform price-structure sandbox — Next. Question: why can the same total price produce different platform participation? **Schema:** `side`, `price`, `demand_elasticity`, `cross_side_effect`, `participation`, `marginal_cost`, `rule`. **Renderer:** linked two-sided participation curves and a balance panel. **Interaction:** move a fixed total charge between sides; observe participation and platform surplus under transparent parameters. **Tooltip/focus:** per-side price, participation, cross-side effect, and welfare components. **Access/static:** worked scenarios and equations. **Mobile:** plus/minus controls with three presets. **Integrity:** stylized comparative statics, not a legal or empirical judgment; do not infer welfare from user count alone.

Epoch IV — Evidence courtroom

IV.1 Paper-by-gate evidence matrix — MVP. Question: which causal gate does each paper actually establish? **Schema:** `paper_id`, `year`, `evidence_type`, `gate`, `status {direct, modeled, absent}`, `coding_note`, `as_of_date`. **Renderer:** semantic HTML table enhanced as a matrix. **Interaction:** filter by gate/evidence type and open the full paper card. **Tooltip/focus:** status in words, design, sample, estimate where applicable, and limit. **Access/static:** the matrix itself is the exact fallback. **Mobile:** sticky first column and horizontal table, plus row-card view. **Integrity:** heterogeneous effect sizes are never normalized or combined; coding is dated and versioned.

IV.2 One queue, three futures — MVP. Question: how can the same 15% support-agent productivity result lead to expansion, augmentation, or labor saving? **Schema:** `scenario`, `demand_response`, `staffing_rule`, `quality_rule`, `wage_sharing`, `tasks_per_hour`, `headcount_index`, `worker_surplus_index`. **Renderer:** branching queue simulation with three synchronized outcome families. The current interface's four buttons—more output, better service, shorter time, and new tasks—are operating presets nested inside those three economic futures rather than four claims. **Interaction:** hold the observed task estimate fixed while changing explicitly hypothetical organizational rules. **Tooltip/focus:** identify observed input versus scenario assumption in every value. **Access/static:** assumption/outcome table for the three futures and four presets. **Mobile:** one preset per panel. **Integrity:** only the 15% study estimate is observed; headcount, wages, and demand responses are illustrative and must never inherit its evidentiary status.

IV.3 The leaky causal pipeline — Experimental. Question: how much macro impact survives each gate between capability and outcome? **Schema:** `gate`, `user_assumption`, `low_case`, `base_case`, `high_case`, `empirical_anchor`, `unit`, `dependency`. **Renderer:** log-scale attenuation cascade with explicitly bounded sensitivity bands. **Interaction:** readers enter adoption, use, reliability, task share, organizational pass-through, and demand response; deterministic low/base/high calculations update the outcome range without implying an empirical probability distribution. **Tooltip/focus:** assumption source, range, sensitivity contribution, and whether observed or chosen. **Access/static:** low/base/high calculation table with formula. **Mobile:** guided one-gate-at-a-time wizard. **Integrity:** a scenario calculator, never a forecast; prohibit default precision beyond source quality and show sensitivity rather than a single headline number.

IV.4 Jagged task frontier — Next. Question: where does AI capability remain reliable, review-intensive, or unsafe inside a job? **Schema:** `task_id`, `occupation`, `benchmark_score`, `error_cost`, `context_requirement`, `human_review_minutes`, `adoption_status`, `source_date`. **Renderer:** task map with uncertainty/error-cost contours. **Interaction:** choose an occupation, then sort tasks by capability, consequence, or review burden; compare model/version dates. **Tooltip/focus:** benchmark definition, confidence interval, error cost, and last-tested date. **Access/static:** task-level exact table. **Mobile:** ranked task cards before the map. **Integrity:** benchmark performance is not workplace performance; versions expire, error costs are context-specific, and occupation-level aggregation must preserve task weights.

Prioritization matrix




Figure 7. Sixteen visualization concepts organized by epoch and delivery priority. Source: author's synthesis.

Priority	Figures	Why now	Dependency / release gate
MVP	I.1 timeline; I.2 workshop-town; II.1 phase portrait; II.3 four diagnoses; III.1 dependency network; IV.1 evidence matrix; IV.2 three futures	They carry the four distinct story engines, reuse the audited corpus, and provide complete static fallbacks.	Keyboard/touch parity, reduced-motion mode, source-card deep links, and exact-value fallbacks must pass QA.
Next	I.3 comparative-advantage distribution; I.4 assumption X-ray; II.2 moving coefficient; III.3 invisible balance sheet; III.4 platform sandbox; IV.4 jagged frontier	They add tangible examples and modern measurement without changing the central narrative contract.	New coding protocols or source datasets; dual review for matrices; version/date metadata for task benchmarks.
Experimental	II.4 Phillips time machine; III.2 tipping basin; IV.3 leaky causal pipeline	They offer high explanatory value and the most innovative interaction, but carry higher risks of synthetic precision or source-definition drift.	Publish equations/data, scenario badges, deterministic seeds, uncertainty tests, and a prominent “not a forecast” boundary before release.

The launch principle is simple: novelty belongs in the reader’s action, not in an opaque encoding. Each epoch should feel different—restore, decide, trace, judge—while every hovered insight remains available by keyboard, every animation has a static explanation, and every visual claim can be followed back to a work, a dataset, or a plainly labeled editorial assumption.

Selection appendix — influence signals and near misses

How the signals should be read

The canon does not convert influence into a single score. It records four inspectable signals instead: **paradigm effect** (the work changed the problem), **durable mechanism** (its analytical device remains in use), **downstream program** (later research built, tested, or argued through it), and **bridge value** (it performs a non-redundant transition in this four-epoch account). The “why it enters” sentence in each epoch’s canon table is the operative selection claim. It can be challenged paper by paper.

Near misses are therefore not an apology or a reserve ranking. They show where the ten-slot constraint bites. Exclusion usually means that another selected work performs the same narrative function more economically, that the work belongs to a different canon, or that including it would crowd out a mechanism necessary for the transition. A research history centered on institutions, development, finance, behavioral economics, gender, ecology, or the Global South would make different choices.

Epoch I near misses

Work	Influence signal	Why it is outside this ten
John Stuart Mill, <i>Principles of Political Economy</i> (1848)	Major synthesis of classical political economy and the distinction between laws of production and arrangements of distribution	Ricardo already carries distribution and trade, while later entries are needed to show marginal choice, general equilibrium, money, welfare, and aggregate failure.
William Stanley Jevons, <i>The Theory of Political Economy</i> (1871)	Independent marginal revolution and mathematical utility theory	Menger supplies subjective marginal value and Walras supplies the interdependent formal system; a third marginalist would reduce mechanism diversity.
Thorstein Veblen, <i>The Theory of the Leisure Class</i> (1899)	Institutional and evolutionary critique of status consumption	Its omission is consequential, but the selected arc prioritizes the market’s analytical construction and macroeconomic break rather than a full institutional history.
Vilfredo Pareto, <i>Manual of Political Economy</i> (1906)	General equilibrium, ordinal utility, and welfare criteria	Walras and Pigou occupy the system and welfare functions; Pareto would sharpen both but add less transition value than Wicksell’s monetary disequilibrium.
Irving Fisher, <i>The Theory of Interest</i> (1930)	Intertemporal choice and capital theory	Wicksell performs the money–real-system bridge in this canon; a finance-centered history should reverse that choice.

Epoch II near misses

Work	Influence signal	Why it is outside this ten
Wassily Leontief, <i>The Structure of American Economy</i> (1941)	Input–output measurement and sectoral interdependence	Tinbergen carries the estimated macro system and Arrow–Debreu the integrated formal benchmark; Leontief belongs in a measurement-centered canon.
Paul Samuelson, <i>Foundations of Economic Analysis</i> (1947)	Unified comparative statics and formal postwar synthesis	Hicks provides the operational macro panel and Arrow–Debreu the existence benchmark, leaving limited non-redundant space for a general formal synthesis.
Richard Stone and the national-accounts program (1940s–1950s)	Operational measurement of national income and social accounts	The program is indispensable context, but this list is organized around works that alter mechanisms; the omission is explicitly recorded rather than hidden.
Edmund Phelps, “Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time” (1967)	Independent expectations-augmented natural-rate argument	Friedman occupies the expectations slot because of its policy reach; Phelps is credited in the chapter and contextual references.
Finn Kydland and Edward Prescott, “Rules Rather than Discretion” (1977)	Dynamic inconsistency and credible policy rules	Lucas closes the epoch with the broader regime-dependence critique; Kydland–Prescott would be the strongest substitution if credibility, rather than model invariance, were the hinge.

Epoch III near misses

Work	Influence signal	Why it is outside this ten
Ronald Coase, “The Nature of the Firm” (1937)	Transaction costs and the boundary of organization	It is foundational to digital organization but predates the epoch and would displace a work specific to knowledge, networks, tasks, platforms, or intangibles.
Paul David, “Clio and the Economics of QWERTY” (1985)	Historical path dependence and standards lock-in	Arthur supplies a more explicit increasing-returns model of tipping and lock-in for the mechanism-led story.
Timothy Bresnahan and Manuel Trajtenberg, “General Purpose Technologies” (1995)	GPT complementarities and innovation spawning	Romer carries endogenous knowledge and Bresnahan–Brynjolfsson–Hitt carries measured organizational complementarity; the GPT frame remains important context between them.
Eric von Hippel, <i>The Sources of Innovation</i> (1988)	Users and distributed actors as sources of innovation	The canon’s network and platform slots emphasize coordination and pricing; an innovation-organization history should include it.

Work	Influence signal	Why it is outside this ten
Yannis Bakos and Erik Brynjolfsson, “Bundling Information Goods” (1999)	Pricing under near-zero marginal cost and heterogeneous valuation	<i>Information Rules</i> already carries the operating economics of information goods, switching costs, versioning, and standards in one slot.

Epoch III ends in 2015 but its latest selected work is from 2005. That gap is deliberate and visible. The chosen ten build the conceptual stack—ideas, increasing returns, compatibility, lock-in, information-good costs, organizational complements, platforms, tasks, and intangibles. Much of the 2005–2015 literature measures diffusion, mobile and cloud markets, platform conduct, and organizational adoption through those mechanisms. A later empirical-diffusion canon would fill the decade differently; this mechanism-led canon reserves the next break for machine-mediated cognition.

Epoch IV near misses

Work	Influence signal	Why it is outside this ten
Edward Felten, Manav Raj, and Robert Seamans, occupational AI-exposure measures (2021)	Early systematic task–occupation exposure mapping	Eloundou and coauthors provide the GPT-specific exposure bridge used by the story’s seven-gate framework.
Sida Peng and coauthors, “The Impact of AI on Developer Productivity” (2023)	Controlled evidence on coding-assistant speed	Noy–Zhang covers bounded task experimentation across professional writing, while Brynjolfsson–Li–Raymond covers a real workplace deployment.
Fabrizio Dell’Acqua and coauthors, “Navigating the Jagged Technological Frontier” (2023)	Field evidence that gains and errors depend sharply on task location relative to capability	It is a prime candidate for expansion, but its task-experiment function overlaps Noy–Zhang in a ten-slot causal-chain canon.
Paweł Gmyrek, Janine Berg, and David Bescond, ILO generative-AI jobs index (2023)	Global occupational exposure with an augmentation-versus-automation lens	Eloundou supplies the selected exposure method; the ILO study is particularly important for a global labor-policy canon.
Anton Korinek and Donghyun Suh, scenarios for the transition to transformative AI (2024)	Explicit long-run macro scenarios under rapid capability change	Acemoglu occupies the macro slot because its bounded calibration disciplines current evidence; the scenario approach remains a useful contrast, not a competing estimate.

Contextual works cited by the argument

- Ronald H. Coase (1937), "The Nature of the Firm," *Economica* 4(16), 386–405. [DOI record](#).
- Roy F. Harrod (1939), "An Essay in Dynamic Theory," *The Economic Journal* 49(193), 14–33. [DOI record](#).
- Evsey D. Domar (1946), "Capital Expansion, Rate of Growth, and Employment," *Econometrica* 14(2), 137–147. [DOI record](#).
- Simon Kuznets (1934), *National Income, 1929–1932*. US Senate document and foundational national-income estimate. [NBER edition record](#).
- Edmund S. Phelps (1967), "Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time," *Economica* 34(135), 254–281. [DOI record](#).
- Charles R. Hulten (1978), "Growth Accounting with Intermediate Inputs," *The Review of Economic Studies* 45(3), 511–518. [DOI record](#).

These references support the surrounding architecture; they are not silently counted among the forty selected works.

What the sequence changes

The forty works do not converge on one final model of the economy. They produce a more demanding standard for economic explanation.

Smith asks how specialization can be coordinated. Keynes shows why coordinated plans need not produce full employment. The policy tradition makes aggregate management operational, and Lucas demonstrates why the operator changes the system being operated. Digital economics explains why ideas and networks create increasing returns, while organizational research shows that the return depends on complements that conventional accounts often classify as expense. AI economics now measures striking capabilities and bounded productivity gains, but the causal path from a capable model to wages, employment, firm structure, or aggregate growth remains incomplete.

Three conclusions follow.

First, the object of economics has widened faster than any one method can absorb. Choice, equilibrium, institutions, power, money, expectations, networks, tasks, organizations, and cognition are not rival decorations around a fixed core. They are different layers of the system. A useful model earns its abstraction by naming which layers it holds fixed.

Second, the most consequential variable is often the one both sides treat as given. In Malthus it is behavioral response; in Wicksell the neutral rate; in Lucas the policy-dependent coefficient; in the digital firm the organizational complement; in AI the translation from task productivity into staffing, quality, demand, and new work. The disagreement begins as a contest over outcomes and ends as a question about an endogenous parameter.

Third, modern evidence becomes stronger when it becomes more modest about scope. A randomized experiment can identify a bounded effect more credibly than a grand historical analogy, yet it may say less about equilibrium adjustment. Administrative data can reveal realized outcomes, yet only for the covered population and time window. A structural model can connect levels, but its aggregate result inherits its calibration. Rigor does not come from forcing all evidence into one scale. It comes from preserving what each design can and cannot establish.

This is also the reason for four different interactive stories. A repeated timeline would flatten the argument. The reader should restore a missing object in Epoch I, make a decision in Epoch II, trace a dependency in Epoch III, and judge an evidence claim in Epoch IV. Those are not interface themes. They are the intellectual actions each literature requires.

The current AI debate is tempting precisely because its first link is vivid. A system writes, codes, recommends, or answers. The conclusion appears close enough to touch. Yet the economic outcome sits several gates away, behind adoption, actual use, reliability, workflow redesign, demand, ownership, bargaining, and institutional response. Technical progress shortens the

distance between possibility and demonstration. It does not abolish the distance between demonstration and consequence.

Economics, in other words, keeps evolving because the economy keeps revealing that yesterday's residual was a mechanism. The inconvenient part is that the mechanism usually becomes visible only after somebody has already called the model complete.

Primary bibliography

Epoch 1 · 1776–1936

- Adam Smith (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. Two-volume political-economy treatise; linked Cannan edition is based on Smith's 1789 final lifetime edition. [Source / edition record](#).
- Thomas Robert Malthus (1798). *An Essay on the Principle of Population, as It Affects the Future Improvement of Society. With Remarks on the Speculations of Mr. Godwin, M. Condorcet, and Other Writers*. Book-length polemical essay; first edition published anonymously. [Source / edition record](#).
- David Ricardo (1817). *On the Principles of Political Economy and Taxation*. Systematic treatise; linked Sraffa–Dobb critical edition reflects a work revised after 1817. [Source / edition record](#).
- Karl Marx (1867). *Das Kapital. Kritik der politischen Oekonomie. Erster Band: Buch I. Der Produktionsprozess des Kapitals*. Critical treatise, Volume I; the only volume published in Marx's lifetime. [Source / edition record](#).
- Carl Menger (1871). *Grundsätze der Volkswirtschaftslehre*. Foundational monograph; linked source contains the original German text and facsimile. [Source / edition record](#).
- Marie-Esprit-Léon Walras (1874–1877). *Éléments d'économie politique pure, ou Théorie de la richesse sociale*. Mathematical treatise issued in two instalments, 1874 and 1877. [Source / edition record](#).
- Alfred Marshall (1890). *Principles of Economics*. Textbook-treatise; linked source is the 1890 first edition. [Source / edition record](#).
- Johan Gustaf Knut Wicksell (1898). *Geldzins und Güterpreise: Eine Studie über die den Tauschwert des Geldes bestimmenden Ursachen*. Monetary monograph; linked source is the original German edition. [Source / edition record](#).
- Arthur Cecil Pigou (1920). *The Economics of Welfare*. Welfare-economics treatise; linked source is the 1920 first edition. [Source / edition record](#).
- John Maynard Keynes (1936). *The General Theory of Employment, Interest and Money*. Book-length macroeconomic treatise; linked authorized reissue records the original 1936 Macmillan publication. [Source / edition record](#).

Epoch 2 · 1936–1980

- John R. Hicks (1937). *Mr. Keynes and the 'Classics'; A Suggested Interpretation*. Journal article · *Econometrica* 5(2), 147–159. [Source / edition record](#).
- Jan Tinbergen (1939). *Statistical Testing of Business-Cycle Theories. II: Business Cycles in the United States of America, 1919–1932*. League of Nations Economic Intelligence Service monograph, 244 pages. [Source / edition record](#).
- Trygve Haavelmo (1944). *The Probability Approach in Econometrics*. Monograph supplement · *Econometrica* 12, 1–115. [Source / edition record](#).
- Friedrich A. Hayek (1945). *The Use of Knowledge in Society*. Journal article · *American Economic Review* 35(4), 519–530. [Source / edition record](#).
- Kenneth J. Arrow and Gérard Debreu (1954). *Existence of an Equilibrium for a Competitive Economy*. Journal article · *Econometrica* 22(3), 265–290. [Source / edition record](#).

- Robert M. Solow (1956). *A Contribution to the Theory of Economic Growth*. Journal article · Quarterly Journal of Economics 70(1), 65–94. [Source / edition record](#).
- A. W. Phillips (1958). *The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957*. Empirical article · Economica 25(100), 283–299. [Source / edition record](#).
- Milton Friedman (1968). *The Role of Monetary Policy*. Presidential address · American Economic Review 58(1), 1–17. [Source / edition record](#).
- George A. Akerlof (1970). *The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism*. Journal article · Quarterly Journal of Economics 84(3), 488–500. [Source / edition record](#).
- Robert E. Lucas Jr. (1976). *Econometric Policy Evaluation: A Critique*. Journal article · Carnegie-Rochester Conference Series on Public Policy 1, 19–46. [Source / edition record](#)80003-6).

Epoch 3 · 1980–2015

- Michael L. Katz and Carl Shapiro (1985). *Network Externalities, Competition, and Compatibility*. American Economic Review 75(3), 424–440; journal article. [Source / edition record](#).
- Paul M. Romer (1986). *Increasing Returns and Long-Run Growth*. Journal of Political Economy 94(5), 1002–1037; journal article. [Source / edition record](#).
- W. Brian Arthur (1989). *Competing Technologies, Increasing Returns, and Lock-In by Historical Events*. The Economic Journal 99(394), 116–131; journal article. [Source / edition record](#).
- Paul M. Romer (1990). *Endogenous Technological Change*. Journal of Political Economy 98(5, Part 2), S71–S102; journal article. [Source / edition record](#).
- Erik Brynjolfsson (1993). *The Productivity Paradox of Information Technology*. Communications of the ACM 36(12), 66–77; research synthesis. [Source / edition record](#).
- Carl Shapiro and Hal R. Varian (1998). *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press; book. [Source / edition record](#).
- Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt (2002). *Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence*. The Quarterly Journal of Economics 117(1), 339–376; empirical journal article. [Source / edition record](#).
- Jean-Charles Rochet and Jean Tirole (2003). *Platform Competition in Two-Sided Markets*. Journal of the European Economic Association 1(4), 990–1029; journal article. [Source / edition record](#).
- David H. Autor, Frank Levy, and Richard J. Murnane (2003). *The Skill Content of Recent Technological Change: An Empirical Exploration*. The Quarterly Journal of Economics 118(4), 1279–1333; theoretical and empirical journal article. [Source / edition record](#).
- Carol Corrado, Charles Hulten, and Daniel Sichel (2005). *Measuring Capital and Technology: An Expanded Framework*. Chapter in *Measuring Capital in the New Economy*, University of Chicago Press/NBER; capital-measurement framework. [Source / edition record](#).

Epoch 4 · 2015–present

- David H. Autor (2015). *Why Are There Still So Many Jobs? The History and Future of Workplace Automation*. Journal of Economic Perspectives 29(3), synthesis and interpretive essay. [Source / edition record](#).

- Daron Acemoglu and Pascual Restrepo (2018). *The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment*. American Economic Review 108(6), structural growth-and-tasks theory. [Source / edition record](#).
- Daron Acemoglu and Pascual Restrepo (2020). *Robots and Jobs: Evidence from US Labor Markets*. Journal of Political Economy 128(6), local-labor-market empirical study. [Source / edition record](#).
- Erik Brynjolfsson, Daniel Rock, and Chad Syverson (2021). *The Productivity J-Curve: How Intangibles Complement General Purpose Technologies*. American Economic Journal: Macroeconomics 13(1), theory plus accounting application. [Source / edition record](#).
- Shakked Noy and Whitney Zhang (2023). *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*. Science 381(6654), preregistered randomized online experiment. [Source / edition record](#).
- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock (2024). *GPTs are GPTs: Labor Market Impact Potential of LLMs*. Science 384(6702), task-exposure framework and Policy Forum. [Source / edition record](#).
- Tania Babina, Anastassia Fedyk, Alex He, and James Hodson (2024). *Artificial Intelligence, Firm Growth, and Product Innovation*. Journal of Financial Economics 151, firm-level panel study. [Source / edition record](#).
- Daron Acemoglu (2025). *The Simple Macroeconomics of AI*. Economic Policy 40(121), task-based macro model and calibration. [Source / edition record](#).
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond (2025). *Generative AI at Work*. Quarterly Journal of Economics 140(2), staggered workplace deployment study. [Source / edition record](#).
- Anders Humlum and Emilie Vestergaard (2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. NBER Working Paper 33777, May 2025, revised March 2026; linked surveys and administrative data. [Source / edition record](#).

Bibliographic and publication notes

- First-publication dates are kept separate from the digital edition or publisher page linked in the package.
- Source links prioritize DOI records, publishers, institutional repositories, archival scans, author pages, and official working-paper series.
- Titles preserve the publication form and language recorded in the project data. Typographic quotation marks may differ across catalog systems.
- The forty-work list is a curated chronological canon, not a citation leaderboard or claim of universal disciplinary consensus.
- “Influence” describes intellectual reach, not empirical validation or endorsement.
- Contemporary AI findings are dated evidence. Benchmark capability, adoption, workplace use, labor-market outcomes, and macroeconomic calibration are kept analytically separate.
- Generated interactive scenarios are labeled illustrative. They are not estimates, forecasts, financial advice, or policy prescriptions.