Firms: Natural Units of Technology Augmentation

Gary H. Jefferson
Brandeis University
jefferson@brandeis.edu

August 4, 2026

DRAFT: FOR REVIEW AND COMMENT ONLY

Submitted to the NBER Conference on Firms and Technology,
September 25, 2026, Cambridge MA

Abstract

Firms are the locus of technology adoption, and this paper argues that they are also the natural unit in which technology augmentation should be counted. The argument begins from a measurement result. The direction of technical change –whether it augments capital, labor, or both – is not determined by the national accounts. The same path of output, factor payments and factor shares admits Hicks-, Harrod- and Solow-neutral readings, agreeing on every value the accounts record and differing only in how each factor's value is split into a quantity and a unit price. That split is settled by deflator convention rather than recovered from data, and the accounts make it once and asymmetrically: capital arrives pre-augmented, its quality change absorbed into a deflated stock, while labor arrives raw, counted in bodies and hours. That pairing *is* Harrod-neutrality, and Samuelson's (1965) "factual assumption" – that capital in natural units accumulates relative to labor in natural units – is stated in units of two different kinds.

We define natural units by three properties and show that firms and households satisfy them while deflated value aggregates do not. This yields a hierarchy: *technology systems* (machines, workers) are the loci at which augmentation is realized; *economic systems* (firms, households) are the agents that direct it, mobilize the investment it requires, and bear its cost and return. Firms and households exist because the corresponding contracts cannot be written – configuration-specific improvement for the firm, capability-specific improvement for the household – and the same non-contractibility determines what the accounts can see. Capital augmentation is largely embodied in purchased goods and therefore priced; labor augmentation is largely internal and therefore not. An allocation of augmentation investment across the two margins by agents responding to factor prices converges to equalized returns, making balanced augmentation an attractor rather than a knife-edge. The framework generates an institutional designation of the production function, $Y = F(A_K n_F, A_L n_H)$, in which capital is counted in firms and labor in households. This decomposes growth into an intensive margin – better firms, better households – and an extensive margin, the counts themselves, which are directly observable where the augmentation indices are not. We execute a test on U.S. data, 1989-2019.

1. Introduction: the question being studied

1.1 Samuelson's fork

Samuelson (1965, p. 350) embraces what he calls a factual assumption: that capital, in *natural units*, is accumulating relative to labor in natural units. Given Samuelson's characterization of dollars and workers and the resulting ratios in the steady state – *K/Y* fixed and *K/L* and *w/r* rising – he asserted (p. 348):

> Unless the production function of that model is Cobb-Douglas, we know that realism requires that the technical change be Harrod-neutral and nothing else.

Hence, any credible alternative to the established Harrod must attend to these three ratio as well as to the well-established neoclassical assumption of the natural units of capital and labor on which they are founded.

Samuelson's "factual assumption" is doing heavy work. It is the empirical premise behind the presumption of capital deepening, behind the stationary capital–output ratio that Uzawa's (1961) theorem converts into purely labor-augmenting technical change, and behind a research program that has treated Harrod-neutrality as a finding rather than a convention.

The term *natural units* is not defined in that paper, nor, so far as we have found, anywhere else in the literature on factor augmentation. This paper supplies a definition and finds that the two sides of Samuelson's comparison are not measured in units of the same kind. Labor is counted in bodies and hours – referents fixed independently of how productively the labor is used. Capital, as the accounts supply it, is a deflated value aggregate: a unit of constant-price capital is defined as the quantity of equipment equivalent in productive service to a base-year dollar's worth. That individuation is by efficiency in performing the function rather than by the function itself, so improvement raises the count by construction.

The consequence is direct. To whatever extent capital is measured in the second way, the factual assumption holds in part by construction rather than by observation. The rise in measured K/L is not evidence about the technology; it is partly the deflator doing its job.

1.2 The indeterminacy

The result behind this is general. Where a quantity is observed only as a product – the *engine* $A_K K$ – any statement about one factor separately is a statement about the chosen decomposition. Let $\phi_K(t)$, $\phi_L(t)$ be positive and differentiable and set

$$K' = \phi_K K, \quad A_K{}' = A_K/\phi_K, \qquad L' = \phi_L L, \quad A_L{}' = A_L/\phi_L.$$

The redesignated system generates an identical path of output, engines, marginal products with respect to the engines, factor payments and factor shares, while $g_{A_K{}'} = g_{A_K} - g_{\phi_K}$. For any target pair of augmentation rates there exist ϕ_K, ϕ_L attaining it. Hicks-, Harrod- and Solow-neutrality are therefore not three hypotheses among which evidence might select; they are three coordinates in which one path can be written.

Four objects are invariant across all designations: output, the two engines, and the factor shares. These are the identified content of the technology. What is *not* invariant is the measured capital–output ratio, the measured capital–labor ratio, and the classification itself – and the standard argument against symmetric factor augmentation runs entirely through the non-invariant objects.

1.3 Inside and outside anatomy

The distinction that organizes the paper is Arthur's (2009). Viewing technologies from the outside is adequate for asking how they enter an economy and spread within it, but not for asking what they are and how they originate, for which the internal anatomy has to be opened up. The national accounts are the outside view: what was purchased, at what price, growing how fast.

The direction of technical change is an inside-anatomy question, because it concerns how an observed value divides between a quantity of instruments and an index of their quality.
The firm is where the two views meet. It is the entity the accounts can see – counted, registered, filing – and it is also the entity within which the internal reconfiguration happens. That is why it is the natural unit for this question, and why a firm-level framing is not an application of the measurement argument but the form the argument takes.

2. What the evidence establishes, and what it does not

Four bodies of evidence are usually taken to settle the direction of technical change. Each constrains the *product* $A_K K$; none resolves its *composition*.

Solow (1960) and embodiment. Solow's vintage framework makes the point that inventions are useless until they are embodied in new capital goods. This is a statement about the capital side of the technology being where improvement is realized – which is an argument *for* capital augmentation being present, not against it. Jorgenson (1966) then established that embodied and disembodied technical change are observationally equivalent, precisely because the allocation between "more efficiency units per machine" and "higher efficiency index on a given machine" is settled by the deflator rather than discovered. The embodiment literature therefore identifies the mechanism and simultaneously demonstrates that the accounts cannot resolve it.

Greenwood, Hercowitz and Krusell (1997). The falling relative price of equipment is the strongest evidence for investment-specific technical change. It is fully consistent with either reading: the improvement may be booked as an increase in the number of efficiency units embodied in the measured stock, or as an increase in a capital-augmenting efficiency index applied to an unchanged stock. Evidence of this kind constrains the product; it does not resolve the composition. Its importance for this paper is different and larger: the wedge between the return on saving and the rental per machine, on the order of the investment-specific rate, is what prevents the Euler equation from pinning the capital unit and thereby making the Harrod classification substantive.

The facts of growth. The stylized regularities – stationary K^m/Y, stationary r, rising K^m/L^h, constant shares – contain fewer independent facts than they appear to. With constant shares, $s_X = p_X X/Y$, so a stationary ratio and a stationary price are the same statement. And for $\sigma \neq 1$, constant shares, a constant engine ratio, and stationarity of both engine–output ratios are equivalent. The invariant content of the first two Kaldor facts is one fact, and symmetric augmentation satisfies it. What symmetric augmentation does not deliver is stationarity of K/Y in natural units –which was never the invariant fact.

Solow (1957) and the residual. Growth accounting decomposes $g_Y = s_K g_{A_K K} + s_L g_{A_L L}$, an exact Divisia identity that holds for every constant-returns technology. An expression valid

across the whole class cannot discriminate within it. The residual measures the joint movement of the engines; the attribution of that movement to one factor rests on the maintained hypothesis, not on the arithmetic.

Implication. Neither the composition of technical change nor the curvature of the technology is recoverable from balanced-path data (Diamond, McFadden and Rodriguez, 1978). Neutrality is therefore not an empirical finding in any growth model. It is a maintained hypothesis, and every model contains one. The relevant question is which maintained hypothesis is better justified, and that question is settled on theoretical coherence and on the observable implications each generates — which is the business of Sections 5 through 7.

---

3. Properties of the natural unit

A **natural unit** is a system individuated by a fixed referent – the function performed, in the case of artifacts and institutions; personhood or time, in the case of labor; area, in the case of land – counted by a rule that may consult that referent and, where necessary, its scale, but never its efficiency.

N1 — Fixed referent. *The unit is individuated by a referent that does not move with productive performance, specified at a stated level of detail and held fixed thereafter.* Within that referent, form and interior composition may both change: a workshop re-tooled and enlarged remains the same workshop, and a worker who retrains remains the same worker. The two cases invert — the machine holds its function and varies its form, the worker holds its form and varies its function — and both are augmentation, because neither event moves the count. Where the referent itself is replaced, the unit has been replaced rather than improved.

N2 — Interior variation. *All variation in productive capacity is interior.* The unit's interior is a population of elements, constant in typology and variable in configuration, and augmentation consists in recombining them into arrangements that perform the fixed function at a higher rate (Weitzman, 1998; Arthur, 2009). Improvement changes the arrangement, not the number of units.

N3 — Counting is by referent, not by efficiency. *Membership in the count is determined by the referent of N1 — what the unit is for, who it is, how long it lasts, how much ground it covers — at a stated level of detail and scale, and not by how well the unit performs.* A

rule may set a scale threshold, provided the threshold is fixed and not itself productivity-denominated. Classification of this kind is standard in the statistical systems themselves – asset taxonomies, industrial and occupational classifications – and is not performance-graded. Given N2, a count so determined cannot register improvement, so $D_X \equiv 1$: all technical change in factor X is carried by A_X.

Aggregation. Any factor quantity aggregates heterogeneous units and so requires weights. The aggregation is admissible where improvement raises the efficiency of the class's members in common, so that a single index applies and factors out of the aggregate. What disqualifies a class is divergence in rates rather than dispersion in levels: persons and machines are heterogeneous and this is no obstacle, whereas an aggregate spanning distinct substrates admits no common index.

Exclusion. Deflated value aggregates fail N3, and fail it uniquely among the measures in common use. The national accounts classify assets functionally and then measure their quantity by deflated value; the two principles are not the same, and only the first is performance-blind. Writing $K^m = D_K K$, stationarity of K^m/Y is equivalent to the conjunction of engine balance with $g_{D_K} = g_{A_K}$ — and the second conjunct is a statement about the deflator, not about technology. Requiring a constant measured capital–output ratio of a growth model therefore imposes $g_{A_K} = 0$ in the measured designation: the premise contains the conclusion.

---

4. Hierarchical outside–inside levels of technical change

The criteria admit two kinds of system, and the argument needs both.

	Technology systems	**Economic systems**
Instances	machines, workers, acres	firms, households
Individuated by	function; personhood or time; area	contractual boundary; residence and shared budget
Function	perform a job, at improving rates	integrate and direct augmentation across subsystems
Interior	one substrate, recombined	technology systems of several substrates

	Technology systems	Economic systems
Role	locus: where augmentation is realized	locus and agent: allocates augmentation investment
Commonly augmentable in	its substrate	its coordinating function

Technology systems are the counted units whose interiors recombine and in which augmentation is realized. They are arguments of the production function, and the aggregation condition applies to them within a substrate.

Economic systems are bounded systems that direct the augmentation of the technology systems they contain, mobilize the investment it requires, and bear its cost and capture its return. They satisfy N1–N3 and are natural units, and they are commonly augmentable in the function that individuates them: firms as a class improve at coordinating production, which is what the management-practice and organizational-capital literatures measure. They can therefore be arguments of F, and Section 7 designates them so.

What separates the two categories is not eligibility but role. A technology system is where augmentation is realized; an economic system additionally chooses it, finances it, and bears its return. Machines do not allocate. The integration across substrates that brings the firm into existence is therefore not a disqualification but a determination of what the firm's index measures: **the firm is the level at which the direction of technical change is chosen rather than merely recorded.**

The two are alternative designations in the sense of Proposition 1, not levels of one hierarchy, and the augmentation index means something different in each — machine quality where the count is machines, coordinating efficiency where the count is firms. Normalizing both to unity at a base year, they diverge thereafter at exactly the rate of capital deepening per firm. Capital per firm is not constant, so the two designations bracket the technological rate rather than agreeing on it.

5. Firms and households as the agents of augmentation

5.1 Why the institutions exist

Both factor augmentations fail the conditions for market provision, and for the same reason: the improvement is inseparable from the asset it improves, and the asset cannot be sold.
Human capital is embodied in a person and cannot be collateralized, so no lender finances the investment against the returns it generates; the household undertakes it because the household internalizes the return (Becker, 1964). Process improvement is specific to a configuration of equipment, workers and routines, so a contract with an outside supplier is incomplete and exposed to hold-up; the firm undertakes it because the firm internalizes the return (Coase, 1937). The two are not analogous institutions but two margins of one problem, divided by two distinct contracting failures.

The role is active rather than custodial. Each institution is a bounded system holding technology systems of both kinds, and its function is to integrate them: to direct the augmentation of its several subsystems along a coordinated vector, to mobilize the investment that requires, and to bear the cost and capture the return.

Neither institution is confined to one margin. The firm invests in its workforce as well as in its equipment; the household holds claims on the capital stock as well as investing in capability. What the two contracting failures determine is not which factor each institution augments but which portion of each margin each is able to capture.

5.2 Why the accounts see one margin and not the other

Measurement absorbs change into a quantity only where the change crosses a transaction boundary. What is not transacted has no price to be adjusted, and remains outside the measured quantity.

Coase's internalization is partial. The firm buys the machine on a market; that transaction has a price, the price index is quality-adjusted, and the embodied component of capital augmentation therefore enters D_K. What the firm internalizes is the complementary configuration investment, which remains outside the measured quantity. Becker's internalization is near-total: child-rearing time, household investment and on-the-job learning cross no boundary, so there is no transaction to quality-adjust and D_L captures almost nothing. Formal schooling is the exception precisely because it is institutionally separated and priced.

Proposition. $g_{D_K} > g_{D_L}$ follows from the position of the transaction boundary rather than from the construction of any particular index. The gap is not closed by improved statistics: the non-transacted component is invisible in principle, not merely in practice.

This is a stronger foundation than the observation that equipment deflators are hedonic while human-capital indices are thin. That observation is contingent and invites the reply "so build a better index." The claim here is that a better index cannot reach the non-transacted portion — which makes the Harrod designation a structural feature of measurement rather than a defect of current practice.

5.3 Balanced augmentation as an attractor

Let a fraction sY of output be allocated between augmenting the two subsystems in shares θ and $1-\theta$. The elasticity of output with respect to each engine is its factor share, so the marginal return to augmentation investment in factor i is proportional to $s_i Y/(A_i X_i)$. Equalizing returns gives $\theta^* = s_K$. Writing $x \equiv A_K K/A_L L$,

$$g_x = \frac{sY\, s_L}{A_L L}\left[\frac{\alpha}{1-\alpha}x^{\rho-1} - 1\right].$$

Since $\rho < 1$ the bracket is strictly decreasing in x: the system has a unique interior rest point and converges to it globally. Two forces produce this — diminishing returns to augmenting an already-large engine, and the share response, since a lagging engine's share rises and with it the return to relieving it. The result does not require gross complementarity.

The mechanism is decentralized, which is what makes it a claim about firms and households rather than about a planner. Neither institution observes the other's allocation, and neither needs to: the factor price of the subsystem each augments is a sufficient statistic for relative scarcity, and it is public. The firm sees the return to effective capital; the household sees the wage on effective labor.

The measure-zero objection reverses. Symmetric augmentation is standardly charged with requiring exact and permanent equality of two independently determined rates. The rates are not independently determined: they are jointly set by one allocation responding to relative scarcity, and equality is where that allocation settles. Meanwhile $\theta^* = s_K$ is interior whenever $s_K \in (0,1)$, while the Harrod restriction requires $\theta = 0$ – no augmentation investment in a factor

earning roughly a third of output, permanently. That is the corner an allocation responding to returns never selects. The knife-edge is Harrod's.

6. Illustration: Artificial Intelligence

AI is the case in which the framework's predictions are currently being generated, and it separates the two margins unusually cleanly.

The firm's margin. A firm adopting AI purchases compute, models and licensed services. These are transacted, they have prices, and the price indices for them are – increasingly – quality-adjusted. The capability improvement is therefore absorbed into measured capital: more efficiency units per dollar of equipment and software. What the firm additionally invests in is the reconfiguration that makes the purchased capability productive – workflow redesign, data infrastructure, retraining, reorganization of task boundaries. That investment is configuration-specific, is not separately transacted, and is Coase's margin. It remains outside the measured quantity, i.e., the firm.

The household's margin. Workers acquiring facility with AI tools do so largely through use, informal instruction, and on-the-job learning. Almost none of this crosses a transaction boundary. It is Becker's margin, and D_L captures essentially none of it.

The prediction. Measured capital deepening will accelerate – and much of the acceleration will be the deflator recording capability improvement as quantity. Measured labor quality will move very little, because the index has no channel through which to see the corresponding capability. Under the framework, the two augmentation margins are being invested in jointly by agents equalizing returns; under the accounting convention, one is visible and one is not. The AI episode is therefore predicted to *widen* the gap between measured and natural-unit factor ratios, and to do so without any change in the underlying symmetry of the technology.

This yields a testable contrast at the firm level. If AI-adopting firms are simultaneously raising both margins – capital-embodied capability and workforce capability – then the observed division of augmentation investment should track $\theta^* = s_K$ rather than concentrating on the capital side alone. Firm-level intangible investment data, and the growing body of AI-adoption survey evidence, bear on this directly.

7. The Hick-Harrod Test and Result

7.1 The institutional designation

Sections 3 and 4 identify firms and households as natural units – economic systems, individuated by contractual boundary and by residence and shared budget, neither of which moves with productivity. Designating the factors in these units gives $Y = F(A_K n_F, A_L n_H)$, where n_F is the number of firms and n_H the number of households, and A_K and A_L are the augmentation indices those institutions carry. This is an admissible designation under Proposition 1: output, the engines, factor payments and factor shares are unchanged. What has changed is which physical thing counts as one unit of each factor.

Two features make this designation the one to test in. Its counting rule is fixed outside the accounts, so – unlike the deflated capital stock – it is not constructed to deliver any particular classification. And the counts are maintained by the statistical system for reasons having nothing to do with growth accounting, so they are available without a deflator, an index, or an assumption about σ. The augmentation indices remain latent. The counts are not. The test below runs entirely through the counts.

Under constant returns the designation can equivalently be written $Y = A\,F(n_F, n_H)$, since $F(An_F, An_H) = A\,F(n_F, n_H)$. Symmetric factor augmentation and a Hicks-neutral shift are one hypothesis written twice, and the null tested below is exactly the claim that the aggregate technology admits that representation. The augmentation index in this designation is the efficiency with which economic units organize production – coordinating and organizational efficiency — rather than the quality of machines. The two coincide only if capital per firm is constant; it is not, so the two designations bracket the technological rate rather than agreeing on it.

7.2 The two hypotheses, in observables

Under constant returns and competitive payment, constant factor shares are equivalent to a constant engine ratio, so along a steady state $g_{AK} + g_{nF} = g_{AL} + g_{nH} = g_Y$. Subtracting, $g_{AK} - g_{AL} = g_{nH} - g_{nF} \equiv d$. The augmentation gap equals the gap between household and firm formation. Both terms on the right are observed. This is the whole of the identification strategy: the difference of two latent indices is exactly a difference of two counts.

Hicks. $g_{AK} = g_{AL}$ implies $d = 0$: firms and households must form at a common rate. The common augmentation rate is then $g_A = g_Y - g_{nF}$.

Harrod. $g_{AK} = 0$ implies $g_{nF} = g_Y$ – with capital counted in firms, the firm population must grow at the rate of output – and therefore $d^{\text{Harrod}} = g_{nH} - g_Y$, with $g_{AL} = g_Y - g_{nH}$.

Harrod is not the complement of the null. It is a point prediction, located a determinate distance from zero, and the data discriminate between the two only to the extent that distance is large relative to the standard error. Reporting the null alone would leave the test's power unstated. Uzawa's theorem is not violated by the Hicks representation here: the theorem's premise is stationarity of the capital –output ratio in value units, which Section 3 shows to be a property of the deflator rather than of the technology.

7.3 Data

The test requires three series and no augmentation index.

Firms. U.S. Census Bureau employer-firm counts. The employer-firm register is roughly an order of magnitude more accurate than the OECD business-register sources used in the companion cross-country panel, which is why the U.S. case is reported separately and given the weight it is.

Households. FRED total household counts.

Output and capital. Penn World Table 11.0: real GDP at constant national prices (rgdpna) and the capital stock at constant national prices (rnna), the latter used only for the comparison designation of Table 7C.

Factor shares. FRED labor income share, on the labor-income-over-GDP construction that is internally consistent with $Y \equiv$ GDP. Measures constructed over national income are not consistent with that denominator and move the estimates materially; the choice is made explicitly rather than inherited.

Window selection. Quasi-steady-state windows are identified by near-constancy of the factor-share ratio. Two appear in the U.S. series: 1988–2003 and 2009–2019, separated by a decade and by the financial crisis. Treating them separately provides two effectively independent episodes. Both exhibit small residual drift in the share ratio (−0.51 and −0.13 pp/yr), so they are quasi- rather than strict steady states; the d-statistic carries the observed drift explicitly.

7.4 Result

Table 7B. Three tests of neutrality against one statistic, $d = g_{AK} - g_{AL}$, in percentage points per year. Rows 1 and 2 report the same observed statistic against two different point predictions. U.S. quasi-steady-state windows.

Test — units (capital / labor)	Predicted d	1988–2003	2009–2019	Pooled	Pooled deviation
Hicks — firms / households	0	+0.32	+0.31	**+0.31**	1.0 SE, p = 0.32
Harrod — firms / households	−1.56	+0.32	+0.31	**+0.31**	**6.0 SE**, $p <$ 0.01
Hicks, robustness — firms / employment	0	−0.04	+0.80	+0.29	—
Memo: accounts' units — deflated stock / employment	—	−1.76	−0.14	−1.11	reproduces Harrod

Newey–West (lag 2) standard errors on the observed statistic: 0.91, 0.33 and 0.31 for the two windows and pooled. Underlying growth rates in Table 7C. The memo row is not a test but the same statistic computed in the units of the national accounts.

Two implied magnitudes follow. Under Harrod the labor-augmentation rate that would have to absorb the difference is $g_{AL} = 1.56$ percent a year; under Hicks the common rate is $g_A = 1.88$. And Harrod requires the firm population to grow at the rate of output: 2.98 against 1.01 observed in the first window, 2.37 against 0.62 in the second — a factor of 2.9 and 3.8.

The two hypotheses are not equally supported. The observed statistic sits one standard error from the Hicks prediction and six from the Harrod prediction. Hicks-neutrality is not rejected in either window or in the pooled sample; Harrod-neutrality is rejected at better than one percent in the second window and in the pooled sample, and the point estimate is on the opposite side of zero from the Harrod prediction in both.

The stability of the estimate is worth noting independently of the test. Two windows separated by a decade, a financial crisis and a substantial change in the composition of the firm population return point estimates differing by one one-hundredth of a percentage point.

The result stated without a statistic. For capital-augmenting technical change to be absent, the firm population would have to grow at the rate of output. Over 1988–2003 output grew at 3.0 percent a year and employer firms at 1.0 percent; over 2009–2019, 2.4 and 0.6. Harrod-neutrality requires the number of U.S. firms to have roughly tripled over a period in

which it grew by about half. The implied labor-augmentation rate needed to absorb the difference, 1.6 percent a year, is not implausible in itself – what is implausible is the firm-count path it must be paired with.

7.5 What the result establishes

Harrod-neutrality is rejected in a designation that was not built to reject it. This is the substance of the finding. The measured capital stock is constructed so that the Harrod classification follows – the deflator's function is to absorb quality change into the quantity, which sets $g_{AK} = 0$ by construction. Firm counts are not constructed that way; the register exists for administrative reasons and is blind to productivity. In units the accounts did not shape, the restriction fails.

Hicks-neutrality survives, but the test's power against it is limited. The pooled interval is roughly [−0.30, +0.92]. It excludes the Harrod prediction comfortably and would fail to exclude a range of small capital-augmenting biases. The correct statement is that the data are consistent with symmetric augmentation and inconsistent with the received restriction, not that symmetry has been confirmed.

The two readings of the same period. Over 1988–2019 the measured capital–labor ratio rose steadily, which is the observation Samuelson's factual assumption rests on. Over the same period, factors counted in the institutions that augment them show an augmentation gap of a third of a percentage point per year. Both statements describe one economy. They differ in what was counted as a unit.

What the test does not establish is the level of augmentation on either margin separately: it constrains the difference $g_{AK} - g_{AL}$ and is silent on the common component. The latency of the augmentation indices established in Sections 2 and 5 is not removed by this design – it is the reason the design works through counts instead.

7.6 Three designations, one economy

Harrod-neutrality is a restriction on the capital side alone, $g_{AK} = 0$. It makes no reference to how labor is counted, and it is therefore computed once. Hicks-neutrality is a restriction relating the two sides, $g_{AK} = g_{AL}$, and so depends on both units. The table below computes each restriction where it lives, and varies only what is being tested. Under engine balance $g_{AX} = g_Y - g_X$ for whichever unit X designates; every entry is that subtraction, and nothing else is estimated.

Table 7C. Growth rates by unit of measurement, percent per year. Augmentation rates are g_Y less the relevant row: capital augmentation is 0.48 percent a year counted as a deflated stock and 1.88 counted as firms; labor augmentation is 1.59 counted as persons and 1.56 counted as households. Output and capital from PWT 11.0.

	1988–2003	2009–2019	Pooled
Output, g_Y	2.98	2.37	2.73
Capital as deflated stock, g_{Km}	2.73	1.55	2.26
Capital as firms, g_{nF}	1.01	0.62	0.85
Labor as persons employed, g_L	0.97	1.41	1.15
Labor as households, g_{nH}	1.33	0.93	1.17
Capital share, s_K	0.381	0.410	0.393
Observations	15	10	25

The pooled column is an observation-weighted average of the two windows rather than a growth rate over 1988–2019; the intervening years are excluded by the window criterion and do not enter it.

The capital rows are the finding. Measured as a deflated stock, capital augmentation runs at a quarter of a percentage point a year in the first window and half a point pooled – near enough to zero that Harrod-neutrality appears to hold. Measured as firms, the same fifteen and ten years give 1.97 and 1.88. The technology has not changed between the two rows; the unit has. And the first row is not independent evidence for Harrod, because stationarity of K^{m}/Y is engine balance together with zero residual capital augmentation, the second conjunct being a property of the deflator. The accounts return approximately zero because they are constructed to.

The labor rows show that the choice of labor unit is second-order. Households and persons employed give 1.56 and 1.59 pooled — a difference of three hundredths of a point. Both are natural units in the sense of Section 3: neither individuation consults productivity, and no deflator stands between the worker and the count. The asymmetry predicted in Section 5.2 is visible here as an absence. On the labor side the choice of unit barely matters; on the capital side it moves the answer by a factor of four. That is the transaction boundary appearing in the arithmetic.

The bias rows separate the two hypotheses. With capital counted in firms, d is +0.31 against households and +0.29 against employment — the same answer from two labor series, and within one standard error of the Hicks restriction in both cases. With capital counted as a

deflated stock, d is −1.11, which is Harrod's neighborhood. The three rows do not disagree about the economy. They disagree about what counts as one unit of capital.

Growth accounting in the institutional designation. Writing $Y = A\,F(n_F, n_H)$ and decomposing, $g_A = g_Y - s_K g_{nF} - s_L g_{nH}$ gives 1.77, 1.56 and 1.69 percent a year across the two windows and pooled — against standard TFP of 1.34, 0.90 and 1.16 computed with the deflated stock and employment. The difference of roughly half a point a year, widening from 0.43 to 0.67 between the windows, is the quality change the deflators move out of the residual and into the measured quantities. Standard growth accounting does not find less technical change; it finds the same technical change with part of it relabeled as capital.

Two qualifications. The 2009–2019 firms/employment entry of +0.80 reflects employment rebounding from a deep trough rather than a change in augmentation, and the window's start date should be tested for sensitivity. And the deflated-stock row weakens in the second window, where capital augmentation reaches 0.81 – measured K/Y fell over the post-crisis decade, so even in the accounts' own units the Harrod restriction sits less comfortably after 2009 than before.

7.7 Extensions

Two further tests follow from the same framework and are reported in companion work rather than here.

The allocation restriction. $\theta^* = s_K$ is a cross-equation restriction in observables: the division of augmentation investment across the two subsystems should equal the capital share. Capital-side measurement is partly available — equipment and software investment, R&D, intangibles; household-side measurement is largely not, for the reason given in Section 5.2. The shortfall is itself an estimate of the non-transacted margin rather than a residual to be discarded.

The quality-accounting gap. Under symmetric augmentation with common natural-unit factor growth, $g_{Km/Lm} = g_{DK} - g_{DL}$. Symmetric augmentation is falsified if the documented quality-accounting gap is materially larger or smaller than observed growth in measured K/L. Two complications govern execution: structures dominate the aggregate capital stock and receive minimal quality adjustment, so an equipment-only version is likely the informative one; and the appropriate g_{DL} depends on how much true human-capital quality growth the schooling index omits, which by Section 5.2 is substantial and structural.

Extending the U.S. window backward. The employer-firm register begins in 1988. Reaching 1950 requires splicing IRS Statistics of Income business-return counts and *Statistical Abstract* series to the Census register across an overlap — a data-construction task in its own right, and not attempted here.

8. Conclusions

The direction of technical change is not a property of a technology but of a technology together with a designation of factor units. Of the Kaldor regularities, only constant factor shares — and, equivalently for $\sigma \neq 1$, stationarity of the engine ratios — survives redesignation. The standard argument against symmetric factor augmentation runs entirely through the regularities that do not survive.

Requiring a stationary measured capital–output ratio of a growth model is equivalent to requiring engine balance together with zero residual capital augmentation. The second conjunct is a measurement convention, and the convention arises because the transaction boundary falls asymmetrically across the two augmentation institutions: the firm's improvement is partly embodied in purchased goods and therefore priced, the household's is almost wholly internal and therefore not.

The firm is thus doubly central to the question. It is the agent that allocates augmentation investment across the two margins, and it is the entity whose purchases determine which margin the statistical system can see. Placing firms and households in the count — the institutional designation of Section 7 — makes the extensive margin observable where the augmentation indices are not, and supplies a test of symmetric augmentation that the standard designation cannot support.

Counted in those institutions, U.S. capital augmentation over 1988–2019 runs at 1.88 percent a year and the augmentation gap at +0.31 percentage points — one standard error from the Hicks restriction and six from the Harrod restriction. Counted as a deflated stock, over the same years, capital augmentation runs at 0.48 percent and the gap at −1.11. The technology did not change between those two sentences. The unit did.

The claim of this essay is not that technical change is symmetric. It is that the question has not been settled by the evidence usually taken to settle it, that the institutions of augmentation are where the answer is determined, and that their counts are where it can be looked for.

References

Arthur, W. B. (2009). *The Nature of Technology: What It Is and How It Evolves*. Free Press.

Becker, G. S. (1964). *Human Capital*. NBER.

Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16).

Cummins, J. G., & Violante, G. L. (2002). Investment-specific technical change in the United States. *Review of Economic Dynamics*, 5(2).

Diamond, P., McFadden, D., & Rodriguez, M. (1978). Measurement of the elasticity of factor substitution and bias of technical change. In M. Fuss & D. McFadden (eds.), *Production Economics*.

Gordon, R. J. (1990). *The Measurement of Durable Goods Prices*. University of Chicago Press.

Greenwood, J., Hercowitz, Z., & Krusell, P. (1997). Long-run implications of investment-specific technological change. *American Economic Review*, 87(3).

Hulten, C. R. (1992). Growth accounting when technical change is embodied in capital. *American Economic Review*, 82(4).

Jones, C. I. (2016). The facts of economic growth. In J. B. Taylor & H. Uhlig (eds.), *Handbook of Macroeconomics*, vol. 2A. [Verify: outline cites Jones (2015); confirm intended reference.]

Jones, C. I., & Scrimgeour, D. (2008). A new proof of Uzawa's steady-state growth theorem. *Review of Economics and Statistics*, 90(1).

Jorgenson, D. W. (1966). The embodiment hypothesis. *Journal of Political Economy*, 74(1).

Kaldor, N. (1961). Capital accumulation and economic growth. In F. A. Lutz & D. C. Hague (eds.), *The Theory of Capital*.

Klump, R., McAdam, P., & Willman, A. (2007). Factor substitution and factor-augmenting technical progress in the United States. *Review of Economics and Statistics*, 89(1).

León-Ledesma, M. A., McAdam, P., & Willman, A. (2010). Identifying the elasticity of substitution with biased technical change. *American Economic Review*, 100(4).

Samuelson, P. A. (1965). A theory of induced innovation along Kennedy–Weizsäcker lines. *Review of Economics and Statistics*, 47(4).

Solow, R. M. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics*, 39(3).

Solow, R. M. (1960). Investment and technical progress. In K. Arrow, S. Karlin & P. Suppes (eds.), *Mathematical Methods in the Social Sciences*.

Uzawa, H. (1961). Neutral inventions and the stability of growth equilibrium. *Review of Economic Studies*, 28(2).

Weitzman, M. L. (1998). Recombinant growth. *Quarterly Journal of Economics*, 113(2).