

Revealed Preference Analysis of Hicks-Neutral Productivity

Wilfried Youmbi Fotso* Samson M'boueke† Roland Pongou‡

October 24, 2026

Abstract

This paper develops a nonparametric framework for assessing whether a proposed measure of total factor productivity is consistent with cost-minimizing behavior under a common concave production technology. Allowing for observation-specific Hicks-neutral productivity, we derive necessary and sufficient conditions for rationalizability and characterize the identified set of relative productivity levels. We show that the existence of a rationalizing technology depends only on observed input choices, while relative productivity is partially identified by sharp, finite bounds that can be computed using linear programming. We also provide an exact joint test of a candidate productivity series and obtain a graph-theoretic characterization under constant returns to scale. We apply the framework to 364 U.S. manufacturing industries over 1958–2018. Approximately half satisfy the revealed-preference restrictions on input choices, but constant returns to scale is rejected in every industry over the full sample. Among industries satisfying these restrictions, the published NBER–CES productivity index violates the implied marginal bounds in nearly three-quarters of industry-year observations, and no complete industry-level series passes the joint rationalizability test. These findings raise questions about the consistency of conventional productivity measures with the maintained assumptions of cost minimization, Hicks-neutral technical change, and a stable production technology.

JEL Classification: C14, D24, E23, O47.

Keywords: Hicks-neutral productivity; revealed preference; Afriat's theorem; partial identification; production GARP.

*Department of Economics, Wilfrid Laurier University. E-mail: wyoumbifotso@wlu.ca.

†Department of Economics, University of Notre Dame. E-mail: smboueke@alumni.nd.edu.

‡Department of Economics, University of Ottawa. E-mail: rpongou@uottawa.ca.

1 Introduction

Total factor productivity (TFP) plays a central role in empirical research on economic growth, resource misallocation, firm dynamics, and development, although it is not directly observed. In growth accounting, productivity is measured as the portion of output growth unexplained by changes in a share-weighted index of inputs (Solow, 1957; Jorgenson and Griliches, 1967). Firm-level studies typically recover productivity as a residual from an estimated production function (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Akerberg et al., 2015; Gandhi et al., 2020). Both approaches rely on assumptions about the underlying production technology. This paper examines whether a proposed productivity series can be rationalized by cost-minimizing behavior under a common concave production technology, without imposing a parametric functional form.

Existing productivity estimators rely on assumptions concerning technology, market structure, and the timing of firms' decisions. We develop a revealed-preference framework for evaluating the consistency of a given Hicks-neutral productivity series with observed input choices. The framework provides necessary and sufficient conditions under which a common concave technology can rationalize the data and the proposed productivity series jointly. Rejection implies that at least one of the maintained assumptions is inconsistent with the observations, although the test does not identify which assumption fails. When productivity is unrestricted, the existence of a rationalizing technology depends entirely on observed input prices and quantities. A specified productivity series introduces additional restrictions that can be tested against the data.

Our analysis builds on the classical production revealed-preference framework of Afriat (1967) and Varian (1984). We consider a production environment in which output satisfies $y_t = \theta_t F(x_t)$, where F denotes a common production technology and θ_t is an observation-specific Hicks-neutral productivity shifter. We characterize the identified set of relative productivity levels and establish an exact joint test for the rationalizability of a complete candidate series. The identified set yields sharp marginal bounds that can be computed using linear programming. These bounds are finite and strictly positive for every observation, a property that follows from the supporting-hyperplane restrictions and the nonnegativity of technology at the origin. Under constant returns to scale, we obtain a graph-theoretic characterization

of rationalizability and derive sharp bounds on relative productivity using longest-path distances. The resulting cycle diagnostics and ratio bounds can be computed without linear programming. We use these theoretical results to evaluate a widely used measure of productivity in U.S. manufacturing.

The framework distinguishes three aspects of productivity identification. First, the existence of a Hicks-neutral rationalization is equivalent to satisfaction of the production generalized axiom of revealed preference (GARP) by observed input prices and quantities. Second, the common scale of productivity is not identified, since proportional changes in the productivity shifters can be offset by an appropriate rescaling of the production function. A normalization is therefore required to express productivity levels relative to a reference observation. Third, relative productivity is partially identified. The Afriat inequalities imply restrictions on productivity ratios that can be characterized by sharp bounds, which generally become tighter as additional observations are incorporated. Under constant returns to scale, these bounds admit a longest-path representation, with the familiar pairwise restrictions corresponding to paths of length one.

The restrictions derived in the paper hold exactly under the maintained assumptions. Their empirical implementation can therefore be sensitive to measurement error in input prices and quantities, particularly as the number of observations increases and the number of inequalities grows quadratically. We assess the discriminatory power of the tests using a Bronars-type random-behavior benchmark. We also examine the implications of relaxing the cost comparisons while maintaining exact interpolation of normalized output levels. Such a relaxation provides necessary conditions for approximate rationalization but does not establish sufficiency, limiting its usefulness as a general measure of approximation.

We apply the framework to 364 U.S. manufacturing industries over the period 1958–2018. Approximately half of the industries satisfy production GARP, indicating that their observed input choices can be rationalized by a common concave technology. This result concerns the consistency of input choices with cost minimization and does not establish Hicks-neutrality. We then examine the published NBER–CES productivity index using the marginal bounds and the exact joint test. Among industries satisfying GARP, the index violates at least one marginal bound in nearly three-quarters of industry-year observations, with violations predominantly occurring

above the upper bound. No industry has a complete productivity series that satisfies the joint rationalizability conditions, including two industries for which every annual productivity value lies within its respective marginal interval. We also find that constant returns to scale is rejected in every industry over the full sample. The restriction is consistent with approximately three-quarters of adjacent-year windows, fewer than one-fifth of five-year windows, and almost no windows spanning a decade or more. These findings indicate substantial tensions between the observed data and the joint assumptions underlying conventional productivity measurement. Rejection may reflect failures of cost minimization, a stable common technology, Hicks-neutral technical change, accurate measurement, or valid aggregation at the industry level.

1.1 Related Literature

The framework descends from Afriat (1967) and Varian (1984), who characterize rationalizability of finite production data, including under constant returns; Hanoch and Rothschild (1972) and Diewert and Parkan (1983) established the linear-programming approach to testing production regularity that our implementation follows. That work treats every observation as generated by one technology; we allow an observation-specific multiplicative productivity shifter.

The closest antecedents are the nonparametric production frameworks with latent productivity of Cherchye et al. (2018) and, allowing factor-biased technical change, Cherchye et al. (2021). Like us, they maintain cost minimization, treat productivity as unobserved, and use linear programming. Their productivity term is a latent factor recovered jointly with the technology across a cross-section of firms, and the biased-change extension allows input-specific shifters, which is more general than our single multiplicative scalar. We instead fix a *candidate* series and ask whether it is admissible. The exercises are complementary: theirs recovers productivity when no candidate is in hand and neutrality is in doubt; ours adjudicates a series that already exists and is silent about factor bias by construction. Aguiar et al. (2023) identify production sets and counterfactual bounds for firms ranked by productivity when quantities and prices may be latent, by duality and an observed optimized value. We take input prices and quantities as observed and work directly with the Afriat inequalities, which delivers bounds on individual ratios comparable with a published index, at the cost of saying nothing when prices or quantities are unobserved.

Our analysis also relates to partial identification (Manski, 2003; Tamer, 2010). In revealed-preference analysis, Blundell et al. (2008) derive sharp nonparametric bounds on consumer demand responses. In production analysis, the methods of Olley and Pakes (1996), Levinsohn and Petrin (2003), Wooldridge (2009), Akerberg et al. (2015), and Gandhi et al. (2020) recover production functions and productivity under maintained assumptions about technology and firms’ input decisions. Using revealed-preference restrictions, we characterize the identified set of relative Hicks-neutral productivity consistent with cost-minimizing behavior under a common concave technology. We obtain sharp marginal bounds and an exact joint test of a proposed productivity series. A series can satisfy every marginal bound without being jointly rationalizable. The framework does not estimate a production frontier or classify observations as inefficient, as in data envelopment analysis.

The index-number tradition (Solow, 1957; Jorgenson and Griliches, 1967; Caves et al., 1982b; Good et al., 1997; Hulten, 2001) measures productivity as a residual after share-weighted input growth, the joint hypothesis our constant-returns results test; Basu and Fernald (1997) study returns to scale on comparable industry data, and our upward violations connect to the markup literature (Hall, 1988; De Loecker and Warzynski, 2012; De Loecker et al., 2020). Because our observations are industry aggregates, we inherit the warning that an aggregate production function need not exist (Fisher, 1969; Blackorby and Schworm, 1988). We follow Bronars (1987) in evaluating test power; goodness-of-fit indices (Afriat, 1973; Varian, 1990) and inference on the bounds (Chernozhukov et al., 2007) are left for future work.

The remainder of the paper is organized as follows. Section 2 presents the production environment and the Afriat system. Section 3 establishes the conditions for Hicks-neutral rationalizability and characterizes the identified set of relative productivity. Section 4 examines the case of constant returns to scale. Section 5 provides a numerical illustration, and Section 6 presents the empirical application to U.S. manufacturing industries. Section 7 concludes.

2 Setup

This section formalizes the production environment, the revealed-feasibility relation it induces, and the Afriat system whose feasibility characterizes Hicks-neutral ratio-

nalizability. Table 1 collects the notation for reference before the formal development begins.

Table 1: Notation

Symbol	Meaning
$\mathbb{T} = \{1, \dots, T\}$	index set of production observations
y_t, x_t, w_t	output, input bundle, and input prices at t
$c_t := w_t \cdot x_t$	observed production cost at t
F	common technology
θ_t	observation-specific Hicks-neutral productivity
$U_t(\theta) := y_t/\theta_t$	normalized output implied by profile θ
λ_t	Afriat multiplier at t
r	reference observation, normalized $\theta_r = 1$
Θ_r^{HN}	identified set of relative productivity profiles
$\overline{U}_t, \underline{U}_t$	marginal upper and lower bounds on $U_t(\theta)$
$\mathcal{F}_r, \overline{\mathcal{F}}_r$	feasible set of (U, λ) and its closure
$v_t := w_t/c_t$	unit-cost share vector (constant-returns section)
$\succeq_D$	direct revealed feasibility

2.1 Production Environment

Let $\mathbb{T} = \{1, \dots, T\}$, $T \geq 2$, index a finite collection of production observations. For each $t \in \mathbb{T}$, the econometrician observes output $y_t \in \mathbb{R}_{++}$, an input bundle $x_t \in \mathbb{R}_{++}^L$, and an input-price vector $w_t \in \mathbb{R}_{++}^L$. The observed production dataset is $S = \{(y_t, x_t, w_t)\}_{t \in \mathbb{T}}$ and observed production cost is $c_t := w_t \cdot x_t$.

Assumption 1 (Interior observations). *Every observed input bundle is strictly positive, $x_t \in \mathbb{R}_{++}^L$ for all $t \in \mathbb{T}$.*

Assumption 1 guarantees that a concave F admits a finite supergradient at x_t ; it is used only in the necessity direction of Theorem 1, and every observation in Section 6 satisfies it. Production is governed by a common technology $F : \mathbb{R}_+^L \rightarrow \mathbb{R}_+$ and observation-specific Hicks-neutral productivity shifters $\theta_t > 0$.

Definition 1 (Admissible technology). *A function $F : \mathbb{R}_+^L \rightarrow \mathbb{R}_+$ is admissible if it is continuous, concave, nonnegative on $\mathbb{R}_+^L$, and strictly increasing in the strong componentwise order, meaning that $x' \geq x$ with $x' \neq x$ implies $F(x') > F(x)$.*

Strict monotonicity delivers local nonsatiation, which yields a strictly positive Afriat multiplier in the necessity direction of Theorem 1; it is not restrictive, because the envelope (6) constructed in the sufficiency direction is itself strictly increasing. The Hicks-neutral specification is

$$y_t = \theta_t F(x_t), \quad t \in \mathbb{T}. \quad (1)$$

Each observed bundle minimizes cost given the output level generated by the technology:

$$x_t \in \arg \min \{w_t \cdot x : F(x) \geq F(x_t)\}, \quad t \in \mathbb{T}. \quad (2)$$

The productivity shifters affect output multiplicatively but do not alter the shape of the common production technology.

Definition 2 (Hicks-neutral cost rationalization). *The dataset S is Hicks-neutral cost-rationalizable if there exist a productivity profile $\theta \in \mathbb{R}_{++}^T$ and an admissible production function $F : \mathbb{R}_+^L \rightarrow \mathbb{R}_+$ such that (1) and (2) hold for every $t \in \mathbb{T}$.*

For any candidate productivity profile $\theta = (\theta_t)$, define the normalized output levels

$$U_t(\theta) := \frac{y_t}{\theta_t}, \quad t \in \mathbb{T}.$$

Under any Hicks-neutral rationalization, $U_t(\theta) = F(x_t)$, so the normalized outputs represent the technology levels assigned to the observed input bundles by a common production function.

2.2 Revealed Feasibility and the Afriat System

The production data induce a revealed-feasibility relation across observations.

Definition 3 (Direct revealed feasibility). *Observation r directly reveals feasibility of observation t , written $r \succeq_D t$, if $w_r \cdot x_t \leq c_r$. Equivalently, the input bundle chosen in observation t was affordable at the prices and cost level observed in observation r .*

Remark 1 (Relation to Varian’s terminology). *The condition $w_r \cdot x_t \leq c_r$ is equivalent to $w_r \cdot x_r \geq w_r \cdot x_t$, which Varian (1984) calls “ x_r is directly revealed preferred to x_t ” and denotes $x_r R^D x_t$. The present paper organizes the relation across observations rather than bundles and uses the term “reveals feasibility” to emphasize that the firm could*

have produced at the technology level $F(x_t)$ using the budget available at observation r ; the mathematical content is identical.

Definition 4 (Revealed-feasibility graph). *The revealed-feasibility graph $\mathcal{G}(S)$ has vertex set $\mathbb{T}$ and a directed edge $r \rightarrow t$ whenever $r \succeq_D t$.*

Definition 5 (Afriat system). *Given $\theta \in \mathbb{R}_{++}^T$, the Afriat system $\mathcal{A}(S, \theta)$ is the collection of inequalities*

$$U_t(\theta) \leq U_s(\theta) + \lambda_s(w_s \cdot x_t - w_s \cdot x_s), \quad \forall s, t \in \mathbb{T}, \quad (3)$$

together with the nonnegativity conditions

$$U_s(\theta) \geq \lambda_s w_s \cdot x_s = \lambda_s c_s, \quad \forall s \in \mathbb{T}, \quad (4)$$

in the unknown multipliers $(\lambda_t)_{t \in \mathbb{T}} \in \mathbb{R}_{++}^T$. The system is feasible if there exists a strictly positive vector (λ_t) that satisfies (3) and (4). Condition (4) is what guarantees that the rationalizing technology constructed in Theorem 1 below is nonnegative on all of $\mathbb{R}_+^L$, not merely at the observed bundles. It is not a technical convenience. Because it caps every multiplier by $\lambda_s \leq U_s(\theta)/c_s$, it is also the source of the two-sided productivity bounds of Proposition 3, and hence of the finding in Section 6.4 that every industry-year receives a finite interval bounded away from zero. Remark 6 below sets out exactly which ingredients that conclusion rests on.

Remark 2 (Intercept condition and Varian (1984)). *The intercept condition (4) does not appear separately in Varian (1984) Theorem 1, but it is implicit in his construction of the rationalizing technology: the Afriat envelope $\tilde{F}$ is nonnegative throughout $\mathbb{R}_+^L$ only when (4) holds, which Varian derives by evaluating the supporting hyperplane at $x = 0$ and using $F(0) \geq 0$. The present paper includes it as an explicit condition because it plays an active role in establishing the upper bound on $\bar{U}_t$ in Proposition 3, which holds for every observation and not only for those connected to the reference, and in guaranteeing that every identified set Θ_r^{HN} delivers a nonnegative rationalizing technology.*

2.3 Identified Set

Since the common scale of productivity is not identified (Section 3.2), we fix a reference observation $r \in \mathbb{T}$ and normalize $\theta_r = 1$, as is standard in index-number analysis.

Definition 6 (Identified set). *For a fixed reference observation $r \in \mathbb{T}$, define*

$$\Theta_r^{HN} = \{\theta \in \mathbb{R}_{++}^T : \theta_r = 1, \mathcal{A}(S, \theta) \text{ is feasible}\}.$$

The set Θ_r^{HN} is the identified set of relative Hicks-neutral productivity profiles. As discussed in Section 3.2 below, part of this set is already restricted before any normalization is imposed: only the common scale of θ , not its ratios, is fixed by $\theta_r = 1$.

2.4 Standard Production Rationalizability

The benchmark treats output levels as free: can the observed input choices be rationalized by a common concave technology?

Definition 7 (Production GARP). *The dataset $\{(x_t, w_t)\}_{t \in \mathbb{T}}$ satisfies production GARP if there exist numbers $(U_t, \lambda_t)_{t \in \mathbb{T}} \in \mathbb{R}_{++}^T \times \mathbb{R}_{++}^T$ such that*

$$U_t \leq U_s + \lambda_s(w_s \cdot x_t - w_s \cdot x_s), \quad \forall s, t \in \mathbb{T}, \quad (5)$$

*and $U_s \geq \lambda_s c_s$ for all $s \in \mathbb{T}$.*¹

By Afriat's theorem for production (Afriat, 1967; Varian, 1984), production GARP is equivalent to rationalization of the input choices by a continuous, concave, monotone technology, nonnegative on $\mathbb{R}_+^L$. Lemma 2 in the Appendix records the elementary direction of the graph characterization: a revealed-feasibility cycle with a strict affordability relation makes the Afriat system infeasible for every productivity profile.

¹Following Afriat (1967) and the computational tradition of Hanoich and Rothschild (1972) and Diewert and Parkan (1983), this definition states production GARP directly as the feasibility of the Afriat numerical system. Varian (1984) instead defines GARP as the graph-theoretic cyclic-consistency condition on the data (if the transitive closure of the direct revealed-preference relation R^D satisfies $x_s R x_t$, then $\neg(x_t P^0 x_s)$, i.e. $w_t \cdot x_s \leq w_t \cdot x_t$) and then proves in his Theorem 1 that this is equivalent to the existence of Afriat numbers satisfying the inequalities above. The two formulations are equivalent by that theorem; the present paper uses the Afriat characterization as the definition because it is the object manipulated directly in proofs and in the linear programs of Section 3.

3 Main Results

3.1 An Afriat Characterization for Fixed Productivity

For a given $\theta \in \mathbb{R}_{++}^T$, the normalized outputs $U_t(\theta) = y_t/\theta_t$ are observed, and the question is whether a common concave technology can generate them. Cost minimization at x_s is a convex program whose Karush–Kuhn–Tucker conditions place a positive multiple of the price vector in the superdifferential, $\lambda_s w_s \in \partial F(x_s)$; strict monotonicity supplies the Slater point, and Assumption 1 makes the nonnegativity constraints inactive. The supergradient inequality then yields the Afriat inequalities. An arbitrary supporting hyperplane to the upper contour set would not do, since the normal cone at x_s need not consist of multiples of w_s . Conversely, the lower envelope of the supporting hyperplanes, the Afriat envelope, is a piecewise-linear, concave, strictly monotone technology, nonnegative on $\mathbb{R}_+^L$, that recovers the normalized outputs.

Theorem 1 (Afriat characterization). *Let $\theta \in \mathbb{R}_{++}^T$. The following statements are equivalent.*

- (i) *The dataset S is Hicks-neutral cost-rationalizable by some pair (θ, F) , where F is admissible.*
- (ii) *The Afriat system $\mathcal{A}(S, \theta)$, including the intercept condition (4), is feasible.*

Moreover, if (ii) holds, then writing $U_t := U_t(\theta)$ and letting $(\lambda_t)_{t \in \mathbb{T}}$ be any multiplier vector feasible for $\mathcal{A}(S, \theta)$, the function

$$\tilde{F}(x) = \min_{s \in \mathbb{T}} \{U_s(\theta) + \lambda_s(w_s \cdot x - w_s \cdot x_s)\} \quad (6)$$

is admissible, in particular satisfies $\tilde{F}(x) \geq 0$ for all $x \in \mathbb{R}_+^L$ by (4), satisfies $\tilde{F}(x_t) = U_t(\theta)$ for all $t \in \mathbb{T}$, and rationalizes the dataset together with the productivity profile θ .

Theorem 1 is, read directly, a membership test for a *complete* candidate productivity series, and it is worth recording in the form used in the empirical application.

Corollary 1 (Joint test of a candidate productivity series). *Let $\hat{\theta} \in \mathbb{R}_{++}^T$ be a proposed productivity series and set $\hat{U}_t := y_t/\hat{\theta}_t$. Write $v_s := w_s/c_s$ and $\gamma_s := \lambda_s c_s$. Then $\hat{\theta}$ is*

Hicks-neutral cost-rationalizable if and only if there exists $\gamma \in \mathbb{R}^T$ with

$$\widehat{U}_t - \widehat{U}_s \leq \gamma_s (v_s \cdot x_t - 1) \quad \forall s, t \in \mathbb{T}, \quad 0 < \gamma_s \leq \widehat{U}_s \quad \forall s \in \mathbb{T}. \quad (7)$$

Equivalently, letting κ^* be the optimal value of

$$\max_{\gamma, \kappa} \kappa \quad \text{s.t.} \quad \widehat{U}_t - \widehat{U}_s \leq \gamma_s (v_s \cdot x_t - 1) \quad \forall s, t, \quad \gamma_s \leq \widehat{U}_s \quad \forall s, \quad \gamma_s \geq \kappa \quad \forall s, \quad \kappa \geq 0, \quad (8)$$

the series is rationalizable if and only if (8) is feasible with $\kappa^* > 0$.

Proof. Dividing the Afriat inequality (3) by $c_s > 0$ and substituting $\gamma_s = \lambda_s c_s$ gives $\lambda_s (w_s \cdot x_t - c_s) = \gamma_s (v_s \cdot x_t - 1)$, and the intercept condition (4) becomes $\gamma_s \leq \widehat{U}_s$; the correspondence $\lambda \mapsto \gamma$ is a bijection of $\mathbb{R}_{++}^T$ onto itself. Theorem 1 then gives the equivalence with (7). For the second statement, if some $\gamma \gg 0$ satisfies (7) then $\kappa := \min_s \gamma_s > 0$ is feasible for (8), so $\kappa^* > 0$; conversely if $\kappa^* > 0$ then any optimal γ has $\gamma_s \geq \kappa^* > 0$. The feasible set of (8) is compact, since $0 \leq \gamma_s \leq \widehat{U}_s$, so the maximum is attained. $\square$

In fact no linear program is needed. Because $\widehat{U}$ is fixed, the multiplier γ_s appears only in the inequalities whose source observation is s , so (7) separates into T independent one-dimensional interval problems.

Corollary 2 (Closed-form joint test and rejection certificate). *Write $d_{st} := v_s \cdot x_t - 1$ and $a_{st} := \widehat{U}_t - \widehat{U}_s$, and for each s define*

$$\underline{\gamma}_s := \max \left\{ 0, \max_{t: d_{st} > 0} \frac{a_{st}}{d_{st}} \right\}, \quad \overline{\gamma}_s := \min \left\{ \widehat{U}_s, \min_{t: d_{st} < 0} \frac{a_{st}}{d_{st}} \right\},$$

with the conventions $\max \emptyset = 0$ and $\min \emptyset = +\infty$. Then $\widehat{\theta}$ is Hicks-neutral cost-rationalizable if and only if, for every $s \in \mathbb{T}$,

$$a_{st} \leq 0 \quad \text{whenever } d_{st} = 0, \quad \overline{\gamma}_s > 0, \quad \underline{\gamma}_s \leq \overline{\gamma}_s. \quad (9)$$

The test requires $O(T^2)$ arithmetic operations and no optimization.

Proof. Each inequality in (7) reads $a_{st} \leq \gamma_s d_{st}$ and involves γ_s alone, so the system separates across s . For fixed s , the inequality is $\gamma_s \geq a_{st}/d_{st}$ when $d_{st} > 0$, is

$\gamma_s \leq a_{st}/d_{st}$ when $d_{st} < 0$, and reduces to the γ_s -free requirement $a_{st} \leq 0$ when $d_{st} = 0$. Together with $0 < \gamma_s \leq \widehat{U}_s$, the admissible set for γ_s is $[\underline{\gamma}_s, \overline{\gamma}_s] \cap (0, \infty)$, which is nonempty if and only if $\overline{\gamma}_s > 0$ and $\underline{\gamma}_s \leq \overline{\gamma}_s$. Note that $s = t$ gives $d_{ss} = v_s \cdot x_s - 1 = 0$ and $a_{ss} = 0$, so the diagonal is always satisfied. $\square$

The three failure modes in (9) are economically distinct: an equal-cost target assigned a strictly higher normalized output ($a_{st} > 0$, $d_{st} = 0$); a strictly cheaper target assigned a weakly higher normalized output, so that no positive multiplier survives ($\overline{\gamma}_s \leq 0$, the fixed-profile counterpart of Theorem 4); and one target forcing γ_s above a bound imposed by another target or by the intercept ($\underline{\gamma}_s > \overline{\gamma}_s$). The argmax and argmin name the years generating the contradiction, which a solver status does not.

Remark 3 (Scale invariance of the joint test). *The test is invariant to the units of the candidate series. Replacing $\widehat{\theta}$ by $a\widehat{\theta}$ for $a > 0$ replaces $\widehat{U}$ by $\widehat{U}/a$, hence a_{st} by a_{st}/a , and dividing every γ_s by a preserves (7) and the intercept bound. The verdict therefore does not depend on the base year or the units in which `tfp5` is reported, consistent with Proposition 2. For numerical conditioning one may rescale so that $\widehat{U}_r = 1$ at the reference year.*

The program (8) has $T + 1$ variables, far fewer than the $2(T - 1)$ marginal-bound programs of Theorem 2, and $v_s = w_s/c_s$ imposes no constant-returns restriction. Strict positivity must be imposed through κ , not a fixed floor such as $\lambda_s \geq 1$: that floor is legitimate in the unrestricted GARP test, where (U, λ) can be rescaled jointly, but once $\widehat{U}$ is fixed it would test a stricter condition.

The intercept condition can be strengthened to pin the technology at the origin, which gives it its natural economic reading.

Corollary 3 (The technology can be chosen with no free output). *Suppose $\mathcal{A}(S, \theta)$ is feasible. Then the dataset admits a Hicks-neutral rationalization by an admissible F satisfying additionally $F(0) = 0$.*

Proof. Let (λ_t) be feasible and $\widetilde{F}$ the envelope (6). Choose $M \geq \max_t U_t(\theta)/(\mathbf{1} \cdot x_t)$, which is finite because $x_t \gg 0$ by Assumption 1, put $\mu := M\mathbf{1} \gg 0$, and define

$$\widehat{F}(x) := \min\{\widetilde{F}(x), \mu \cdot x\}, \quad x \in \mathbb{R}_+^L.$$

Being a minimum of finitely many affine functions with strictly positive slopes, $\widehat{F}$ is continuous, concave and strictly increasing in the strong componentwise order. It is nonnegative on $\mathbb{R}_+^L$ because $\widetilde{F}$ is, by (4), and because $\mu \cdot x \geq 0$. At the origin, $\widehat{F}(0) = \min\{\widetilde{F}(0), 0\} = 0$, since $\widetilde{F}(0) \geq 0$. Interpolation is preserved: $\mu \cdot x_t = M(\mathbf{1} \cdot x_t) \geq U_t(\theta)$ by the choice of M , so $\widehat{F}(x_t) = \min\{U_t(\theta), \mu \cdot x_t\} = U_t(\theta)$. Finally, $\widehat{F} \leq \widetilde{F}$ implies $\{x : \widehat{F}(x) \geq U_t(\theta)\} \subseteq \{x : \widetilde{F}(x) \geq U_t(\theta)\}$, and x_t minimizes $w_t \cdot x$ on the larger set by Theorem 1 while belonging to the smaller one; hence it minimizes cost on the smaller set as well. $\square$

Corollary 3 matters for interpretation, not for the bounds: positive output at the null bundle is an artifact of a particular envelope, never of the data, so the intercept condition reads as a *no free output* restriction. Under constant returns the envelope of Theorem 5 already vanishes at the origin. The rationalization is generally not unique: different feasible multipliers give different envelopes, so even for a fixed profile the technology is set-identified rather than point-identified.

3.2 Vacuity of the Existence Test, and Scale versus Ratio Non-identification

When the profile is left unrestricted, two questions must be kept apart: whether *some* Hicks-neutral rationalization exists, and whether productivity *ratios* are restricted. The first has a vacuous answer; the second does not.

Existence is vacuous in outputs

One might expect outputs to discipline admissible productivity, since $y_t = \theta_t F(x_t)$. For existence they do not: any Afriat solution (U_t^*, λ_t^*) of the standard production problem yields a decomposition $\theta_t = y_t/U_t^*$, so outputs are entirely absorbed.

Proposition 1 (Observable equivalence with production GARP). *The following statements are equivalent.*

- (i) *There exists a productivity profile $\theta \in \mathbb{R}_{++}^T$ such that the dataset S is Hicks-neutral cost-rationalizable.*
- (ii) *The input-price data $\{(x_t, w_t)\}_{t \in \mathbb{T}}$ satisfy production GARP.*

Hence the output observations $(y_t)_{t \in \mathbb{T}}$ impose no testable restriction, beyond production GARP, on whether some Hicks-neutral rationalization of the dataset exists.

Remark 4 (The existence test is not a test of neutrality). *Proposition 1 is easily over-read, and the misreading is the natural one. What the existence test examines is whether the observed input choices are consistent with stable concave isoquants and cost minimization. It does not examine neutrality, because the productivity profile is left unrestricted and therefore absorbs any output pattern whatsoever: two datasets with identical (x_t, w_t) and wildly different (y_t) receive the same verdict. A pass therefore says that the input data admit a common concave technology; it says nothing about whether technical change was neutral, and nothing about any particular productivity series. Neutrality acquires testable content only once a candidate profile is fixed, as in Theorem 1, or once ratios are bounded and compared against a proposed series, as in Theorem 2 and Section 6.4. The same caution applies to the constant-returns existence result after Theorem 5, where the output terms again telescope out of the cyclical condition.*

Proposition 1 does not say that every profile is admissible or that ratios are unrestricted; ratios are restricted even before any normalization.

Scale is not identified, but ratios can already be restricted

If (θ, F) rationalizes the data, so does $(a\theta, F/a)$ for every $a > 0$: outputs are unchanged and (2) depends only on the level sets of F . The *scale* of productivity is therefore never identified, and a normalization $\theta_r = 1$ is a choice of units. Ratios, however, are generally restricted. Whenever $r \succeq_D t$, the Afriat inequality with $s = r$ gives, for *any* rationalizing profile, $U_t(\theta) \leq U_r(\theta) + \lambda_r(w_r \cdot x_t - c_r) \leq U_r(\theta)$, since $w_r \cdot x_t \leq c_r$ and $\lambda_r > 0$; that is,

$$\frac{\theta_t}{\theta_r} \geq \frac{y_t}{y_r}, \quad (10)$$

which involves only the ratio and so holds at any scale. The normalization does not create such restrictions; it is a coordinate convention for reporting them, and Theorem 2 below gives genuine bounds on ratios.

Proposition 2 (Scale normalization is a unit choice). *For any reference observation $r \in \mathbb{T}$, $\Theta_r^{HN} \neq \emptyset$ if and only if the input-price data satisfy production GARP. Moreover, every scale-equivalence class of rationalizing productivity profiles, $[\theta] = \{a\theta :$*

$a > 0\}$, contains exactly one element satisfying $\theta_r = 1$; consequently, Θ_r^{HN} is in bijection with the set of attainable productivity-ratio vectors $\{(\theta_t/\theta_r)_{t \in \mathbb{T}} : \theta \text{ rationalizes } S\}$, so that $\theta_r = 1$ selects one representative from each ray of observationally equivalent profiles without altering which relative productivity ratios are attainable. The map $\theta \mapsto (\theta_t/\theta_r)_t$ from the full, unnormalized set of rationalizing profiles is of course not injective on its own, since every profile on a given ray maps to the same ratio vector; it is the quotient by this ray equivalence, not the unnormalized set itself, that is in bijection with Θ_r^{HN} .

We maintain throughout the distinction between (a) non-testability of unrestricted existence, (b) nonidentification of scale, and (c) partial identification of ratios.

3.3 Partial Identification of Relative Productivity

The normalization pins $U_r = y_r$, and the Afriat inequalities propagate this information to every observation, generating the sharp marginal bounds below.

Definition 8 (Feasible set of normalized technology levels). *Fix a reference observation $r \in \mathbb{T}$ satisfying production GARP. The feasible set is*

$$\mathcal{F}_r = \{(U_t, \lambda_t)_{t \in \mathbb{T}} : U_r = y_r, \text{ (3) and (4) hold, } U_t > 0, \lambda_t > 0\},$$

and its closure $\overline{\mathcal{F}}_r$ is obtained by relaxing the strict positivity constraints to $U_t \geq 0, \lambda_t \geq 0$.

Definition 9 (Marginal bounds on normalized technology levels). *For each $t \neq r$, define*

$$\overline{U}_t := \sup\{U_t(\theta) : \theta \in \Theta_r^{HN}\}, \quad \underline{U}_t := \inf\{U_t(\theta) : \theta \in \Theta_r^{HN}\}, \quad (11)$$

with the convention $y_t/\overline{U}_t := 0$ if $\overline{U}_t = +\infty$ and $y_t/\underline{U}_t := +\infty$ if $\underline{U}_t = 0$.

Theorem 2 (Linear-programming characterization of marginal bounds). *Suppose the dataset satisfies production GARP and fix a reference observation $r \in \mathbb{T}$. The closure $\overline{\mathcal{F}}_r$ is a compact polytope. For every $t \neq r$, $\overline{U}_t$ and $\underline{U}_t$ are respectively the maximum and minimum of the linear programs*

$$\begin{aligned} \overline{U}_t &= \max\{U_t : U_r = y_r, \text{ (3) and (4) hold}\}, \\ \underline{U}_t &= \min\{U_t : U_r = y_r, \text{ (3) and (4) hold}\}. \end{aligned} \quad (12)$$

taken over $\overline{\mathcal{F}}_r$. These closure optima equal the supremum and infimum over the strictly positive feasible set $\mathcal{F}_r$. Consequently, the closure of the sharp marginal identified set for the ratio is

$$\frac{\theta_t}{\theta_r} \in \left[\frac{y_t}{\overline{U}_t}, \frac{y_t}{\underline{U}_t} \right]. \quad (13)$$

A finite endpoint belongs to the actual identified set if and only if the corresponding optimizing face of $\overline{\mathcal{F}}_r$ intersects $\mathcal{F}_r$.

The closure programs attain their optima, but an endpoint need not be generated by strictly positive multipliers (Section 6.2 discusses implementation). Theorem 3 explains why the product of marginal intervals need not be sharp.

Theorem 3 (Geometry of the normalized identified set). *Suppose the dataset satisfies production GARP and fix a reference observation $r \in \mathbb{T}$. Then:*

(i) $\mathcal{F}_r$ is convex and dense in its closure $\overline{\mathcal{F}}_r$, and $\overline{\mathcal{F}}_r$ is a compact polytope.

(ii) Write

$$v_s := \frac{w_s}{c_s}, \quad z_t := \frac{\theta_t}{y_t} = \frac{1}{U_t}, \quad d_{st} := v_s \cdot x_t - 1,$$

and, for each source observation s , define $P_s := \{p : d_{sp} > 0\}$, $N_s := \{n : d_{sn} < 0\}$, and $Z_s := \{t : d_{st} = 0\}$. The identified set is exactly the positive vectors z satisfying $z_r = 1/y_r$ and, for every s ,

$$z_t \geq z_s \quad \forall t \in Z_s, \quad (14)$$

$$z_n > z_s \quad \forall n \in N_s, \quad (15)$$

$$z_s \leq (1 + d_{sp})z_p \quad \forall p \in P_s, \quad (16)$$

$$z_s \leq \frac{d_{sp} - d_{sn}}{(-d_{sn})/z_p + d_{sp}/z_n} \quad \forall p \in P_s, n \in N_s. \quad (17)$$

The productivity profile is recovered by $\theta_t = y_t z_t$.

(iii) Consequently, Θ_r^{HN} is convex, and its closure is compact and convex. It need not be polyhedral: the right-hand side of (17) is a weighted harmonic mean, so a nonredundant constraint of this form generates a curved boundary.

The auxiliary set in (U, λ) -space is polyhedral, which is why marginal bounds are

linear programs; eliminating the multipliers preserves convexity but replaces linear faces by weighted-harmonic-mean boundaries.

Remark 5 (Marginal versus joint identification). *Theorem 2 delivers marginal bounds: for each t separately, θ_t/θ_r lies in the interval $[y_t/\bar{U}_t, y_t/\underline{U}_t]$. The Cartesian product of these intervals is a valid outer bound for the joint identified set Θ_r^{HN} , since each interval is by construction the projection of Θ_r^{HN} onto the corresponding coordinate and a set is always contained in the product of its projections. What fails is sharpness: the product is in general strictly larger than Θ_r^{HN} , because achieving the extreme value for θ_t/θ_r and the extreme value for θ_s/θ_r simultaneously may require conflicting configurations of the common multiplier vector $(\lambda_s)_{s \in \mathbb{T}}$. In the three-observation example of Section 5, the identified intervals are $\theta_2 \in [14/15, 2.1]$ and $\theta_3 \in [0.8, 1.2]$ individually, but the joint computation in Appendix A.1 shows that $\theta_2 = 14/15$ and $\theta_3 = 1.2$ are not simultaneously attained by the same feasible solution. The pair $(14/15, 1.2)$ therefore lies in the product of the marginal intervals but outside Θ_r^{HN} , which is precisely a demonstration that the product is a loose outer bound rather than a sharp characterization. The sharp joint object is the convex, generally nonpolyhedral set characterized in Theorem 3. Marginal bounds are a valid outer bound, but membership of a complete candidate series in the product of its marginal intervals is only a necessary condition for membership in Θ_r^{HN} . The corresponding necessary-and-sufficient test is the single linear program (8) of Corollary 1, or equivalently the closed-form source-by-source test of Corollary 2.*

Theorem 4 (Sufficient condition for a nontrivial bound). *Suppose the dataset satisfies production GARP and fix a reference observation $r \in \mathbb{T}$. If $r \succeq_D t$ for some $t \neq r$, then*

$$\bar{U}_t \leq y_r, \quad \frac{\theta_t}{\theta_r} \geq \frac{y_t}{y_r}. \quad (18)$$

As discussed in Section 3.2, this bound holds for the ratio θ_t/θ_r independently of the normalization and is not, in general, the only source of nontrivial bounds: chains of revealed feasibility not involving a direct edge from r can also restrict θ_t/θ_r , so direct affordability $r \succeq_D t$ is sufficient, but this paper does not establish that it is necessary, for a nontrivial marginal bound.

The bound combines the input data, through $r \succeq_D t$, with the output data, through y_t/y_r ; neither alone suffices, and the interaction does not depend on the

normalization.

Proposition 3 (Universal finite two-sided bounds). *Suppose the dataset satisfies production GARP and fix a reference observation r . For every $t \in \mathbb{T}$,*

$$\frac{y_r}{\max\{1, v_t \cdot x_r\}} \leq \underline{U}_t \leq \overline{U}_t \leq y_r \max\{1, v_r \cdot x_t\}, \quad (19)$$

where $v_t = w_t/c_t$. Consequently the marginal identified interval for the ratio is nonempty and finite for every t , with a strictly positive lower endpoint:

$$\frac{y_t}{y_r \max\{1, v_r \cdot x_t\}} \leq \frac{\theta_t}{\theta_r} \leq \frac{y_t \max\{1, v_t \cdot x_r\}}{y_r}, \quad (20)$$

so the lower bound is strictly positive and the upper bound is finite.

Remark 6 (What the two-sided bound rests on). *Proposition 3 is not a connectivity result, and it is worth being explicit about its ingredients, since the natural reading attributes it to the revealed-feasibility graph. The derivation uses four elements. Concavity of F yields the supergradient inequality and hence the Afriat inequalities. Exact cardinal interpolation, $F(x_t) = U_t(\theta)$, ties those inequalities to the observed outputs rather than to an arbitrary ordinal ranking. Nonnegativity of F at the origin, through (4), caps each multiplier by $\lambda_s \leq U_s(\theta)/c_s$. And the reference normalization $U_r = y_r$ converts that cap into a numerical bound on the remaining coordinates. In the present derivation, the intercept condition supplies the multiplier cap; without exact interpolation, the observed outputs drop out of the existence problem, as Proposition 1 shows. These observations explain the role of the maintained assumptions without asserting that no alternative set of assumptions could also deliver finite bounds. What the revealed-feasibility graph contributes is the sharpness of the bounds through Theorem 2, not their existence.*

Whereas Theorem 4 bounds only observations reached by a direct edge from the reference, Proposition 3 bounds θ_t/θ_r on both sides for *every* t , connected or not: each bound follows from one Afriat inequality and one intercept cap, with the roles of r and t exchanged. Partial identification of relative productivity is therefore never vacuous once a reference is normalized.

3.4 Identification and Information Accumulation

Two further results complete the characterization. The first confirms that an edge from the reference excludes an interval of ratio values.

Proposition 4 (Nontrivial identification). *Suppose there exists an observation $t \neq r$ such that $r \succeq_D t$. Then every productivity profile $\theta \in \Theta_r^{HN}$ satisfies $\theta_t/\theta_r \geq y_t/y_r$, while any candidate value $\theta_t/\theta_r < y_t/y_r$ is excluded from the marginal identified interval for this ratio. This is a sufficient, not a necessary, condition for a nontrivial marginal bound; ratios θ_t/θ_r for which no direct edge $r \succeq_D t$ exists may nonetheless be restricted through indirect chains of revealed feasibility, a possibility this paper does not rule out.*

Remark 7 (Three distinct sources of identification). *Direct revealed feasibility $r \succeq_D t$ is sufficient for a nontrivial lower bound on θ_t/θ_r but, as Theorem 4 notes, it is not necessary. It is worth separating three mechanisms, because they are easily conflated and only the first two involve the revealed-feasibility graph at all.*

Direct affordability. *If $r \succeq_D t$, the single Afriat inequality with $s = r$ gives $U_t \leq U_r = y_r$ and hence $\theta_t/\theta_r \geq y_t/y_r$. This is Theorem 4.*

Directed paths. *If $r \succeq_D s \succeq_D t$ but r does not directly afford t , chaining gives*

$$U_t \leq U_r + \lambda_r(w_r \cdot x_s - c_r) + \lambda_s(w_s \cdot x_t - c_s) \leq y_r + \underbrace{\lambda_r(w_r \cdot x_s - c_r)}_{\leq 0} + \underbrace{\lambda_s(w_s \cdot x_t - c_s)}_{\leq 0} \leq y_r,$$

so the same bound propagates along any directed path. In the general concave model a path cannot do better than this: setting $U_s = y_r$ and $\lambda_s = 0$ for every s satisfies the closure constraints, so $\bar{U}_t \geq y_r$, and an affordability path from r to t therefore gives exactly $\bar{U}_t = y_r$. Under constant returns homogeneity removes the free multipliers, longer paths can bind, and the mechanism becomes the path maximization of Theorem 7.

Cross-cost inequalities and intercepts, absent any path. *The Afriat system contains an inequality for every ordered pair, including pairs with $w_s \cdot x_t > c_s$, which are not edges of $\mathcal{G}(S)$. Such an inequality still restricts U_t once the intercept condition caps $\lambda_s \leq U_s/c_s$, and this is the mechanism behind the universal two-sided bounds of Proposition 3. It operates whether or not t is reachable from r .*

The example of Section 5 illustrates the third mechanism, not the second. Its

only edges are $1 \rightarrow 3$ and $2 \rightarrow 3$, so there is no directed path from the reference $r = 1$ to observation 2; nevertheless both bounds on θ_2 are nontrivial. Solving the linear programs of Appendix A.1 and inspecting which constraints bind, $\bar{U}_2 = 15$ is determined by the Afriat inequality for the ordered pair $(1, 2)$, whose cross-cost $w_1 \cdot x_2 - c_1 = 4$ is strictly positive, together with the intercept condition at observation 1, which caps $\lambda_1 \leq y_1/c_1 = 5/4$; symmetrically, $\underline{U}_2 = 20/3$ is determined by the pair $(2, 1)$, again a non-edge, together with the intercept at observation 2. No revealed-feasibility relation involving observation 2 and the reference is used. This is why Proposition 3 is a global statement rather than a connectivity statement, and it illustrates that the linear program of Theorem 2 extracts identifying information that Theorem 4 and path arguments alone do not reach.

The second shows that identification strengthens as observations are added. Since new observations add variables, the comparison is with the *projection* of the augmented feasible set onto the original coordinates.

Proposition 5 (Monotonicity of identification). *Let $S' \supseteq S$ augment S by one or more additional observations, retaining the same reference observation r , and let $\text{proj}_S \mathcal{F}_r(S')$ denote the projection of $\mathcal{F}_r(S')$ onto the coordinates $(U_t, \lambda_t)_{t \in \mathbb{T}}$ corresponding to the original observations. If $\mathcal{F}_r(S')$ is nonempty, then*

$$\text{proj}_S \mathcal{F}_r(S') \subseteq \mathcal{F}_r(S).$$

Consequently,

$$\bar{U}_t(S') \leq \bar{U}_t(S), \quad \underline{U}_t(S') \geq \underline{U}_t(S),$$

for every original observation t , and the marginal identified interval for relative productivity under S' is contained in the corresponding interval under S .

The conditioning matters: a new observation can also empty $\mathcal{F}_r(S')$, falsifying the maintained model, a stronger form of information than a tightened bound.

4 Constant Returns to Scale

Constant returns to scale (CRS) removes the linear program: it fixes the radial scaling law $F(\alpha x) = \alpha F(x)$, while the common cardinal scale of F and θ remains indetermi-

nate, and turns the question into closed-form inequalities checkable pairwise, without a reference normalization.

Let

$$v_t := \frac{w_t}{w_t \cdot x_t} = \frac{w_t}{c_t},$$

so that $v_t \cdot x_t = 1$ for all t : input prices per unit of expenditure. The following theorem extends to the Hicks-neutral setting Varian (1984)'s characterization of CRS rationalizability, under which a dataset admits a continuous, concave, monotone, degree-one-homogeneous rationalization if and only if the sums of $\log(v_s \cdot x_t)$ around every cycle are nonnegative.

Theorem 5 (CRS Hicks-neutral rationalizability). *Fix a productivity profile $\theta \in \mathbb{R}_{++}^T$ and define normalized outputs $U_t(\theta) = y_t/\theta_t$. The following statements are equivalent.*

(i) *There exists a function $G : \mathbb{R}_+^L \rightarrow \mathbb{R}_+$, homogeneous of degree one, not assumed a priori to be continuous, concave, or monotone, such that $y_t = \theta_t G(x_t)$ and x_t minimizes $w_t \cdot x$ subject to $G(x) \geq G(x_t)$, for every $t \in \mathbb{T}$.*

(ii) *For every $s, t \in \mathbb{T}$,*

$$v_s \cdot x_t \geq \frac{U_t(\theta)}{U_s(\theta)} = \frac{y_t \theta_s}{y_s \theta_t}. \quad (21)$$

(iii) *The dataset S is Hicks-neutral cost-rationalizable, in the sense of Definition 2, by a common admissible production function F that is additionally homogeneous of degree one; that is, an admissible, continuous, concave, monotone representative can always be chosen for the function G in part (i).*

Since part (i) does not presuppose regularity, the equivalence of (i) and (iii) shows that the testable inequality (21) already forces a fully regular homogeneous representative, the lower envelope $F(x) = \min_s U_s(\theta)(v_s \cdot x)$. In logs, (21) is a system of difference constraints on $\varphi_t := \log \theta_t$; output terms telescope around cycles, so existence of *some* CRS rationalization is, as in Proposition 1, a condition on input data alone (Section 4.4). The inequality also yields explicit bounds without a linear program.

Corollary 4 (Pairwise CRS bounds on relative productivity). *Fix a productivity profile $\theta \in \mathbb{R}_{++}^T$. Suppose the dataset S is Hicks-neutral cost-rationalizable by a common*

constant-returns technology. Then, for every pair $s, t \in \mathbb{T}$,

$$\frac{\theta_t}{\theta_s} \geq \frac{y_t}{y_s(v_s \cdot x_t)}.$$

Consequently, for each t ,

$$\theta_t \geq \max_{s \in \mathbb{T}} \left\{ \theta_s \frac{y_t}{y_s(v_s \cdot x_t)} \right\}.$$

If a reference observation r is normalized by $\theta_r = 1$, then $\theta_t \geq y_t / (y_r(v_r \cdot x_t))$ for all $t \in \mathbb{T}$. These are pairwise bounds: each is valid, and their intersection over s is valid, but the resulting interval is not in general the sharp marginal identified set for θ_t/θ_s once more than two observations are present, since chains of restrictions through intermediate observations may tighten it further. The sharp bounds are the path-based ones of Theorem 7, which take the maximum over all directed paths rather than over single edges; the present corollary is the special case of paths of length one.

Under CRS the general bound of Theorem 4 holds for every pair, not only edges, and is tightened by the factor $v_s \cdot x_t \leq 1$ whenever $s \succeq_D t$.

4.1 Productivity Ranking

The bounds give a revealed-preference criterion for ranking observations.

Corollary 5 (CRS productivity ranking). *Suppose the dataset is Hicks-neutral cost-rationalizable by a common constant-returns technology. If $y_t > y_s(v_s \cdot x_t)$, then every rationalizing productivity profile satisfies $\theta_t > \theta_s$.*

The ranking identifies the *sign* of the productivity difference for every element of the identified set (*ordinal* identification, as opposed to identification of magnitudes), from a fully observable condition requiring neither a normalization nor a linear program.

4.2 Productivity Dominance

The criterion extends to a global notion of productivity superiority.

Proposition 6 (Productivity dominance). *Suppose there exists an observation t such that $y_t > y_s(v_s \cdot x_t)$ for all $s \neq t$. Then every Hicks-neutral CRS rationalization satisfies $\theta_t > \theta_s$ for all $s \neq t$.*

Global productivity leadership can thus be revealed without any normalization, independently of the rationalizing technology.

4.3 A Refutation Test for Constant Returns

The pairwise inequality of Theorem 5(ii) also yields a falsification test requiring neither a profile nor a normalization; what it rejects is stronger than the natural reading suggests.

Proposition 7 (Cross-expenditure refutation of CRS). *Suppose the dataset S is Hicks-neutral cost-rationalizable by a common constant-returns technology at some productivity profile $\theta \in \mathbb{R}_{++}^T$. Then for every pair $s, t \in \mathbb{T}$,*

$$(v_s \cdot x_t)(v_t \cdot x_s) \geq 1. \quad (22)$$

The proof is immediate and is worth displaying, because it shows why the productivity profile plays no role. Theorem 5(ii) applied in both directions gives

$$\frac{\theta_t}{\theta_s} \geq \frac{y_t}{y_s(v_s \cdot x_t)}, \quad \frac{\theta_s}{\theta_t} \geq \frac{y_s}{y_t(v_t \cdot x_s)}.$$

Multiplying the two inequalities, the left-hand sides cancel to one and the output levels cancel as well, leaving $1 \geq 1/[(v_s \cdot x_t)(v_t \cdot x_s)]$, which is (22). Neither θ nor y survives the multiplication.

Corollary 6 (Refutation). *If $(v_s \cdot x_t)(v_t \cdot x_s) < 1$ for some pair s, t , then no Hicks-neutral constant-returns rationalization of S exists at any productivity profile.*

First, (22) is the two-cycle case of the cycle condition (Theorem 7), hence necessary but not sufficient. Second, it does *not* test equal productivity: every profile implies it, because θ cancels, so a violation licenses the much stronger conclusion of Corollary 6, that the data are incompatible with constant returns whatever the productivities. The product measures the cost of a round trip between the two bundles at each

other's normalized prices, which must not be cost-reducing; it is also what fails when the interval of Theorem 6 is empty.

Lemma 1 (Product bound under CRS rationalizability). *Suppose the dataset is Hicks-neutral cost-rationalizable by a common constant-returns technology. Then, for every pair $s, t \in \mathbb{T}$, $(v_s \cdot x_t)(v_t \cdot x_s) \geq 1$.*

Lemma 1 restates Proposition 7 for use in the proof of Theorem 6; it is a restriction on the technology, not on the productivities.

Closed-Form Identification without Normalization

Applying Corollary 4 to (s, t) and (t, s) bounds relative productivity from both sides without a reference or a linear program; Lemma 1 makes the interval nonempty.

Theorem 6 (CRS identification interval). *Suppose the dataset is Hicks-neutral cost-rationalizable by a common constant-returns technology. Then, for every pair s, t ,*

$$\frac{y_t}{y_s(v_s \cdot x_t)} \leq \frac{\theta_t}{\theta_s} \leq \frac{y_t(v_t \cdot x_s)}{y_s}, \quad (23)$$

and this interval is nonempty because $(v_s \cdot x_t)(v_t \cdot x_s) \geq 1$ under rationalizability. This pairwise interval is valid but is not in general the sharp range of θ_t/θ_s ; the sharp two-sided bounds are the path bounds of Theorem 7, of which this interval is the length-one case.

This is the closed-form CRS analogue of Section 3.3.

Connection to the Classical Productivity-Index Literature

Since $v_s \cdot x_t = (w_s \cdot x_t)/(w_s \cdot x_s)$, the interval (23) says that the productivity ratio lies between the Laspeyres-type index $(y_t/y_s)/(w_s \cdot x_t/w_s \cdot x_s)$ and its Paasche counterpart $(y_t/y_s) \cdot (w_t \cdot x_s/w_t \cdot x_t)$, which deflate the output ratio by a base-weighted and a current-weighted input quantity index. The geometric mean of the endpoints, the midpoint of the interval in logs, is

$$\sqrt{\frac{y_t/y_s}{v_s \cdot x_t} \cdot \frac{y_t}{y_s} (v_t \cdot x_s)} = \frac{y_t/y_s}{\sqrt{Q_{st}^L Q_{st}^P}}, \quad Q_{st}^L := \frac{w_s \cdot x_t}{w_s \cdot x_s}, \quad Q_{st}^P := \frac{w_t \cdot x_t}{w_t \cdot x_s},$$

the output ratio deflated by the *Fisher ideal* input quantity index: the midpoint of the bilateral interval is a Fisher-type productivity index. Relating it to Malmquist and Törnqvist constructions requires additional structure: under translog distance functions with common second-order coefficients, [Caves et al. \(1982a\)](#) show that a geometric mean of Malmquist indices equals a Törnqvist index, and [Caves et al. \(1982b\)](#) develop the multilateral counterpart; the identification result requires none of these assumptions. What Theorem 6 adds is a revealed-preference reading of the two *length-one* bounds, which extends [Varian \(1984\)](#)'s bilateral connection to the Hicks-neutral setting. With two observations they are sharp; with more, longer paths can tighten either endpoint, so a ratio between them, the Fisher midpoint included, need not be consistent with the complete dataset. The Fisher measure summarizes the pairwise outer interval; full-data consistency is determined by the sharp all-path bounds of Theorem 7.

4.4 Cycle Characterization and Sharp Path Bounds

Writing $\varphi_t := \log \theta_t$, condition (21) is a system of difference constraints,

$$\varphi_t - \varphi_s \geq \log \frac{y_t}{y_s} - \log(v_s \cdot x_t), \quad \forall s, t \in \mathbb{T}, \quad (24)$$

For a directed path $\pi = (s_0, s_1, \dots, s_k)$ write $w(\pi) := \prod_{i=1}^k (v_{s_{i-1}} \cdot x_{s_i})$, and let $\Pi(s \rightarrow t)$ denote the set of directed paths from s to t . The following theorem consolidates the results above and extends them to arbitrary directed paths.

Theorem 7 (CRS cycle characterization and sharp path bounds). *A Hicks-neutral constant-returns rationalization of S exists for some $\theta \in \mathbb{R}_{++}^T$ if and only if*

$$\prod_{i=1}^k (v_{s_i} \cdot x_{s_{i+1}}) \geq 1 \quad (25)$$

for every directed cycle $(s_1, \dots, s_k, s_1)$ in $\mathbb{T}$. Whenever this holds, fix a reference r and normalize $\theta_r = 1$. The sharp identified interval for θ_t/θ_r is

$$\frac{y_t}{y_r} \max_{\pi \in \Pi(r \rightarrow t)} \frac{1}{w(\pi)} \leq \frac{\theta_t}{\theta_r} \leq \frac{y_t}{y_r} \min_{\pi' \in \Pi(t \rightarrow r)} w(\pi'). \quad (26)$$

Both endpoints are attained, and both are computable from all-pairs longest-path dis-

tances on the weights of (24) in $O(T^3)$ operations.

Theorem 7 asks whether *some* productivity profile admits a constant-returns rationalization. The corresponding question for a *given* profile is settled without optimization by Theorem 5(ii) read directly.

Corollary 7 (Fixed-profile CRS test). *Let $\hat{\theta} \in \mathbb{R}_{++}^T$ be a proposed productivity series and $\hat{U}_t := y_t/\hat{\theta}_t$. Then $\hat{\theta}$ admits a Hicks-neutral constant-returns rationalization if and only if*

$$v_s \cdot x_t \geq \frac{\hat{U}_t}{\hat{U}_s} = \frac{y_t \hat{\theta}_s}{y_s \hat{\theta}_t} \quad \text{for every ordered pair } s, t \in \mathbb{T}. \quad (27)$$

The test requires $O(T^2)$ arithmetic and no optimization, and the largest value of $\hat{U}_t/\hat{U}_s - v_s \cdot x_t$ together with its argmax identifies the pair of years responsible for any rejection.

Unlike the cycle condition (25), in which $\hat{\theta}$ cancels, (27) restricts the series itself. The tests are nested: (27) implies (25) and, since every homogeneous rationalization is admissible, the joint test of Corollary 1.

Cycles of length two give Proposition 7, and paths of length one give Corollary 4 and Theorem 6. The content beyond those cases is that longer paths can bind and that, once all are used, nothing more can be extracted: (26) is the exact range of θ_t/θ_r . Existence again depends on input data alone, since output terms telescope around cycles, and a single Floyd–Warshall pass returns every bound and the cycle condition, a positive cycle appearing as a positive diagonal entry.

5 An Illustrative Example

A three-observation dataset shows how revealed feasibility, outputs, and the intercept condition (4) combine to bound relative productivity.

Example 1. *Let $L = 3$, $T = 3$, and choose observation $r = 1$ as the reference observation.*

t	w_t	x_t	y_t	c_t
1	(3, 1, 2)	(1, 3, 1)	10	8
2	(1, 3, 2)	(3, 1, 1)	14	8
3	(2, 2, 1)	(1, 1, 2)	8	6

Step 1: Revealed-feasibility graph

The relevant cross-costs are

$$\begin{aligned}
w_1 \cdot x_2 - c_1 &= 12 - 8 = 4, & w_1 \cdot x_3 - c_1 &= 8 - 8 = 0, \\
w_2 \cdot x_1 - c_2 &= 12 - 8 = 4, & w_2 \cdot x_3 - c_2 &= 8 - 8 = 0, \\
w_3 \cdot x_1 - c_3 &= 9 - 6 = 3, & w_3 \cdot x_2 - c_3 &= 9 - 6 = 3.
\end{aligned}$$

Hence $1 \succeq_D 3$ and $2 \succeq_D 3$, and no other revealed-feasibility relations hold. The revealed-feasibility graph contains the two edges $1 \rightarrow 3$ and $2 \rightarrow 3$. Since the graph is acyclic, production GARP holds and Proposition 2 implies that the identified set Θ_1^{HN} is nonempty.

Step 2: Productivity normalization

Normalize the reference observation by setting $\theta_1 = 1$, which fixes $U_1 = y_1 = 10$. The edge $1 \succeq_D 3$ yields $w_1 \cdot x_3 - c_1 = 0$. Applying the Afriat inequality with $(s, t) = (1, 3)$:

$$U_3 \leq U_1 + \lambda_1(w_1 \cdot x_3 - c_1) = 10.$$

Consequently,

$$\theta_3 = \frac{y_3}{U_3} \geq \frac{8}{10} = 0.8,$$

exactly the bound of Theorem 4, which holds for the ratio θ_3/θ_1 even without setting $\theta_1 = 1$.

Step 3: A feasible productivity profile

Consider $U_1 = 10$, $U_2 = 10$, $U_3 = 8$, and $\lambda_1 = \lambda_2 = \lambda_3 = 1$. The Afriat inequalities are satisfied. It remains to check the intercept condition (4), $U_s \geq \lambda_s c_s$:

$$U_1 - \lambda_1 c_1 = 10 - 8 = 2, \quad U_2 - \lambda_2 c_2 = 10 - 8 = 2, \quad U_3 - \lambda_3 c_3 = 8 - 6 = 2,$$

so all three are satisfied with slack 2. The corresponding productivity profile is

$$\theta = \left(1, \frac{14}{10}, \frac{8}{8}\right) = (1, 1.4, 1).$$

By Theorem 1, the associated Afriat envelope $\tilde{F}(x) = \min_{s \in \mathbb{T}} \{U_s + \lambda_s(w_s \cdot x - w_s \cdot x_s)\}$ is an admissible, nonnegative production function that rationalizes the dataset.

Step 4: The complete marginal LP bounds

Appendix A.1 solves both linear programs in (12) for θ_2 and θ_3 . The complete marginal identified intervals are

$$\theta_2 \in \left[\frac{14}{15}, 2.1\right] \approx [0.933, 2.1], \quad \theta_3 \in [0.8, 1.2],$$

both reported as closures of the sharp marginal bounds in the sense of Section 3.3.

Discussion

Existence is governed by the input data alone (Step 1). The edge $1 \succeq_D 3$ restricts θ_3/θ_1 to $[0.8, \infty)$ whatever the normalization (Step 2), but the complete system bounds both θ_2 and θ_3 on *both* sides (Step 4), even though observation 2 has no edge to or from the reference.

The marginal intervals do not describe the joint set. Figure 1 displays it exactly: eliminating the multipliers (Appendix A.1) gives $0.8 \leq \theta_3 \leq 1.2$ and $(7/6)\theta_3 \leq \theta_2 \leq (7/4)\theta_3$, a parallelogram strictly inside the product of the marginal intervals. The corner $(\theta_2, \theta_3) = (14/15, 1.2)$ satisfies both marginal bounds but is jointly inadmissible, which is the failure described in Remark 5. Here the joint set happens to be polyhedral; Theorem 3 shows that in general it need not be.




Figure 1: Joint and marginal identification in Example 1, with $\theta_1 = 1$. The dashed rectangle is the product of the marginal intervals; the shaded parallelogram is the exact joint identified set. The cross satisfies each marginal bound but violates the joint restrictions; the dot is the feasible profile of Step 3.

6 An Empirical Application

We take the framework to the NBER-CES Manufacturing Industry Database: we test production GARP, test the complete published productivity series of each industry, compute the marginal bounds of Theorem 2 against a reference year, and confront the constant-returns restrictions of Section 4 with six decades of data. Throughout, a rejection is a rejection of the full maintained model (a common technology, concavity, monotonicity, cost minimization, accurate measurement, valid aggregation, and, for the CRS specification, constant returns), not evidence isolating any one of these assumptions (Section 6.7). The online appendix contains the variable construction, numerical details, supplementary tables, and the complete industry-level results.

6.1 Data

We use version 2021a of the NBER-CES database (Becker et al., 2021), whose deflators follow Bartelsman and Gray (1996): 364 six-digit NAICS-2012 industries observed annually from 1958 to 2018. Three industries enter only in 1997; their 117 empty earlier records are the only observations lost, leaving 22,087 usable industry-years. Each industry’s annual series is one dataset S , so the profile (θ_t) traces Hicks-

neutral technical change within the industry.

Output is real gross shipments, the output concept of the database’s five-factor TFP index `tfp5`. The four variable inputs are production-worker hours, non-production employment, real energy, and real non-energy materials, priced so that cost $c_t = \sum_i w_{it}x_{it}$ equals payroll plus total materials cost exactly (Online Appendix B.1). Only relative prices within a year matter, since a common rescaling is absorbed by the Afriat multipliers.

Capital is treated as quasi-fixed and excluded from the baseline input vector. The common technology should then be read as a short-run technology conditional on capital, $F(x; K_t)$, and representing it by a single capital-free $F(x)$ is an additional maintained assumption: capital differences across years must not shift the variable-input technology, or any such shift must be absorbed by the Hicks-neutral scalar. Section 6.6 reports an alternative specification that prices capital as a fifth input.

6.2 Implementation

The input-side and candidate tests are run on a common support,

$$\mathbb{T}_i^{\text{tfp5}} := \{t : y_t, \hat{\theta}_t, x_t, w_t \text{ are all observed and strictly positive}\}, \quad (28)$$

so that their verdicts are drawn from the same observations. The published index ends in 2016, so this support is 1958–2016 for 361 industries and 1997–2016 for the other three. The input-side test is production GARP, evaluated by the transitive closure of revealed affordability and cross-checked by an Afriat feasibility linear program. The candidate test is the joint test of Corollary 1 at $\hat{\theta}_t = \text{tfp5}_t$, evaluated in closed form by Corollary 2 and cross-checked by the slack program (8); strict positivity is imposed through κ , not through a fixed multiplier floor. We also run the fixed-profile CRS test of Corollary 7. The tests are nested: the fixed-profile CRS test implies the joint test, which implies production GARP, and the code verifies these implications for every industry.

For localization we normalize $\theta_{2012} = 1$, matching the database’s normalization of `tfp5`, and solve the two closure programs (12) for every other year on each industry’s full input-data support, which runs through 2018. By Theorem 2 the closure optima are the supremum and infimum over the strictly positive set; we report

closure intervals without claiming that every endpoint is attained. The CRS cycle condition is checked by a max-plus Floyd–Warshall pass, on the full sample and on non-overlapping windows of 2, 5, 10, and 15 years. Solver tolerances, scaling, residual checks, and the handling of numerical failures, which are never counted as economic rejections, are described in Online Appendix B.2.

6.3 The Joint Test of the Published Series

On the common support the tests partition industries into three outcomes. An *input-side rejection* is a failure of production GARP: no productivity profile can rationalize the industry, and by Proposition 1 the failure concerns input prices and quantities alone. A *candidate-specific rejection* occurs when production GARP passes but the published profile fails the joint test: some Hicks-neutral profile is admissible, and the one in use is not. *Joint acceptance* means that a common admissible technology reconciles the complete published profile with the observed input choices.

Table 2 reports the partition. Production GARP fails in 172 industries and passes in 192, and the complete `tfp5` profile is rejected in every one of the 192. The closed-form verdicts agree with the linear program for all 364 industries. Two rejections are *joint-only*: in household refrigerator and home freezer manufacturing (NAICS 335222) and household laundry equipment manufacturing (NAICS 335224), every annual value lies inside its own marginal interval, yet no common multiplier vector reconciles them. They are the empirical counterpart of Remark 5: the joint test has content beyond the pointwise comparison.

Figure 2 shows the geometry for NAICS 335222, projecting the joint set onto the 2001 and 2006 coordinates with every other year free. The published point (0.763, 1.077) is interior to each marginal interval but lies outside the projected set. In both joint-only industries the conditions of Corollary 2 fail at most source years (53 and 51 of 59), so the conflict is not confined to a single pair of dates.

6.4 Marginal Bounds and Pointwise Localization

Over the full input-data support, 186 of 364 industries (51.1%) pass production GARP. By Proposition 1 this is a statement about the input data, not a test of neutrality; it differs from the 192 above only because the full support adds 2017 and

Table 2: Exact joint specification test of the complete **tfp5** profile

Outcome	Industries	Share of common support	Share conditional on GARP pass
Input-side rejection	172	47.3%	–
Candidate-specific rejection	192	52.7%	100.0%
Joint acceptance	0	0.0%	0.0%
Joint-only rejection	2	0.5%	1.0%

Notes: Industries with a common support $\mathbb{T}_i^{\text{tfp5}}$ of at least two years including the 2012 reference ($N = 364$). “Input-side rejection” is failure of production GARP on $\mathbb{T}_i^{\text{tfp5}}$; “candidate-specific rejection” is production GARP passing while the published profile fails the joint test of Corollary 1; “joint acceptance” is the profile satisfying the full system. “Joint-only rejection” is the subset of candidate-specific rejections whose every annual **tfp5** value nonetheless lies inside its marginal interval. The last column conditions on the GARP-passing industries.




Figure 2: Joint versus marginal identification for a joint-only industry, household refrigerator and home freezer manufacturing (NAICS 335222), projected onto $(\theta_{2001}/\theta_{2012}, \theta_{2006}/\theta_{2012})$ with the remaining annual levels free. The light rectangle is the product of the two marginal intervals; the shaded region is the projection of the closure of the joint identified set, computed on a 90×90 grid of closure feasibility programs. The published **tfp5** point lies inside the rectangle but outside the projected set.

2018. Pass rates vary across sectors, from 11% in nonmetallic minerals and 17% in primary metals to all seven textile mills (Online Appendix Table B.2). The concentration of failures in energy- and materials-intensive industries is consistent with, but does not establish, factor-biased responses to the post-1973 relative-price shocks (Acemoglu, 2002); the test cannot separate this from failures of cost minimization, omitted quasi-fixed inputs, or aggregation (Fisher, 1969).

In the 186 passing industries, every one of the 11,043 non-reference industry-years² obtains a finite marginal interval with a strictly positive lower endpoint, as Proposition 3 requires. The 2012 reference directly affords the bundle of 52.6% of years, and the median interval width is 0.139. Median widths fall from 0.26 in the 1950s to 0.04 in the 2010s (Online Appendix Table B.3). Because this compares different target years within one dataset, it is an empirical regularity rather than an application of Proposition 5.

Figure 3 illustrates with two industries. In light trucks (NAICS 336112) the published index lies inside the marginal interval in 93% of years; the 2000 interval is $[0.858, 1.031]$ against $\mathbf{tfp5} = 0.876$. In turbines (NAICS 333611) it lies above the upper bound in every year; the 2000 interval is $[0.633, 0.638]$ against $\mathbf{tfp5} = 1.002$. Across the 10,671 industry-years with a published value in passing industries, $\mathbf{tfp5}$ lies inside the marginal interval in 25.7% of cases, below it in 19.2%, and above it in 55.1%. These pointwise comparisons localize where a series departs from the identified set; by Remark 5 they are necessary but not sufficient for joint admissibility, and the industry-level verdict is the joint test above. Violations concentrate where intervals are narrow: food, chemicals, and primary metals combine median widths of at most 0.13 with above-bound shares of 69% to 74%, whereas apparel and leather, with median widths above 2, almost never reject.

6.5 Constant Returns to Scale

Over the full sample, the CRS cycle condition fails in all 364 industries. Because outputs telescope out of the cycle condition, this is an input-side verdict: at least one of constant returns, a common technology, or the other maintained assumptions fails at that horizon. Table 3 shows how the verdict depends on the horizon. The

²183 industries contribute 60 non-reference years each and the three late entrants 21 each: $183 \times 60 + 3 \times 21 = 11,043$.




Figure 3: Marginal identified intervals and the published TFP5 index: two industries. The shaded band is the marginal identified interval for θ_t/θ_{2012} , a coordinate projection of the joint set and not the joint set itself; the dashed line is the published five-factor TFP index. The dot marks the reference year 2012, where $\theta_{2012} = 1$ by normalization and $\text{tfp5} = 1$ by construction. A series may lie inside every band and nevertheless fail the joint test. Online Appendix B.6 shows every GARP-passing industry.

general concave model passes in at least 93% of windows up to fifteen years, whereas CRS passes in 74.5% of two-year windows, 19.2% of five-year windows, and almost no window of a decade or more. The windows are identical across columns, so the gap isolates the homogeneity restriction, conditional on the other maintained assumptions.

Table 3: Rationalizability by window length and model

Window length	General concave	CRS
2 years	100.0%	74.5%
5 years	99.5%	19.2%
10 years	96.4%	2.3%
15 years	93.4%	0.9%
Full sample (61 years)	51.1%	0.0%

Notes: Share of windows passing each test, all 364 industries, non-overlapping windows of consecutive years anchored at the start of each industry’s series. “General concave” is production GARP (Definition 7); “CRS” is the cycle condition of Theorem 7. The windows are identical across columns. Longer windows test more inequalities, which the horizon comparison does not net out.

Where the cycle condition holds, Theorem 7 can be taken to the data. Table 4 compares the bilateral Laspeyres–Paasche interval with the sharp all-path interval

on CRS-passing windows. With two years the two coincide, as the theory requires. On longer windows, intermediate years strictly tighten the interval for 66% to 89% of ordered pairs, and the bilateral Fisher midpoint falls outside the sharp interval for 12% to 20% of them, so the Fisher index summarizes the pairwise outer bound rather than a rationalizable ratio (Online Appendix Figure B.1 shows one window). The published profile essentially never satisfies the fixed-profile test of Corollary 7: it passes in 1.6% of CRS-passing two-year windows, in none of the longer ones, and in no industry on the common support.

Table 4: Bilateral outer bounds and sharp CRS path bounds

Window	CRS pass	Strict path tightening	Median Δ^P (tightened)	Fisher excluded	Fixed tfp5 pass
2 years	74.5%	0.0%	—	0.0%	1.6%
5 years	19.2%	65.9%	0.001	12.1%	0.0%
10 years	2.3%	83.0%	0.003	18.6%	0.0%
15 years	0.9%	88.7%	0.005	20.1%	0.0%

Notes: Non-overlapping windows anchored at each industry’s series start, identical to Table 3; “CRS pass” reproduces its constant-returns column. The remaining statistics use only CRS-passing windows and are pooled over all ordered year pairs (r, t) within them. “Strict path tightening” is the share of pairs whose sharp all-path interval of Theorem 7 is strictly narrower than the length-one bilateral interval; Δ^P is the reduction in log width, whose median is reported over the strictly tightened pairs; “Fisher excluded” is the share of pairs whose bilateral Fisher midpoint lies outside the sharp interval. “Fixed **tfp5** pass” is the share of CRS-passing windows (with a complete **tfp5** series) that also satisfy the fixed-profile test of Corollary 7.

Descriptive pairwise statistics, such as the share of year pairs ranked by Corollary 5, are reported in Online Appendix B.3.5.

6.6 Robustness and Test Power

Pricing capital. The published index credits output growth to five inputs including capital services, whereas the baseline disciplines a four-input short-run technology. Adding the real capital stock as a fifth input, priced at $w_K = \text{piinv} \cdot \kappa$ with $\kappa = r + \delta \in \{0.08, 0.10, 0.15\}$, on the sample through 2016, raises the GARP pass count to 219, 229, and 232. This is an empirical finding rather than a logical necessity, since an added priced input can create or destroy affordability relations. The share of published values above the upper bound moves only to 54.8%–58.2%, against 55.1% in the baseline (Online Appendix Table B.4). Pricing capital changes the maintained

technology, so this is an alternative specification rather than a test that leaves the baseline unchanged. Within the grid considered, the upper-bound violations are not an artifact of treating capital as quasi-fixed.

Test power. A pass rate is uninterpretable without knowing how demanding the test is on the data at hand. Following Bronars (1987), we hold observed prices and total cost fixed, allocate cost across inputs by a uniform draw on the simplex, and record how often 1,000 artificial datasets pass at each horizon. For production GARP, random data pass in 70.0% of five-year windows, 5.0% of twenty-year windows, and none at forty or sixty-one years, against observed rates of 99.5%, 87.9%, and 51.1%. For the CRS cycle test, random data pass in 56.3% of two-year windows, 1.0% of five-year windows, and none at ten or fifteen years, against observed rates of 74.5%, 19.2%, 2.3%, and 0.9%. At two years the CRS test therefore has limited power; from five years on it is demanding, and the observed data are far closer to rationalizability than random behavior even where they fail. The two experiments use independent seeds and different horizon grids; Online Appendix B.3.3 reports both.

6.7 Interpretation

Each rejection is a rejection of a joint maintained model. The upward violations are consistent with, though not identified as, markups absorbed by revenue-share weighting (Hall, 1988; De Loecker and Warzynski, 2012; De Loecker et al., 2020); the bounds use cost data only and locate the wedge without attributing it. The joint model behind growth accounting survives at adjacent-year frequencies but not over the decadal horizons at which it is used, and adding constant returns is what turns most passing windows into failures, although Basu and Fernald (1997) find approximately constant returns once utilization, aggregation, and imperfect competition are accounted for, precisely the confounds our test cannot separate. Industry aggregation (Fisher, 1969; Blackorby and Schworm, 1988), deflator-based prices, measurement error, and adjustment costs can each generate violations; Online Appendix B.4 discusses these limitations.

7 Conclusion

This paper develops a nonparametric framework for assessing the consistency of Hicks-neutral productivity measures with cost-minimizing behavior under a common concave technology. We establish necessary and sufficient conditions for rationalizability and characterize the identified set of relative productivity. Unrestricted rationalizability is equivalent to production GARP, while a proposed productivity series can be assessed using an exact joint test and sharp, finite, strictly positive marginal bounds computed by linear programming. Under constant returns to scale, rationalizability is characterized by the absence of positive cycles in a weighted graph, with sharp ratio bounds obtained from longest-path distances. Laspeyres- and Paasche-type bounds are special cases, although their Fisher midpoint need not be jointly rationalizable.

Applying the framework to 364 U.S. manufacturing industries over 1958–2018, we find that approximately half satisfy production GARP. Among these industries, the published NBER–CES productivity index violates marginal bounds in nearly three-quarters of industry-year observations, and no complete industry-level series passes the joint test. In two industries, all annual observations satisfy their marginal bounds despite joint rejection. Constant returns to scale is rejected in every industry over the full sample but is frequently consistent with shorter windows. These results indicate tensions between the data and the joint assumptions of cost minimization, Hicks-neutral technical change, and a stable common technology. Random-behavior benchmarks indicate that the tests are demanding over longer horizons. The violations do not quantify departures from rationalizability: relaxing cost comparisons while maintaining exact output interpolation provides necessary but insufficient conditions for approximate rationalization.

Future work could incorporate productivity dynamics, measurement error, and inference for the identified bounds ([Chernozhukov et al., 2007](#)), as well as address industry aggregation and factor-augmenting technical change. The framework provides a basis for evaluating existing productivity measures against the restrictions of optimizing production behavior.

Appendix A: Proofs

Proof of Theorem 1. Write $U_t := U_t(\theta)$.

(i) implies (ii). Suppose (θ, F) rationalizes the dataset, with F admissible: concave, strictly increasing, continuous, and nonnegative on $\mathbb{R}_+^L$. Strict monotonicity implies local nonsatiation at every x_s : for any $\varepsilon > 0$ the point $x_s + \varepsilon \mathbf{1}$ satisfies $F(x_s + \varepsilon \mathbf{1}) > F(x_s)$ and lies in any ball around x_s for small ε . Fix $s \in \mathbb{T}$ and write the cost-minimization problem (2) at observation s as the convex program

$$\min_{x \in \mathbb{R}_+^L} w_s \cdot x \quad \text{subject to} \quad g_s(x) := U_s - F(x) \leq 0, \quad (29)$$

where g_s is convex because F is concave and the objective is linear. The program satisfies Slater's condition: by strict monotonicity of F , the point $x_s + \varepsilon \mathbf{1}$ satisfies $F(x_s + \varepsilon \mathbf{1}) > F(x_s) = U_s$, hence $g_s(x_s + \varepsilon \mathbf{1}) < 0$, and it lies in $\mathbb{R}_+^L$ for every $\varepsilon > 0$ since $x_s \geq 0$. By Assumption 1, $x_s \in \mathbb{R}_{++}^L$, so the nonnegativity constraints are inactive at x_s and the normal cone $N_{\mathbb{R}_+^L}(x_s)$ is $\{0\}$.

The Karush–Kuhn–Tucker conditions for a convex program under Slater's condition (Rockafellar, 1970, Theorem 28.3) therefore yield a multiplier $\eta_s \geq 0$ and a supergradient $g_s^* \in \partial F(x_s)$ with

$$w_s = \eta_s g_s^*.$$

Necessarily $\eta_s > 0$: if $\eta_s = 0$ the displayed equation would give $w_s = 0$, contradicting $w_s \gg 0$. Setting $\lambda_s := 1/\eta_s > 0$ gives $g_s^* = \lambda_s w_s \in \partial F(x_s)$, and concavity of F gives the supergradient inequality

$$F(x) \leq F(x_s) + \lambda_s w_s \cdot (x - x_s) \quad \text{for all } x \in \mathbb{R}_+^L.$$

Two remarks on this argument. It does not proceed by taking an arbitrary supporting normal to the upper contour set $V_s = \{x : F(x) \geq F(x_s)\}$ and asserting proportionality to w_s : that step would not be valid, because the normal cone to V_s at x_s may be multidimensional and an arbitrary element of it need not be a positive multiple of the observed price vector. For instance, if $F(x) = \min\{2x_1 + x_2, x_1 + 2x_2\}$ and $x_s = (1, 1)$, the normal cone at the kink is generated by $(2, 1)$ and $(1, 2)$, and neither generator is proportional to a price vector such as $w_s = (3/2, 3/2)$ under which x_s

minimizes cost. The KKT route selects the correct element of the normal cone rather than an arbitrary one. Second, Assumption 1 is used here to make the normal cone of $\mathbb{R}_+^L$ vanish at the observed bundle. At a boundary observation the KKT stationarity condition contains an additional normal-cone term, so the observed price vector need not be proportional to a supergradient without further complementary-slackness conditions. Using the already defined multiplier $\lambda_s = 1/\eta_s > 0$ and evaluating at $x = x_t$ yields $U_t \leq U_s + \lambda_s(w_s \cdot x_t - w_s \cdot x_s)$, the Afriat inequality. Evaluating at $x = 0$ and using $F(0) \geq 0$ gives $0 \leq F(x_s) - \lambda_s w_s \cdot x_s = U_s - \lambda_s c_s$, the intercept condition (4).

(ii) *implies (i)*. Suppose the Afriat system, including (4), is feasible and define $\tilde{F}$ by (6). Since $\tilde{F}$ is the minimum of finitely many affine functions with strictly positive slopes $\lambda_s w_s \gg 0$, it is concave and continuous; it is in fact strictly increasing in the strong componentwise order, not merely weakly increasing: if $x' \geq x$ and $x' \neq x$, then $\lambda_s w_s \cdot (x' - x) > 0$ for every s , since every component of $\lambda_s w_s$ is strictly positive and $x' - x \geq 0$ with at least one strictly positive component; taking $\varepsilon := \min_s \lambda_s w_s \cdot (x' - x) > 0$ over the finitely many $s \in \mathbb{T}$ gives $\tilde{F}(x') \geq \tilde{F}(x) + \varepsilon > \tilde{F}(x)$. For each t , the Afriat inequalities state $U_t \leq U_s + \lambda_s w_s \cdot (x_t - x_s)$ for every s , and hence $U_t \leq \min_s \{U_s + \lambda_s w_s \cdot (x_t - x_s)\} = \tilde{F}(x_t)$; the $s = t$ term of the same minimum gives $\tilde{F}(x_t) \leq U_t$ directly, so $\tilde{F}(x_t) = U_t$. For nonnegativity on all of $\mathbb{R}_+^L$: for any $x \geq 0$ and any s , $U_s + \lambda_s w_s \cdot (x - x_s) = (U_s - \lambda_s c_s) + \lambda_s w_s \cdot x \geq 0 + 0 = 0$ by (4), $w_s > 0$, and $x \geq 0$; taking the minimum over s preserves this lower bound, so $\tilde{F}(x) \geq 0$ on $\mathbb{R}_+^L$. For any bundle with $w_t \cdot x \leq w_t \cdot x_t$, the t -th term gives $\tilde{F}(x) \leq U_t + \lambda_t(w_t \cdot x - w_t \cdot x_t) \leq U_t = \tilde{F}(x_t)$, so x_t minimizes cost on the upper contour set of $\tilde{F}$ at x_t . Finally, $\theta_t \tilde{F}(x_t) = \theta_t U_t = y_t$. $\square$

Proof of Proposition 1. Suppose (i) holds. By Theorem 1, the Afriat system $\mathcal{A}(S, \theta)$, including (4), is feasible. The normalized outputs $U_t(\theta) = y_t/\theta_t$ provide a feasible solution to the standard production Afriat inequalities of Definition 7, so production GARP holds.

Conversely, suppose (ii) holds. Let (U_t^*, λ_t^*) be any feasible solution to the standard production Afriat system, including $U_t^* \geq \lambda_t^* c_t$. Define $\theta_t := y_t/U_t^*$. Since $U_t^* > 0$, we have $\theta \in \mathbb{R}_{++}^T$ and $U_t(\theta) = U_t^*$. Hence (U_t^*, λ_t^*) satisfies $\mathcal{A}(S, \theta)$, including the intercept condition, and Theorem 1 yields a Hicks-neutral rationalization. $\square$

Proof of Proposition 2. If $\Theta_r^{HN} \neq \emptyset$, production GARP holds by Proposition 1. Con-

versely, let (U_t^*, λ_t^*) be feasible for the standard production Afriat system. Define $\alpha := y_r/U_r^* > 0$ and $\widehat{U}_t := \alpha U_t^*$, $\widehat{\lambda}_t := \alpha \lambda_t^*$. The Afriat inequalities and the intercept condition (4) are homogeneous of degree one jointly in (U, λ) : multiplying $U_t^* \leq U_s^* + \lambda_s^*(w_s \cdot x_t - w_s \cdot x_s)$ and $U_s^* \geq \lambda_s^* c_s$ by $\alpha > 0$ gives $\widehat{U}_t \leq \widehat{U}_s + \widehat{\lambda}_s(w_s \cdot x_t - w_s \cdot x_s)$ and $\widehat{U}_s \geq \widehat{\lambda}_s c_s$, so $(\widehat{U}_t, \widehat{\lambda}_t)$ remains feasible, and $\widehat{U}_r = \alpha U_r^* = y_r$. Setting $\theta_t := y_t/\widehat{U}_t$ gives $\theta_r = y_r/\widehat{U}_r = 1$, so $\theta \in \Theta_r^{HN}$.

For the second statement, if $\theta, \theta' \in \Theta_r^{HN}$ satisfy $\theta_t/\theta_r = \theta'_t/\theta'_r$ for all t , then since $\theta_r = \theta'_r = 1$ we have $\theta = \theta'$, so the map $\theta \mapsto (\theta_t/\theta_r)_t$ is injective on Θ_r^{HN} ; it is a bijection onto the set of ratio vectors attainable by any rationalizing profile because, by the scaling argument above, every rationalizing profile θ has a representative $a\theta \in \Theta_r^{HN}$ with $a = 1/\theta_r$ attaining the same ratios. $\square$

Proof of Theorem 2. Non-emptiness of $\mathcal{F}_r$ follows from Proposition 2; fix once and for all a point $z^0 = (U^0, \lambda^0) \in \mathcal{F}_r$, so that $U^0 \gg 0$ and $\lambda^0 \gg 0$. We first show that $\mathcal{F}_r$ is dense in $\overline{\mathcal{F}}_r$, where the latter is the nonnegative relaxation obtained by replacing $U_t > 0, \lambda_t > 0$ with $U_t \geq 0, \lambda_t \geq 0$.

Let $z \in \overline{\mathcal{F}}_r$ and, for $\mu \in (0, 1]$, set $z_\mu := (1 - \mu)z + \mu z^0$. Every constraint other than the positivity restrictions is either a weak linear inequality or the linear equality $U_r = y_r$, and both z and z^0 satisfy them; convex combinations therefore satisfy them as well. For the positivity restrictions, each coordinate of z_μ is a sum of a nonnegative term and a strictly positive term, so $z_\mu \gg 0$. Hence $z_\mu \in \mathcal{F}_r$ for every $\mu \in (0, 1]$, and $z_\mu \rightarrow z$ as $\mu \downarrow 0$, which establishes density.

It is worth noting why the more direct perturbation, adding a common $\varepsilon > 0$ to every coordinate, does not work. In the Afriat inequality the multiplier λ_s is multiplied by $w_s \cdot x_t - c_s$, which is *negative* precisely when $s \succeq_D t$. Raising λ_s therefore tightens rather than relaxes the constraint, and a uniform increase in all coordinates can leave the feasible set; it also violates the equality $U_r = y_r$. The convexity argument above avoids both difficulties.

Given density, the supremum of any continuous objective over $\mathcal{F}_r$ equals the supremum over $\overline{\mathcal{F}}_r$, and similarly for the infimum. The quantities $\overline{U}_t$ and $\underline{U}_t$ are therefore the optimal values of the linear programs stated in (12), taken over the closed polyhedron $\overline{\mathcal{F}}_r$; the maximum and minimum are attained on the compact closure. A closure endpoint is attained by a strictly positive rationalization if and only if the

corresponding optimizing face intersects $\mathcal{F}_r$, which need not hold in general. The interval (13) follows from the monotone identity $\theta_t = y_t/U_t$, together with the stated conventions for the degenerate cases $\overline{U}_t = +\infty$ and $\underline{U}_t = 0$. $\square$

Proof of Theorem 3. Convexity of $\mathcal{F}_r$ is immediate from its linear constraints. Density in $\overline{\mathcal{F}}_r$ follows from the convex-combination argument in the proof of Theorem 2. Proposition 3 bounds every U_t above and away from zero on the closure. The intercept condition gives $0 \leq \lambda_t \leq U_t/c_t$, so every multiplier is also bounded. Hence $\overline{\mathcal{F}}_r$ is a compact polytope.

It remains to eliminate the source-specific multipliers. Fix s and write $d_{st} = v_s \cdot x_t - 1$. The inequalities originating at s are

$$U_t - U_s \leq \gamma_s d_{st} \quad \forall t, \quad 0 < \gamma_s \leq U_s,$$

where $\gamma_s = \lambda_s c_s$. If $d_{st} = 0$, feasibility requires $U_t \leq U_s$, equivalently $z_t \geq z_s$. If $d_{sn} < 0$, the upper bound $(U_n - U_s)/d_{sn}$ on γ_s must be strictly positive, so $U_n < U_s$, equivalently $z_n > z_s$.

For $p \in P_s$, comparing the lower bound $(U_p - U_s)/d_{sp}$ with the intercept upper bound U_s gives

$$U_p \leq (1 + d_{sp})U_s \iff z_s \leq (1 + d_{sp})z_p.$$

For $p \in P_s$ and $n \in N_s$, comparing the same lower bound with the upper bound $(U_n - U_s)/d_{sn}$ gives

$$(-d_{sn})U_p + d_{sp}U_n \leq (d_{sp} - d_{sn})U_s.$$

Substituting $U_j = 1/z_j$ and rearranging yields

$$z_s \leq \frac{d_{sp} - d_{sn}}{(-d_{sn})/z_p + d_{sp}/z_n}.$$

These pairwise comparisons, together with the sign conditions above, are necessary and sufficient for the interval of admissible γ_s to be nonempty. Since the blocks separate across s , they give the exact representation in part (ii).

The right-hand side of the last display is a weighted harmonic mean. For $a, b > 0$, the function

$$H(x, y) = \frac{a + b}{a/x + b/y} = \frac{(a + b)xy}{ay + bx}$$

is concave on $\mathbb{R}_{++}^2$: its Hessian is

$$-\frac{2ab(a+b)}{(ay+bx)^3} \begin{pmatrix} y^2 & -xy \\ -xy & x^2 \end{pmatrix},$$

which is negative semidefinite. Therefore each restriction $z_s \leq H(z_p, z_n)$ is a convex hypograph constraint. The remaining restrictions are linear halfspaces, some strict. Their intersection is convex, and the positive diagonal map $z_t \mapsto \theta_t = y_t z_t$ preserves convexity. Replacing strict inequalities by weak ones gives the compact convex closure.

Finally, nonpolyhedrality can arise from a nonredundant harmonic-mean boundary. For example, take one input, three observations with $w_t = 1$, $y_t = 1$, $x_1 = 2$, $x_2 = 3$, $x_3 = 1$, and normalize $\theta_1 = 1$. Eliminating the multipliers yields

$$\frac{1}{2} \leq U_3 < 1, \quad 1 < U_2 \leq 2 - U_3.$$

In productivity coordinates this includes the curved boundary $1/\theta_2 + 1/\theta_3 = 2$, so the identified set is not a polyhedron. $\square$

Proof of Theorem 4. Suppose $r \succeq_D t$. Then $w_r \cdot x_t - c_r \leq 0$. For any θ with $\mathcal{A}(S, \theta)$ feasible (not necessarily normalized), the Afriat inequality with $s = r$ gives $U_t(\theta) \leq U_r(\theta) + \lambda_r(w_r \cdot x_t - c_r) \leq U_r(\theta)$, hence $y_t/\theta_t \leq y_r/\theta_r$, i.e., $\theta_t/\theta_r \geq y_t/y_r$; this derivation nowhere uses $\theta_r = 1$. Restricting to $\theta \in \Theta_r^{HN}$, where $U_r(\theta) = y_r$, gives $U_t(\theta) \leq y_r$ for every such θ , hence $\overline{U}_t \leq y_r$. Necessity of direct affordability for a nontrivial bound is not claimed and is not established by this argument, which only shows sufficiency. $\square$

Proof of Proposition 3. Let $(U, \lambda) \in \overline{\mathcal{F}}_r$, so that $U_r = y_r$.

Upper bound on U_t . The Afriat inequality with $s = r$ gives $U_t \leq y_r + \lambda_r(w_r \cdot x_t - c_r)$, and the intercept condition at $s = r$ gives $y_r = U_r \geq \lambda_r c_r$, so $\lambda_r \leq y_r/c_r$. If $w_r \cdot x_t \leq c_r$ the second term is nonpositive and $U_t \leq y_r$. If $w_r \cdot x_t > c_r$, the coefficient of λ_r is positive, so substituting its upper bound gives

$$U_t \leq y_r + \frac{y_r}{c_r}(w_r \cdot x_t - c_r) = \frac{y_r(w_r \cdot x_t)}{c_r} = y_r(v_r \cdot x_t).$$

In both cases $U_t \leq y_r \max\{1, v_r \cdot x_t\} < \infty$.

Lower bound on U_t . Apply the same two restrictions with the roles of r and t exchanged. The Afriat inequality with $s = t$ and target r gives $y_r = U_r \leq U_t + \lambda_t(w_t \cdot x_r - c_t)$, and the intercept condition at $s = t$ gives $\lambda_t \leq U_t/c_t$. If $w_t \cdot x_r \leq c_t$ the second term is nonpositive and $U_t \geq y_r$. If $w_t \cdot x_r > c_t$, the coefficient of λ_t is positive, so substituting its upper bound gives

$$y_r \leq U_t + \frac{U_t}{c_t}(w_t \cdot x_r - c_t) = \frac{U_t(w_t \cdot x_r)}{c_t} = U_t(v_t \cdot x_r),$$

whence $U_t \geq y_r/(v_t \cdot x_r)$. In both cases $U_t \geq y_r/\max\{1, v_t \cdot x_r\} > 0$.

Taking the infimum and supremum over $\overline{\mathcal{F}}_r$ gives (19). Since $\theta_t/\theta_r = y_t/U_t$ is strictly decreasing in U_t , the bounds invert to (20), with the lower endpoint strictly positive and the upper endpoint finite. $\square$

Proof of Proposition 4. By Theorem 4, $r \succeq_D t$ implies $\theta_t/\theta_r \geq y_t/y_r$ for every $\theta \in \Theta_r^{HN}$, so every candidate ratio value below y_t/y_r is excluded from the marginal identified interval for θ_t/θ_r . This establishes only a sufficient condition for exclusion of a nontrivial region of ratio values; the argument does not address, and this paper does not otherwise establish, whether the absence of a direct edge $r \succeq_D t$ implies the marginal interval is unrestricted, since indirect chains of revealed feasibility could still generate a bound through a sequence of Afriat inequalities not reducible to the single inequality used here. $\square$

Proof of Proposition 5. Let $(U, \lambda) \in \text{proj}_S \mathcal{F}_r(S')$, so there exist values $(U_{t'}, \lambda_{t'})$ for the added observations such that $(U, \lambda, U_{t'}, \lambda_{t'}) \in \mathcal{F}_r(S')$. By definition of $\mathcal{F}_r(S')$, this point satisfies $U_r = y_r$, $U_t > 0$, $\lambda_t > 0$ for all original t , and the Afriat and intercept inequalities (3)–(4) for every pair of observations drawn from S' , which in particular includes every pair drawn from S alone. Hence (U, λ) satisfies every constraint defining $\mathcal{F}_r(S)$, so $(U, \lambda) \in \mathcal{F}_r(S)$. This shows $\text{proj}_S \mathcal{F}_r(S') \subseteq \mathcal{F}_r(S)$ whenever $\mathcal{F}_r(S') \neq \emptyset$; the argument uses no property of the added observations beyond membership in S' and in particular does not require $r \succeq_{D, S'} t'$ for any added t' . Maximizing a linear objective in U_t over the smaller set $\text{proj}_S \mathcal{F}_r(S')$ can only weakly decrease its optimum relative to the same objective over $\mathcal{F}_r(S)$, and minimizing can only weakly increase it, which gives the stated inequalities for $\overline{U}_t$ and $\underline{U}_t$. $\square$

Proof of Theorem 5. (iii) implies (i). Immediate: an admissible representative is in

particular a function G satisfying the requirements of part (i).

(i) *implies* (ii). Suppose G is homogeneous of degree one and x_t minimizes $w_t \cdot x$ subject to $G(x) \geq G(x_t)$, for every t . Fix $s, t \in \mathbb{T}$ and suppose $G(x_t) > 0$ (the case $G(x_t) = 0$ is immediate since the right side of (21) is then not positive). Consider the rescaled bundle $\hat{x} := (G(x_s)/G(x_t))x_t$. By homogeneity of degree one, $G(\hat{x}) = (G(x_s)/G(x_t))G(x_t) = G(x_s)$, so $\hat{x}$ lies in observation s 's upper contour set $\{x : G(x) \geq G(x_s)\}$. Since x_s minimizes $w_s \cdot x$ over this set, $w_s \cdot \hat{x} \geq w_s \cdot x_s = c_s$, i.e., $(G(x_s)/G(x_t))(w_s \cdot x_t) \geq c_s$. Dividing by $c_s G(x_s) > 0$ and using $v_s = w_s/c_s$ gives $v_s \cdot x_t \geq G(x_t)/G(x_s) = U_t(\theta)/U_s(\theta)$, which is (21); the argument uses only homogeneity and cost minimization, not continuity, concavity, or monotonicity of G .

(ii) *implies* (iii). We give an explicit construction rather than appealing to an external characterization. Write $U_t := U_t(\theta) > 0$ and define

$$F(x) := \min_{s \in \mathbb{T}} U_s(v_s \cdot x), \quad x \in \mathbb{R}_+^L. \quad (30)$$

This is the constant-returns analogue of the Afriat envelope (6): a lower envelope of the finitely many linear forms $x \mapsto U_s(v_s \cdot x)$, one for each observation.

Regularity. As a pointwise minimum of finitely many linear functions, F is concave, continuous, and positively homogeneous of degree one, so $F(\kappa x) = \kappa F(x)$ for every $\kappa > 0$. Each coefficient vector $U_s v_s = U_s w_s / c_s$ is strictly positive because $U_s > 0$, $w_s \gg 0$ and $c_s > 0$; hence $F(x) \geq 0$ on $\mathbb{R}_+^L$, and F is strictly increasing in the strong componentwise order: if $x' \geq x$ with $x' \neq x$, then $U_s v_s \cdot (x' - x) > 0$ for every s , and taking $\varepsilon := \min_s U_s v_s \cdot (x' - x) > 0$ over the finitely many $s \in \mathbb{T}$ gives $F(x') \geq F(x) + \varepsilon > F(x)$. Thus F is admissible in the sense of Definition 2 and homogeneous of degree one.

Interpolation. Fix t . The $s = t$ term of the minimum equals $U_t(v_t \cdot x_t) = U_t(w_t \cdot x_t)/c_t = U_t$, so $F(x_t) \leq U_t$. Conversely, condition (ii) states $v_s \cdot x_t \geq U_t/U_s$ for every s , that is, $U_s(v_s \cdot x_t) \geq U_t$; taking the minimum over s gives $F(x_t) \geq U_t$. Hence $F(x_t) = U_t$ for every t , and this is the only step at which (ii) is used.

Cost minimization. Fix t and let $x \in \mathbb{R}_+^L$ satisfy $F(x) \geq F(x_t) = U_t$. By (30) the $s = t$ term of the minimum is at least $F(x)$, so $U_t(v_t \cdot x) \geq U_t$, and dividing by $U_t > 0$ gives $v_t \cdot x \geq 1$, that is, $w_t \cdot x \geq c_t = w_t \cdot x_t$. Therefore x_t minimizes $w_t \cdot x$ on the upper contour set $\{x \in \mathbb{R}_+^L : F(x) \geq F(x_t)\}$.

Finally $\theta_t F(x_t) = \theta_t U_t = y_t$, so F Hicks-neutral cost-rationalizes S together with θ , which is (iii). Because F is continuous, concave, monotone, and homogeneous, the construction also establishes the claim made after the theorem, that imposing these regularity properties in addition to homogeneity is without loss of generality once (ii) holds. $\square$

Proof of Corollary 4. Theorem 5 gives $v_s \cdot x_t \geq y_t \theta_s / (y_s \theta_t)$. Rearranging gives $\theta_t / \theta_s \geq y_t / (y_s (v_s \cdot x_t))$. The remaining inequalities follow by fixing $s = r$ or taking the maximum over all $s \in \mathbb{T}$; joint sharpness across more than two observations simultaneously is not claimed. $\square$

Proof of Corollary 5. By Corollary 4, $\theta_t / \theta_s \geq y_t / (y_s (v_s \cdot x_t)) > 1$ under the maintained hypothesis, so $\theta_t > \theta_s$; this holds for every rationalizing θ , establishing the sign, not the magnitude, of the productivity difference. $\square$

Proof of Proposition 6. Apply Corollary 5 to each pair (s, t) . $\square$

Proof of Proposition 7. Let $\theta \in \mathbb{R}_{++}^T$ be any profile at which S is Hicks-neutral cost-rationalizable by a common constant-returns technology. Corollary 4 gives $\theta_t / \theta_s \geq y_t / (y_s (v_s \cdot x_t))$ and, interchanging s and t , $\theta_s / \theta_t \geq y_s / (y_t (v_t \cdot x_s))$. Both right-hand sides are strictly positive, so multiplying the two inequalities is legitimate and yields

$$1 = \frac{\theta_t}{\theta_s} \cdot \frac{\theta_s}{\theta_t} \geq \frac{y_t}{y_s (v_s \cdot x_t)} \cdot \frac{y_s}{y_t (v_t \cdot x_s)} = \frac{1}{(v_s \cdot x_t)(v_t \cdot x_s)},$$

which is (22). Both the productivity ratio and the output levels cancel, so the conclusion holds for every rationalizing profile and does not require $\theta_t = \theta_s$. The condition is necessary but not sufficient for CRS rationalizability, being the two-cycle case of the cyclical condition following Theorem 5. $\square$

Proof of Corollary 6. Contrapositive of Proposition 7. $\square$

Proof of Lemma 1. Theorem 5 gives $v_s \cdot x_t \geq y_t \theta_s / (y_s \theta_t)$ and, interchanging s and t , $v_t \cdot x_s \geq y_s \theta_t / (y_t \theta_s)$. Multiplying the two inequalities, the factor θ_t / θ_s cancels against its reciprocal, giving $(v_s \cdot x_t)(v_t \cdot x_s) \geq 1$. $\square$

Proof of Theorem 6. Corollary 4 gives $\theta_t/\theta_s \geq y_t/(y_s(v_s \cdot x_t))$. Interchanging s and t gives $\theta_s/\theta_t \geq y_s/(y_t(v_t \cdot x_s))$. Taking reciprocals of the second inequality gives $\theta_t/\theta_s \leq y_t(v_t \cdot x_s)/y_s$. Combining establishes (23). Nonemptiness of the interval, $y_t/(y_s(v_s \cdot x_t)) \leq y_t(v_t \cdot x_s)/y_s$, is equivalent to $(v_s \cdot x_t)(v_t \cdot x_s) \geq 1$, which holds under rationalizability by Lemma 1; joint sharpness with more than two observations present is not claimed, as in Corollary 4. $\square$

Lemma 2. *If $\mathcal{G}(S)$ contains a cycle $t_1, \dots, t_K = t_1$ with $w_{t_k} \cdot x_{t_{k+1}} \leq c_{t_k}$ for all k and $w_{t_j} \cdot x_{t_{j+1}} < c_{t_j}$ for some j , then $\mathcal{A}(S, \theta)$ is infeasible for every $\theta \in \mathbb{R}_{++}^T$.*

Proof of Lemma 2. Suppose (λ_t) is feasible. Summing the Afriat inequalities around the cycle, the left side telescopes to zero and the right side is ≤ 0 , so each term is zero. But the j -th term equals $\lambda_{t_j}(w_{t_j} \cdot x_{t_{j+1}} - c_{t_j}) < 0$ since $\lambda_{t_j} > 0$. Contradiction. $\square$

Proof of Theorem 7

Write $\omega(s \rightarrow t) := \log(y_t/y_s) - \log(v_s \cdot x_t)$, so that (24) reads $\varphi_t - \varphi_s \geq \omega(s \rightarrow t)$, and for a path $\pi = (s_0, \dots, s_k)$ put $\omega(\pi) := \sum_{i=1}^k \omega(s_{i-1} \rightarrow s_i)$.

Cycle characterization. By Theorem 5, a CRS rationalization at θ exists if and only if (21) holds, equivalently if and only if φ solves (24). A finite system of difference constraints admits a solution if and only if the associated weighted digraph has no cycle of strictly positive total weight. Around a cycle $(s_1, \dots, s_k, s_1)$ the terms $\log(y_{s_{i+1}}/y_{s_i})$ telescope to zero, so the total weight is $-\sum_i \log(v_{s_i} \cdot x_{s_{i+1}})$ and non-positivity is $\prod_i (v_{s_i} \cdot x_{s_{i+1}}) \geq 1$, which is (25). Since θ has dropped out, existence for *some* θ is equivalent to the same condition.

Lower bound. Suppose (25) holds and normalize $\varphi_r = 0$. Chaining the constraints along any $\pi \in \Pi(r \rightarrow t)$ gives $\varphi_t \geq \omega(\pi)$, hence $\varphi_t \geq D(r, t) := \max_{\pi \in \Pi(r \rightarrow t)} \omega(\pi)$. The maximum is finite and attained because the absence of positive cycles implies no optimal path repeats a vertex. It is achieved by the potential $\varphi_u := D(r, u)$, which satisfies every constraint since $D(r, u) + \omega(u \rightarrow u') \leq D(r, u')$ by maximality.

Upper bound. Applying the same argument to paths into r , $\varphi_r - \varphi_t \geq \omega(\pi')$ for every $\pi' \in \Pi(t \rightarrow r)$, so $\varphi_t \leq -D(t, r)$, attained by the potential $\varphi_u := -D(u, r)$.

Conversion. Along any path from r to t the output terms telescope to $\log(y_t/y_r)$, so $\omega(\pi) = \log(y_t/y_r) - \log w(\pi)$. Exponentiating $D(r, t) \leq \varphi_t \leq -D(t, r)$ therefore

gives (26). A single Floyd–Warshall pass on ω returns $D(u, u')$ for all ordered pairs, hence both endpoints and the cycle condition, in $O(T^3)$ operations. Restricting both extrema to paths of length one recovers Corollary 4 and Theorem 6, and restricting (25) to cycles of length two recovers Proposition 7. $\square$

A.1 Detailed Linear-Programming Bounds for Example 1

This subsection carries out the full linear-programming computation summarized in Step 4 of Section 5, over the closure with $U_1 = 10$ fixed and $U_2, U_3, \lambda_1, \lambda_2, \lambda_3 \geq 0$.

For $\bar{U}_3$ (equivalently $\underline{\theta}_3$): the binding constraint is $U_3 \leq U_1 + \lambda_1(w_1 \cdot x_3 - c_1) = 10$ from the edge $1 \succeq_D 3$, attained at $\lambda_1 = 0$, so $\bar{U}_3 = 10$ and $\underline{\theta}_3 = y_3/\bar{U}_3 = 8/10 = 0.8$, matching Step 2.

For $\underline{U}_3$ (equivalently $\bar{\theta}_3$): the constraints $U_3 \geq 10 - 3\lambda_3$ (from $U_1 \leq U_3 + \lambda_3(w_3 \cdot x_1 - c_3)$) and $U_3 \geq 6\lambda_3$ (the intercept condition) together imply $U_3 \geq \max\{10 - 3\lambda_3, 6\lambda_3\}$, minimized by equating the two bounds, $10 - 3\lambda_3 = 6\lambda_3$, giving $\lambda_3 = 10/9$ and $\underline{U}_3 = 6 \cdot (10/9) = 20/3$. One checks directly that $U_2 = 20/3$, $\lambda_1 = 0$, $\lambda_2 = 5/6$ completes a feasible point of the closure together with this U_3 and λ_3 , so $\underline{U}_3 = 20/3$ and $\bar{\theta}_3 = y_3/\underline{U}_3 = 8/(20/3) = 1.2$.

The same reasoning, minimizing and maximizing U_2 subject to all Afriat and intercept constraints with U_3 and the λ 's free, gives $\bar{U}_2 = 15$ (attained at $\lambda_1 = 5/4$, $U_3 = 10$, $\lambda_3 = 5/3$) and $\underline{U}_2 = 20/3$ (attained at the same point used for $\underline{U}_3$ above, since $U_2 \geq U_3 \geq 20/3$ binds simultaneously). This yields $\underline{\theta}_2 = y_2/\bar{U}_2 = 14/15 \approx 0.933$ and $\bar{\theta}_2 = y_2/\underline{U}_2 = 14/(20/3) = 2.1$.

Which constraints bind. It is worth recording the active set at the two optima for observation 2, since Remark 7 refers to it. At $\bar{U}_2 = 15$ the binding constraints are the Afriat inequalities for the ordered pairs $(1, 2)$, $(1, 3)$ and $(3, 2)$, with cross-costs +4, 0 and +3 respectively, together with all three intercept conditions; the optimum is $U = (10, 15, 10)$ and $\lambda = (5/4, 15/8, 5/3)$. At $\underline{U}_2 = 20/3$ the binding constraints are the pairs $(2, 1)$, $(2, 3)$ and $(3, 1)$, with cross-costs +4, 0 and +3, again with all three intercepts; the optimum is $U = (10, 20/3, 20/3)$ and $\lambda = (5/4, 5/6, 10/9)$. In both cases the pair carrying the binding restriction between the reference and observation 2 has strictly positive cross-cost and is therefore not an edge of $\mathcal{G}(S)$.

The two marginal minima are not jointly attainable. This is the computation referred

to in Remark 5. Attaining $\theta_2 = 14/15$ requires $U_2 = 15$, and attaining $\theta_3 = 1.2$ requires $U_3 = 20/3$. The Afriat inequality for the pair $(s, t) = (3, 2)$ reads $U_2 \leq U_3 + \lambda_3(w_3 \cdot x_2 - c_3) = U_3 + 3\lambda_3$, so these two values together force $\lambda_3 \geq (15 - 20/3)/3 = 25/9$. The intercept condition at observation 3 reads $U_3 \geq \lambda_3 c_3 = 6\lambda_3$, which with $U_3 = 20/3$ forces $\lambda_3 \leq 10/9$. Since $25/9 > 10/9$, no feasible point attains both, so the pair $(14/15, 1.2)$ lies in the Cartesian product of the two marginal intervals but outside Θ_r^{HN} .

The exact joint set. The system also admits a complete description. With $U_1 = 10$, the Afriat and intercept inequalities are $U_2 \leq 10 + 4\lambda_1$, $U_3 \leq 10$, $8\lambda_1 \leq 10$; $10 \leq U_2 + 4\lambda_2$, $U_3 \leq U_2$, $8\lambda_2 \leq U_2$; and $10 \leq U_3 + 3\lambda_3$, $U_2 \leq U_3 + 3\lambda_3$, $6\lambda_3 \leq U_3$. Eliminating the multipliers through their intercept caps gives the necessary conditions

$$\frac{20}{3} \leq U_3 \leq 10, \quad U_3 \leq U_2 \leq \frac{3}{2}U_3.$$

They are also sufficient: for any such (U_2, U_3) the strictly positive choice $\lambda_1 = 10/8$, $\lambda_2 = U_2/8$, $\lambda_3 = U_3/6$ satisfies every inequality, since $U_2 \leq \frac{3}{2}U_3 \leq 15$ and $U_2 \geq U_3 \geq 20/3$. Substituting $U_2 = 14/\theta_2$ and $U_3 = 8/\theta_3$,

$$\frac{4}{5} \leq \theta_3 \leq \frac{6}{5}, \quad \frac{7}{6}\theta_3 \leq \theta_2 \leq \frac{7}{4}\theta_3,$$

the parallelogram with vertices $(14/15, 4/5)$, $(7/5, 4/5)$, $(21/10, 6/5)$, and $(7/5, 6/5)$ plotted in Figure 1.

The complete marginal identified intervals are therefore

$$\theta_2 \in \left[\frac{14}{15}, 2.1 \right] \approx [0.933, 2.1], \quad \theta_3 \in [0.8, 1.2],$$

reported as sharp marginal bounds. The strictly positive multiplier vectors displayed above verify that each of the four marginal endpoints is individually attained. What fails is only the simultaneous attainment of the particular pair $(14/15, 1.2)$.

References

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies*, 69(4):781–809.

- Akerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.
- Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review*, 8(1):67–77.
- Afriat, S. N. (1973). On a system of inequalities in demand analysis: An extension of the classical method. *International Economic Review*, 14(2):460–472.
- Aguiar, V. H., Kashaev, N., and Allen, R. (2023). Prices, profits, proxies, and production. *Journal of Econometrics*, 235(2):666–693.
- Bartelsman, E. J. and Gray, W. B. (1996). The NBER manufacturing productivity database. NBER Technical Working Paper 205, National Bureau of Economic Research.
- Basu, S. and Fernald, J. G. (1997). Returns to scale in U.S. production: Estimates and implications. *Journal of Political Economy*, 105(2):249–283.
- Becker, R. A., Gray, W. B., and Marvakov, J. (2021). NBER-CES manufacturing industry database (1958–2018, version 2021a). Technical report, National Bureau of Economic Research.
- Blackorby, C. and Schworm, W. (1988). The existence of input and output aggregates in aggregate production functions. *Econometrica*, 56(3):613–643.
- Blundell, R., Browning, M., and Crawford, I. (2008). Best nonparametric bounds on demand responses. *Econometrica*, 76(6):1227–1262.
- Bronars, S. G. (1987). The power of nonparametric tests of preference maximization. *Econometrica*, 55(3):693–698.
- Caves, D. W., Christensen, L. R., and Diewert, W. E. (1982a). The economic theory of index numbers and the measurement of input, output, and productivity. *Econometrica*, 50(6):1393–1414.
- Caves, D. W., Christensen, L. R., and Diewert, W. E. (1982b). Multilateral comparisons of output, input, and productivity using superlative index numbers. *Economic Journal*, 92(365):73–86.

- Cherchye, L., Demuynck, T., De Rock, B., Duprez, C., Magerman, G., and Verschelde, M. (2021). Structural identification of productivity under biased technological change. Working Paper 2021-28, ECARES, Université Libre de Bruxelles.
- Cherchye, L., Demuynck, T., De Rock, B., and Verschelde, M. (2018). Nonparametric production analysis with unobserved heterogeneity in productivity. Working Paper 2018-25, ECARES, Université Libre de Bruxelles.
- Chernozhukov, V., Hong, H., and Tamer, E. (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284.
- De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2):561–644.
- De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.
- Diewert, W. E. and Parkan, C. (1983). Linear programming tests of regularity conditions for production functions. In Eichhorn, W., Henn, R., Neumann, K., and Shephard, R. W., editors, *Quantitative Studies on Production and Prices*, pages 131–158. Physica-Verlag, Würzburg.
- Fisher, F. M. (1969). The existence of aggregate production functions. *Econometrica*, 37(4):553–577.
- Gandhi, A., Navarro, S., and Rivers, D. A. (2020). On the identification of gross output production functions. *Journal of Political Economy*, 128(8):2973–3016.
- Good, D. H., Nadiri, M. I., and Sickles, R. C. (1997). Index number and factor demand approaches to the estimation of productivity. In Pesaran, M. H. and Schmidt, P., editors, *Handbook of Applied Econometrics, Volume II: Microeconomics*, pages 14–80. Blackwell, Oxford.
- Hall, R. E. (1988). The relation between price and marginal cost in U.S. industry. *Journal of Political Economy*, 96(5):921–947.
- Hanoch, G. and Rothschild, M. (1972). Testing the assumptions of production theory: A nonparametric approach. *Journal of Political Economy*, 80(2):256–275.

- Hulten, C. R. (2001). Total factor productivity: A short biography. In Hulten, C. R., Dean, E. R., and Harper, M. J., editors, *New Developments in Productivity Analysis*, volume 63 of *Studies in Income and Wealth*, pages 1–52. University of Chicago Press, Chicago.
- Jorgenson, D. W. and Griliches, Z. (1967). The explanation of productivity change. *Review of Economic Studies*, 34(3):249–283.
- Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.
- Manski, C. F. (2003). *Partial Identification of Probability Distributions*. Springer Series in Statistics. Springer, New York.
- Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.
- Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.
- Solow, R. M. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics*, 39(3):312–320.
- Tamer, E. (2010). Partial identification in econometrics. *Annual Review of Economics*, 2:167–195.
- Varian, H. R. (1984). The nonparametric approach to production analysis. *Econometrica*, 52(3):579–597.
- Varian, H. R. (1990). Goodness-of-fit in optimizing models. *Journal of Econometrics*, 46(1-2):125–140.
- Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters*, 104(3):112–114.