The Impact of Electrification on Quality and Productivity: Evidence from Japanese Cotton Spinning

Chun Pang Chow* Hiroyuki Kasahara† Atsuki Kotani‡ Yoichi Sugita§

January 30, 2026

Abstract

Conventional productivity estimation treats unobserved heterogeneity as a single efficiency term, conflating product quality with physical output efficiency. This paper shows that unobserved quality and quantity-based productivity (TFPQ) can be separately identified from panel data on prices, quantities, and inputs when the two latent variables exhibit differential persistence. A constructive nonparametric identification argument based on transformation models is developed, alongside a parametric estimator for implementation. Applied to Japanese cotton spinning (1914–1920), instrumenting for electric motor adoption with geographic variation in electricity grid access, quality increased by 1.5 percent per percentage point increase in electric motor share while TFPQ remained unchanged. Full electrification raised quality by nearly 60 percent relative to no adoption.

JEL Classification: D24, L11, N65, O33 **Keywords:** Production function, quality, productivity, electrification

*Vancouver School of Economics, University of British Columbia. Email: alexccp@student.ubc.ca
†Vancouver School of Economics, University of British Columbia. Email: hkasahar@mail.ubc.ca
‡Department of Economics, Boston University.
§Faculty of Business and Commerce, Keio University. Email: ysugita@fbc.keio.ac.jp

1 Introduction

Standard productivity measures miss product quality. When a firm improves reliability, uniformity, or other attributes that consumers value, conventional estimates attribute the resulting revenue gains to Hicks-neutral efficiency—or miss them entirely. This mismeasurement has aggregate consequences: it confounds the evaluation of new technologies, distorts measured resource misallocation, and obscures whether firm success reflects supply-side efficiency or demand-side capabilities (**?**). This paper establishes conditions under which unobserved quality and quantity-based productivity (TFPQ) can be separately identified from panel data on prices, quantities, and inputs. The identification argument exploits differential persistence in the stochastic processes governing quality and productivity: quality is more persistent than productivity, and this difference is sufficient for separate recovery. Because quality and productivity load differently on lagged observables, their distinct autoregressive structures can be recovered from the conditional distribution of current prices and quantities given past inputs and outputs. Applied to Japanese cotton spinning (1914–1920), with electric motor adoption instrumented using geographic variation in electricity grid access, the estimates indicate that quality increased substantially while TFPQ remained unchanged.

The simultaneity problem in production function estimation, first formalized by **?**, has given rise to a large literature estimating production functions from firm-level data. The proxy-variable approach of **?** and **?**, refined by **?** to address collinearity and by **?** to exploit nonparametric identification from first-order conditions, constitutes the dominant framework. **?** and **?** extend the framework to allow finite mixture technologies across latent firm types, and **?** address bias from using deflated expenditure as a proxy for input quantities. These extensions accommodate richer heterogeneity but treat unobserved firm differences as a single Hicks-neutral productivity term.

A separate strand emphasizes the distinction between revenue-based and quantity-based productivity. **?** showed that using revenue as output biases estimates when firms have market power. **?** exploited rare plant-level price data to construct both TFPR and TFPQ, finding that TFPQ is negatively correlated with prices while TFPR is positively correlated—confounding efficiency with demand. **?** demonstrate further pitfalls of revenue-based measures when quality and marginal cost are correlated. **?** build on this distinction in their influential misallocation framework, where dispersion in TFPR reflects distortions while TFPQ captures true efficiency differences. The revenue-quantity distinction addresses price variation but leaves quality variation unresolved. A firm may achieve high revenue by producing goods that consumers value more highly rather than by producing more units at lower cost. **?**, **?**, **?**, and **?** document the empirical importance of quality upgrading but do

not provide methods to separately identify quality from productivity when quality is unobserved. **?** demonstrates that revenue production functions are generically non-identified, underscoring the need for structural decomposition of firm heterogeneity.

Existing approaches to incorporating quality in production function estimation face important limitations. Some control for demand shocks without isolating quality as a persistent latent state variable: **?** introduces an explicit demand system to control for unobserved prices, and **?** extend this framework to estimate markups, but in both cases quality remains subsumed in the composite residual. **?** propose markups as a control function for unobserved prices and quality, yet the composite is not decomposed. **?** simultaneously recover demand, productivity, and markups but require restricting the demand specification to achieve identification. **?** decompose firm size into appeal, cost, and product scope but operate at the product level rather than within the production function framework. Others require quality to be directly observed: **?** examine quality using Chinese steel-making data with an observed quality index, demonstrating that quality-adjusted productivity differs substantially from traditional measures—but the approach cannot be applied when quality is unobserved. Still others decompose revenue growth margins without identifying quality and productivity as separate state variables: **?** find that import competition in India primarily drives quality upgrading rather than efficiency gains, but their method recovers growth margins rather than levels. The closest methodological comparator is **?**, who use quantity data and external instruments to address quality and variety bias in production function estimation; however, they do not separately identify quality as a latent variable with its own dynamic process.

This paper establishes conditions under which unobserved quality and productivity can be separately identified as distinct persistent latent variables. Quality captures demand-side attributes—product uniformity, reliability, and characteristics valued by consumers—while productivity captures supply-side efficiency. This represents the first result showing that quality and productivity can be separately identified without observing quality directly. The identification result is general: it applies to any industry where panel data on prices, quantities, and inputs are available and quality and productivity follow processes with differential persistence.

The identification argument proceeds in two stages. In the first stage, two control functions recover latent composites from observables. Revenue monotonicity permits inversion of the demand system, recovering quality from the price-quantity relationship. Material demand monotonicity permits inversion of the input demand function, recovering a quality-productivity composite from input choices. In the second stage, panel dynamics with differential persistence separate the two latent variables. Quality and productivity each follow first-order Markov processes; because they load differently on lagged observables, their dis-

tinct autoregressive structures can be disentangled from the joint distribution of current and past observables. Identification requires that the total persistence of quality differs from that of productivity, a rank condition that is testable via overidentifying restrictions on the reduced-form dynamics.

The identification argument builds on the transformation model literature. Transformation models express an observed variable as an unknown monotone function of a latent index and additive noise; **?** establish nonparametric identification under weak conditions. **?** extend this framework to the analysis of firm heterogeneity, demonstrating that production technologies and consumer demand functions can be jointly identified from revenue data under monopolistic competition. The present paper provides an additional identification argument: **?** establish joint identification of production and demand but do not decompose demand-side heterogeneity into a persistent quality state variable with its own law of motion. Exploiting differential persistence enables this decomposition. The approach also connects to the nonparametric identification literature: **?** establish identification of latent variable models using measurement equations, **?** develops identification in nonseparable models with unobserved heterogeneity, and **?** show how panel structure can identify latent heterogeneity through mixture models. Based on the nonparametric argument, the paper develops a parametric estimator for implementation.

The empirical setting is Japanese cotton spinning between 1914 and 1920, drawing on the plant-level panel compiled by **?** and extended in **?**. Three features of the data enable identification: plant-level observations of output prices, physical output quantities, and input quantities. **?** documented that firm heterogeneity in this industry was driven not by supply-side technical efficiency but by demand-side capabilities, including product quality and customer retention; the present paper provides a structural framework to quantify these demand-side capabilities. The technology shock is the transition from steam to electric power. Steam engines transmitted power through oscillating line shafts that introduced speed variability and uneven yarn tension, while electric motors delivered the constant torque required for uniform output; contemporary engineering accounts emphasize that this stability was critical for fine yarns, where speed fluctuations caused costly thread breakage (**?**). This mechanism is consistent with the productivity paradox described by **?**: the initial benefits of electrification took the form of unmeasured quality upgrading rather than immediate improvements in measured physical efficiency. **?** and **?** provide complementary evidence on electrification and industrial outcomes. To identify the causal effect of electric motor adoption, the paper instruments using a 1909 indicator of electricity grid access (**?**). Because Japan's early electrification relied on hydroelectric generation, municipalities closer to mountainous water resources gained grid access earlier, independent of local demand conditions.

The contribution is both methodological and empirical. The paper provides the first identification result for separately recovering unobserved quality and productivity from panel data, extending the production function literature to accommodate demand-side heterogeneity as a distinct persistent latent variable. The key identifying restriction—differential persistence—is testable via overidentifying restrictions on the reduced-form dynamics. On the empirical side, electrification raised quality by approximately 1.5 percent per percentage point increase in electric motor share while leaving TFPQ unchanged; counterfactual simulations indicate that full electrification raised quality by nearly 60 percent relative to no adoption. These findings have three implications. First, standard productivity measures miss an important channel through which technology affects performance. Second, the results provide micro-level evidence for the **?** hypothesis that general-purpose technology benefits may operate through channels not captured by standard physical productivity measures. Third, they help explain the **?** finding that successful acquirers excelled at demand management: quality improvements may have been the primary margin of adjustment.

The paper proceeds as follows. Section 2 presents the model. Section 3 establishes identification. Section 4 develops the estimator. Section 5 presents data and results. Section 6 reports counterfactuals. Section 7 concludes.

2 Model

Firms differ in productive efficiency and product quality. Both are unobserved by the econometrician but known to the firm. The model generates two control functions that, combined with dynamic restrictions, permit separate identification.

The economy consists of firms indexed by $i = 1, \ldots, N$ observed over periods $t = 1, \ldots, T$. Firm i produces a single differentiated product using materials M_{it}, capital K_{it}, and labor L_{it}. Let $\Delta_{it} = \exp(\delta_{it})$ denote product quality and $\Omega_{it} = \exp(\omega_{it})$ denote productive efficiency. Lowercase letters denote logarithms: y_{it} for output, p_{it} for price, $r_{it} = p_{it} + y_{it}$ for revenue, and $(m_{it}, k_{it}, \ell_{it})$ for inputs. The production technology is

$$y_{it} = f_t(m_{it}, k_{it}, \ell_{it}) + \omega_{it}, \tag{1}$$

where $f_t(\cdot)$ is the production function and ω_{it} is Hicks-neutral TFPQ. The function f_t may vary over time to accommodate industry-level technical change.

Product quality enters demand by augmenting effective output. Under CES preferences with elasticity of substitution $\sigma > 1$, consumers maximize utility over quality-adjusted quan-

tities $\tilde{Y}_{it} = Y_{it} \cdot \Delta_{it}$. The first-order conditions yield inverse demand

$$p_{it} = \alpha_t - \frac{1}{\sigma}(y_{it} + \delta_{it}) + \delta_{it}, \tag{2}$$

where α_t captures aggregate demand conditions. Log revenue is

$$r_{it} = p_{it} + y_{it} = \alpha_t + \frac{\sigma - 1}{\sigma}(y_{it} + \delta_{it}). \tag{3}$$

Define the markup $\mu \equiv \sigma/(\sigma - 1) > 1$, so that $(\sigma - 1)/\sigma = 1/\mu$. The demand specification admits a more general representation:

$$p_{it} = \psi_t(y_{it} + \delta_{it}) + \delta_{it}, \tag{4}$$

where $\psi_t(\cdot)$ is strictly decreasing. Log revenue is $r_{it} \equiv \varphi_t(y_{it} + \delta_{it})$ where $\varphi_t(\cdot) \equiv \psi_t(\cdot) + (\cdot)$.

Assumption 1 (Revenue Monotonicity)**.** $\varphi_t : \mathbb{R} \to \mathbb{R}$ *is strictly increasing with* $\varphi_t'(x) > 0$.

Under CES with $\sigma > 1$, $\varphi_t'(x) = 1/\mu > 0$, so Assumption **??** is satisfied. Monotonicity permits inversion:

$$\delta_{it} = \varphi_t^{-1}(r_{it}) - y_{it}. \tag{5}$$

This constitutes the first control function. Under CES, $\varphi_t^{-1}(r) = \mu(r - \alpha_t)$, so $\delta_{it} = \mu(r_{it} - \alpha_t) - y_{it}$.

Quality and productivity evolve jointly according to a first-order vector autoregression. Let $z_{it} \in \mathbb{R}$ denote an observable state variable that may affect the evolution of both quality and productivity. The joint dynamics are

$$\begin{pmatrix} \omega_{it} \\ \delta_{it} \end{pmatrix} = \begin{pmatrix} \rho_{\omega\omega} & \rho_{\omega\delta} \\ \rho_{\delta\omega} & \rho_{\delta\delta} \end{pmatrix} \begin{pmatrix} \omega_{it-1} \\ \delta_{it-1} \end{pmatrix} + \begin{pmatrix} \lambda_\omega \\ \lambda_\delta \end{pmatrix} z_{it} + \begin{pmatrix} \epsilon_{it} \\ \eta_{it} \end{pmatrix}, \tag{6}$$

where $\rho_{\omega\omega}, \rho_{\omega\delta}, \rho_{\delta\omega}, \rho_{\delta\delta}$ are persistence and cross-effect parameters, λ_ω and λ_δ capture the effects of z_{it} on productivity and quality, and $(\epsilon_{it}, \eta_{it})$ are mean-zero innovations. The off-diagonal terms permit spillovers: $\rho_{\omega\delta} \neq 0$ allows lagged quality to affect current productivity, while $\rho_{\delta\omega} \neq 0$ allows lagged productivity to affect current quality.

Assumption 2 (Timing)**.** *Capital* k_{it}*, labor* ℓ_{it}*, and the state variable* z_{it} *are determined at* $t - 1$ *before* $(\epsilon_{it}, \eta_{it})$ *are realized. Materials* m_{it} *are chosen at* t *after the firm observes* $(\omega_{it}, \delta_{it})$.

Assumption 3 (Orthogonality)**.** $(\epsilon_{it}, \eta_{it}) \perp\!\!\!\perp (k_{it}, \ell_{it}, z_{it}, r_{it-1}, y_{it-1}, m_{it-1}, k_{it-1}, \ell_{it-1}, z_{it-1})$.

Consider the firm's material choice problem. Conditional on (k_{it}, ℓ_{it}) determined at $t-1$, the firm chooses m_{it} to maximize profit $\pi_{it} = R_{it} - P_t^M M_{it}$. Substituting the production function (**??**) into revenue yields $r_{it} = \varphi_t(f_t(m_{it}, k_{it}, \ell_{it}) + \omega_{it} + \delta_{it})$, which depends on $(\omega_{it}, \delta_{it})$ only through the sum $\omega_{it} + \delta_{it}$. The first-order condition

$$\varphi_t'(y_{it} + \delta_{it}) \cdot \frac{\partial f_t}{\partial m} = \frac{P_t^M M_{it}}{R_{it}} \equiv s_{it}^m \tag{7}$$

inherits this property. By the implicit function theorem, optimal material demand takes the form

$$m_{it} = \mathbb{M}_t(\omega_{it} + \delta_{it}, k_{it}, \ell_{it}). \tag{8}$$

Assumption 4 (Monotonicity)**.** *$\mathbb{M}_t(\cdot, k, \ell)$ is strictly increasing in its first argument.*

Monotonicity permits inversion:

$$\omega_{it} + \delta_{it} = \mathbb{M}_t^{-1}(m_{it}, k_{it}, \ell_{it}). \tag{9}$$

This is the second control function. The model generates two control functions: $\delta_{it} = \varphi_t^{-1}(r_{it}) - y_{it}$ from revenue monotonicity and $\omega_{it} + \delta_{it} = \mathbb{M}_t^{-1}(m_{it}, k_{it}, \ell_{it})$ from material demand monotonicity. The functions φ_t^{-1} and $\mathbb{M}_t^{-1}$ are unknown, and cross-sectional data cannot identify them. Panel data introduce additional restrictions through the VAR(1) dynamics. As shown in the next section, separate identification requires that quality and productivity respond differently to their lagged values.

3 Identification

The identification approach exploits the transformation model structure arising from the joint evolution of quality and productivity.

Adding the two equations in (**??**) yields the law of motion for the composite:

$$\omega_{it} + \delta_{it} = (\rho_{\omega\omega} + \rho_{\delta\omega})\omega_{it-1} + (\rho_{\omega\delta} + \rho_{\delta\delta})\delta_{it-1} + (\lambda_\omega + \lambda_\delta)z_{it} + \xi_{it}, \tag{10}$$

where $\xi_{it} \equiv \epsilon_{it} + \eta_{it}$ is the composite innovation. Define

$$\rho_\omega \equiv \rho_{\omega\omega} + \rho_{\delta\omega}, \qquad \rho_\delta \equiv \rho_{\omega\delta} + \rho_{\delta\delta}, \qquad \lambda \equiv \lambda_\omega + \lambda_\delta. \tag{11}$$

The parameter ρ_ω captures the total effect of lagged productivity on the current composite, while ρ_δ captures the total effect of lagged quality. Rewriting (**??**) using $\omega_{it-1} = (\omega_{it-1} +$

$\delta_{it-1}) - \delta_{it-1}$:

$$\omega_{it} + \delta_{it} = \rho_\omega(\omega_{it-1} + \delta_{it-1}) + \nu\delta_{it-1} + \lambda z_{it} + \xi_{it}, \tag{12}$$

where

$$\nu \equiv \rho_\delta - \rho_\omega = (\rho_{\omega\delta} + \rho_{\delta\delta}) - (\rho_{\omega\omega} + \rho_{\delta\omega}) \tag{13}$$

is the differential persistence parameter. When $\nu \neq 0$, lagged quality has a distinct effect on the current composite beyond its contribution to the lagged composite, enabling separate identification.

Let $G_{m|\mathbf{u}}$ denote the conditional distribution of m_{it} given predetermined variables $\mathbf{u}_{it} \equiv (k_{it}, \ell_{it}, z_{it}, r_{it-1}, y_{it-1}, m_{it-1}, k_{it-1}, \ell_{it-1}, z_{it-1})'$. Substituting the control functions into (**??**):

$$\mathbb{M}_t^{-1}(m_{it}, k_{it}, \ell_{it}) = \nu[\varphi_{t-1}^{-1}(r_{it-1}) - y_{it-1}] + \rho_\omega \mathbb{M}_{t-1}^{-1}(m_{it-1}, k_{it-1}, \ell_{it-1}) + \lambda z_{it} + \xi_{it}. \tag{14}$$

Define the index function $h_t(\mathbf{u}_{it}) \equiv \nu[\varphi_{t-1}^{-1}(r_{it-1}) - y_{it-1}] + \rho_\omega \mathbb{M}_{t-1}^{-1}(m_{it-1}, k_{it-1}, \ell_{it-1}) + \lambda z_{it}$, so that

$$\mathbb{M}_t^{-1}(m_{it}, k_{it}, \ell_{it}) = h_t(\mathbf{u}_{it}) + \xi_{it}. \tag{15}$$

Assumption 5 (Regularity)**.** *(a) The distribution of ξ_{it} is absolutely continuous with density g_ξ that is continuous and strictly positive on its support. (b) The predetermined variables $\mathbf{u}_{it}$ are continuously distributed. (c) The functions $\mathbb{M}_t^{-1}$ and φ_t^{-1} are continuously differentiable.*

Proposition 1. *Let Assumptions* **??**–**??** *hold with $\nu \neq 0$. Fix reference points $(m_0^*, m_1^*, k^*, \ell^*)$ in the interior of the support with $m_0^* < m_1^*$, and impose the normalizations $\mathbb{M}_t^{-1}(m_0^*, k^*, \ell^*) = 0$ and $\mathbb{M}_t^{-1}(m_1^*, k^*, \ell^*) = 1$. Then the differential persistence ν, the composite persistence ρ_ω, the composite effect λ, the inverse material demand $\mathbb{M}_t^{-1}$, the production function f_t, the inverse revenue function φ_t^{-1}, and the latent variables $(\delta_{it}, \omega_{it})$ are identified.*

Proof. The idea extends from **?** and **?**. Under Assumption **??**, $\xi_{it} \perp \mathbf{u}_{it}$. By monotonicity of $\mathbb{M}_t^{-1}$ in m, the conditional distribution of m_{it} given $\mathbf{u}_{it}$ satisfies

$$G_{m|\mathbf{u}}(m \mid \mathbf{u}_{it}) = G_\xi\big(\mathbb{M}_t^{-1}(m, k_{it}, \ell_{it}) - h_t(\mathbf{u}_{it})\big). \tag{16}$$

Differentiating with respect to m yields

$$\frac{\partial G_{m|\mathbf{u}}}{\partial m} = g_\xi\big(\mathbb{M}_t^{-1} - h_t\big) \cdot \frac{\partial \mathbb{M}_t^{-1}}{\partial m}. \tag{17}$$

Differentiating with respect to y_{it-1}, noting from the definition of h_t that $\partial h_t/\partial y_{it-1} = -\nu$:

$$\frac{\partial G_{m|\mathbf{u}}}{\partial y_{it-1}} = g_\xi\big(\mathbb{M}_t^{-1} - h_t\big) \cdot \nu. \tag{18}$$

The density g_ξ is evaluated at the same argument in both expressions. Taking the ratio eliminates the unknown density:

$$\frac{\partial \mathbb{M}_t^{-1}(m,k,\ell)}{\partial m} = \nu \cdot \frac{\partial_m G_{m|\mathbf{u}}(m \,|\, \mathbf{u})}{\partial_{y_{t-1}} G_{m|\mathbf{u}}(m \,|\, \mathbf{u})}. \tag{19}$$

Integrating (**??**) over $[m_0^*, m_1^*]$ at reference values $\mathbf{z}^*$ and applying the normalization yields $\nu = S_{y_{t-1}}^{-1}$, where $S_{y_{t-1}} \equiv \int_{m_0^*}^{m_1^*} (\partial_m G_{m|\mathbf{u}}/\partial_{y_{t-1}} G_{m|\mathbf{u}})\, dm$ is identified from data. Substituting back into (**??**) identifies $\partial_m \mathbb{M}_t^{-1}$. Analogous arguments identify $\partial_k \mathbb{M}_t^{-1}$ and $\partial_\ell \mathbb{M}_t^{-1}$. Integrating from the reference point recovers $\mathbb{M}_t^{-1}$.

For the remaining parameters, differentiate (**??**) with respect to m_{it-1} and z_{it}, take ratios with (**??**), and solve for ρ_ω and λ. For the production function, define $\phi_t(m,k,\ell) \equiv \mathbb{E}[r_{it} | m_{it} = m, k_{it} = k, \ell_{it} = \ell]$. Differentiating and combining with (**??**) identifies $\partial f_t/\partial m$. The inverse revenue function follows from $\varphi_t^{-1}(\phi_t) = f_t + \mathbb{M}_t^{-1}$. Quality and TFPQ follow from the control functions. □ □

The rank condition $\nu \neq 0$ is equivalent to $\rho_\delta \neq \rho_\omega$: the total effect of lagged quality on the composite must differ from the total effect of lagged productivity. When $\nu = 0$, the index function $h_t(\mathbf{u}_{it})$ depends only on the lagged composite, and any decomposition with common persistence is observationally equivalent.

Point identification of the full VAR(1) structure requires recovering the individual elements of the persistence matrix.

Assumption 6 (Stationarity)**.** *The VAR(1) process* (**??**) *is covariance stationary with* $\mathbb{E}[\delta_{it}] = \mathbb{E}[\omega_{it}] = 0$ *and eigenvalues of the persistence matrix inside the unit circle.*

Proposition 2. *Let Assumptions* **??**–**??** *hold with* $\nu \neq 0$*. Then the persistence matrix, the coefficient vector* $(\lambda_\omega, \lambda_\delta)$*, and the latent series* $(\delta_{it}, \omega_{it})$ *are point identified.*

Proof. Proposition **??** establishes identification of $\mathbb{M}_t^{-1}$, φ_t^{-1}, ν, ρ_ω, and λ up to scale and location normalizations. Let $(\tilde{\delta}_{it}, \tilde{\omega}_{it})$ denote the latent series constructed under an arbitrary normalization, related to the true series by $\tilde{\delta}_{it} = a + b\delta_{it}$ and $\tilde{\omega}_{it} = c + b\omega_{it}$ for unknown constants (a, b, c) with $b > 0$.

The constructed series satisfy the VAR(1) dynamics up to the affine transformation.

Substituting into (**??**) for quality:

$$\tilde{\delta}_{it} = \tilde{a}_\delta + \rho_{\delta\omega}\tilde{\omega}_{it-1} + \rho_{\delta\delta}\tilde{\delta}_{it-1} + b\lambda_\delta z_{it} + b\eta_{it}, \tag{20}$$

where $\tilde{a}_\delta = a(1-\rho_{\delta\delta}) - c\rho_{\delta\omega}$. Regressing $\tilde{\delta}_{it}$ on $(\tilde{\omega}_{it-1}, \tilde{\delta}_{it-1}, z_{it})$ identifies $(\rho_{\delta\omega}, \rho_{\delta\delta})$ as slope coefficients, regardless of (a, b, c). Analogously, regressing $\tilde{\omega}_{it}$ identifies $(\rho_{\omega\omega}, \rho_{\omega\delta})$.

Stationarity implies $\mathbb{E}[\delta_{it}] = \mathbb{E}[\omega_{it}] = 0$, so $a = \mathbb{E}[\tilde{\delta}_{it}]$ and $c = \mathbb{E}[\tilde{\omega}_{it}]$ are identified as sample means. For scale, the production function must be deterministic, which requires $b = 1$. The coefficients on z_{it} then identify λ_δ and λ_ω directly. □

4 Estimation

We develop a parametric estimator based on the identification results. The parametric specification imposes CES demand, Cobb-Douglas production, and VAR(1) dynamics for the latent variables. Estimation proceeds in four steps:

Step 1: Reduced-form material demand. Under CES preferences with elasticity of substitution $\sigma > 1$, inverse demand is

$$p_{it} = \alpha_t - \frac{1}{\sigma}(y_{it} + \delta_{it}) + \delta_{it}, \tag{21}$$

where α_t captures time-varying aggregate demand conditions. Log revenue follows as

$$r_{it} = \alpha_t + \frac{\sigma - 1}{\sigma}(y_{it} + \delta_{it}). \tag{22}$$

Define the markup $\mu \equiv \sigma/(\sigma-1) > 1$, so that $(\sigma-1)/\sigma = 1/\mu$. Inverting (**??**) yields the quality control function:

$$\delta_{it} = \mu(r_{it} - \alpha_t) - y_{it}. \tag{23}$$

The production technology is Cobb-Douglas:

$$y_{it} = \beta_t + \beta_m m_{it} + \beta_k k_{it} + \beta_\ell \ell_{it} + \omega_{it}, \tag{24}$$

where $(\beta_m, \beta_k, \beta_\ell)$ are output elasticities and β_t captures industry-level technical change. The first-order condition for materials equates marginal revenue product to input price:

$$\frac{1}{\mu} \cdot \beta_m \cdot \frac{R_{it}}{M_{it}} = P_t^M. \tag{25}$$

The materials share in revenue $s_{it}^m \equiv P_t^M M_{it}/R_{it}$ satisfies

$$s_{it}^m = \frac{\beta_m}{\mu}, \qquad \text{equivalently} \qquad \beta_m = \mu \cdot s^m. \tag{26}$$

Taking logarithms of (**??**) yields $r_{it} = m_{it} + \gamma_r$, where $\gamma_r \equiv p_t^m + \ln \mu - \ln \beta_m$ absorbs input prices and structural parameters.

Substituting production (**??**) into revenue (**??**) and using the first-order condition, material demand satisfies

$$m_{it} = \gamma_{0t} + \gamma_\omega \beta_k k_{it} + \gamma_\omega \beta_\ell \ell_{it} + \gamma_\omega(\omega_{it} + \delta_{it}), \tag{27}$$

where

$$\gamma_\omega \equiv \frac{1}{\mu - \beta_m} = \frac{1}{\mu(1 - s^m)} \tag{28}$$

and γ_{0t} absorbs time-varying constants. Quality and productivity follow the VAR(1) process (**??**). Substituting the composite dynamics (**??**) and the lagged control functions into (**??**):

$$\begin{aligned} m_{it} &= \gamma_{0t} + \gamma_\omega \beta_k k_{it} + \gamma_\omega \beta_\ell \ell_{it} + \gamma_\omega \big[\rho_\omega(\omega_{it-1} + \delta_{it-1}) + \nu \delta_{it-1} + \lambda z_{it} + \xi_{it}\big] \\ &= \gamma_{0t} + \gamma_\omega \beta_k k_{it} + \gamma_\omega \beta_\ell \ell_{it} + \rho_\omega(m_{it-1} - \gamma_{0,t-1} - \gamma_\omega \beta_k k_{it-1} - \gamma_\omega \beta_\ell \ell_{it-1}) \\ &\quad + \gamma_\omega \nu[\mu(m_{it-1} + \gamma_r - \alpha_{t-1}) - y_{it-1}] + \gamma_\omega \lambda z_{it} + \gamma_\omega \xi_{it}, \end{aligned} \tag{29}$$

where the second line substitutes $\omega_{it-1} + \delta_{it-1} = \gamma_\omega^{-1}(m_{it-1} - \gamma_{0,t-1} - \gamma_\omega \beta_k k_{it-1} - \gamma_\omega \beta_\ell \ell_{it-1})$ from (**??**) at $t-1$, and the third line substitutes $\delta_{it-1} = \mu(r_{it-1} - \alpha_{t-1}) - y_{it-1} = \mu(m_{it-1} + \gamma_r - \alpha_{t-1}) - y_{it-1}$ from (**??**) at $t-1$.

Collecting terms yields the reduced-form material demand equation:

$$m_{it} = \phi_{0t} + \phi_k k_{it} + \phi_\ell \ell_{it} + \phi_m m_{it-1} + \phi_{k,-1} k_{it-1} + \phi_{\ell,-1} \ell_{it-1} + \phi_y y_{it-1} + \phi_z z_{it} + \tilde{\xi}_{it}, \tag{30}$$

where $\tilde{\xi}_{it} = \gamma_\omega \xi_{it}$ and the reduced-form coefficients are

$$\phi_k = \gamma_\omega \beta_k, \qquad \phi_\ell = \gamma_\omega \beta_\ell, \tag{31}$$

$$\phi_{k,-1} = -\rho_\omega \gamma_\omega \beta_k = -\rho_\omega \phi_k, \qquad \phi_{\ell,-1} = -\rho_\omega \gamma_\omega \beta_\ell = -\rho_\omega \phi_\ell, \tag{32}$$

$$\phi_y = -\gamma_\omega \nu, \qquad \phi_z = \gamma_\omega \lambda, \tag{33}$$

with $\rho_\omega = \rho_{\omega\omega} + \rho_{\delta\omega}$, $\nu = \rho_\delta - \rho_\omega$, and $\lambda = \lambda_\omega + \lambda_\delta$. The coefficient on m_{it-1} is $\phi_m = \rho_\omega + \gamma_\omega \nu \mu$, which combines contributions from the lagged composite and lagged quality.

The key cross-equation restriction from (**??**) is

$$\frac{\phi_{k,-1}}{\phi_k} = \frac{\phi_{\ell,-1}}{\phi_\ell} = -\rho_\omega. \tag{34}$$

Estimate (**??**) by OLS with time fixed effects. If z_{it} is endogenous (correlated with ξ_{it}), instrument with w_{it} (1909 grid access). This yields the reduced-form coefficients $(\hat{\phi}_k, \hat{\phi}_\ell, \hat{\phi}_m, \hat{\phi}_{k,-1}, \hat{\phi}_{\ell,-1}, \hat{\phi}_y, \hat{\phi}_z)$. From the cross-equation restriction (**??**), the composite persistence is

$$\hat{\rho}_\omega = -\frac{\hat{\phi}_{k,-1}}{\hat{\phi}_k} = -\frac{\hat{\phi}_{\ell,-1}}{\hat{\phi}_\ell}. \tag{35}$$

Step 2: Production function and structural parameters. To identify γ_ω, quasi-difference the production function. From (**??**) and the VAR(1) dynamics for productivity, $\omega_{it} = \rho_{\omega\omega}\omega_{it-1} + \rho_{\omega\delta}\delta_{it-1} + \lambda_\omega z_{it} + \epsilon_{it}$. Substituting $\omega_{it} = y_{it} - \beta_t - \beta_m m_{it} - \beta_k k_{it} - \beta_\ell \ell_{it}$ and defining the quasi-difference $\Delta_\rho x_{it} \equiv x_{it} - \rho x_{it-1}$:

$$\Delta_{\rho_{\omega\omega}} y_{it} = \kappa_t + \beta_m \Delta_{\rho_{\omega\omega}} m_{it} + \beta_k \Delta_{\rho_{\omega\omega}} k_{it} + \beta_\ell \Delta_{\rho_{\omega\omega}} \ell_{it} + \tilde{\epsilon}_{it}, \tag{36}$$

where $\kappa_t = \beta_t - \rho_{\omega\omega}\beta_{t-1}$ and $\tilde{\epsilon}_{it}$ absorbs VAR residual terms. Using $\hat{\rho}_\omega$ from Step 1 as a proxy for $\rho_{\omega\omega}$, define $\Delta_{\hat{\rho}} x_{it} \equiv x_{it} - \hat{\rho}_\omega x_{it-1}$. The production function (**??**) becomes

$$\Delta_{\hat{\rho}} y_{it} = \iota_t + \beta_m \Delta_{\hat{\rho}} m_{it} + \beta_k \Delta_{\hat{\rho}} k_{it} + \beta_\ell \Delta_{\hat{\rho}} \ell_{it} + \tilde{\epsilon}_{it}. \tag{37}$$

The quasi-differenced material $\Delta_{\hat{\rho}} m_{it}$ is endogenous (correlated with $\tilde{\epsilon}_{it}$). To address this, substitute the reduced-form material demand (**??**). Applying (**??**):

$$\begin{aligned}
\Delta_{\hat{\rho}} m_{it} &= m_{it} - \hat{\rho}_\omega m_{it-1} \\
&= \phi_{0t} + \phi_k k_{it} + \phi_\ell \ell_{it} + (\phi_m - \hat{\rho}_\omega) m_{it-1} + \phi_{k,-1} k_{it-1} + \phi_{\ell,-1} \ell_{it-1} + \phi_y y_{it-1} + \dots \\
&= \phi_{0t} + \phi_k \Delta_{\hat{\rho}} k_{it} + \phi_\ell \Delta_{\hat{\rho}} \ell_{it} + (\phi_m - \hat{\rho}_\omega) m_{it-1} + \phi_y y_{it-1} + \dots,
\end{aligned} \tag{38}$$

where the last line uses $\phi_{k,-1} = -\hat{\rho}_\omega \phi_k$ and $\phi_{\ell,-1} = -\hat{\rho}_\omega \phi_\ell$.

Multiplying (**??**) by γ_ω and using $\gamma_\omega \beta_m = s^m/(1 - s^m)$, $\gamma_\omega \beta_k = \phi_k$, and $\gamma_\omega \beta_\ell = \phi_\ell$:

$$\gamma_\omega \Delta_{\hat{\rho}} y_{it} = \gamma_\omega \iota_t + \frac{s^m}{1 - s^m} \Delta_{\hat{\rho}} m_{it} + \phi_k \Delta_{\hat{\rho}} k_{it} + \phi_\ell \Delta_{\hat{\rho}} \ell_{it} + \gamma_\omega \tilde{\epsilon}_{it}. \tag{39}$$

Define the composite regressor by substituting (**??**) into the right-hand side of (**??**):

$$\Upsilon_{it} \equiv \frac{\hat{\phi}_k}{1-s^m}\Delta_{\hat{\rho}}k_{it} + \frac{\hat{\phi}_\ell}{1-s^m}\Delta_{\hat{\rho}}\ell_{it} + \frac{s^m}{1-s^m}(\hat{\phi}_m - \hat{\rho}_\omega)m_{it-1} + \frac{s^m}{1-s^m}\hat{\phi}_y y_{it-1}. \tag{40}$$

The coefficient $1/(1-s^m)$ on the quasi-differenced inputs arises from combining the direct effect $\phi_k\Delta_{\hat{\rho}}k$ with the indirect effect through $[s^m/(1-s^m)]\phi_k\Delta_{\hat{\rho}}k$ from (**??**). The lagged terms m_{it-1} and y_{it-1} enter because $\Delta_{\hat{\rho}}m_{it}$ depends on them through the reduced-form dynamics.

Regressing $\Delta_{\hat{\rho}}y_{it}$ on Υ_{it} with time fixed effects:

$$\Delta_{\hat{\rho}}y_{it} = \iota_t + \frac{1}{\gamma_\omega}\Upsilon_{it} + \tilde{\epsilon}_{it}, \tag{41}$$

identifies $\hat{\gamma}_\omega$ as the inverse of the coefficient on Υ_{it}. If z_{it} is endogenous, include it as a regressor and instrument with w_{it}.

The markup and structural elasticities follow:

$$\hat{\mu} = \frac{1}{\hat{\gamma}_\omega(1-s^m)}, \qquad \hat{\sigma} = \frac{\hat{\mu}}{\hat{\mu}-1}, \tag{42}$$

$$\hat{\beta}_m = s^m \cdot \hat{\mu}, \qquad \hat{\beta}_k = \frac{\hat{\phi}_k}{\hat{\gamma}_\omega}, \qquad \hat{\beta}_\ell = \frac{\hat{\phi}_\ell}{\hat{\gamma}_\omega}. \tag{43}$$

Step 3: Latent series construction. Stationarity (Assumption **??**) implies $\mathbb{E}[\omega_{it}] = \mathbb{E}[\delta_{it}] = 0$. Taking expectations of the quality control function (**??**):

$$0 = \mu(\mathbb{E}[r_{it}] - \alpha_t) - \mathbb{E}[y_{it}] \quad \Longrightarrow \quad \hat{\alpha}_t = \bar{r}_t - \frac{\bar{y}_t}{\hat{\mu}}, \tag{44}$$

where bars denote sample means at time t. Taking expectations of the production function (**??**):

$$0 = \mathbb{E}[y_{it}] - \beta_t - \beta_m\mathbb{E}[m_{it}] - \beta_k\mathbb{E}[k_{it}] - \beta_\ell\mathbb{E}[\ell_{it}] \quad \Longrightarrow \quad \hat{\beta}_t = \bar{y}_t - \hat{\beta}_m\bar{m}_t - \hat{\beta}_k\bar{k}_t - \hat{\beta}_\ell\bar{\ell}_t. \tag{45}$$

Construct the latent series from the control functions:

$$\hat{\delta}_{it} = \hat{\mu}(r_{it} - \hat{\alpha}_t) - y_{it}, \tag{46}$$

$$\hat{\omega}_{it} = y_{it} - \hat{\beta}_t - \hat{\beta}_m m_{it} - \hat{\beta}_k k_{it} - \hat{\beta}_\ell \ell_{it}. \tag{47}$$

Step 4: VAR estimation. Estimate the VAR(1) system by regressing the constructed series on their lags and the exogenous shifter:

$$\hat{\omega}_{it} = \hat{\rho}_{\omega\omega}\hat{\omega}_{it-1} + \hat{\rho}_{\omega\delta}\hat{\delta}_{it-1} + \hat{\lambda}_{\omega}z_{it} + \hat{\epsilon}_{it}, \tag{48}$$

$$\hat{\delta}_{it} = \hat{\rho}_{\delta\omega}\hat{\omega}_{it-1} + \hat{\rho}_{\delta\delta}\hat{\delta}_{it-1} + \hat{\lambda}_{\delta}z_{it} + \hat{\eta}_{it}. \tag{49}$$

The diagonal elements $(\hat{\rho}_{\omega\omega}, \hat{\rho}_{\delta\delta})$ capture own persistence; the off-diagonal elements $(\hat{\rho}_{\omega\delta}, \hat{\rho}_{\delta\omega})$ capture spillovers between productivity and quality. The coefficients $(\hat{\lambda}_{\omega}, \hat{\lambda}_{\delta})$ measure the causal effects of the technology shifter on productivity and quality.

Remark 1 (Endogeneity of z_{it})**.** *If z_{it} is endogenous, due to selection on unobservables, the reduced-form coefficient $\hat{\phi}_z$ and VAR coefficients $(\hat{\lambda}_{\omega}, \hat{\lambda}_{\delta})$ are inconsistent. An instrument IV_{it} is required to provide exogenous variation in all steps of the estimation procedure.*

5 Empirical Application: Japanese Cotton Spinning

This section applies the estimation framework to Japanese cotton spinning mills during 1913–1920, a period of rapid electrification and industrial expansion.

5.1 Data and Institutional Background

The analysis draws on plant-level data from **?**. The data covers the universe of member mills and records physical output (bales of yarn), input quantities (raw cotton consumption, spindle capacity, worker counts), output prices, input prices, and power source composition (steam, water, electricity). The final sample in our analysis comprises 493 plant-year observations from mills observed in at least two consecutive years during 1913–1920, and Table **??** reports summary statistics.

Japanese cotton spinning experienced rapid growth and technological transformation during the Meiji and Taishō eras. By 1913, Japan had emerged as the world's fourth-largest cotton textile producer, with an industry structure characterized by large, vertically integrated mills producing standardized yarn counts for domestic and export markets. The industry association facilitated coordination on output standards, quality grading, and information sharing, generating the detailed mill-level records that underpin this analysis.

Cotton yarn quality is determined by uniformity, strength, and cleanliness—attributes that command price premia in downstream weaving and export markets. Uniformity refers to consistency in yarn thickness along the length of the thread; non-uniform yarn produces fabric with visible defects. Strength determines breakage rates during weaving, affecting

Table 1: Summary Statistics

Variable	Mean	SD	Median	Min	Max
Log output	13.42	0.78	13.43	10.01	15.01
Log revenue	18.15	0.78	18.17	14.49	19.97
Log capacity	15.96	0.73	15.91	13.64	17.58
Log workers	11.75	0.73	11.77	8.79	13.42
Log cotton	13.56	0.79	13.58	10.20	15.16
Materials share	0.711	0.182	0.742	0.247	0.989
Electricity engine share	0.240	0.377	0.000	0.000	1.000
Plant age (years)	14.8	7.5	15.7	1	36
1909 grid access	0.688	0.464	1.000	0	1

Notes: $N = 493$ plant-year observations, 1913–1920.

both productivity and final cloth quality. Cleanliness captures the absence of fiber impurities (neps, short fibers, foreign matter) that create surface blemishes.

These quality attributes are partially observable to buyers through standardized grading systems, but considerable within-grade heterogeneity remains. High-quality yarn enables downstream weavers to operate at higher speeds with fewer stoppages, generating willingness-to-pay that manifests in the revenue data as quality premia. The econometric framework identifies quality δ_{it} as the component of revenue not explained by physical output, capturing precisely these premia.

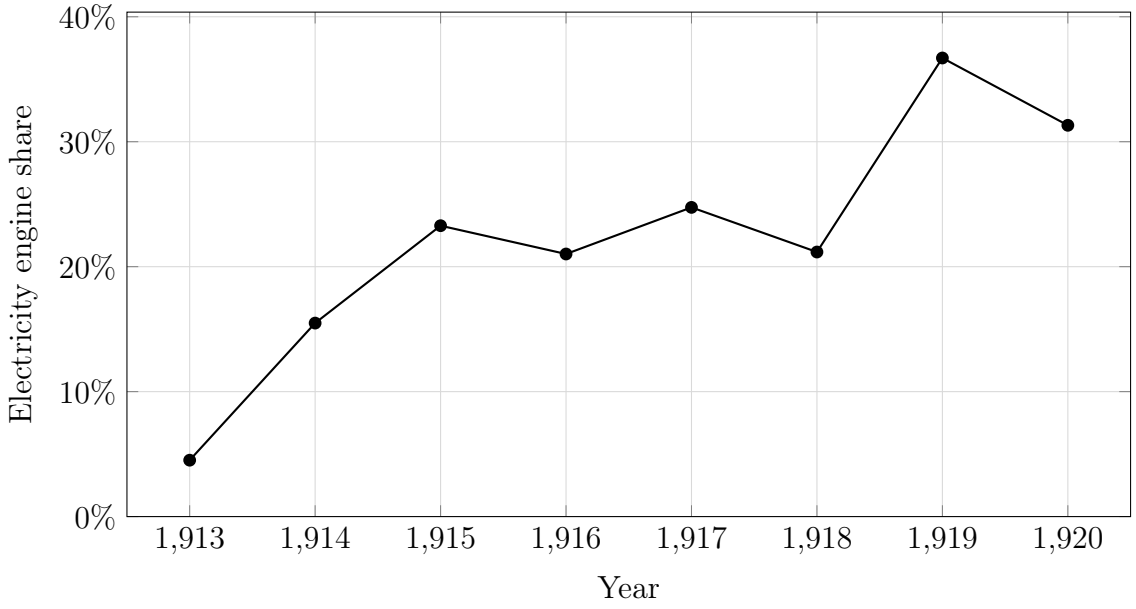


Figure 1: Electricity engine share in Japanese Cotton Spinning, 1913–1920

Figure **??** documents the diffusion of electric power in Japanese cotton spinning. The

Electricity engine share—the fraction of motive power derived from electric motors rather than steam engines or water wheels—rose from under 10 percent in 1913 to over 50 percent by 1920. This transition reflects both grid expansion and mill-level adoption decisions conditional on access.

The hypothesized mechanism linking electrification to quality operates through speed uniformity. Steam engines transmit power via overhead shafts and belt drives, creating speed fluctuations as belts slip and stretch. These fluctuations cause yarn thickness variation: faster spindle rotation produces thinner yarn, slower rotation produces thicker yarn. Electric motors, by contrast, maintain constant rotational speed, eliminating this source of non-uniformity. The engineering literature documents improvements in yarn evenness of 15–25 percent following electrification (**?**).

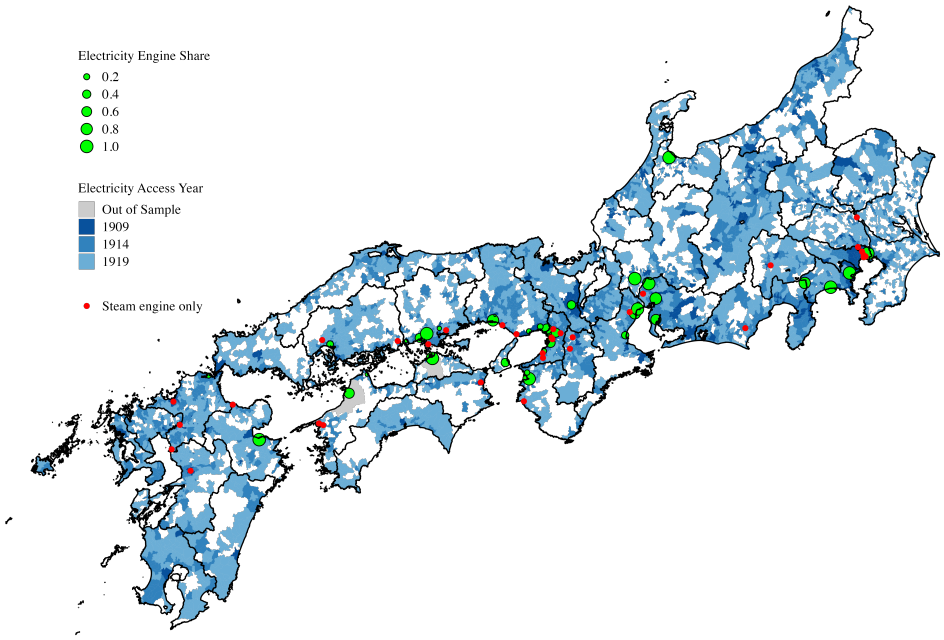
Electrification may also affect productivity through reduced downtime (electric motors require less maintenance than steam engines), flexible factory layout (eliminating the constraint of proximity to the central power shaft), and improved working conditions (reduced heat, noise, and air pollution from steam boilers). The VAR framework allows the data to determine the relative magnitudes of quality versus productivity effects without imposing restrictions on the signs or relative magnitudes.

5.2 Instrumental Variables Strategy

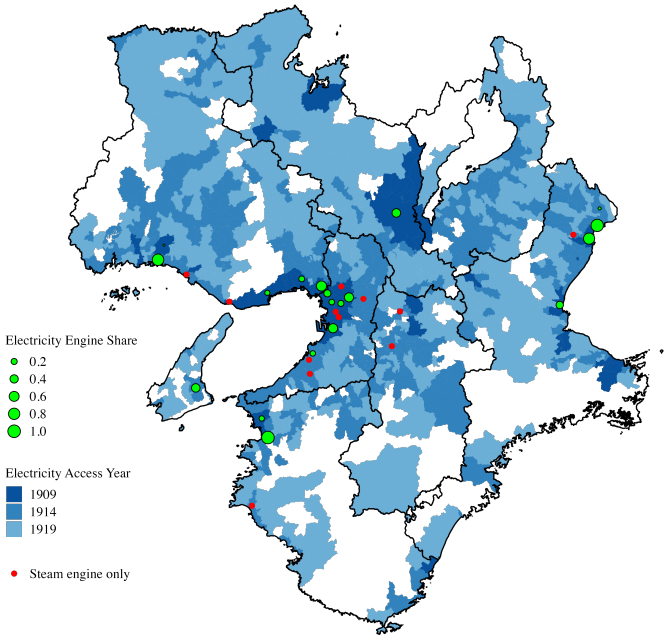
Electricity engine adoption is not randomly assigned. Mills with higher latent productivity or quality may select into early electrification if returns to the technology are heterogeneous and correlated with unobserved mill characteristics. In the notation of Section **??**, z_{it} (Electricity engine share) may be correlated with the VAR innovations $(\epsilon_{it}, \eta_{it})$, rendering OLS estimates of $(\lambda_\omega, \lambda_\delta)$ inconsistent.

The instrument IV_i equals one if mill i is located in a municipality with grid access as of 1909, and zero otherwise. Approximately 69 percent of sample mills had 1909 grid access (Table **??**). The identifying assumption is that 1909 grid access affects mill outcomes during 1913–1920 only through its effect on electricity adoption, which requires that 1909 grid access has no direct effect on productivity or quality conditional on controls. Several features of the setting support this assumption. First, grid access reflects geographic fundamentals (hydroelectric potential) rather than industrial concentration: early grids served diverse customers including residential, commercial, and agricultural users. Second, the four-year lag between the instrument date (1909) and sample start (1913) attenuates concerns about reverse causality from anticipated mill productivity to grid construction. Third, the instrument operates at the municipality level while the outcome variation is at the plant

Figure 2: Geographic Distribution of Electricity Grid Access (1909)



(a) Japan



(b) Kansai Region

Notes: Blue lines indicate high-voltage transmission lines operational by 1909. Red dots represent cotton spinning mills. The instrument is defined as a dummy variable equal to one if a municipality had grid access in 1909.

level, reducing the scope for omitted mill-specific confounders. Figure **??** displays the geographic distribution of electricity adoption during the sample period, showing substantial within-region variation in adoption rates conditional on access.

Table **??** reports first-stage estimates. The coefficient on 1909 grid access is 0.101 (SE = 0.039) in the VAR first stage, indicating that grid access increases electricity engine share by approximately 10 percentage points. Standard errors throughout are computed via wild bootstrap with Rademacher weights to accommodate heteroskedasticity and the multi-step estimation procedure.

Table 2: First-Stage Estimates

	Electricity engine share
1909 grid access	0.101^{***}
	(0.039)
Lagged productivity ($\hat{\omega}_{it-1}$)	0.112
	(0.167)
Lagged quality ($\hat{\delta}_{it-1}$)	0.165^{***}
	(0.049)
Observations	493

Notes: Additional controls include plant age tercile indicators and year. Wild bootstrap (500 replications) standard errors in parentheses.
$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

5.3 Results

Table **??** reports structural parameter estimates from Steps 1–2 of the estimation procedure. The materials, capital , and labor elasticities are $\hat{\beta}_m = 0.808$, $\hat{\beta}_k = 0.035$, and $\hat{\beta}_\ell = 0.096$, implying a returns-to-scale parameter of $\hat{\beta}_m + \hat{\beta}_k + \hat{\beta}_\ell = 0.939$, close to constant returns. The estimated markup is $\hat{\mu} = 1.138$. These estimates are consistent with a competitive industry facing elastic demand.

Table **??** reports the VAR(1) coefficient estimates from Step 4. The diagonal elements reveal substantial persistence in both productivity ($\hat{\rho}_{\omega\omega} = 0.729$) and quality ($\hat{\rho}_{\delta\delta} = 0.308$). Productivity is considerably more persistent than quality. The off-diagonal elements capture dynamic spillovers between the two latent dimensions. The coefficient $\hat{\rho}_{\omega\delta} = 0.010$ (SE = 0.027) is economically small and statistically indistinguishable from zero, indicating that lagged quality does not imply current productivity. In contrast, $\hat{\rho}_{\delta\omega} = -0.405$ (SE = 0.300) is negative, where the sign suggesting a quantity–quality tradeoff.

Table 3: Production Function and Demand Estimates

Parameter	Estimate
Materials elasticity (β_m)	0.808
	(0.038)
Capital elasticity (β_k)	0.035
	(0.088)
Labor elasticity (β_ℓ)	0.096
	(0.022)
Returns to scale ($\sum \beta$)	0.939
	(0.108)
Markup (μ)	1.138
	(0.053)

Notes: $N = 493$. Wild bootstrap (500 replications) standard errors in parentheses.

Table 4: VAR(1) Estimates: Productivity and Quality Dynamics

	Productivity (ω_{it})	Quality (δ_{it})
Lagged productivity (ω_{it-1})	0.729^{***}	-0.405
	(0.066)	(0.300)
Lagged quality (δ_{it-1})	0.010	0.308^{**}
	(0.027)	(0.136)
Electricity engine share (z_{it})	-0.227^{*}	1.488^{**}
	(0.138)	(0.713)
AR statistic	1.43	16.76
AR p-value	0.233	< 0.001

Notes: $N = 493$. IV estimates with 1909 grid access as instrument.
Wild bootstrap standard errors in parentheses.
AR statistic is the Anderson-Rubin test for weak-IV robust inference.
$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The central empirical finding concerns the differential effects of electrification on productivity versus quality. The coefficient $\hat{\lambda}_\omega = -0.227$ (SE = 0.138) indicates that electrification has no detectable effect on productivity, with the robust Anderson-Rubin p-value of 0.233 fails to reject zero at conventional significance levels. The coefficient $\hat{\lambda}_\delta = 1.488$ (SE = 0.713), by contrast, indicates a large and statistically significant effect on quality. The Anderson-Rubin p-value is below 0.001, providing weak-IV robust evidence that electrification causally improves quality.

These findings are consistent with the engineering mechanism: electric motors improve yarn uniformity through speed stability, raising quality, but do not substantially affect the physical quantity of output per unit of inputs. An analysis that conflated quality and productivity into a single "revenue TFP" measure would attribute the quality gains to productivity, overstating the efficiency improvements from electrification.

The VAR specification includes plant age tercile indicators as additional exogenous shifters to capture learning-by-doing effects. In the present context, older plants may exhibit different quality and productivity levels due to accumulated operational experience, vintage of installed equipment, or organizational routines that evolve with tenure. The estimates indicate that older plants exhibit higher quality relative to young plants ($\hat{\gamma}_{\delta,\text{age2}} = 0.522$, SE = 0.240; $\hat{\gamma}_{\delta,\text{age3}} = 0.494$, SE = 0.230), suggesting that cumulative experience improves product quality through refined operational practices. Plant age effects on productivity are negative but statistically insignifcant ($\hat{\gamma}_{\omega,\text{age2}} = -0.058$, SE = 0.042; $\hat{\gamma}_{\omega,\text{age3}} = -0.084$, SE = 0.045), potentially reflecting vintage effects whereby older equipment embodies less advanced technology. These patterns are consistent with learning-by-doing operating primarily through the quality dimension.

6 Counterfactual Analysis

6.1 Methodology

The counterfactual analysis quantifies electrification's contribution to quality, productivity, and firm performance by simulating plant trajectories under alternative technology adoption scenarios. We restrict attention to plants observed continuously from 1915 through 1920, yielding a balanced panel of fourteen plants with complete adoption histories. This restriction ensures that counterfactual comparisons are not confounded by compositional changes in the sample. For each plant i, initial conditions in the base year t_0 are computed from observed data using the estimated structural parameters. TFPQ is recovered by $\omega_{i,t_0} = y_{i,t_0} - \hat{\beta}_k k_{i,t_0} - \hat{\beta}_\ell \ell_{i,t_0} - \hat{\beta}_m m_{i,t_0} - \hat{\beta}_t$, and Quality is recovered by inverting the CES demand

system:

$$\delta_{i,t_0} = \mu\,(r_{i,t_0} - \hat{\alpha}) - y_{i,t_0}. \tag{50}$$

Given initial conditions, quality and TFPQ evolve forward according to the estimated VAR(1) system from Table **??**. Letting $\mathbf{s}_{it} = (\omega_{it}, \delta_{it})'$ denote the state vector, the law of motion under scenario $\mathcal{S}$ is $\mathbf{s}_{it}^{\mathcal{S}} = \hat{\mathbf{P}}\,\mathbf{s}_{i,t-1}^{\mathcal{S}} + \hat{\boldsymbol{\lambda}}\,\tilde{z}_{i,t-1}^{\mathcal{S}} + \hat{\boldsymbol{\gamma}}\,\tilde{\mathbf{x}}_{i,t-1} + \hat{\boldsymbol{\rho}}_t\,\tilde{t} + \boldsymbol{\varepsilon}_{it}$, where $\hat{\mathbf{P}}$ is the persistence matrix, $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}_\omega, \hat{\lambda}_\delta)'$ captures the direct effects of electrification, $\tilde{z}_{it}^{\mathcal{S}} = z_{it}^{\mathcal{S}} - \bar{z}$ is the demeaned electric engine share under scenario $\mathcal{S}$, $\tilde{\mathbf{x}}_{it}$ contains demeaned equipment age tercile indicators capturing learning-by-doing effects, and $\boldsymbol{\varepsilon}_{it} = (\eta_{it}, \epsilon_{it})'$ captures idiosyncratic shocks to productivity and quality. A key feature of the methodology is that these idiosyncratic shocks are held fixed across scenarios. For each plant-year observation, residuals are recovered from the estimated VAR and then added to counterfactual state trajectories, ensuring that any differences across scenarios are attributable solely to counterfactual variation in electric engine shares rather than to differences in unobserved shocks.

The analysis considers three scenarios spanning the range of possible electrification outcomes. Under *Actual Adoption*, electric engine shares follow their observed historical path, $z_{it}^{A} = z_{it}^{\text{obs}}$, serving as the baseline for comparison. Under *No Adoption*, electric engine shares are set to zero for all plants in all periods, $z_{it}^{B} = 0$, representing a counterfactual world in which electrification technology was never adopted. Under *Full Grid Expansion*, electric engine shares are predicted under the assumption that all plants had access to the electricity grid in 1909, with the predicted share following from the first-stage regression using the recovered residuals and setting the grid access indicator to unity for all plants.

A crucial distinction of this counterfactual exercise is that firms are permitted to reoptimize their material input choices in response to counterfactual changes in quality and productivity. This equilibrium adjustment is derived from the firm's static profit maximization problem under CES demand and Cobb-Douglas production. Solving for the equilibrium material input under counterfactual states $(\omega_{it}^{\mathcal{S}}, \delta_{it}^{\mathcal{S}})$ requires substituting the structural equations into the first-order condition. Defining $c_{yit}^{\mathcal{S}} \equiv \beta_t + \beta_k k_{it} + \beta_\ell \ell_{it} + \omega_{it}^{\mathcal{S}}$ as the production capacity excluding materials, the equilibrium log material input satisfies:

$$m_{it}^{\mathcal{S}} = \frac{\alpha + \frac{\sigma-1}{\sigma}\delta_{it}^{\mathcal{S}} + \left(1 - \frac{1}{\sigma}\right) c_{yit}^{\mathcal{S}} + \log\left(\frac{\sigma-1}{\sigma}\right) + \log\beta_m - w_{Mit}}{1 - \beta_m\left(1 - \frac{1}{\sigma}\right)}. \tag{51}$$

This expression captures the key economic intuition: higher quality raises the marginal revenue product of materials, inducing firms to expand material usage, and higher productivity has a similar effect through the production capacity term. The denominator ensures a well-defined interior equilibrium. Given counterfactual states and equilibrium material input,

output follows from the production function $y_{it}^{\mathcal{S}} = \beta_t + \beta_k k_{it} + \beta_\ell \ell_{it} + \beta_m m_{it}^{\mathcal{S}} + \omega_{it}^{\mathcal{S}}$, the equilibrium price satisfies $p_{it}^{\mathcal{S}} = \alpha + [(\sigma-1)/\sigma]\delta_{it}^{\mathcal{S}} - (1/\sigma)y_{it}^{\mathcal{S}}$, and revenue is $r_{it}^{\mathcal{S}} = \alpha + \mu^{-1}(y_{it}^{\mathcal{S}} + \delta_{it}^{\mathcal{S}})$. Capital and labor quantities are held fixed at observed values.

6.2 Results

Table **??** reports average outcomes for 1915–1920 under each scenario, using the parameter estimates from Tables **??** and **??**. The results are aggregated over the post-initialization period to allow the VAR dynamics and equilibrium adjustments to accumulate.

Table 5: Counterfactual Analysis: Average Outcomes, 1915–1920

	Actual	No Adoption	Full Grid
Electric Engine Share	0.203	0.000	0.257
Quality (δ)	1.125	0.665 [−0.460]	1.249 [+0.124]
TFPQ (ω)	0.672	0.778 [+0.106]	0.644 [−0.028]
Log Materials	12.89	12.74 [−0.15]	12.94 [+0.05]
Log Output	13.30	13.22 [−0.08]	13.34 [+0.04]
Log Price	5.73	5.32 [−0.41]	5.85 [+0.12]
Log Revenue	19.04	18.54 [−0.50]	19.19 [+0.15]

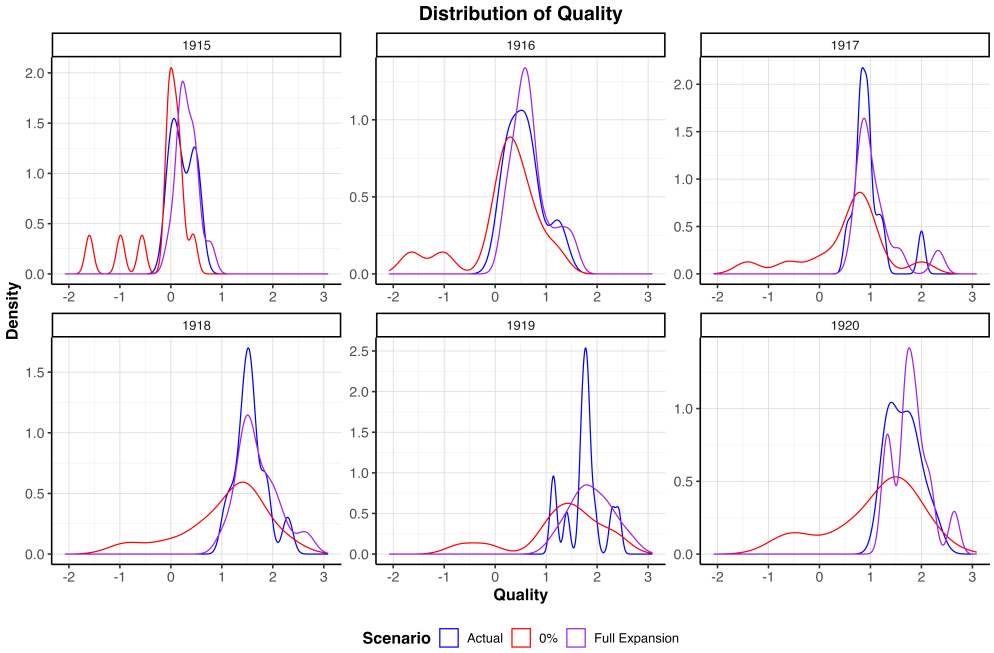
Notes: Averages across plant-years, 1915–1920. Differences from Actual in brackets. Material inputs reoptimized via equation (**??**). Sample: 14 plants with consecutive observations 1915–1920.

The counterfactual reveals substantial quality effects from electrification. Under no adoption, average quality would be 0.665 versus 1.125 under actual adoption. Full grid expansion would raise quality further to 1.249. These large quality effects reflect the cumulative impact of electrification operating through the VAR dynamics: electrification raises quality directly through $\hat{\lambda}_\delta > 0$, and quality improvements partially persist across periods through $\hat{\rho}_{\delta\delta} > 0$. TFPQ exhibits the opposite pattern, with slightly higher values under no adoption than under actual, reflecting the quality-productivity trade-off embedded in the VAR dynamics through $\hat{\rho}_{\delta\omega} < 0$. Plants that invest in quality accept lower physical throughput in exchange for premium prices.

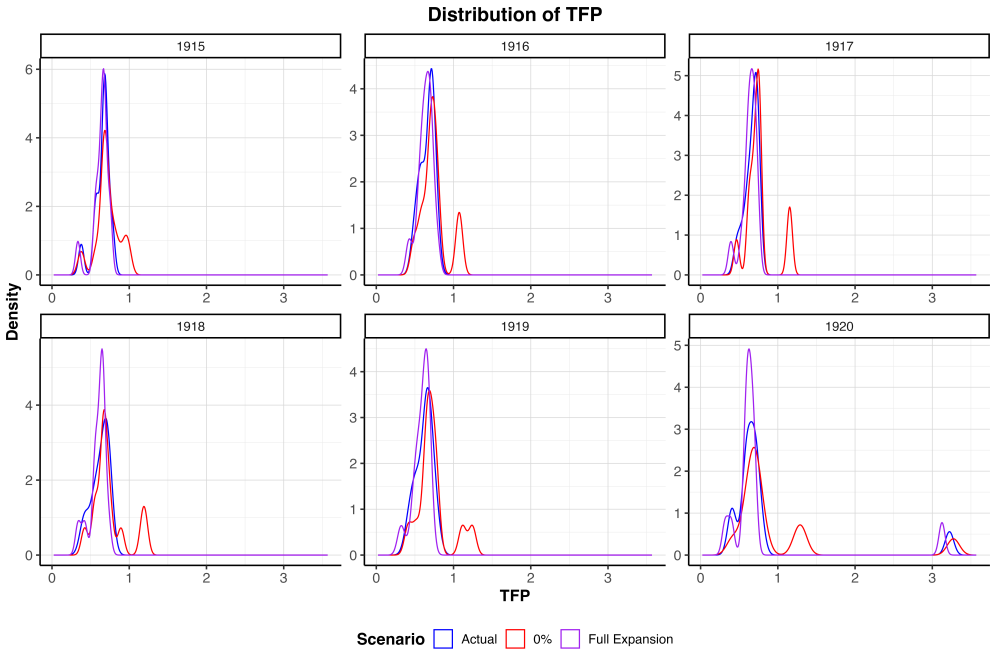
The equilibrium material adjustment amplifies the revenue effects of electrification. Under no adoption, firms would use 0.15 log points less material than under actual adoption, reflecting the lower marginal revenue product of inputs when quality is diminished. This input contraction contributes to output that is 0.08 log points lower than actual—in contrast to a naive counterfactual holding materials fixed, which would show higher output under no adoption due to the higher TFPQ. Prices are substantially lower under no adoption: 5.32 versus 5.73 log points, a difference of 0.41 log points reflecting quality commanding higher willingness to pay under CES demand. Combining the output and price effects with equilibrium material adjustment, log revenue is 0.50 log points higher under actual adoption than under no adoption, corresponding to 65% in levels. Full grid expansion would raise revenue by an additional 0.15 log points above actual.

Figures **??**, **??**, and **??** display the distribution of counterfactual outcomes across the 1915–1920 period, illustrating how differences accumulate over time and how firms respond through equilibrium adjustment.

Figure 3: Counterfactual Distributions: Quality and TFPQ (1915–1920)



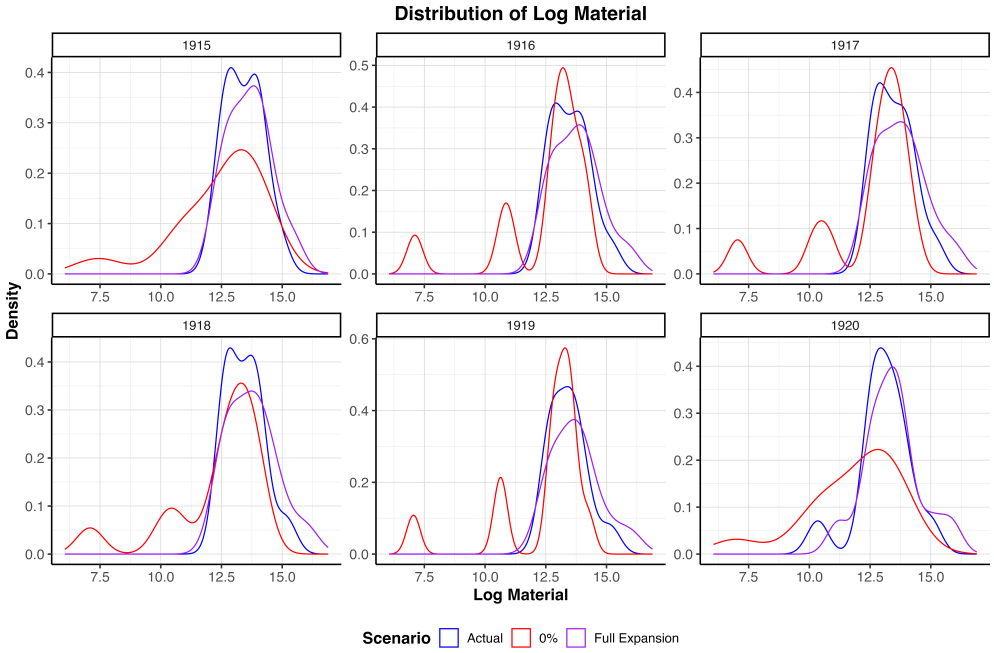
(a) Quality (δ)



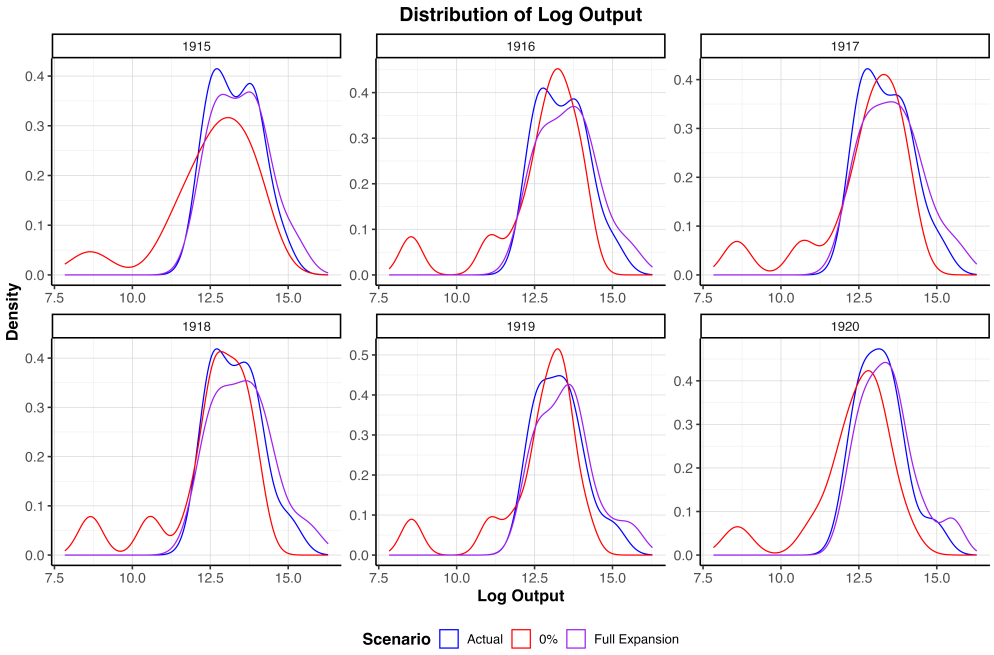
(b) TFPQ (ω)

Notes: Kernel density estimates under three scenarios: actual adoption (blue), no adoption (red), and full grid expansion (purple). Continued in Figures **??** and **??**.

Figure 4: Counterfactual Distributions: Materials and Output (1915–1920)

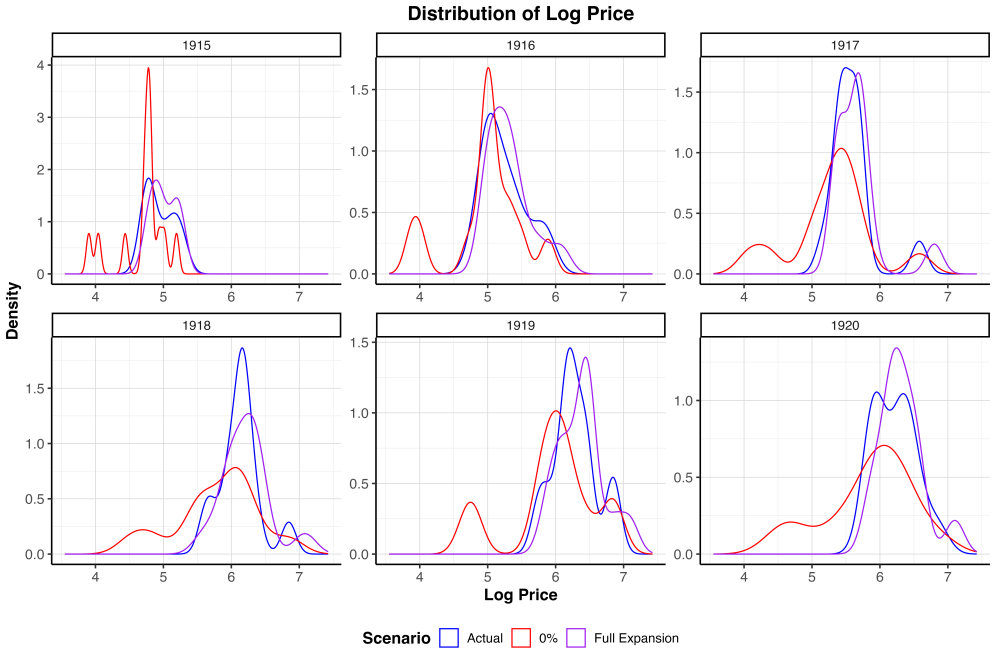


(a) Log Materials

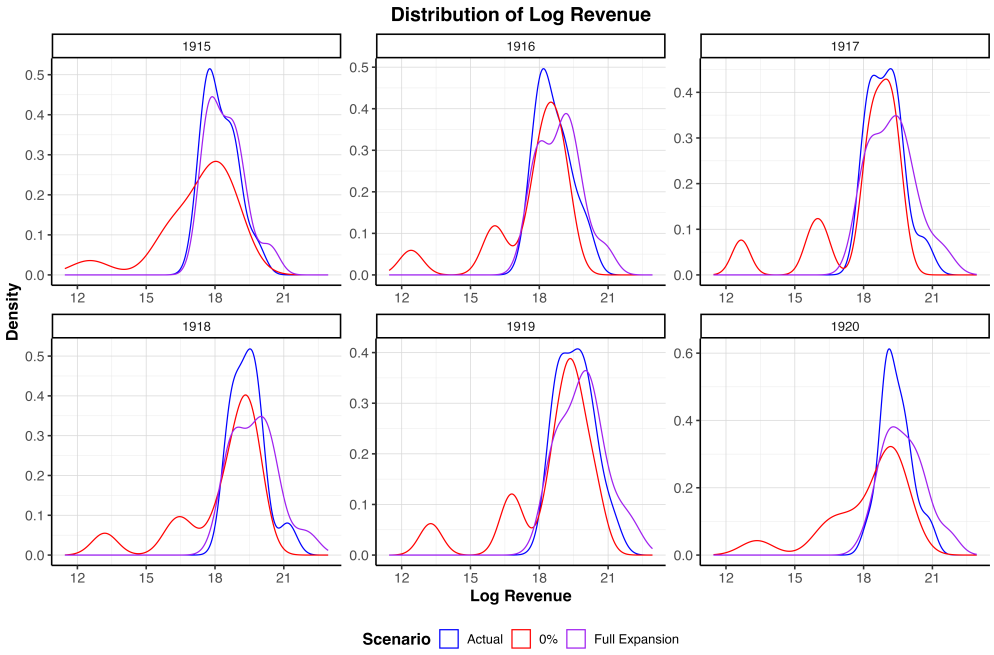


(b) Log Output

Figure 5: Counterfactual Distributions: Prices and Revenue (1915–1920)



(a) Log Price



(b) Log Revenue

Notes: Material inputs adjust endogenously according to equation (**??**).

In Figures **??**, panel (a) displays the distribution of quality under each scenario. The quality gap between scenarios widens over time because quality improvements partially per-

sist and electrification raises quality each period. By 1920, the distributions are clearly separated, with the no-adoption scenario concentrated at substantially lower quality levels. Panel (b) shows that TFPQ distributions largely overlap across scenarios, confirming that electrification's primary effect operated through quality rather than physical productivity.

In Figures **??**, panel (a) illustrates the equilibrium material adjustment: under no adoption, the distribution of log materials shifts leftward, reflecting firms' optimal response to lower quality through reduced variable input usage. Panel (b) displays output dynamics incorporating the equilibrium material adjustment, showing that output is lower under no adoption once material contraction is accounted for.

In Figures **??** ,Panel (a) shows pronounced separation in price distributions that mirrors the quality distributions, reflecting quality commanding higher willingness to pay under CES demand. Panel (b) displays log revenue, reflecting both the direct price premium from quality and the indirect output effects from equilibrium material adjustment.

7 Conclusion

This paper develops methods to separately identify unobserved product quality and quantity-based total factor productivity from panel data on prices, quantities, and inputs. The identification strategy exploits differential persistence in the stochastic processes governing quality and productivity. When quality and productivity follow autoregressive processes with different persistence parameters, this differential persistence appears in the conditional distribution of optimally chosen inputs, permitting nonparametric identification of both latent variables. Based on this identification argument, the paper develops a parametric estimator that recovers the joint dynamics of quality and productivity and identifies causal effects of technology adoption on each dimension separately.

The empirical application to Japanese cotton spinning during 1913–1920 yields a striking finding: electrification raised product quality substantially while leaving physical productivity essentially unchanged. A percentage point increase in electric engine share raises quality by 1.5 percent, with the effect precisely estimated and robust to weak-instrument inference. The effect on TFPQ is economically small and statistically indistinguishable from zero.

The counterfactual analysis quantifies the cumulative contribution of electrification to firm performance. Simulating plant trajectories under a no-adoption counterfactual reveals that product quality would have been nearly sixty percent lower in the absence of electrification, with the quality gap widening over time as persistent quality dynamics compound the direct effects.

These findings carry implications for productivity measurement and for understanding

technology adoption. Revenue-based productivity measures conflate output quantity with product quality. Technologies that improve what firms produce without raising how much they produce appear in standard productivity statistics as modest gains, understating the true benefits of adoption. The productivity paradox documented by **?**—that general purpose technologies often fail to generate immediate measured productivity gains—may partly reflect quality improvements that standard measures miss. The present findings provide micro-level evidence consistent with this interpretation: electrification delivered substantial benefits that operated primarily through quality rather than efficiency channels.