Parameter Stability as a Diagnostic Tool for Production Function Specification

Daniel Miles-Touya
Universidad de Vigo

July 2, 2025

Abstract

The control function approach remains the dominant method for estimating production functions. In this framework, unobserved productivity is recovered by inverting a proxy function that relates it to observed state variables and flexible inputs. A key specification test proposed by Levinsohn and Petrin (2003) relies on the invariance of parameter estimates across different nonparametric proxies for flexible inputs. Given the widespread use of Cobb-Douglas production functions in empirical industrial organization, we propose a related test within the parametric inversion framework of Doraszelski and Jaumandreu (2013) (hereafter DJ). Specifically, we demonstrate that under the identification assumptions of DJ, the recovered productivity process—and hence the structural parameter estimates—should be invariant to the choice of flexible input used for inversion. Consequently, systematic discrepancies across inversion strategies can serve as a diagnostic tool for detecting violations of the model's underlying assumptions.

1.- Introduction

The dominant approach to estimating production functions in the presence of unobserved productivity in empirical industrial organization is the control function framework introduced by Olley and Pakes (1996) (hereafter OP). This class of methods relies on four key assumptions (see Ackerberg et al., 2015; Ackerberg, 2023; De Loecker and Syverson, 2021). First, productivity is the only unobserved state variable (the scalar unobservable assumption). Second, a proxy variable—investment in original OP—exists that responds to productivity shocks and observed state variables, thereby allowing productivity to be expressed as a function of observables through the inversion of the firm's input decision rule. Third, a timing assumption ensures that capital is predetermined with respect to current productivity shocks but may correlate with persistent productivity. Fourth, productivity evolves according to a first-order Markov process.

Building on the framework of Olley and Pakes (1996), Levinsohn and Petrin (2003) (hereafter LP) propose using intermediate input demand—rather than investment—as the proxy variable in the control function for unobserved productivity. By relying on intermediate inputs, which are typically non-zero and more smoothly varying, the LP method enhances the feasibility and robustness of the estimation procedure, while still maintaining the scalar unobservable and monotonicity assumptions.

The existence of multiple intermediate inputs also enables specification testing based on the internal consistency of the model. In this spirit, Levinsohn and Petrin (2003) propose a simple test grounded in the invariance of parameter estimates across proxy functions constructed from different intermediate inputs. Substantial variation in estimates obtained using distinct proxies may signal model misspecification. However, because LP's estimation relies on nonparametric approximations to the proxy functions, observed discrepancies may also stem from systematic differences in approximation error rather than genuine violations of the identification assumptions.

In this note, following the logic of Levinsohn and Petrin (2003), we propose building the stability test on the parametric inversion framework developed by Doraszelski and Jaumandreu (2013) (hereafter DJ).[1] While DJ framework retains the scalar unobservable assumption it exploits the parametric structure of the Cobb-Douglas production function to recover productivity directly from the first-order conditions implied by the firm's optimal flexible input decisions. Specifically, DJ assume a Cobb-Douglas production function, one of the most widely used functional forms in empirical industrial organization (see De Loecker and Syverson, 2021). Under this specification, the static profit maximization problem yields closed-form expressions for the flexible input FOCs, allowing productivity to be expressed as a parametric function of observables and structural parameters.

A direct implication of the DJ approach is that the number of candidate productivity inversion equations corresponds to the number of flexible inputs in the production function. However, the estimates of the structural parameters should remain invariant to the specific inversion equation used to recover productivity. In other words, assuming the model is correctly specified, parameter estimates should be consistent across different inversion strategies. Consequently, any divergence in estimates obtained from alternative inversion approaches can serve as a diagnostic tool for evaluating the validity of the underlying assumptions—an idea closely aligned with the specification test proposed by Levinsohn and Petrin (2003).

To illustrate the proposed test, we draw on the empirical application presented in DJ. Our findings show that DJ's estimated production function coefficients vary significantly depending on the input demand equation used to recover productivity. For example, in the metals and metal products sector, the estimated labor coefficient increases from 0.111 to 0.247 when intermediate inputs, rather than labor, are used to invert productivity. A similar pattern emerges in the food, drink, and tobacco sector, where the labor coefficient rises from 0.129 to 0.211. These shifts are not limited to labor inputs: the materials coefficient declines from 0.684 to 0.543 in the metals sector and from 0.766 to 0.608 in the food sector. Meanwhile, the capital coefficient in the latter sector increases from 0.068 to 0.147. These discrepancies are not isolated cases; rather, comparable divergences are observed across other Spanish manufacturing sectors analyzed in DJ's study.

Taken together, the dependence of parameter estimates on the choice of inversion procedure

[1] DJ is a widely cited contribution in empirical industrial organization, offering a transparent and tractable alternative to nonparametric control function methods by leveraging first-order conditions associated with flexible input choices.

raises concerns about potential violations of the underlying specification assumptions. In what follows, we identify two plausible sources of misspecification that may account for the observed variation in DJ's estimates: optimization errors and measurement error—beyond the more straightforward possibility that the production function itself is misspecified. As we demonstrate, these issues introduce additional unobservables into the estimation process, thereby violating the scalar unobservable assumption that underpins the DJ framework.

Beyond statistical concerns, these discrepancies may also carry important economic implications. For instance, in the context of markup estimation via the production approach (e.g., De Loecker and Warzynski, 2012), the implied markup based on the DJ-estimated labor coefficient more than doubles when switching from a labor-based to a materials-based inversion. Conversely, markups constructed using materials fall by nearly 20% relative to those based on labor. In short, the choice of inversion strategy can dramatically alter empirical conclusions, including those central to questions of market power (see also Raval, 2023).

The note is organized as follows. Section 2 presents the parametric inversion approach to show that the estimates should be invariant to the inversion approach. In section 3 we give two possible reasons for the variability of the estimates. In section 4 we conclude.

2.- Parameter Estimates Independent of Inversion Approach

In this section, we show that, under the identifying assumptions of the DJ framework, the recovered productivity process is invariant to the choice of flexible input used in the inversion. Consequently, parameter estimates should not depend on the specific inversion strategy employed.

Let the production function for firm j be given by

$$Q_{jt}^{*} = e^{\omega_{jt}} F(M_{jt}, L_{jt}, K_{jt}; \theta),$$

where $F(\cdot)$ is a known functional form, M_{jt} and L_{jt} are flexible inputs (materials and labor), K_{jt} is the dynamic input (capital), and θ denotes the vector of parameters. The scalar unobservable assumption implies that ω_{jt}, the firm-specific productivity shock, is the sole unobserved variable to the econometrician. Productivity is assumed to follow a first-order Markov process:

$$\omega_{jt} = g(\omega_{jt-1}, rd_{jt-1}) + \xi_{jt},$$

where $g(\cdot)$ is an unknown function, rd_{jt-1} denotes lagged R&D expenditure, and ξ_{jt} is an exogenous productivity shock. The firm faces an inverse demand curve given by

$$P_{jt} = Q_{jt}^{*\,-1/\eta},$$

where $\eta > 1$ is the elasticity of demand.

The empirical production function is derived from $Q_{jt} = Q_{jt}^{*} e^{u_{jt}}$, where u_{jt} captures measurement error. Taking logarithms yields:

$$\ln Q_{jt} = \ln F(M_{jt}, L_{jt}, K_{jt}; \theta) + g(\omega_{jt-1}, rd_{jt-1}) + \xi_{jt} + u_{jt}. \tag{1}$$

Here, ω_{jt} has been substituted using its law of motion. [2]

To control for unobserved productivity in (1), DJ propose using the firm's optimal input choices derived from first-order conditions. At each t, the firm chooses M_{jt} and L_{jt} to maximize profits:

$$\max_{M_{jt}, L_{jt}} \left(e^{\omega_{jt}} F(M_{jt}, L_{jt}, K_{jt}; \theta)\right)^{1-1/\eta} - P_{M_{jt}} M_{jt} - P_{L_{jt}} L_{jt}.$$

Inverting the first-order conditions and taking logarithms yields the following expressions for productivity:

$$h_{Ljt} := \omega_{jt} = \frac{1}{1-\frac{1}{\eta}} \left[p_{Ljt} - \ln\left(1 - \frac{1}{\eta}\right) + \frac{1}{\eta} \ln F(K_{jt}, L_{jt}, M_{jt}; \theta) - \ln \frac{\partial F}{\partial L}(K_{jt}, L_{jt}, M_{jt}; \theta) \right], \tag{2}$$

[2] DJ assume a Cobb-Douglas production function, so their empirical specification becomes:

$$q_{jt} = \alpha_0 + \alpha_k k_{jt} + \alpha_l l_{jt} + \alpha_m m_{jt} + g(\omega_{jt-1}, rd_{jt-1}) + \xi_{jt} + u_{jt}.$$

$$h_{Mjt} := \omega_{jt} = \frac{1}{1-\frac{1}{\eta}}\left[p_{Mjt} - \ln\left(1-\frac{1}{\eta}\right) + \frac{1}{\eta}\ln F(K_{jt}, L_{jt}, M_{jt};\theta) - \ln\frac{\partial F}{\partial M}(K_{jt}, L_{jt}, M_{jt};\theta)\right]. \quad (3)$$

A direct implication of the DJ approach is that the number of candidate inversion equations for recovering unobserved productivity equals the number of flexible inputs in the production function. Each inversion strategy yields an expression for the same underlying productivity term. Stated differently, (2) $\equiv$ (3). [3]

Consequently, under correct model specification, the parameter estimates should be the same regardless of the specific flexible input used to recover productivity. This invariance becomes evident when substituting the recovered productivity term into the empirical production function (1). Whether productivity is recovered via labor—i.e., $\omega_{jt-1} = h_{Ljt-1}$ from equation (2)—or via materials—i.e., $\omega_{jt-1} = h_{Mjt-1}$ from equation (3)—the resulting empirical specification remains unchanged:

$$\begin{aligned}\ln Q_{jt} &= \ln F(M_{jt}, L_{jt}, K_{jt};\theta) + g(\omega_{jt-1}, rd_{jt-1}) + \xi_{jt} + u_{jt}\\ &= \ln F(M_{jt}, L_{jt}, K_{jt};\theta) + g(h_{Ljt-1}, rd_{jt-1}) + \xi_{jt} + u_{jt}\\ &= \ln F(M_{jt}, L_{jt}, K_{jt};\theta) + g(h_{Mjt-1}, rd_{jt-1}) + \xi_{jt} + u_{jt}.\end{aligned}$$

Thus, DJ parametric framework implies that the empirical production function—after controlling for unobserved productivity—should be invariant to the specific inversion strategy employed. Any observed divergence in parameter estimates across inversion strategies may therefore reflect violations of the underlying assumptions required for identification.

In the following section, we examine two potential sources of misspecification that may explain the discrepancies observed in DJ's empirical implementation.

[3] Under the Cobb-Douglas assumption, the inversion of the labor first-order condition at $t-1$ is given by:

$$\begin{aligned}\omega_{jt-1} &= h_{Ljt-1} = h_L(k_{jt-1}, l_{jt-1}, p_{L_{jt-1}}, p_{M_{jt-1}}, p_{jt-1}, d_{jt-1};\theta)\\ &= \lambda_l - \alpha_k k_{jt-1} + (1-\alpha_l-\alpha_m)l_{jt-1} + (1-\alpha_m)p_{Ljt-1} + \alpha_m p_{Mjt-1} - p_{jt-1} - \mu(p_{jt-1}, d_{jt-1}),\end{aligned}$$

where λ_l is a constant depending on production parameters, and $\mu(p_{jt-1}, d_{jt-1}) = \ln\left(1-\frac{1}{\eta(p_{jt-1}, d_{jt-1})}\right)$ with $\eta(\cdot)$ modeled as a function of price and demand shifters. A similar expression holds for materials:

$$\begin{aligned}\omega_{jt-1} &= h_{Mjt-1} = h_M(k_{jt-1}, m_{jt-1}, p_{L_{jt-1}}, p_{M_{jt-1}}, p_{jt-1}, d_{jt-1};\theta)\\ &= \lambda_m - \alpha_k k_{jt-1} + (1-\alpha_l-\alpha_m)m_{jt-1} + (1-\alpha_l)p_{Mjt-1} + \alpha_l p_{Ljt-1} - p_{jt-1} - \mu(p_{jt-1}, d_{jt-1}).\end{aligned}$$

3.-Parameter Estimate Discrepancies as a Diagnostic for Misspecification

In the previous section, we established that, under the identifying assumptions of the DJ framework, parameter estimates should be invariant to the choice of flexible input used to invert and control for unobserved productivity. Consequently, any systematic variation in estimates across inversion strategies may signal a violation of these underlying assumptions.

This implication is put to the test in the empirical application by Doraszelski and Jaumandreu (2013), who estimate production functions using firm-level data from the Spanish manufacturing sector. Assuming a Cobb-Douglas specification, they employ both labor and intermediate materials—the two flexible inputs in the firm's static optimization problem—to recover unobserved productivity. However, the resulting parameter estimates vary substantially depending on which input demand equation is used for inversion. For instance, in the timber and furniture sector, the estimated capital coefficient falls from 0.131 to 0.037, the labor coefficient rises from 0.176 to 0.266, and the materials coefficient declines slightly from 0.697 to 0.666. As discussed earlier, similar discrepancies are observed across other sectors in DJ's empirical implementation.

Taken together, these discrepancies suggest a failure of one or more key assumptions underlying the DJ framework. In the following section, we outline two potential sources of such misspecification—beyond the standard concern of functional form misspecification—each corresponding to a distinct violation of the identification strategy implicit in inversion-based estimation.

Optimization Errors

Discrepancies in parameter estimates across inversion strategies may stem from optimization errors—that is, deviations from the first-order conditions implied by profit maximization or cost minimization. Such deviations occur when firms make input choices that do not exactly satisfy the necessary conditions for optimality (see (Ackerberg et al., 2015; Reiss and Wolak, 2007)). Crucially, these violations undermine the scalar unobservable assumption that underlies the control function approach of Olley and Pakes (OP) and, by extension, the DJ framework.

Suppose, specifically, that firm j's choice of flexible input $Z_{jt} \in L_{jt}, M_{jt}$ is subject to an optimization error. Then, the recovered productivity term becomes:

$$\omega_{jt} = \frac{1}{1-\frac{1}{\eta}} \left[\ln p_{Zjt} - \ln\left(1 - \frac{1}{\eta}\right) + \frac{1}{\eta} \ln F(K_{jt}, L_{jt}, M_{jt}; \theta) - \ln\left(\frac{\partial F(K_{jt}, L_{jt}, M_{jt}; \theta)}{\partial Z_{jt}} \right) + \nu_{jt}^{Z} \right],$$

where ν_{jt}^{Z} denotes the optimization error associated with input Z. Because ω_{jt} is no longer a deterministic function of observables and parameters, the inversion introduces an additional unobservable into the estimation problem.

When this expression for productivity is substituted back into the empirical production function, the optimization error ν_{jt}^{Z} enters the unknown function $g(\cdot)$. Since these errors are difficult to instrument, they violate the conditions required for consistent estimation of the structural parameters. In other words, failure to account for optimization errors may result in biased parameter estimates.

Measurement Errors

Measurement error in the quantities or prices of flexible inputs can generate discrepancies in parameter estimates across inversion strategies. This issue is particularly relevant in DJ's empirical application, given the characteristics of the data source—the *Encuesta Sobre Estrategias Empresariales* (ESEE), a Spanish firm-level survey. While the ESEE reports firm-specific rates of price change, it does not provide firm-specific price levels. As shown in González et al. (2025), this limitation introduces measurement error in deflated input quantities, especially when nominal expenditures must be converted into real quantities.

To illustrate, consider the case of intermediate materials. Let the log of nominal expenditures be given by:

$$e_{jt} = m_{jt} + p_{Mjt} = m_{jt} + p_{Mjb} + \Delta_{Mjb}^{t},$$

where m_{jt} is the log of the quantity of materials input, p_{Mjt} the log of the materials price level, p_{Mjb} the base-year price level, and Δ_{Mjb}^{t} the observed rate of price change between the base year b and period t. [4]

Since only Δ^t_{Mjb} is observed, the deflated input measure becomes:

$$m^*_{jt} = e_{jt} - \Delta^t_{Mjb} = m_{jt} + p_{Mjb},$$

which implies that m^*_{jt} contains an additive measurement error equal to the unknown base-year price level p_{Mjb}. This error is transmitted directly into the inversion when intermediate materials are used as the flexible input. In contrast, labor is typically measured in hours or number of employees, which are directly reported in the dataset and thus less prone to this type of error. [5]

Moreover, in the absence of firm-specific output price levels, the base-year output price p_{jb} is also unobserved. As a result, output is measured with error when constructed from deflated revenues, i.e., $r^*_{jt} = q_{jt} + p_{jb}$, due to the presence of p_{jb}, a time-invariant, firm-specific price component. Although this measurement error enters additively, it is *non-classical* in nature, as it may be correlated with input choices. If ignored, this correlation violates the standard assumptions required for consistent estimation.

In sum, measurement error—particularly when it differentially affects specific input or output variables—can induce variation in the recovered productivity process across inversion strategies, thereby offering a plausible explanation for the divergence in parameter estimates.

4.- Conclusion

In this note, we propose a specification test for production function estimation based on the stability of parameter estimates across alternative inversion strategies within the parametric framework of Doraszelski and Jaumandreu (2013). Inspired by the logic of Levinsohn and Petrin (2003), who

[4]Any log price level can be decomposed as:

$$p_{jt} = p_{jb} + (p_{jt} - p_{jb}) = p_{jb} + \Delta^t_{jb},$$

where Δ^t_{jb} is the observed price rate of change and p_{jb} is the base-year price level, which is unobserved because price levels are not reported in the dataset.

[5]Under a Cobb-Douglas production function, intermediate materials affect only the inversion equation based on materials FOCs, not the labor-based inversion. See footnote 4.

suggest testing model validity by comparing estimates across nonparametric proxies, our approach exploits the parametric structure arising from the Cobb-Douglas production function behind the of the DJ method, where unobserved productivity is recovered from the first-order conditions of flexible inputs under a Cobb-Douglas production function.

A key implication of the DJ framework is that, under correct specification, the recovered productivity—and thus the structural parameters—should be invariant to the choice of input used for inversion. We show that this prediction fails in DJ's empirical application: estimates for the same parameter differ markedly across inversion strategies, with economically meaningful consequences, particularly for markup estimation. For example, the labor parameter estimate increases from 0.158 to 0.239 in the Spanish transport and equipment sector, or from 0.122 to 0.240 in the chemical sector, when switching from the labor to the materials inversion strategy. These discrepancies suggest internal inconsistency and motivate the use of parameter variation as a simple, model-implied diagnostic for misspecification.

We outline two potential sources of such variation: optimization errors and measurement errors in prices or quantities. These violate the scalar unobservable assumption fundamental to identification.

References

Ackerberg, D. A. (2023). Timing assumptions and efficiency: Empirical evidence in a production function context. *The Journal of Industrial Economics*, 71(3):644–674.

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

De Loecker, J. and Syverson, C. (2021). An industrial organization perspective on productivity. In *Handbook of Industrial Organization*, volume 4, pages 141–223.

De Loecker, J. D. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Doraszelski, U. and Jaumandreu, J. (2013). R&D and productivity: Estimating endogenous productivity. *Review of Economic Studies*, 80(4):1338–1383.

González, X., Lach, S., and Miles-Touya, D. (2025). Revisiting the omitted price bias in the estimation of production function. *Journal of Industrial Economics*, forthcoming.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies*, 70(2):317–341.

Olley, S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Raval, D. (2023). Testing the production approach to markup estimation. *Review of Economic Studies*, 90(5):2592–2611.

Reiss, P. C. and Wolak, F. A. (2007). Structural econometric modeling: Rationales and examples from industrial organization. *Handbook of econometrics*, 6:4277–4415.