

Production Function Estimation without Invertibility: Imperfectly Competitive Environments and Demand Shocks^{*}

Ulrich Doraszelski

University of Pennsylvania[†]

Lixiong Li

Johns Hopkins University[‡]

June 25, 2025

Abstract

We advance the proxy variable approach to production function estimation. We show that the invertibility assumption at its heart is testable. We characterize what goes wrong if invertibility fails and what can still be done. We show that rethinking how the estimation procedure is implemented either eliminates or mitigates the bias that arises if invertibility fails. In particular, a simple change to the first step of the estimation procedure provides a first-order bias correction for the GMM estimator in the second step. Furthermore, a modification of the moment condition in the second step ensures Neyman orthogonality, thereby enhancing efficiency and robustness by rendering the asymptotic distribution of the GMM estimator invariant to estimation noise from the first step.

^{*}We thank Dan Akerberg and Steve Berry for helpful discussions and comments and Alexander Whitefield, Vivian Zhang, and Wanxi Zhou for excellent research assistance.

[†]Wharton School, Email: doraszelski@wharton.upenn.edu

[‡]Department of Economics, Email: lixiong.li@jhu.edu

1 Introduction

Production functions are one of the oldest concepts in economics (von Thünen 1842, Wicksteed 1894, Cobb & Douglas 1928). As a description of the relationship between inputs and output, they are of interest in and of themselves and also a vehicle for measuring productivity and technical change (Solow 1957). More recently, production functions have become a key input into estimating markups and markdowns in order to assess firms' market power (Hall 1988, De Loecker & Warzynski 2012). The production approach to markup estimation recognizes that, given an estimate of the production function, the markup can be recovered from the firm's cost minimization problem. To estimate the production function, this large and rapidly growing literature relies on the procedure initially developed by Olley & Pakes (1996) and further advanced by Levinsohn & Petrin (2003) and Akerberg, Caves & Frazer (2015) (henceforth OP, LP, and ACF).

The OP/LP/ACF procedure starts from the observation that there is an endogeneity problem in estimating production functions. As first pointed out by Marschak & Andrews (1944), this problem arises because the decisions that the firm makes regarding its inputs depend on its productivity, which is unobserved by the econometrician. To resolve this problem, OP turn it on its head and use the firm's decisions that are observed by the econometrician to infer its productivity. The resulting proxy variable approach to production function estimation has quickly become standard practice.

The inversion from observables to productivity at the heart of the OP/LP/ACF procedure requires that firms with the same productivity make the same choices. Yet, there is no reason to believe this is the case, e.g., if these firms face different demands in the output market. The current paper therefore revisits the OP/LP/ACF procedure if invertibility fails. It characterizes what goes wrong in the OP/LP/ACF procedure and what can still be done.

The point that the OP/LP/ACF procedure cannot accommodate unobserved demand heterogeneity has been made before. Foster, Haltiwanger & Syverson (2008) put it as follows:

... idiosyncratic demand shocks make the proxies functions of both technology and demand shocks, thereby inducing a possible omitted variable bias. Put simply, proxy methods require a one-to-one mapping between plant-level productivity and the observables used to proxy for productivity. This mapping breaks down if other unobservable plant-level factors besides productivity drive changes in the observable proxy. (p. 403)

Indeed, to establish invertibility, OP rule out that firms face different demands and abstract from competition between firms.¹ LP assume a perfectly competitive industry where firms act as price

¹OP assume that any profitability differences across firms are due to differences in their capital stocks and productivities (p. 1273), thereby ruling out that firms face different demands. Limiting the state variables in the firm's investment policy to its own capital stock and productivity moreover abstracts from competition between firms (see also Lemma 3 and Theorem 1 in Pakes 1994).

takers and thus face the same horizontal demand curve (p. 322 and Appendix A). Based on their results, ACF start from scalar unobservable and strict monotonicity assumptions, and almost all subsequent literature simply imposes invertibility as a high-level assumption. For this reason, it may not have fully appreciated the demanding nature of the invertibility assumption.

Unobserved demand heterogeneity is likely to be empirically important. The large literatures on demand estimation and productivity analysis highlight the considerable heterogeneity in demand that remains even after controlling for detailed product attributes (Berry, Levinsohn & Pakes 1995) or honing in on (nearly) homogenous products (Foster et al. 2008). The problem is compounded by the fact that, in imperfectly competitive environments, a firm’s decisions in equilibrium depend on its own productivity as well as on the productivities of its rivals. Hence, imperfectly competitive environments require jointly inverting the decisions of all firms for the productivities of all firms. This many-to-many inverse may not exist (Biondi 2022) or it may be too high-dimensional to be practical (Akerberg & De Loecker 2024). Moreover, a firm’s rivals are partially or completely unobserved in typical datasets used for production function estimation. The decisions of its rivals have thus to be thought of as shocks to the demand the focal firm faces.

Invertibility can fail for reasons other than unobserved demand heterogeneity. Changes in firm conduct due to mergers and acquisitions or switches from competition to collusion are best thought of as changes to the ownership matrix. Invertibility can fail if these changes in firm conduct are partially or completely unobserved by the econometrician. Finally, invertibility can fail if there is unobserved variation across firms or time in input prices, investment opportunities, or financial constraints.

In light of the demanding nature of the invertibility assumption, this paper makes five contributions. First, we propose tests for invertibility. Our tests exploit that if invertibility fails, then productivity becomes a hidden state in a Markov model. The literature on Kalman and particle filters in dynamic systems shows that the best guess for the hidden state uses the entire history of the observables as opposed to just their current value. Invertibility can therefore be tested by including lags of the observables in the regression in the first step of the OP/LP/ACF procedure. Implementing our tests on an unbalanced panel of Spanish manufacturing firms and a balanced panel of US manufacturing industries, we strongly reject invertibility.

We therefore characterize the consequences of a failure of invertibility for the OP/LP/ACF procedure. The first step of the OP/LP/ACF procedure regresses output on observables. If invertibility fails, then the prediction of output contains an error. Because the prediction is used to control for lagged productivity in the GMM estimation in the second step, the lagged prediction error enters into the conditional moment. Building on Doraszelski & Jaumandreu (2024), we show that this invalidates capital as an instrument and results in biased estimates.

Turning from what goes wrong if invertibility fails to what can still be done, our second

contribution is to provide a necessary and sufficient condition for the moment condition in the second step of the OP/LP/ACF procedure to hold for the true production function and some (not necessarily the true) law of motion for productivity. This condition calls for rethinking how the OP/LP/ACF procedure is implemented. In particular, it compels us to ensure that any instrument used in the second step is appropriately included in the regression in the first step. Due to the lag structure of the model, this means including the lead of capital in addition to its current value in the regression. We provide a series of examples where this simple expedient suffices to satisfy our necessary and sufficient condition.

Beyond these examples, our necessary and sufficient condition may be violated. While this results in biased estimates, our third contribution is to show that rethinking how the OP/LP/ACF procedure is implemented mitigates the bias. Specifically, we show that the moment condition in the second step implicitly incorporates a first-order bias correction provided any instrument used in the second step is appropriately included in the regression in the first step.

Our fourth contribution is to explicitly incorporate a bias correction into the second step of the OP/LP/ACF procedure. We show that the modified moment condition has a property known as Neyman orthogonality (Neyman 1959). While our modification remains as straightforward to implement as the original OP/LP/ACF procedure, Neyman orthogonality renders the asymptotic distribution of the GMM estimator in the second step invariant to estimation noise and the quality of the prediction of output from the first step. This is particularly advantageous if the regression in the first step includes a large number of covariates. Neyman orthogonality facilitates the use of a wide range of estimation methods in the first step, including traditional nonparametric methods as well as modern machine learning techniques such as neural networks and random forests. A Monte Carlo exercise shows that Neyman orthogonality can substantially improve the performance of the GMM estimator in finite samples.

Our fifth and final contribution is to provide a diagnostic to assess the sensitivity of the estimates to the size of the prediction error in the first step. While our diagnostic has some similarities to the sensitivity measure in Andrews, Gentzkow & Shapiro (2017), it is not local to the true model and directly informative about the estimated model. Our diagnostic is neither necessary nor sufficient for no bias. However, a small value of the diagnostic provides assurance that small changes in the prediction error do not dramatically alter the estimates.

In sum, this paper examines the OP/LP/ACF procedure if the invertibility assumption fails. Invertibility fails if there are demand shocks or in imperfectly competitive environments with partially or completely unobserved rivals or changes in firm conduct. Whether invertibility fails can be tested. A failure of invertibility can have a substantial, detrimental impact on the estimates.

Fortunately, much can still be done. We provide a necessary and sufficient condition for the moment condition in the second step of the OP/LP/ACF procedure to hold for the true production

function. This condition compels us to ensure that any instrument used in the second step is appropriately included in the regression in the first step. We show that this simple change either eliminates or mitigates the bias. Going a step further, we modify the moment condition in the second step to endow it with Neyman orthogonality. Finally, we provide a diagnostic to assess the sensitivity of the estimates to the size of the prediction error that arises in the first step of the OP/LP/ACF procedure if invertibility fails.

The remainder of the paper is organized as follows. In Section 2, we recall the setup and the OP/LP/ACF procedure. In Section 3, we develop tests for invertibility. In Section 4, we develop a necessary and sufficient condition for the moment condition in the second step to hold for the true production function. We show that ensuring that any instrument used in the second step is appropriately included in the regression in the first step either eliminates or mitigates the bias that arises if invertibility fails. In Section 5, we modify the moment condition in the second step to endow it with Neyman orthogonality. In Section 6, we conduct a Monte Carlo exercise to illustrate what goes wrong in the OP/LP/ACF procedure if invertibility fails and what can still be done. In Section 7, we provide a diagnostic to assess the sensitivity of the estimates to the size of the prediction error. We conclude in Section 8.

2 Setup and OP/LP/ACF procedure

Firm i in period t produces output Q_{it} with inputs K_{it} and V_{it} according to the production function

$$q_{it} = f(k_{it}, v_{it}) + \omega_{it} + \varepsilon_{it}, \quad (1)$$

where lower case letters denote logs. Capital k_{it} is a predetermined input that is chosen in period $t-1$ whereas v_{it} is freely variable and decided on in period t after the firm observes its productivity ω_{it} .² Productivity follows a first-order Markov process with law of motion

$$\omega_{it} = E[\omega_{it} | \omega_{it-1}] + \xi_{it} = g(\omega_{it-1}) + \xi_{it}, \quad (2)$$

where the productivity innovation ξ_{it} is by construction mean independent of lagged productivity ω_{it-1} and further assumed to be mean independent of any variable included in the firm's information set in period $t-1$. The disturbance ε_{it} sits between the firm's output q_{it} as recorded in the data and the output $q_{it}^* = q_{it} - \varepsilon_{it} = f(k_{it}, v_{it}) + \omega_{it}$ that the firm planned on when it decided on the variable input v_{it} . It can be interpreted as measurement error or as an unanticipated shock to output (OP, pp. 1273–1274) and is assumed to be mean independent of the inputs and

²While k_{it} and v_{it} may be vectors, we think of them as scalars for simplicity. The variable input v_{it} may accordingly be interpreted as a composite of labor and materials such as cost of goods sold.

other included variables as formalized below.³ While the econometrician observes actual output q_{it} and the inputs k_{it} and v_{it} , productivity ω_{it} , the disturbance ε_{it} , and planned output q_{it}^* remain unobserved.

Invertibility. The literature following OP relies on invertibility. Invertibility assumes that there exists a function $\omega_{it} = h(x_{it})$ that maps observables $x_{it} = (k_{it}, v_{it}, \dots)$ into productivity ω_{it} . This is equivalently to

$$E[\omega_{it}|x_{it}] = \omega_{it}.$$

The observables included in x_{it} depend on the decision of the firm that is being inverted.⁴ We remain agnostic about which variables are included in x_{it} , aside from imposing, without loss of generality, that the inputs k_{it} and v_{it} are included.

OP/LP/ACF procedure. Estimation proceeds in two steps. Step 1 assumes $E[\varepsilon_{it}|x_{it}] = 0$ and flexibly or nonparametrically estimates the conditional expectation

$$E[q_{it}|x_{it}] = E[f(k_{it}, v_{it}) + \omega_{it} + \varepsilon_{it}|x_{it}] = f(k_{it}, v_{it}) + E[\omega_{it}|x_{it}] = f(k_{it}, v_{it}) + \omega_{it} = q_{it}^*, \quad (3)$$

where the first equality uses equation (1), the second-to-last equality uses invertibility, and the last equality uses the definition of planned output q_{it}^* . Because $E[q_{it}|x_{it}] = q_{it}^*$, step 1 separates actual output q_{it} into planned output q_{it}^* and the disturbance $\varepsilon_{it} = q_{it} - q_{it}^*$.

Step 2 of the OP/LP/ACF procedure assumes $E[\xi_{it} + \varepsilon_{it}|z_{it}] = 0$ for instruments $z_{it} = (k_{it}, k_{it-1}, v_{it-1}, \dots)$ and estimates the parameters $\theta = (\theta_f, \theta_g)$ in the production function f and the law of motion g by GMM.⁵ Capital k_{it} is a valid instrument because it is a predetermined input, and the lagged inputs k_{it-1} and v_{it-1} are valid instruments because the productivity innovation ξ_{it} is mean independent of any variable included in the firm's information set in period $t - 1$. We remain agnostic regarding additional instruments in z_{it} .

Using equations (1) and (2) and the definition of planned output q_{it}^* , the assumption $E[\xi_{it} + \varepsilon_{it}|z_{it}] = 0$ implies

$$E[q_{it} - f(k_{it}, v_{it}) - g(\omega_{it-1})|z_{it}] = E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1}^* - f(k_{it-1}, v_{it-1}))|z_{it}] = 0. \quad (4)$$

Moment condition (4) is infeasible for estimation because lagged planned output q_{it-1}^* is unobserved. Substituting $E[q_{it-1}|x_{it-1}]$ from step 1 for q_{it-1}^* , step 2 therefore proceeds by estimating

³Mundlak & Hoch (1965) refer to ω_{it} and ε_{it} as the transmitted, respectively, untransmitted component of productivity. The untransmitted component may include machine breakdowns, labor actions, supply chain disruptions, and power outages that are not anticipated by the firm. The interpretation as measurement error accommodates serial correlation in the disturbance ε_{it} .

⁴OP invert the firm's demand for investment whereas LP and ACF invert its demand for materials.

⁵For notational convenience we suppress θ in much of what follows.

the parameters θ from the moment condition

$$\mathbb{E}[q_{it} - f(k_{it}, v_{it}) - g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) | z_{it}] = 0. \quad (5)$$

Throughout the remainder of the paper, we follow OP, LP, and ACF and maintain $\mathbb{E}[\varepsilon_{it}|x_{it}] = 0$ for observables $x_{it} = (k_{it}, v_{it}, \dots)$ and $\mathbb{E}[\xi_{it} + \varepsilon_{it}|z_{it}] = 0$ for instruments $z_{it} = (k_{it}, k_{it-1}, v_{it-1}, \dots)$. To avoid cumbersome notation, we adopt the convention that all equalities involving random variables and conditional expectations are understood to hold almost surely.

3 Tests for invertibility

If invertibility fails and $\mathbb{E}[\omega_{it}|x_{it}] \neq \omega_{it}$, then productivity ω_{it} becomes a hidden state in a Markov model. The literature on Kalman and particle filters in dynamic systems shows that the best guess for the unobservables uses the entire history of the observables as opposed to just their current value. Our tests for invertibility are based on this intuition.

Our first proposition shows that invertibility implies a testable mean-independence restriction:

Proposition 1. *If $\mathbb{E}[\omega_{it}|x_{it}] = \omega_{it}$ and $\mathbb{E}[\varepsilon_{it}|x_{it}, x_{it-1}] = \mathbb{E}[\varepsilon_{it}|x_{it}]$, then $\mathbb{E}[q_{it}|x_{it}, x_{it-1}] = \mathbb{E}[q_{it}|x_{it}]$.*

While step 1 of the OP/LP/ACF procedure assumes $\mathbb{E}[\varepsilon_{it}|x_{it}] = 0$ for observables $x_{it} = (k_{it}, v_{it}, \dots)$ and step 2 assumes $\mathbb{E}[\xi_{it} + \varepsilon_{it}|z_{it}] = 0$ for instruments $z_{it} = (k_{it}, k_{it-1}, v_{it-1}, \dots)$, this does not quite imply $\mathbb{E}[\varepsilon_{it}|x_{it}, x_{it-1}] = \mathbb{E}[\varepsilon_{it}|x_{it}]$. However, the interpretation of the disturbance ε_{it} as an unanticipated shock to output that is outside the firm's information set in period t implies $\mathbb{E}[\varepsilon_{it}|x_{it}, x_{it-1}] = \mathbb{E}[\varepsilon_{it}|x_{it}] = 0$.⁶ The proof of Proposition 1 is straightforward:

Proof. If $\mathbb{E}[\omega_{it}|x_{it}] = \omega_{it}$, then

$$q_{it} = f(k_{it}, v_{it}) + \omega_{it} + \varepsilon_{it} = f(k_{it}, v_{it}) + \mathbb{E}[\omega_{it}|x_{it}] + \varepsilon_{it}.$$

$$\mathbb{E}[\varepsilon_{it}|x_{it}, x_{it-1}] = \mathbb{E}[\varepsilon_{it}|x_{it}] \text{ therefore implies } \mathbb{E}[q_{it}|x_{it}, x_{it-1}] = \mathbb{E}[q_{it}|x_{it}]. \quad \square$$

Our second proposition shows that invertibility implies a testable conditional-independence restriction:

Proposition 2. *If $\mathbb{E}[\omega_{it}|x_{it}] = \omega_{it}$ and $\varepsilon_{it} \perp\!\!\!\perp x_{it-1}|x_{it}$, then $q_{it} \perp\!\!\!\perp x_{it-1}|x_{it}$.*

⁶Under the interpretation of the disturbance ε_{it} as measurement error, the assumption $\mathbb{E}[\varepsilon_{it}|x_{it}, x_{it-1}] = \mathbb{E}[\varepsilon_{it}|x_{it}]$ holds if x_{it-1} does not have predictive power for ε_{it} beyond x_{it} .

Going beyond the first-order implication of $E[\omega_{it}|x_{it}] = \omega_{it}$ in Proposition 1 calls for a stronger assumption on the disturbance ε_{it} . The assumption $\varepsilon_{it} \perp\!\!\!\perp x_{it-1}|x_{it}$ in Proposition 2 means that any possible dependence between ε_{it} and x_{it-1} is through x_{it} . It is implied by $\varepsilon_{it} \perp\!\!\!\perp (x_{it}, x_{it-1})$ and implies $E[\varepsilon_{it}|x_{it}, x_{it-1}] = E[\varepsilon_{it}|x_{it}]$. The proof of Proposition 2 parallels the proof of Proposition 1 and is therefore omitted.

We implement our tests for invertibility on an unbalanced panel of Spanish manufacturing firms from 1990 to 2006 (Encuesta Sobre Estrategias Empresariales) and a balanced panel of US manufacturing industries from 1958 to 2018 (NBER-CES). In both cases, we reject invertibility by a wide margin. We provide further details on the data and the tests in Appendix A. In the next section, we turn to the consequences of a failure of invertibility for production function estimation.

4 Failure of invertibility and validity of OP/LP/ACF moment condition

If invertibility fails and $E[\omega_{it}|x_{it}] \neq \omega_{it}$, then a prediction error $\zeta_{it} = \omega_{it} - E[\omega_{it}|x_{it}]$ arises in step 1 of the OP/LP/ACF procedure that is generally non-zero. By construction, the prediction error satisfies $E[\zeta_{it}|x_{it}] = 0$.

To see how the presence of the prediction error ζ_{it} affects the OP/LP/ACF procedure, recall that the conditional expectation estimated in step 1 is

$$E[q_{it}|x_{it}] = f(k_{it}, v_{it}) + E[\omega_{it}|x_{it}] = f(k_{it}, v_{it}) + \omega_{it} - \zeta_{it} = q_{it}^* - \zeta_{it}, \quad (6)$$

where the first equality uses equation (1) and $E[\varepsilon_{it}|x_{it}] = 0$, the second equality uses the definition of the prediction error ζ_{it} , and the last equality uses the definition of planned output q_{it}^* . Equation (6) provides an alternative interpretation of the prediction error as the difference between planned output and its prediction in step 1: $\zeta_{it} = q_{it}^* - E[q_{it}|x_{it}]$.

Substituting $E[q_{it-1}|x_{it-1}]$ from equation (6), the left-hand side of moment condition (5) in step 2 of the OP/LP/ACF procedure becomes

$$E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1})) | z_{it}]. \quad (7)$$

Because of the lagged prediction error ζ_{it-1} , conditional moment (7) is generally non-zero. To see this more clearly, consider the special case of an $AR(1)$ process for productivity. If $g(\omega_{it-1}) = \rho\omega_{it-1}$, then conditional moment (7) evaluates to

$$E[\xi_{it} + \varepsilon_{it} + \rho\zeta_{it-1} | z_{it}] = \rho E[\zeta_{it-1} | z_{it}], \quad (8)$$

where the equality uses $E[\xi_{it} + \varepsilon_{it}|z_{it}] = 0$.

To further see how this affects the GMM estimation, recall that $x_{it-1} = (k_{it-1}, v_{it-1}, \dots)$ and $z_{it} = (k_{it}, k_{it-1}, v_{it-1}, \dots)$. The lagged inputs k_{it-1} and v_{it-1} hence remain valid instruments because $E[\zeta_{it-1}|x_{it-1}] = 0$ by construction. However, capital k_{it} may no longer be a valid instrument as $E[\zeta_{it-1}|k_{it}] \neq 0$, as noted by Doraszelski & Jaumandreu (2024). This is because the firm chooses k_{it} in period $t - 1$ with knowledge of ω_{it-1} . What remains of lagged productivity after controlling for the lagged observables x_{it-1} may therefore be correlated with k_{it} . This results in biased estimates.

The above discussion foreshadows part of the solution to the problem: adding the lead of capital k_{it+1} to the observables x_{it} in step 1 of the OP/LP/ACF procedure ensures that $E[\zeta_{it-1}|k_{it}] = 0$ and hence the validity of capital k_{it} as an instrument in step 2. This simple expedient amounts to rethinking how the OP/LP/ACF procedure is implemented. The remaining question is to what extent it suffices beyond the special case of an $AR(1)$ process for productivity.

The following theorem provides a complete answer to this question in the form of a necessary and sufficient condition.

Theorem 1. *Moment condition (5) in step 2 of the OP/LP/ACF procedure holds for the true production function f^0 and some law of motion \tilde{g} if and only if*

$$E[g^0(\omega_{it-1})|z_{it}] = E[\tilde{g}(E[\omega_{it-1}|x_{it-1}])|z_{it}], \quad (9)$$

where g^0 is the true law of motion.

We prove Theorem 1 before parsing it:

Proof. Recall that the true production function f^0 and the true law of motion g^0 satisfy moment condition (4).

“Only if” part: Suppose moment condition (5) holds for the true production f^0 and some law of motion \tilde{g} . Then we have

$$\begin{aligned} 0 &= E[q_{it} - f^0(k_{it}, v_{it}) - \tilde{g}(E[q_{it-1}|x_{it-1}] - f^0(k_{it-1}, v_{it-1}))|z_{it}] \\ &= E[q_{it} - f^0(k_{it}, v_{it}) - \tilde{g}(E[\omega_{it-1}|x_{it-1}])|z_{it}], \end{aligned}$$

where the last equality uses equation (6). Subtracting from equation (4) implies condition (9).

“If” part: Suppose condition (9) holds. Substituting into equation (4), we have

$$\begin{aligned} 0 &= E[q_{it} - f^0(k_{it}, v_{it}) - \tilde{g}(E[\omega_{it-1}|x_{it-1}])|z_{it}] \\ &= E[q_{it} - f^0(k_{it}, v_{it}) - \tilde{g}(E[q_{it-1}|x_{it-1}] - f^0(k_{it-1}, v_{it-1}))|z_{it}], \end{aligned}$$

where the last equality uses equation (6). Hence, moment condition (5) holds for f^0 and \tilde{g} . \square

Theorem 1 covers the invertibility assumption at the heart of the literature following OP as a special case. If $E[\omega_{it}|x_{it}] = \omega_{it}$, then condition (9) becomes $E[g^0(\omega_{it-1})|z_{it}] = E[\tilde{g}(\omega_{it-1})|z_{it}]$ and is therefore satisfied with $\tilde{g} = g^0$. Importantly, however, condition (9) can be satisfied even if invertibility fails and $E[\omega_{it}|x_{it}] \neq \omega_{it}$.

To understand condition (9), note that it becomes $E[g^0(\omega_{it-1})|z_{it}] = \tilde{g}(E[\omega_{it-1}|x_{it-1}])$ if $x_{it-1} \subsetneq z_{it}$. This is difficult to satisfy as the left-hand side is a function of z_{it} whereas the right-hand side is a function of $x_{it-1} \subsetneq z_{it}$.

We therefore set $x_{it-1} \supseteq z_{it}$ in what follows. This amounts to rethinking how the OP/LP/ACF procedure is implemented, as foreshadowed by the special case of an $AR(1)$ process for productivity. In particular, because k_{it} is included in z_{it} in step 2, we must include k_{it+1} in x_{it} in step 1.

With this choice of $x_{it-1} \supseteq z_{it}$, we further examine condition (9) through a series of examples:

Example 1 (linear law of motion, $AR(1)$ process). If $g^0(\omega_{it-1}) = \rho\omega_{it-1}$, then $E[\rho\omega_{it-1}|z_{it}] = E[\rho E[\omega_{it-1}|x_{it-1}]|z_{it}]$ by the law of iterated expectations. Condition (9) is therefore satisfied with $\tilde{g}(E[\omega_{it-1}|x_{it-1}]) = g^0(E[\omega_{it-1}|x_{it-1}])$. ■

Remarkably, moment condition (5) in step 2 of the OP/LP/ACF procedure holds in Example 1 irrespective of the quality of the estimate of $E[q_{it}|x_{it}]$ in step 1.

Example 2 (quadratic law of motion). If $g^0(\omega_{it-1}) = \rho_1\omega_{it-1} + \rho_2\omega_{it-1}^2$ and $\text{Var}(\zeta_{it-1}|z_{it}) = \sigma^2$, then

$$E[\rho_1\omega_{it-1} + \rho_2\omega_{it-1}^2|z] = E[\rho_1 E[\omega_{it-1}|x_{it-1}] + \rho_2 E[\omega_{it-1}|x_{it-1}]^2 + \rho_2 \zeta_{it-1}^2|z].$$

Condition (9) is therefore satisfied with $\tilde{g}(E[\omega_{it-1}|x_{it-1}]) = g^0(E[\omega_{it-1}|x_{it-1}]) + \rho_2\sigma^2$.⁷ ■

Example 2 is less restrictive on the law of motion but more restrictive on the lagged prediction error than Example 1. The next example pushes this tradeoff further:

Example 3 (analytic law of motion). If the Taylor series of g^0 around $E[\omega_{it-1}|x_{it-1}]$ converges absolutely so that $g^0(\omega_{it-1}) = \sum_{j=0}^{\infty} \frac{g^{0,(j)}(E[\omega_{it-1}|x_{it-1}])}{j!} \zeta_{it-1}^j$, where $g^{0,(j)}$ is the j th derivative of g^0 , then

$$\begin{aligned} E\left[\sum_{j=0}^{\infty} \frac{g^{0,(j)}(E[\omega_{it-1}|x_{it-1}])}{j!} \zeta_{it-1}^j \middle| z_{it}\right] &= E\left[E\left[\sum_{j=0}^{\infty} \frac{g^{0,(j)}(E[\omega_{it-1}|x_{it-1}])}{j!} \zeta_{it-1}^j \middle| x_{it-1}\right] \middle| z_{it}\right] \\ &= E\left[\sum_{j=0}^{\infty} \frac{g^{0,(j)}(E[\omega_{it-1}|x_{it-1}])}{j!} E[\zeta_{it-1}^j|x_{it-1}] \middle| z_{it}\right]. \end{aligned}$$

If $E[\zeta_{it-1}^j|x_{it-1}] = \alpha_j$ for some constant α_j for all j , then condition (9) is satisfied in this example

⁷We can relax $\text{Var}(\zeta_{it-1}|z_{it}) = \sigma^2$ to $\text{Var}(\zeta_{it-1}|z_{it}) = h(E[\omega_{it-1}|x_{it-1}])$ for some function h . In this case, condition (9) holds with $\tilde{g}(E[\omega_{it-1}|x_{it-1}]) = g^0(E[\omega_{it-1}|x_{it-1}]) + \rho_2 h(E[\omega_{it-1}|x_{it-1}])$.

with $\tilde{g}(E[\omega_{it-1}|x_{it-1}]) = \sum_{j=0}^{\infty} \frac{g^{0,(j)}(E[\omega_{it-1}|x_{it-1}])}{j!} \alpha_j$.⁸ ■

Beyond these examples, condition (9) may be violated even if $x_{it-1} \supseteq z_{it}$.⁹ While this results in biased estimates, setting $x_{it-1} \supseteq z_{it}$ mitigates the bias. If $x_{it-1} \supseteq z_{it}$, then we have $E[\omega_{it-1}|z_{it}] = E[E[\omega_{it-1}|x_{it-1}]|z_{it}]$ by the law of iterated expectations and hence $E[\zeta_{it-1}|z_{it}] = E[\omega_{it-1} - E[\omega_{it-1}|x_{it-1}]|z_{it}] = 0$. This mean independence of the lagged prediction error ζ_{it-1} from the instruments z_{it} intuitively eliminates a first-order source of bias. The following theorem formalizes this intuition:

Theorem 2. *If $x_{it-1} \supseteq z_{it}$ and g is differentiable, then moment condition (5) in step 2 of the OP/LP/ACF procedure implicitly incorporates a first-order bias correction:*

$$E[q_{it} - f(k_{it}, v_{it}) - g(E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}))|z_{it}] = E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1})) - g'(q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1}))\zeta_{it-1}|z_{it}], \quad (10)$$

where g' is the first derivative of g .

We note that equation (10) holds for any production function f and any law of motion g for which the conditional moments exist. The proof of Theorem 2 is in Appendix B.

To appreciate Theorem 2, recall that moment condition (4) is valid even if invertibility fails. Because lagged planned output q_{it-1}^* is unobserved, moment condition (4) is infeasible for estimation. Step 2 of the OP/LP/ACF procedure therefore proceeds by substituting $E[q_{it-1}|x_{it-1}]$ from step 1 for q_{it-1}^* . If invertibility fails, then this substitution introduces the lagged prediction error $\zeta_{it-1} = q_{it-1}^* - E[q_{it-1}|x_{it-1}]$ into moment condition (5) (see conditional moment (7)). Theorem 2 shows that setting $x_{it-1} \supseteq z_{it}$ adjusts moment condition (5) toward the valid but infeasible moment condition (4).

Indeed, comparing the term $g(E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}))$ in moment condition (5) to the term $g(q_{it-1}^* - f(k_{it-1}, v_{it-1}))$ in moment condition (4) yields

$$\begin{aligned} & g(E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) - g(q_{it-1}^* - f(k_{it-1}, v_{it-1})) \\ &= g(q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1})) - g(q_{it-1}^* - f(k_{it-1}, v_{it-1})) \\ &= -g'(q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1}))\zeta_{it-1} + o(\zeta_{it-1}), \end{aligned} \quad (11)$$

where the last equality uses a first-order Taylor expansion of g around $q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1})$

⁸Condition (9) remains valid if we relax the condition that $E[\zeta_{it-1}^j|x_{it-1}] = \alpha_j$ for some constant α_j for all j to $\zeta_{it-1} \perp\!\!\!\perp z_{it} | E[\omega_{it-1}|x_{it-1}]$.

⁹Testing condition (9) is empirically challenging. A necessary condition for it to be satisfied is that moment condition (5) in step 2 of the OP/LP/ACF procedure holds at the estimated production function and law of motion. This can be tested with an overidentification test that corrects for the plug-in nature of the OP/LP/ACF procedure and clustering in the data.

and $o(\zeta_{it-1})$ denotes the higher-order remainder. Equation (10) in Theorem 2 shows that if $x_{it-1} \supseteq z_{it}$, then moment condition (5) implicitly accounts for the term $-g'(q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1})) \zeta_{it-1}$. Moment condition (5) thus differs from the valid but infeasible moment condition (4) at most by higher-order terms. Put differently, setting $x_{it-1} \supseteq z_{it}$ provides a first-order bias correction in step 2 of the OP/LP/ACF procedure.

Because setting $x_{it-1} \supseteq z_{it}$ provides a first-order bias correction in step 2 of the OP/LP/ACF procedure, a violation of condition (9) must be of second order relative to size of the prediction error. The following theorem formalizes this intuition:

Theorem 3. *Let \mathcal{G} be a set of functions that includes the true law of motion g^0 . If $x_{it-1} \supseteq z_{it}$ and g^0 is twice continuously differentiable, then*

$$\inf_{\tilde{g} \in \mathcal{G}} \left\| \mathbb{E} [g^0(\omega_{it-1}) | z_{it}] - \mathbb{E} [\tilde{g}(\mathbb{E}[\omega_{it-1} | x_{it-1}]) | z_{it}] \right\|_{L,1} \leq \tau \text{Var}(\zeta_{it-1}),$$

where $\|\cdot\|_{L,1}$ is the L_1 norm, $\tau = \sup_{\omega_{it-1}} |g^{0''}(\omega_{it-1})|$, and $g^{0''}$ is the second derivative of g^0 .

The proof of Theorem 3 is in Appendix B.

Theorem 3 ties the violation of condition (9) to the variance of the lagged prediction error ζ_{it-1} . Because $\text{Var}(\zeta_{it-1}) = \mathbb{E}[(\omega_{it-1} - \mathbb{E}[\omega_{it-1} | x_{it-1}])^2]$, this variance gets smaller as the observables x_{it-1} get richer even beyond z_{it} . Theorem 3 therefore suggests to take a kitchen sink approach to the regression in step 1 of the OP/LP/ACF procedure and add as many relevant covariates as possible. Adding covariates that speak to demand conditions and firm conduct may be especially helpful.¹⁰

Adding covariates may make the assumption $\mathbb{E}[\varepsilon_{it} | x_{it}] = 0$ that underpins step 1 of the OP/LP/ACF procedure more demanding. We advocate replacing $x_{it} = (k_{it}, v_{it}, \dots)$ in the original procedure by $x_{it} = (k_{it+1}, k_{it}, v_{it}, \dots)$. In the original procedure, the assumption $\mathbb{E}[\varepsilon_{it} | x_{it}] = 0$ is justified by stipulating that ε_{it} is outside the firm's information set in period t . This justification extends to our approach as the firm is assumed to choose k_{it+1} in period t . However, adding leads of other covariates to x_{it} may be more challenging if the firm eventually learns ε_{it} . To guard against this possibility, we recommend adding current values or lags of other covariates to x_{it} .

Finally, we caution that adding covariates introduces a curse of dimensionality. The resulting larger mean squared error of the estimate of $\mathbb{E}[q_{it-1} | x_{it-1}]$ may counterbalance the benefit of the smaller $\text{Var}(\zeta_{it-1})$ in finite samples. In the next section, we propose a modification of the OP/LP/ACF procedure that addresses this problem.

¹⁰The recent literature in fact proceeds along this line: “The materials demand function in our setting will take as arguments all state variables of the firm (...), including productivity, and all additional variables that affect a firm's demand for materials. These include firm location (...), output prices (...), product dummies (...), market shares (...), input prices (...), the export status of a firm (...), and the input (...) and output tariffs (...) that the firm faces on the product it produces” (De Loecker, Goldberg, Khandelwal & Pavcnik 2016, p. 466). Whether adding covariates restores invertibility can be assessed using the tests in Section 3.

Summary. Our analysis calls for rethinking how the OP/LP/ACF procedure is implemented. Theorem 1 makes a strong case for ensuring $x_{it-1} \supseteq z_{it}$. For a linear law of motion, the moment condition in step 2 of the OP/LP/ACF procedure holds for the true production function f^0 and the true law of motion g^0 . For a nonlinear law of motion, recovering the true production function f^0 requires condition (9) to be satisfied. A violation of condition (9) results in biased estimates. In this case, Theorem 2 shows that setting $x_{it-1} \supseteq z_{it}$ provides a first-order bias correction.

Theorem 1 further suggests to be flexible in modeling the law of motion in step 2 as condition (9) is easier to satisfy if \tilde{g} can come from a larger set \mathcal{G} . Theorem 3 finally suggests to take a kitchen sink approach to step 1.

Related literature. An alternative to relaxing the invertibility assumption is to forgo step 1 of the OP/LP/ACF procedure. Instead of substituting $E[q_{it-1}|x_{it-1}]$ for q_{it-1}^* in moment condition (4), the dynamic panel approach to production function estimation pioneered by Blundell & Bond (2000) substitutes q_{it-1} and imposes the linear law of motion $g(\omega_{it-1}) = \rho\omega_{it-1}$ to obtain

$$E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1} - f(k_{it-1}, v_{it-1}))|z_{it}] = E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1}^* - f(k_{it-1}, v_{it-1})) - \rho\varepsilon_{it-1}|z_{it}].$$

Besides the assumption $E[\xi_{it} + \varepsilon_{it}|z_{it}] = 0$ maintained in moment condition (4), the dynamic panel approach requires assuming $E[\varepsilon_{it-1}|z_{it}] = 0$. In comparison, our Example 1 requires setting $x_{it-1} \supseteq z_{it}$ and assuming $E[\varepsilon_{it-1}|x_{it-1}] = 0$ in step 1 of the OP/LP/ACF procedure. Either assumption can be justified by appealing to the firm's information set and the timing of decisions, as discussed above.

Hu, Huang & Sasaki (2020), Brand (2020), and Ponder (2021) generalize the dynamic panel approach from a linear to a nonlinear law of motion.¹¹ To illustrate the similarities and differences with our approach, consider the quadratic law of motion $g(\omega_{it-1}) = \rho_1\omega_{it-1} + \rho_2\omega_{it-1}^2$. Substituting q_{it-1} for q_{it-1}^* in moment condition (4) yields

$$\begin{aligned} E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1} - f(k_{it-1}, v_{it-1}))|z_{it}] &= E[q_{it} - f(k_{it}, v_{it}) \\ &\quad - g(q_{it-1}^* - f(k_{it-1}, v_{it-1})) - \rho_1\varepsilon_{it-1} - 2\rho_2\varepsilon_{it-1}(q_{it-1}^* - f(k_{it-1}, v_{it-1})) - \rho_2\varepsilon_{it-1}^2|z_{it}]. \end{aligned}$$

The generalized dynamic panel approach therefore requires assuming $E[\varepsilon_{it-1}|z_{it}] = 0$, $E[\varepsilon_{it-1}\omega_{it-1}|z_{it}] = 0$, and $\text{Var}(\varepsilon_{it-1}|z_{it}) = \sigma^2$. In comparison, our Example 2 requires setting $x_{it-1} \supseteq z_{it}$ and assuming $E[\varepsilon_{it-1}|x_{it-1}] = 0$ and $\text{Var}(\zeta_{it-1}|z_{it}) = \sigma^2$. Both approaches are similar in that they impose exclusion restrictions on second-order moments of latent variables to ensure that the moment condition

¹¹The estimation strategy in Hu et al. (2020) and Brand (2020) differs from their identification strategies. The latter builds on Hu & Schennach (2008) and relies on the existence of two conditionally independent proxies for productivity and the invertibility of an integral operator. The implied semiparametric maximum likelihood estimator involves significant computational challenges that preclude its implementation.

holds for the true production function.

This similarity extends to the polynomial law of motion $g(\omega_{it-1}) = \sum_{j=1}^{\infty} \rho_j \omega_{it-1}^j$. In this case, the generalized dynamic panel approach requires assuming that $E[\varepsilon_{it-1}^j | \omega_{it-1}, z_{it}]$ is constant for all j . This restriction on the joint distribution of ε_{it-1} and ω_{it-1} may be unpalatable if ε_{it-1} and ω_{it-1} are interpreted as the untransmitted, respectively, transmitted component of productivity (see footnote 3).¹² In comparison, our Example 3 requires assuming that $E[\zeta_{it-1}^j | x_{it-1}]$ is constant for all j .

The key difference is that our approach imposes restrictions on the lagged prediction error ζ_{it-1} rather than the lagged disturbance ε_{it-1} . Whereas the magnitude and properties of ε_{it-1} are inherent in the data generating process, the advantage of our approach is that the magnitude of ζ_{it-1} can be reduced by adding covariates to the regression in step 1 of the OP/LP/ACF procedure. Proceeding to add covariates may shrink $\text{Var}(\zeta_{it-1}) = E[(\omega_{it-1} - E[\omega_{it-1} | x_{it-1}])^2]$ towards zero if we approach invertibility in the limit. This makes our approach particularly attractive for datasets with a rich set of observables.

5 Modification of OP/LP/ACF moment condition

Step 1 of the OP/LP/ACF procedure estimates the conditional expectation $E[q_{it-1} | x_{it-1}]$. To understand the impact that a noisy estimate in step 1 has on the GMM estimator in step 2, we use $\hat{e}(x_{it-1})$ to denote the estimate of a nuisance parameter $e(x_{it-1})$ with true value $E[q_{it-1} | x_{it-1}]$. We remain agnostic about the estimation method used in step 1.

Replacing $E[q_{it-1} | x_{it-1}]$ by $\hat{e}(x_{it-1})$ in moment condition (5) that underpins step 2 yields

$$E[q_{it} - f(k_{it}, v_{it}) - g(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1})) | z_{it}] = 0. \quad (12)$$

Because $\hat{e}(x_{it-1})$ converges to $E[q_{it-1} | x_{it-1}]$ as the sample size increases, we write $\hat{e}(x_{it-1}) = E[q_{it-1} | x_{it-1}] + \lambda \delta(x_{it-1})$, where λ converges to zero as sample size increases and $\delta(x_{it-1})$ is the direction of the finite-sample deviation of $\hat{e}(x_{it-1})$ from $E[q_{it-1} | x_{it-1}]$. The pathwise (or Gateaux) derivative of the left-hand side of moment condition (12) with respect to the direction $\delta(x_{it-1})$ at $\hat{e}(x_{it-1}) = E[q_{it-1} | x_{it-1}]$ is

$$\begin{aligned} & \frac{\partial}{\partial \lambda} E[q_{it} - f(k_{it}, v_{it}) - g(E[q_{it-1} | x_{it-1}] + \lambda \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) | z_{it}] \Big|_{\lambda=0} \\ &= -E[g'(E[q_{it-1} | x_{it-1}] - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) | z_{it}]. \end{aligned}$$

Because this derivative is generally non-zero, moment condition (5) is impacted by the deviation

¹²If $x_{it-1} \subseteq z_{it}$ and invertibility holds, then $E[\varepsilon_{it-1}^j | \omega_{it-1}, z_{it}] = E[\varepsilon_{it-1}^j | z_{it}]$. While this facilitates interpreting the restriction, we focus on the case where invertibility fails.

of $\hat{e}(x_{it-1})$ from $E[q_{it-1}|x_{it-1}]$ that arises from estimation noise in step 1.

To eliminate this impact and its adverse consequences for the asymptotic distribution of the GMM estimator, we modify the moment condition in step 2 as follows:

$$E[q_{it} - f(k_{it}, v_{it}) - g(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1})) - g'(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1}))(q_{it-1} - \hat{e}(x_{it-1}))|z_{it}] = 0. \quad (13)$$

Our modification explicitly incorporates a feasible version of the first-order bias correction that is implicit in moment condition (5) if $x_{it-1} \supseteq z_{it}$ (see Theorem 2).

If $\hat{e}(x_{it-1}) = E[q_{it-1}|x_{it-1}]$ and $x_{it-1} \supseteq z_{it}$, then we have $E[q_{it-1} - \hat{e}(x_{it-1})|z_{it}] = 0$ so that moment condition (13) coincides with moment condition (5). Our modification therefore retains the identification power of the original OP/LP/ACF procedure when $x_{it-1} \supseteq z_{it}$.

However, unlike the original moment condition (5), the modified moment condition (13) is not sensitive to small deviations of $\hat{e}(x_{it-1})$ from $E[q_{it-1}|x_{it-1}]$, a property known as *Neyman orthogonality* (Neyman 1959):

Theorem 4 (Neyman orthogonality). *Define the set of functions $\tilde{\mathcal{T}} = \{\delta(x_{it-1}) = e(x_{it-1}) - E[q_{it-1}|x_{it-1}] : e \in \mathcal{T}\}$, where \mathcal{T} is the set of real-valued integrable functions of x_{it-1} . If $x_{it-1} \supseteq z_{it}$ and g has a bounded and continuous derivative, then the pathwise (or Gateaux) derivative of the left-hand side of moment condition (13) with respect to any direction $\delta \in \tilde{\mathcal{T}}$ at $\hat{e}(x_{it-1}) = E[q_{it-1}|x_{it-1}]$ is zero:*

$$\frac{\partial}{\partial \lambda} E[q_{it} - f(k_{it}, v_{it}) - g(E[q_{it-1}|x_{it-1}] + \lambda \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) - g'(E[q_{it-1}|x_{it-1}] + \lambda \delta(x_{it-1}) - f(k_{it-1}, v_{it-1}))(q_{it-1} - E[q_{it-1}|x_{it-1}] - \lambda \delta(x_{it-1}))|z_{it}] \Big|_{\lambda=0} = 0.$$

The proof of Theorem 4 is in Appendix D.^{13,14}

Neyman orthogonality has received renewed attention in the double-debiased machine learning literature. Following the arguments in Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey & Robins (2018) and Chernozhukov, Escanciano, Ichimura, Newey & Robins (2022), as long as $\hat{e}(x_{it-1})$ converges to $E[q_{it-1}|x_{it-1}]$ at a rate faster than $N^{-1/4}$, the GMM estimator based on moment condition (13) is *oracle efficient*. This means that its asymptotic distribution is *as if* the true value $E[q_{it-1}|x_{it-1}]$ of the nuisance parameter $e(x_{it-1})$ is known. The GMM estimator

¹³The assumption that g has a bounded derivative simplifies stating Theorem 4. As shown in Appendix D, it can be replaced by a much weaker assumption.

¹⁴We recommend ensuring $x_{it-1} \supseteq z_{it}$ when implementing our modification of the OP/LP/ACF procedure. Without $x_{it-1} \supseteq z_{it}$, the modified moment condition (13) entails a first-order bias correction toward $E[q_{it-1}|x_{it-1}, z_{it}]$ instead of $E[q_{it-1}|x_{it-1}]$ and the expectation of the correction term conditional on z_{it} is generally non-zero even if $\hat{e}(x_{it-1}) = E[q_{it-1}|x_{it-1}]$.

based on moment condition (13) therefore achieves the efficient asymptotic distribution.

Besides efficiency, the modified moment condition (13) enjoys further advantages over the original moment condition (5). First, because Theorem 4 allows for any integrable deviation, it accommodates a wide range of estimation methods in step 1 of the OP/LP/ACF procedure. This includes traditional nonparametric methods such as sieve, kernel, and LASSO estimators, which produce smooth estimates, as well as modern machine learning techniques such as neural networks and random forests, which may produce non-smooth estimates. Neyman orthogonality means that small deviations of $\hat{e}(x_{it-1})$ from $E[q_{it-1}|x_{it-1}]$ that can arise from estimation noise, functional approximation error in sieve estimation, local smoothing bias in kernel estimation, or regularization bias in LASSO and machine learning methods do not have a first-order impact on the GMM estimator in step 2. As long as $\hat{e}(x_{it-1})$ converges to $E[q_{it-1}|x_{it-1}]$ at a rate faster than $N^{-1/4}$, its asymptotic distribution is robust to the choice of estimation method and its implementation in step 1.¹⁵

Second, moment condition (13) eliminates the need to correct standard errors in step 2 to account for estimation noise in step 1 through techniques such as those in Hahn, Liao & Ridder (2018).

Unlike general approaches that generate Neyman orthogonality using influence functions or score projections as described in Chernozhukov et al. (2022), moment condition (13) is Neyman orthogonal by construction. By avoiding introducing high-dimensional nuisance parameters, our modification remains as straightforward to implement as the original OP/LP/ACF procedure.

Finally, we note that if OLS is used in step 1 of OP/LP/ACF procedure to estimate the conditional expectation $E[q_{it-1}|x_{it-1}]$, then explicitly including the first-order bias correction can be redundant in some cases. Converting moment condition (13) into an unconditional moment condition yields

$$E \left[h(z_{it}) \left(q_{it} - f(k_{it}, v_{it}) - g(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1})) \right. \right. \\ \left. \left. - g'(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1}))(q_{it-1} - \hat{e}(x_{it-1})) \right) \right] = 0,$$

where $h(z_{it})$ is a vector function of the instruments z_{it} . The first-order bias correction is redundant in a finite sample if

$$\frac{1}{NT} \sum_i \sum_t h(z_{it}) g'(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1}))(q_{it-1} - \hat{e}(x_{it-1})) = 0. \quad (14)$$

¹⁵ Achieving this rate of convergence requires properly implementing a nonparametric estimation method to avoid functional form bias and the validity of smoothness assumptions specific to the selected method. Because Neyman orthogonality is a local property, the first-order bias correction in moment condition (13) loses its theoretical justification if $\hat{e}(x_{it-1})$ deviates substantially from $E[q_{it-1}|x_{it-1}]$.

Because $q_{it-1} - \hat{e}(x_{it-1})$ is the residual from the OLS estimation in step 1, equation (14) holds if $h(z_{it})g'(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1}))$ can be expressed as a linear combination of the regressors used in step 1. This is the case if the law of motion g is linear and the regressors include $h(z_{it})$. More generally, if $h(z_{it})g'(\hat{e}(x_{it-1}) - f(k_{it-1}, v_{it-1}))$ is well approximated by a linear combination of the regressors used in step 1, then the left-hand side of equation (14) is almost zero and we expect the GMM estimator in step 2 to attain near-oracle performance. While our Monte Carlo exercise in Section 6 illustrates that this can happen, we recommend basing the GMM estimator in step 2 on the modified moment condition (13), as this ensures oracle efficiency irrespective of whether equation (14) holds or not. Moreover, if estimation methods other than OLS are used in step 1, then equation (14) generally does not hold and the modified moment condition (13) improves efficiency.

6 Monte Carlo exercise

Data generating process. We specify the CES production function

$$f(k_{it}, v_{it}) = \frac{\nu}{\rho} \ln(\alpha \exp(\rho k_{it}) + (1 - \alpha) \exp(\rho v_{it}))$$

and the disturbance $\varepsilon_{it} \sim N(0, \sigma_\varepsilon^2)$, where $\alpha \in (0, 1)$ is a distributional parameter, $\rho \leq 0$ and $\sigma = \frac{1}{1-\rho}$ is the elasticity of substitution, and $\nu > 0$ is returns to scale. We further specify the law of motion

$$g(\omega_{it-1}) = \mu_\omega + \rho_\omega \left((1 - \alpha_\omega) \omega_{it-1} + \frac{\alpha_\omega}{6} \ln(\ln(1 + \exp(6\omega_{it-1}))) \right)$$

and the productivity innovation $\xi_{it} \sim N(0, \sigma_\omega^2)$, where $\rho_\omega \in (0, 1)$ and $\alpha_\omega \in [0, 1]$. Our particular interest lies in comparing the Gaussian $AR(1)$ process for $\alpha_\omega = 0$ with the nonlinear process for $\alpha_\omega = 1$. To facilitate this comparison we calibrate μ_ω , ρ_ω , and σ_ω^2 to hold fixed $E[\omega_{it}]$, $\text{Var}(\omega_{it})$, and $\text{corr}(\omega_{it}, \omega_{it-1})$. Figure 1 illustrates the resulting distribution of productivity and overlays the law of motion for the two processes.

We specify the CES demand

$$q_{it}^* = \delta_{1i} - (1 + \exp(-\delta_{2i}))p_{it},$$

where p_{it} is the output price and $\delta_i = (\delta_{1i}, \delta_{2i})$ captures shocks to the demand the firm faces and unobserved rivals. This avoids having to specify the imperfectly competitive environment that the firm operates in. Abstracting from time-series variation for simplicity, we specify $\delta_{1i} \sim N(\mu_{\delta_1}, \sigma_{\delta_1}^2)$ and $\delta_{2i} \sim N(\mu_{\delta_2}, \sigma_{\delta_2}^2)$.

We denote the price of capital as p_i^K and the price of the variable input as p_i^V . Abstracting from time-series variation, we specify $p_i^K \sim N(\mu_{p^K}, \sigma_{p^K}^2)$ and $p_i^V \sim N(\mu_{p^V}, \sigma_{p^V}^2)$. Assuming short-run

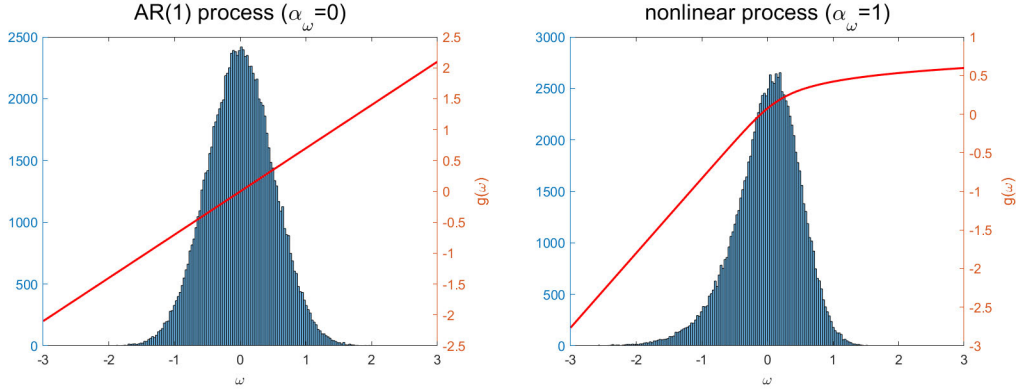


Figure 1: Distribution of productivity (left axis) and law of motion (right axis). Baseline parameterization with Gaussian $AR(1)$ process ($\alpha_\omega = 0$, left panel) and nonlinear process ($\alpha_\omega = 1$, right panel).

parameter	α	ρ	ν	$E[\omega_{it}]$	$\text{Var}(\omega_{it})$	$\text{corr}(\omega_{it}, \omega_{it-1})$...	
value	0.3	-1	0.95	0	0.5^2	0.7	...	
...	μ_{δ_1}	$\sigma_{\delta_1}^2$	μ_{δ_2}	$\sigma_{\delta_2}^2$	μ_{p^K}	$\sigma_{p^K}^2$	μ_{p^V}	$\sigma_{p^V}^2$
...	10	5^2	-1.3543	0.5^2	0	0.5^2	0	0.5^2

Table 1: Baseline parameterization.

profit maximization and equating marginal revenue with marginal cost determines v_{it} along with q_{it}^* , q_{it} , and p_{it} . Note that because v_{it} is a function of k_{it} , p_i^V , ω_{it} , and δ_i , the firm's demand for the variable input cannot be inverted to express ω_{it} as a function of observables: invertibility fails as long as $\sigma_{\delta_1}^2 > 0$ or $\sigma_{\delta_2}^2 > 0$.

Recall that capital k_{it+1} is chosen in period t . We assume that capital fully depreciates between periods and endow the firm with static expectations regarding ω_{it+1} . Hence, the firm chooses k_{it+1} in period t given ω_{it} , δ_i , p_i^K , and p_i^V as if k_{it+1} and v_{it+1} are variable inputs.¹⁶

Table 1 shows our baseline parameterization. The elasticity of substitution is within the range of estimates in the literature (Chirinko 2008), as are the 90:10 percentile ratio and persistence of productivity (De Loecker & Syverson 2021). Short-run profit maximization implies the markup $\mu_{it} = \frac{P_{it}}{MC_{it}} = 1 + \exp(\delta_{2i})$, with $E[\ln \mu_{it}] = 0.25$ and $\text{Var}(\ln \mu_{it}) = 0.0126$. We simulate $S = 1000$ datasets with $N = 5,000$ firms and $T = 20$ periods.¹⁷ We provide further details on the

¹⁶ Assuming constant returns to scale in production and quadratic adjustment costs to capital, the firm's dynamic programming problem that determines investment can be solved in closed form if the firm is a price-taker in the output market (Syverson 2001, Van Biesebroeck 2007, Akerberg et al. 2015). As this is no longer possible if the firm has market power, we opt for a different, tractable specification of the evolution of capital.

¹⁷ We use $T_0 = 5000$ burn-in periods.

specification in Appendix C.

Estimation. We estimate $E[q_{it}|x_{it}]$ in step 1 of the OP/LP/ACF procedure by OLS using the complete set of Hermite polynomials of total degree 4 in the variables in x_{it} as regressors. To illustrate Theorems 1–3, we vary the specification of x_{it} as detailed below. As an alternative to OLS, we use a neural network estimator with two hidden layers of 128 neurons. We provide further details on the neural network estimator in Appendix C.

We estimate the parameters $\theta = (\alpha, \rho, \nu, \mu_\omega, \rho_\omega, \alpha_\omega)$ in step 2 using the complete set of Hermite polynomials of total degree 4 in the variables in $z_{it} = (k_{it}, k_{it-1}, v_{it-1}, p_i^V)$ as instruments.¹⁸ To illustrate Theorem 4 and the advantages of Neyman orthogonality in finite samples, we alternatively base the GMM estimator on moment conditions (5) and (13). We provide further details on the GMM estimator in Appendix C.

We use three statistics to summarize the results. First, we conduct a Lagrange multiplier test for the true parameter value θ^0 . Our test corrects for the plug-in nature of the OP/LP/ACF procedure and clustering in the data. We provide further details on the Lagrange multiplier test in Appendix C.

Second, following the production function approach to markup estimation, we use the estimate of θ_f to estimate the markup $\mu_{it} = \frac{P_{it}}{MC_{it}}$ of firm i in period t as

$$\ln \mu_{it} + \varepsilon_{it} = p_{it} + q_{it} - p_i^V - v_{it} + \ln \frac{\partial f(k_{it}, v_{it})}{\partial v_{it}}.$$

The right-hand side is the log of the output elasticity minus the log of the expenditure share of the variable input.¹⁹ Noting that the disturbance ε_{it} averages out as $E[\varepsilon_{it}] = 0$, we refer to the average of $\ln \mu_{it} + \varepsilon_{it}$ across firms and time simply as the average log markup.

Third, because the law of motion is of primary interest in some applications (Aw, Roberts & Xu 2011, De Loecker 2013, Doraszelski & Jaumandreu 2013), we use the estimate of θ_g to estimate

$$g'(0) = \rho_\omega \left((1 - \alpha_\omega) + \alpha_\omega \frac{1}{2 \ln 2} \right)$$

as a measure of persistence in the productivity process.

¹⁸We treat the price of the variable input p_i^V as observable to focus on demand shocks as the reason for the failure of invertibility. The nonidentification result in Gandhi, Navarro & Rivers (2020) assumes invertibility and therefore does not apply in our setting.

¹⁹De Loecker & Warzynski (2012) isolate $\ln \mu_{it}$ on the left-hand side and use the residual from the regression in step 1 of the OP/LP/ACF procedure to estimate ε_{it} on the right-hand side. However, using equation (6), the residual is $q_{it} - E[q_{it}|x_{it}] = \varepsilon_{it} + \zeta_{it}$. Having ε_{it} on the left-hand side thus avoids that ζ_{it} taints the correlation of the estimated markup with a variable of interest such as the firm's export status or a measure of trade liberalization (Doraszelski & Jaumandreu 2024).

Results. We compare three cases for the observables x_{it} in step 1 of the OP/LP/ACF procedure:

- *Case 1:* $x_{it} = (k_{it}, v_{it}, p_i^V)$;
- *Case 2:* $x_{it} = (k_{it+1}, k_{it}, v_{it}, p_i^V)$;
- *Case 3:* $x_{it} = (k_{it+1}, k_{it}, v_{it}, p_{it}, p_i^V)$.

Case 1 is our baseline and corresponds to how the OP/LP/ACF procedure is implemented in the existing literature. In light of Theorems 1 and 2, case 2 ensures $x_{it-1} \supseteq z_{it}$. Case 3 heeds Theorem 3 and pursues a kitchen sink approach by adding the output price as a covariate that speaks to demand conditions.

Figure 2 shows the results of the OP/LP/ACF procedure with moment condition (5) for the Gaussian $AR(1)$ process ($\alpha_\omega = 0$). In case 1, the failure of invertibility causes a massive bias in the production function and in the average log markup derived from it. Indeed, the average log markup centers around -0.24 whereas short-run profit maximization implies $\ln \mu_{it} = \ln(1 + \exp(\delta_{2i})) > 0$. The failure of invertibility also causes a large bias in the law of motion as summarized by the measure of persistence $g'(0)$. The Lagrange multiplier test for the true parameter value θ^0 rejects by a wide margin. In short, the failure of invertibility can have a substantial, detrimental impact on the OP/LP/ACF procedure as implemented in the existing literature.

In cases 2 and 3, the Lagrange multiplier test does not reject. As in Example 1 in Section 4, ensuring $x_{it-1} \supseteq z_{it}$ eliminates the bias. The average log markup centers around its true value of 0.25 and the measure of persistence $g'(0)$ centers around its true value of 0.7. The first panel of Table 2 summarizes the results. In addition to the bias, it shows the variance and mean squared error of the average log markup and the measure of persistence $g'(0)$.

Turning to the nonlinear process ($\alpha_\omega = 1$), Figure 3 shows the results of the OP/LP/ACF procedure with moment condition (5). In case 1, the failure of invertibility again causes a massive bias in the average log markup and the measure of persistence $g'(0)$. The Lagrange multiplier test for the true parameter value θ^0 rejects by a wide margin.

Ensuring $x_{it-1} \supseteq z_{it}$ greatly reduces the bias in cases 2 and 3. In case 2, the average log markup centers around 0.30 compared to its true value of 0.25. The measure of persistence $g'(0)$ centers around its true value of 0.7. The Lagrange multiplier test rejects by a wide margin. Adding the output price as a covariate that speaks to demand conditions further reduces—and here almost eliminates—the bias. In case 3, the average log markup centers around 0.25 and the measure of persistence $g'(0)$ centers around 0.7. The Lagrange multiplier test rejects for just 72 of the 1000 simulated datasets at a significance level of 5%. The second panel of Table 2 summarizes the results.

As discussed in Section 5, explicitly including the first-order bias correction generally mitigates the adverse impact that a noisy estimate in step 1 has on the GMM estimator in step 2, although

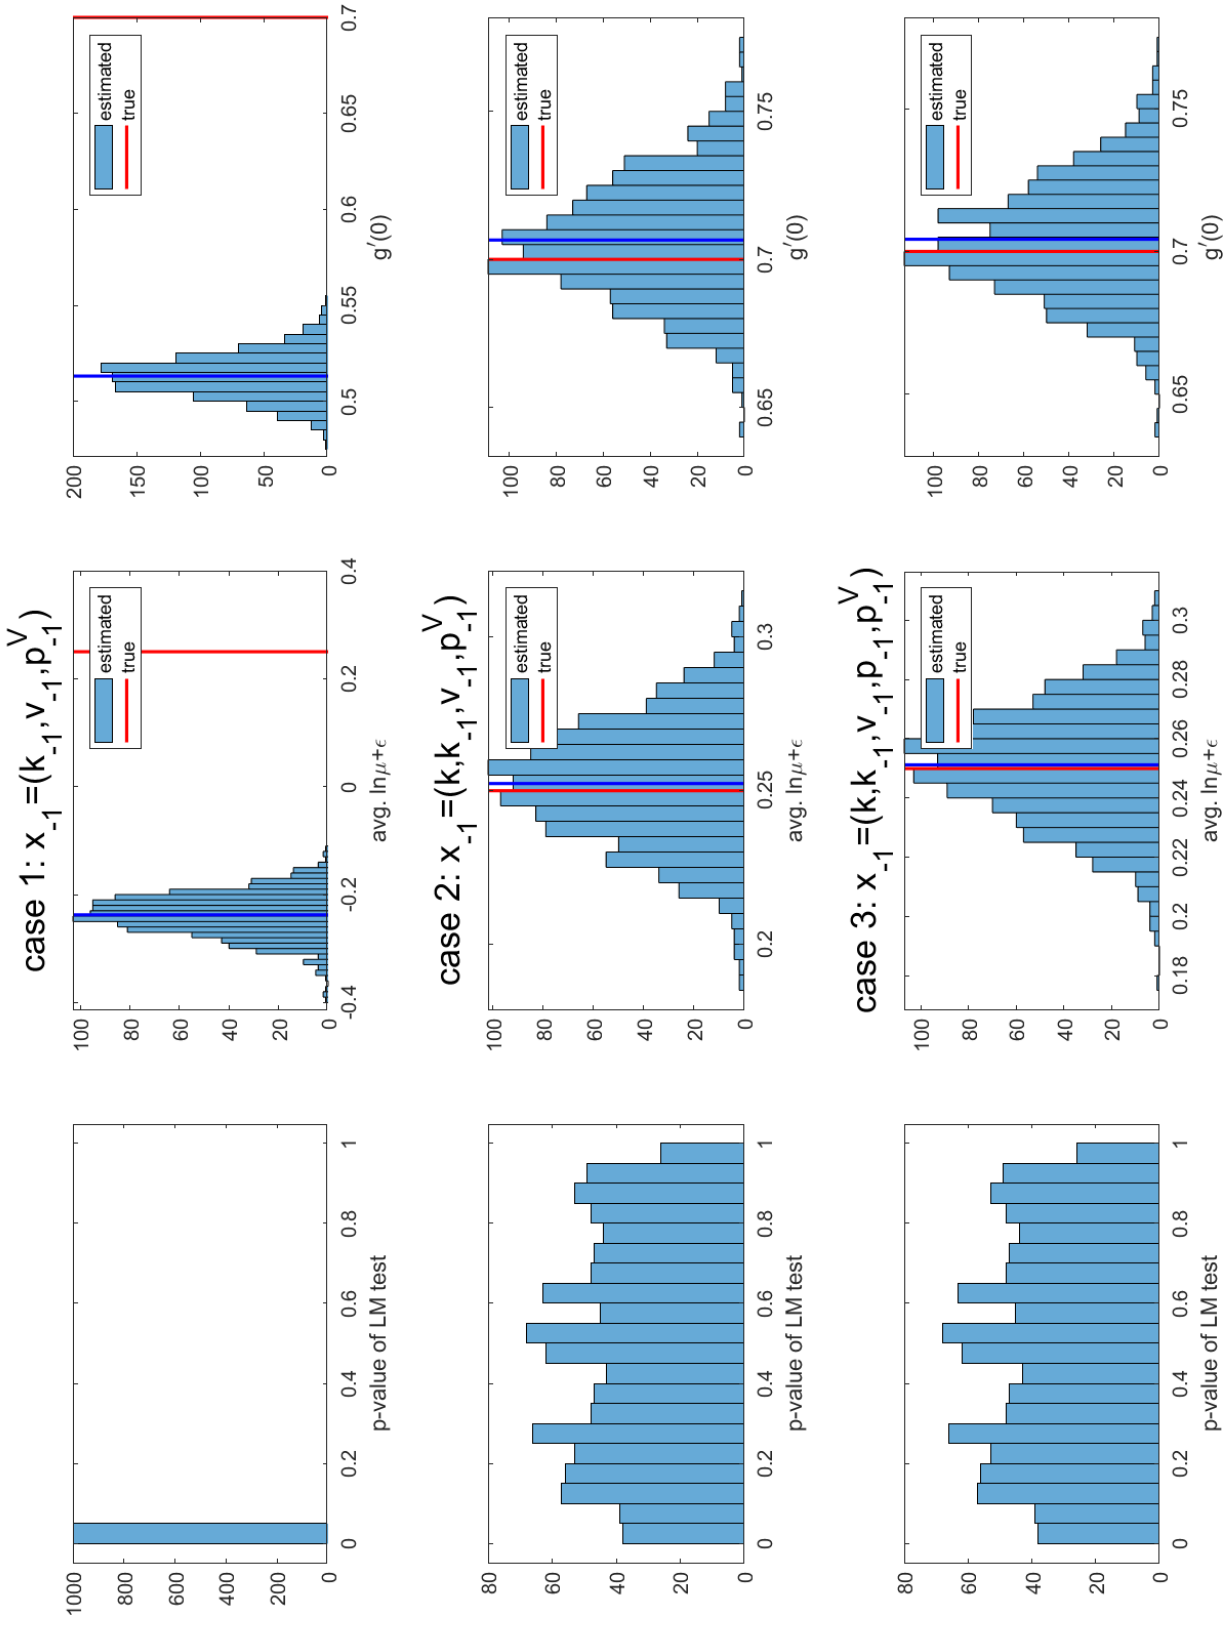


Figure 2: Cases 1, 2, and 3 and moment condition (5). Baseline parameterization with Gaussian $AR(1)$ process ($\alpha_\omega = 0$). Case 1 (upper row), case 2 (middle row), and case 3 (lower row). p -value for Lagrange multiplier test (left column), average log markup (middle column), and measure of persistence $g'(0)$ (right column).

	average log markup			measure of persistence $g'(0)$		
	bias	variance	MSE	bias	variance	MSE
panel 1: baseline parameterization with $AR(1)$ process, moment condition (5) :						
case 1	-0.4875	0.0016	0.2392	-0.1869	0.0001	0.0351
case 2	0.0023	0.0004	0.0004	0.0065	0.0004	0.0005
case 3	0.0012	0.0004	0.0004	0.0042	0.0004	0.0004
panel 2: baseline parameterization with nonlinear process, moment condition (5) :						
case 1	-0.3970	0.0015	0.1591	-0.1524	0.0003	0.0235
case 2	0.0511	0.0005	0.0031	0.0052	0.0003	0.0003
case 3	0.0030	0.0004	0.0004	0.0041	0.0002	0.0002
panel 3: baseline parameterization with nonlinear process, moment condition (13) :						
case 2	0.0415	0.0005	0.0022	-0.0073	0.0003	0.0003
case 3	0.0037	0.0004	0.0004	0.0028	0.0002	0.0002
panel 4: modified parameterization with nonlinear process, moment condition (5) :						
case 2	0.5743	0.0012	0.3309	0.0960	0.0001	0.0093
case 3	-0.0314	0.0033	0.0043	0.0021	0.0000	0.0000
panel 5: modified parameterization with nonlinear process, moment condition (13) :						
case 2	-0.1327	0.0260	0.0436	-0.0098	0.0000	0.0001
case 3	-0.0031	0.0023	0.0023	-0.0012	0.0000	0.0000
panel 6: baseline parameterization with nonlinear process, neural network estimator, moment condition (5) :						
case 2	0.0353	0.0202	0.0214	-0.0209	0.0222	0.0227
case 3	-0.0196	0.0185	0.0189	-0.0165	0.0081	0.0084
panel 7: baseline parameterization with nonlinear process, neural network estimator, moment condition (13) :						
case 2	0.0400	0.0006	0.0022	-0.0054	0.0003	0.0004
case 3	0.0065	0.0005	0.0005	0.0043	0.0002	0.0002

Table 2: Bias, variance, and mean squared error of average log markup and measure of persistence $g'(0)$.

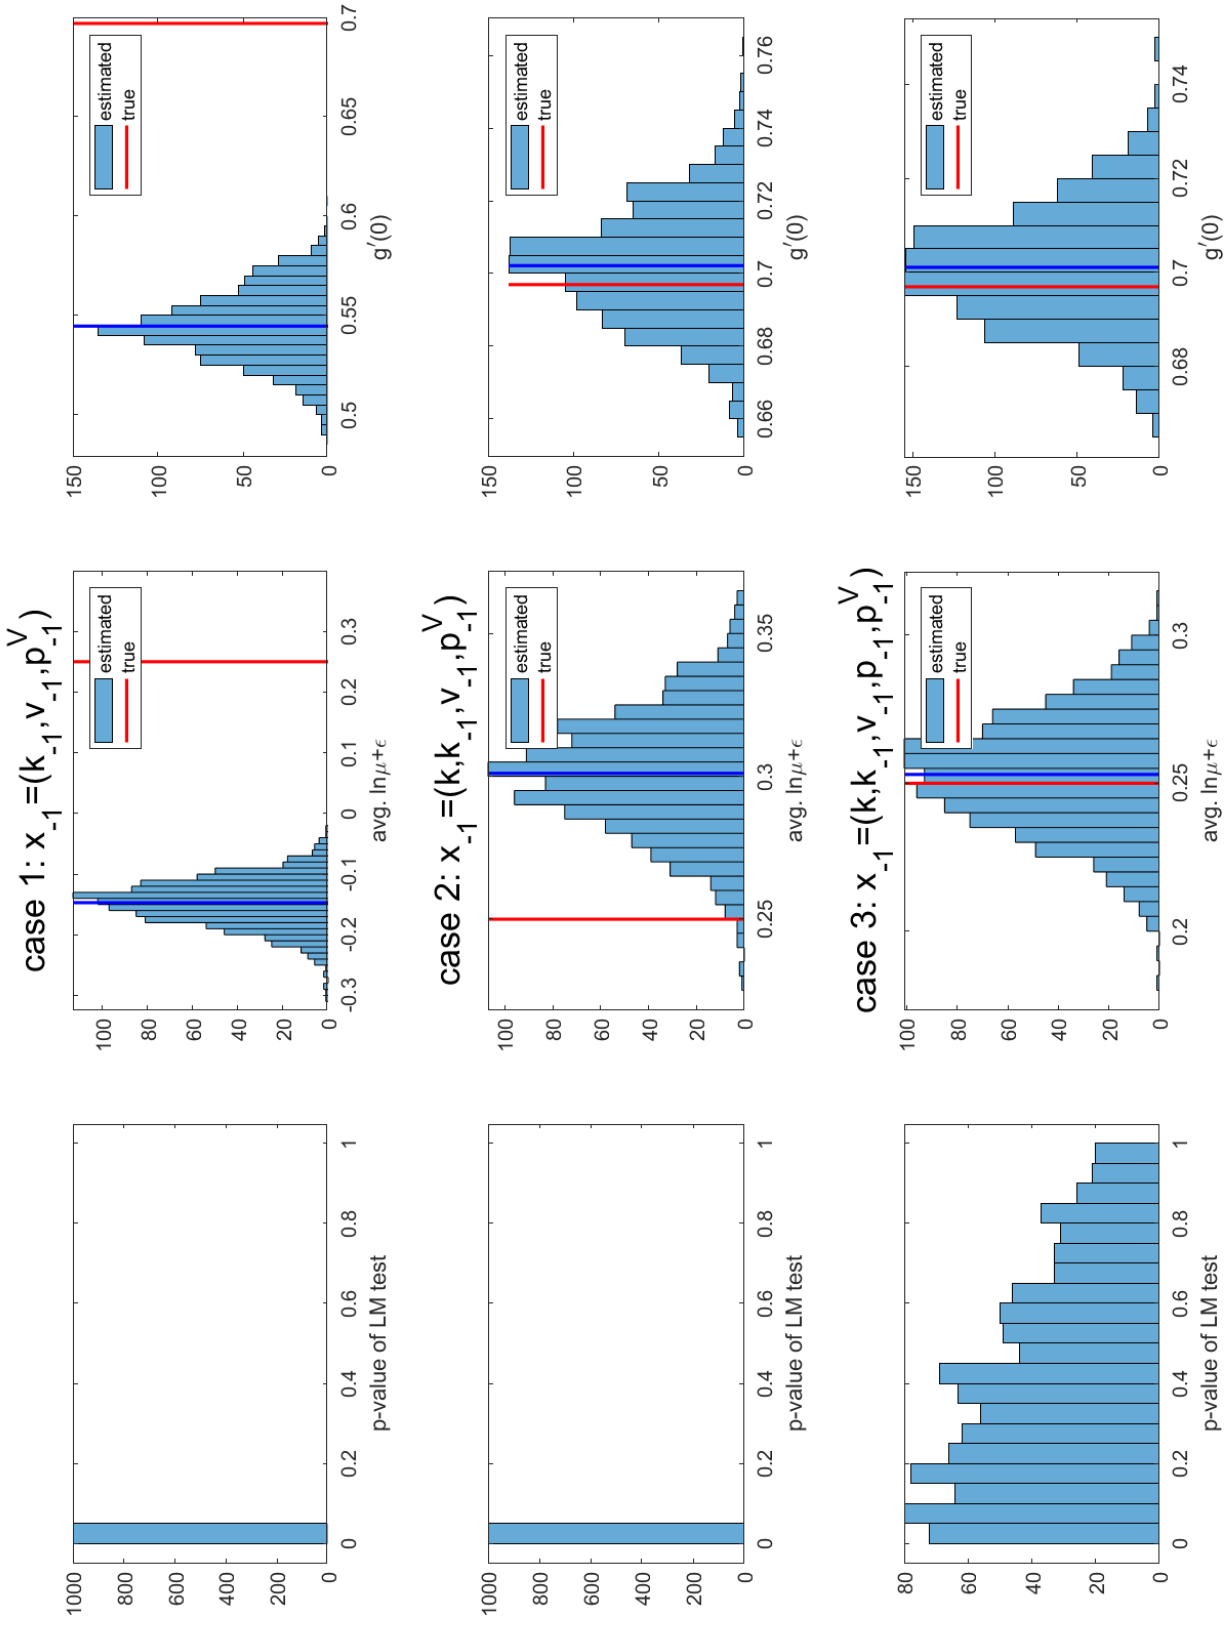


Figure 3: Cases 1, 2, and 3 and moment condition (5). Baseline parameterization with nonlinear process ($\alpha_\omega = 1$). Case 1 (upper row), case 2 (middle row), and case 3 (lower row). p -value for Lagrange multiplier test (left column), average log markup (middle column), and measure of persistence $g'(0)$ (right column).

it can sometimes be redundant if OLS is used in step 1. The third panel of Table 2 summarizes the results of replacing moment condition (5) by moment condition (13) for the nonlinear process ($\alpha_\omega = 1$). The results are indeed similar for the two moment conditions.²⁰

However, explicitly including the first-order bias correction is not always redundant even if OLS is used in step 1. To illustrate the advantages of Neyman orthogonality in finite samples, we modify the parameterization to $E[\omega_{it}] = -1.25$, $\text{Var}(\omega_{it}) = 2^2$, $\text{corr}(\omega_{it}, \omega_{it-1}) = 0.85$, $\sigma_{\delta_1}^2 = 0.5^2$, $\mu_{\delta_2} = -2.5425$, $\sigma_{\delta_2}^2 = 2^2$, and $\sigma_{pK}^2 = 2^2$. This implies an increase in the 90:10 percentile ratio and persistence of productivity and a mean-preserving spread of the log markup ($E[\ln \mu_{it}] = 0.25$ and $\text{Var}(\ln \mu_{it}) = 0.1986$).

Figures 4 and 5 show the results of the OP/LP/ACF procedure with moment conditions (5) and (13) for the nonlinear process ($\alpha_\omega = 1$). Figure 4 pertains to case 2 and Figure 5 to case 3. In case 2, there is a large bias with the original moment condition (5). The average log markup centers around 0.82 compared to its true value of 0.25 and the measure of persistence $g'(0)$ centers around 0.79 compared to its true value of 0.7. The bias is reduced with the modified moment condition (13). The average log markup centers around 0.12 and the measure of persistence $g'(0)$ centers around 0.69. In case 3, the bias is again reduced with the modified moment condition (13). With the original moment condition (5), the average log markup centers around 0.22 and the measure of persistence $g'(0)$ centers around 0.70. The Lagrange multiplier tests rejects by a wide margin. With the modified moment condition (13), the average log markup centers around 0.25 and the measure of persistence $g'(0)$ centers around 0.69. The Lagrange multiplier test rejects for just 205 of the 1000 simulated datasets at a 5% significance level. The fourth and fifth panels of Table 2 summarize the results.

Finally, we use a neural network estimator in step 1 as an alternative to OLS. Reverting to the baseline parameterization with the nonlinear process ($\alpha_\omega = 1$), the sixth and seventh panels of Table 2 summarize the results. Comparing the sixth to the second panel of Table 2 shows that using the neural network estimator increases the variance and thereby the mean squared error of the GMM estimator in step 2 by an order of magnitude if we base the GMM estimator on the original moment condition (5). To the best of our knowledge, there is no theory on using a neural network as a plug-in estimator. In contrast, we achieve much better results if we base the GMM estimator on the modified moment condition (13). Comparing the seventh to the third panel shows that using the neural network estimator is comparable to using OLS. As discussed in

²⁰To confirm, we disrupt the finite-sample relation in equation (14) by running OLS on a separate dataset and using the estimates to construct $\hat{e}(x_{it-1})$ for the focal dataset in step 1. For the original moment condition (5), the mean squared error of the average log markup and the measure of persistence $g'(0)$ nearly doubles in case 3. In case 2, the mean squared error of the average log markup decreases slightly while the mean squared error of the measure of persistence $g'(0)$ increases by an order of magnitude. In contrast, there are no meaningful changes for the modified moment condition (13).

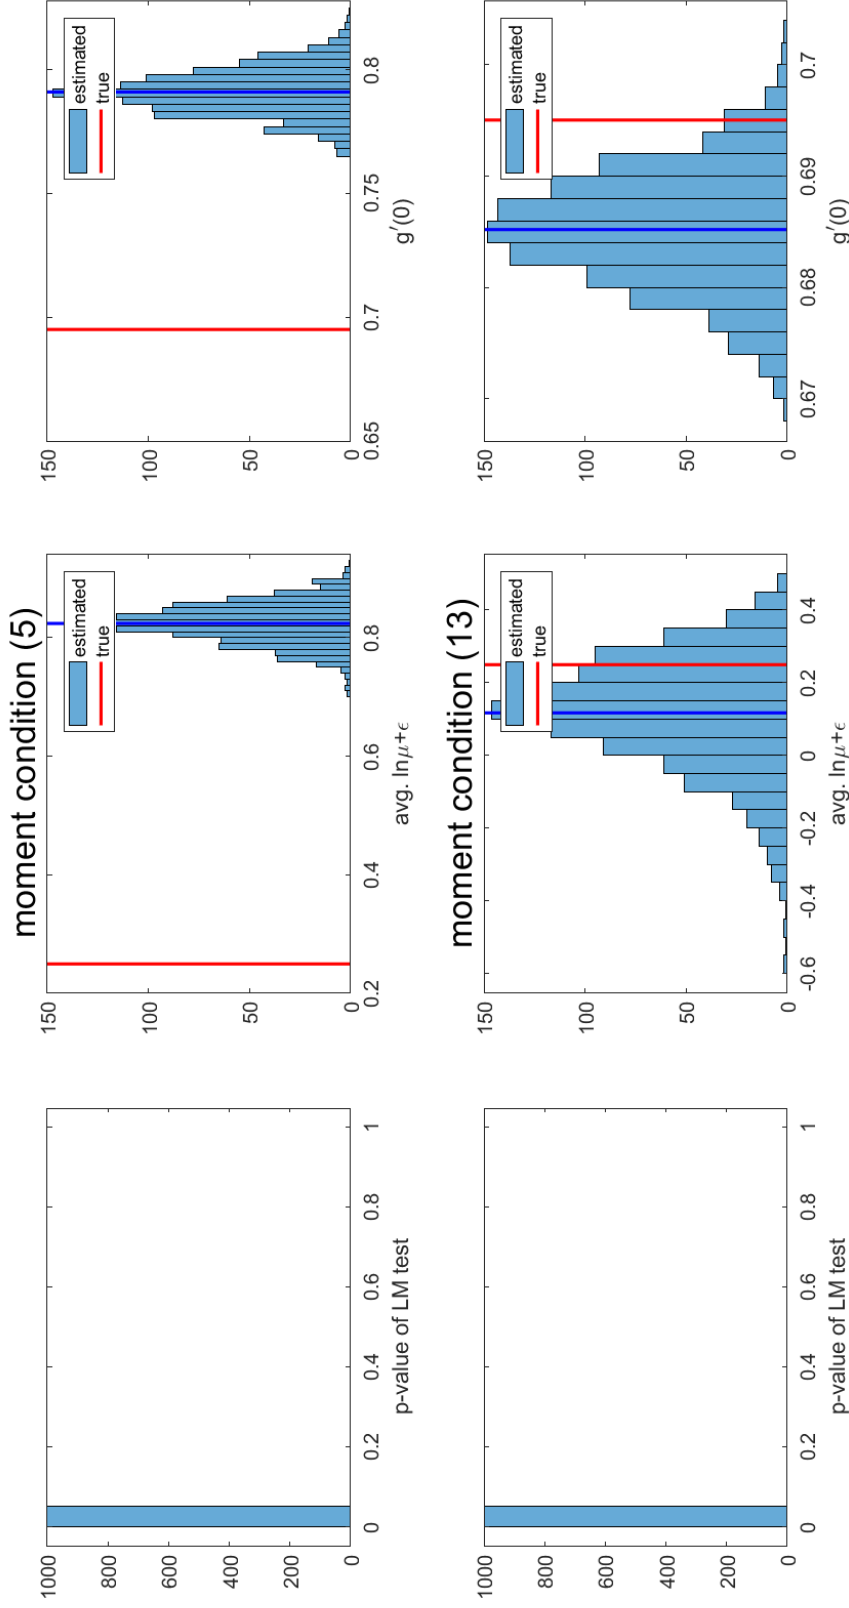


Figure 4: Case 2 and moment conditions (5) and (13). Modified parameterization with nonlinear process ($\alpha_\omega = 1$). Moment condition (5) (upper row) and moment condition (13) (lower row). p -value for Lagrange multiplier test (left column), average log markup (middle column), and measure of persistence $g'(0)$ (right column).

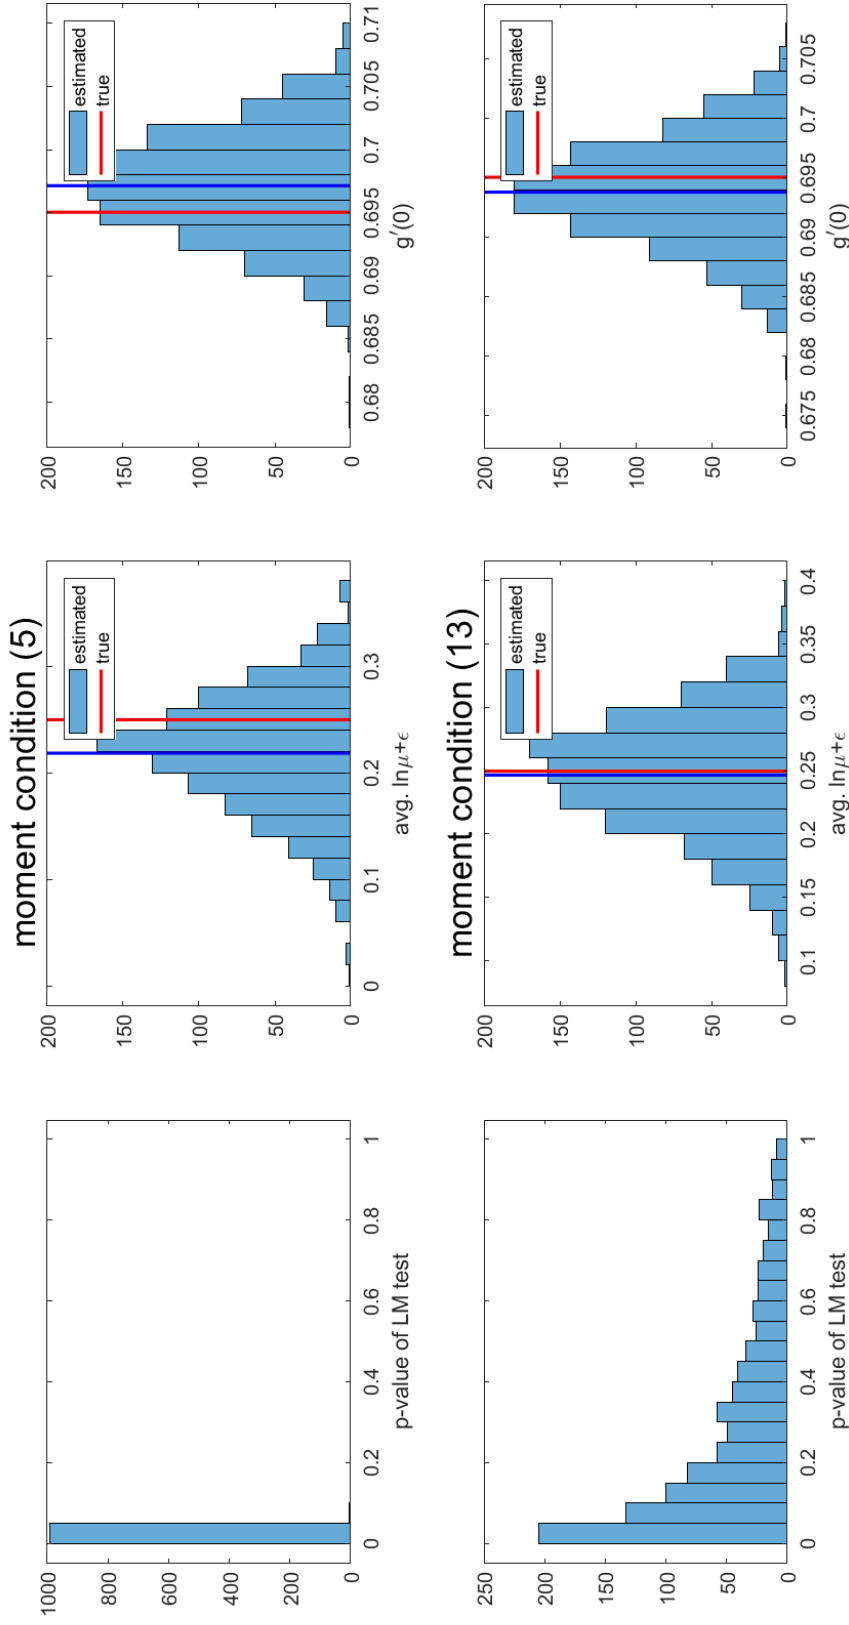


Figure 5: OP/LP/ACF procedure. Case 3 and moment conditions (5) and (13). Modified parameterization with nonlinear process ($\alpha_\omega = 1$). Moment condition (5) (upper row) and moment condition (13) (lower row). p -value for Lagrange multiplier test (left column), average log markup (middle column), and measure of persistence $g'(0)$ (right column).

Section 5, this aligns with the theory on Neyman orthogonality in the double/debiased machine learning literature as our modification ensures that the GMM estimator in step 2 is oracle efficient regardless of the estimation method used in step 1.

7 Sensitivity analysis

While setting $x_{it-1} \supseteq z_{it}$ provides a first-order bias correction, higher-order biases may remain and adversely affect the estimates of $\theta = (\theta_f, \theta_g)$ in step 2 of the OP/LP/ACF procedure. In this section, we develop a diagnostic to assess the sensitivity of the estimates to biases arising from the failure of invertibility. Our diagnostic remains valid without $x_{it-1} \supseteq z_{it}$.

To simplify the analysis, we assume that the nuisance parameter $e(x_{it-1})$ equals its true value $E[q_{it-1}|x_{it-1}]$ and abstract from estimation noise. We consider

$$E[q_{it} - f(k_{it}, v_{it}) - g(q_{it-1}^* - \lambda \zeta_{it-1} - f(k_{it-1}, v_{it-1})) | z_{it}] = 0, \quad (15)$$

where $\lambda \geq 0$ scales the lagged prediction error ζ_{it-1} . If $\lambda = 0$, then moment condition (15) equals the valid but infeasible moment condition (4). If $\lambda = 1$, then moment condition (15) equals the original moment condition (5) because $\zeta_{it-1} = q_{it-1}^* - E[q_{it-1}|x_{it-1}]$ from equation (6). If, moreover, $x_{it-1} \supseteq z_{it}$, then it also equals the modified moment condition (13) with $e(x_{it-1}) = E[q_{it-1}|x_{it-1}]$, as discussed in Section 5.

Define

$$m_{it}(\theta, \lambda) = q_{it} - f(k_{it}, v_{it}; \theta_f) - g(q_{it-1}^* - \lambda \zeta_{it-1} - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g),$$

where we make explicit the parameterization of the production function and law of motion. Moreover, define the pseudo-true value of θ as a function of λ as

$$\theta(\lambda) \in \arg \min_{\theta} E[h(z_{it})m_{it}(\theta, \lambda)]^\top W E[h(z_{it})m_{it}(\theta, \lambda)],$$

where the superscript \top denotes the transpose, $h(z_{it})$ is a vector function of the instruments z_{it} , and W is the GMM weighting matrix. By construction, $\theta(0)$ equals the true parameter value θ^0 and $\theta(1)$ is the probability limit of the GMM estimator.

Although it is not possible to consistently estimate the bias $\theta(1) - \theta(0)$, it is possible to consistently estimate $\frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1}$ and thus to assess the sensitivity at the pseudo-true value $\theta(1)$. A small value of $\frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1}$ provides assurance that small changes in the prediction error that arises in step 1 of the OP/LP/ACF procedure if invertibility fails do not dramatically alter the pseudo-true value $\theta(1)$. Indeed, one can show that $\frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1}$ is zero in the examples in Section 4 where no bias arises despite the failure of invertibility. Conversely, a large value of $\frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1}$ means that the

pseudo-true value $\theta(1)$ is sensitive and thus warns of potentially large bias in the estimates.

To derive our diagnostic $\frac{d\theta(\lambda)}{d\lambda}\big|_{\lambda=1}$, we show in Appendix E that the pseudo-true value $\theta(\lambda)$ under suitable regularity conditions satisfies the first-order condition

$$\mathbb{E} \left[\frac{\partial m_{it}(\theta(\lambda), \lambda)}{\partial \theta} h^\top(z_{it}) \right] \text{WE}[h(z_{it})m_{it}(\theta(\lambda), \lambda)] = 0 \quad (16)$$

where $\frac{\partial m_{it}(\theta(\lambda), \lambda)}{\partial \theta}$ is a $\dim(\theta) \times 1$ vector. The implicit function theorem implies

$$\frac{d\theta(\lambda)}{d\lambda}\bigg|_{\lambda=1} = -\Gamma^{-1}\gamma, \quad (17)$$

where Γ is a $\dim(\theta) \times \dim(\theta)$ matrix and γ a $\dim(\theta) \times 1$ vector. Using $\frac{\partial m_{it}}{\partial \theta}$, $\frac{\partial m_{it}}{\partial \lambda}$, $\frac{\partial^2 m_{it}}{\partial \theta \partial \theta^\top}$, and $\frac{\partial^2 m_{it}}{\partial \theta \partial \lambda}$ to abbreviate $\frac{\partial m_{it}(\theta(\lambda), \lambda)}{\partial \theta}\big|_{\lambda=1}$, $\frac{\partial m_{it}(\theta(\lambda), \lambda)}{\partial \lambda}\big|_{\lambda=1}$, $\frac{\partial^2 m_{it}(\theta(\lambda), \lambda)}{\partial \theta \partial \theta^\top}\big|_{\lambda=1}$, and $\frac{\partial^2 m_{it}(\theta(\lambda), \lambda)}{\partial \theta \partial \lambda}\big|_{\lambda=1}$, respectively, the matrix Γ is defined as

$$\begin{aligned} \Gamma = & \left[\mathbb{E} \left[\frac{\partial^2 m_{it}}{\partial \theta \partial \theta^\top} e_1 h^\top(z_{it}) \right] \text{WE}[h(z_{it})m_{it}(\theta, \lambda)], \dots, \mathbb{E} \left[\frac{\partial^2 m_{it}}{\partial \theta \partial \theta^\top} e_{\dim(\theta)} h^\top(z_{it}) \right] \text{WE}[h(z_{it})m_{it}(\theta, \lambda)] \right] \\ & + \mathbb{E} \left[\frac{\partial m_{it}}{\partial \theta} h^\top(z_{it}) \right] \text{WE} \left[h(z_{it}) \frac{\partial m_{it}}{\partial \theta^\top} \right], \end{aligned} \quad (18)$$

where e_l is a $\dim(\theta) \times 1$ vector with a one in the l th position and zeros elsewhere. The vector γ is defined as

$$\gamma = \mathbb{E} \left[\frac{\partial^2 m_{it}}{\partial \theta \partial \lambda} h^\top(z_{it}) \right] \text{WE}[h(z_{it})m_{it}(\theta, \lambda)] + \mathbb{E} \left[\frac{\partial m_{it}}{\partial \theta} h^\top(z_{it}) \right] \text{WE} \left[h(z_{it}) \frac{\partial m_{it}}{\partial \lambda} \right]. \quad (19)$$

To evaluate our diagnostic $\frac{d\theta(\lambda)}{d\lambda}\big|_{\lambda=1}$, we assume $\mathbb{E}[\varepsilon_{it-1}|x_{it-1}, z_{it}] = 0$. As we detail in Appendix E, consistently estimating Γ is straightforward. Consistently estimating γ is complicated by the fact that

$$\frac{\partial m_{it}}{\partial \lambda} = \frac{\partial g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}} \zeta_{it-1}$$

and

$$\frac{\partial^2 m_{it}}{\partial \theta \partial \lambda} = \begin{pmatrix} \frac{\partial^2 m_{it}}{\partial \theta_f \partial \lambda} \\ \frac{\partial^2 m_{it}}{\partial \theta_g \partial \lambda} \end{pmatrix} = \begin{pmatrix} -\frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}^2} \frac{\partial f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f} \zeta_{it-1} \\ \frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1} \partial \theta_g} \zeta_{it-1} \end{pmatrix}$$

depend on the lagged prediction error ζ_{it-1} . The key insight is that assuming $\mathbb{E}[\varepsilon_{it-1}|x_{it-1}, z_{it}] = 0$ implies

$$\mathbb{E}[\zeta_{it-1}|x_{it-1}, z_{it}] = \mathbb{E}[q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}] | x_{it-1}, z_{it}].$$

Using the law of iterated expectations, we can therefore substitute $q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}]$ for ζ_{it-1}

	α	ρ	ν	μ_ω	ρ_ω	α_ω
<u>case 1:</u>						
bias	0.1306	-0.4932	0.0090	-0.0717	-0.2806	-0.2659
diagnostic	0.1556 (0.0111)	-0.0739 (0.1106)	0.0077 (0.0009)	-0.0013 (0.0088)	-0.2118 (0.0358)	-0.1771 (0.1018)
<u>case 2:</u>						
bias	-0.0237	0.0432	-0.0013	-0.0789	-0.0759	-0.2419
diagnostic	-0.0034 (0.0011)	0.0271 (0.0103)	-0.0012 (0.0002)	-0.0033 (0.0010)	-0.0010 (0.0028)	-0.0541 (0.0110)
<u>case 3:</u>						
bias	-0.0034	-0.0126	-0.0004	-0.0054	-0.0028	-0.0237
diagnostic	-0.0003 (0.0009)	-0.0004 (0.0071)	-0.0002 (0.0001)	-0.0003 (0.0014)	0.0006 (0.0025)	-0.0045 (0.0082)

Table 3: Bias and diagnostic $\frac{d\theta(\lambda)}{d\lambda}\big|_{\lambda=1}$. Average with standard deviation in brackets. Baseline parameterization with nonlinear process ($\alpha_\omega = 1$). Bias and diagnostic use estimate of θ based on moment condition (5).

in the above expressions and rewrite γ as

$$\gamma = \mathbb{E} [\varphi_{it} h^\top(z_{it})] \text{WE}[h(z_{it})m_{it}(\theta, \lambda)] + \mathbb{E} \left[\frac{\partial m_{it}}{\partial \theta} h^\top(z_{it}) \right] \text{WE} [h(z_{it})\psi_{it}], \quad (20)$$

where $\frac{\partial m_{it}}{\partial \lambda}$ and $\frac{\partial^2 m_{it}}{\partial \theta \partial \lambda}$ in (19) are replaced by

$$\psi_{it} = \frac{\partial g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}} (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}])$$

and

$$\varphi_{it} = \left(\begin{array}{c} -\frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}^2} \frac{\partial f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f} (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}]) \\ \frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1} \partial \theta_g} (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}]), \end{array} \right).$$

respectively. We can therefore consistently estimate γ by the finite-sample analog to equation (20) using the estimate of $\mathbb{E}[q_{it-1}|x_{it-1}]$ from step 1 and that of θ from step 2.

Table 3 shows our diagnostic $\frac{d\theta(\lambda)}{d\lambda}\big|_{\lambda=1}$ for the baseline parameterization with the nonlinear process ($\alpha_\omega = 1$). We average the diagnostic over the $S = 1000$ datasets from Section 6. For comparison, we also show the bias $\theta(1) - \theta(0)$.

In case 1, the diagnostic warns of large biases in the production function parameter α and in the law of motion parameters ρ_ω and α_ω . While the diagnostic reliably captures the sign of the bias, it is an order of magnitude smaller than the bias in the production function parameter ρ . In case 2, the bias and the diagnostic are both smaller than in case 1, with the exception of the law of motion parameter α_ω . The diagnostic accurately reflects the remaining bias in the production

function parameter ρ . However, it is an order of magnitude smaller than the bias in α , ρ_ω , and α_ω . In case 3, the small biases go hand-in-hand with small values of the diagnostic. Overall, while the diagnostic does not consistently estimate the bias $\theta(1) - \theta(0)$, it usefully signals the sensitivity of the estimates to a failure of invertibility.

Finally, we emphasize that our sensitivity analysis is conducted around the potentially misspecified model at $\lambda = 1$ as opposed to the true model at $\lambda = 0$ and remains valid regardless of the magnitude of misspecification. Our approach thus differs from the sensitivity analysis in Andrews et al. (2017) that is only valid locally around the true model.

8 Concluding remarks

The OP/LP/ACF procedure for production function estimation hinges on an invertibility assumption. We show that this assumption is testable and strongly rejected in widely used panel data. A failure of invertibility has important consequences: the prediction of planned output q_{it}^* from observables x_{it} in step 1 of the OP/LP/ACF procedure contains an error that invalidates capital k_{it} as an instrument, leading to biased estimates in step 2.

Fortunately, much can still be done. We establish a necessary and sufficient condition for the moment condition in step 2 of the OP/LP/ACF procedure to hold for the true production function. This condition compels a rethinking of the OP/LP/ACF procedure: any instrument used in step 2 must be appropriately included in the regression in step 1 to ensure $x_{it-1} \supseteq z_{it}$. At a minimum, this calls for adding the lead of capital k_{it+1} to x_{it} . Our condition further suggests to flexibly model the law of motion in step 2 and to take a kitchen sink approach to the regression in step 1 by adding as many relevant covariates as possible.

In case our condition is violated, we show that setting $x_{it-1} \supseteq z_{it}$ provides a first-order bias correction. Explicitly incorporating a bias correction achieves Neyman orthogonality in the modified moment condition. Neyman orthogonality ensures that the asymptotic distribution of the GMM estimator in step 2 is invariant to estimation noise from step 1 and is particularly advantageous if the regression in the step 1 includes a large number of covariates or if modern machine learning techniques such as neural networks and random forests are used. Monte Carlo simulations confirm that Neyman orthogonality can substantially enhance the performance of the GMM estimator in step 2.

Finally, we recognize that, despite ensuring $x_{it-1} \supseteq z_{it}$, higher-order biases may remain. To gauge their importance for the GMM estimator in step 2, we introduce a diagnostic that measures the sensitivity of the estimates to the size of the prediction error that arises in step 1 if invertibility fails.

In sum, the invertibility assumption is demanding and can fail because unobserved demand

heterogeneity or in imperfectly competitive environments with partially or fully unobserved rivals or changes in firm conduct and for other reasons. We provide tools for testing the invertibility assumption and propose straightforward modifications of the OP/LP/ACF procedure that either eliminate or mitigate the bias that arises from a failure of invertibility. To address the challenges associated with including a large number of covariates into the regression in step 1, we provide a modified moment condition that achieves Neyman orthogonality and enhances efficiency and robustness. Finally, we hope our diagnostic proves valuable to researchers seeking to assess the potential biases from a failure of invertibility in their estimates.

Appendix

A Tests for invertibility

Tests. Delgado & Manteiga (2001) and Cai, Guo & Zhong (2024) develop nonparametric inference procedures for a mean-independence restriction (Proposition 1), while Shah & Peters (2020) focus on a conditional-independence restriction (Proposition 2). These procedures presume i.i.d. data. Instead of adapting them for panel data, we proceed semiparametrically in what follows.

We test $E[q_{it}|x_{it}, x_{it-1}] = E[q_{it}|x_{it}]$ in Proposition 1 using the partially linear regression model

$$q_{it} = x_{it-1}^\top \beta + \psi(x_{it}) + e_{it}, \quad (21)$$

where the superscript \top denotes the transpose and $\psi(x_{it})$ is a sufficiently flexible function. In practice, we specify $\psi(x_{it}) = r^\top(x_{it})\gamma$, where $r(x_{it})$ is the complete set of polynomials of a given total degree in the variables in x_{it} and γ is a parameter vector. $E[q_{it}|x_{it}, x_{it-1}] = E[q_{it}|x_{it}]$ implies $\beta = 0$. Rejecting $H_0 : \beta = 0$ is therefore evidence against invertibility. To implement this test, we use OLS to estimate the partially linear model (21) and conduct a standard F -test with firm-level clustering.

Because the partially linear model (21) involves a large number of regressors, the standard F -test may perform poorly in finite samples. To address this concern, we use the procedure in Chernozhukov, Hansen & Spindler (2015) (henceforth CHS) instead of OLS. This allows us to use LASSO for dimensionality reduction while achieving oracle efficiency for the estimator of β via Neyman orthogonalization.

We test $q_{it} \perp\!\!\!\perp x_{it-1}|x_{it}$ in Proposition 2 using quantile regression to estimate the partially linear model (21). Let $Q_\tau(q_{it}|x_{it})$ and $Q_\tau(q_{it}|x_{it}, x_{it-1})$ denote the τ th quantile of q_{it} conditional on x_{it} and (x_{it}, x_{it-1}) , respectively. $q_{it} \perp\!\!\!\perp x_{it-1}|x_{it}$ implies $\beta = 0$ and is equivalent to $Q_\tau(q_{it}|x_{it}) = Q_\tau(q_{it}|x_{it}, x_{it-1})$ for all $\tau \in (0, 1)$. We restrict attention to the median ($\tau = 0.5$) and conduct the F -test with firm-level clustering in Parente & Santos Silva (2016).

Data. We implement our tests on two widely used datasets: an unbalanced panel of Spanish manufacturing firms from 1990 to 2006 (Encuesta Sobre Estrategias Empresariales, henceforth ESEE) and the balanced panel of US manufacturing industries from 1958 to 2018 (NBER-CES) from Becker, Gray & Marvakov (2021).

We follow the data cleaning procedures and variable definitions in Doraszelski & Jaumandreu (2018, 2024) for the ESEE data and in Jaumandreu (2022) and Jaumandreu & Lopez (2024) for the NBER-CES data. The available variables include output q_{it} , capital k_{it} , labor l_{it} , materials m_{it} , the price of output p_{it} , the wage w_{it} , and the price of materials p_{it}^M (in logs). In addition, the

industry	#obs	#firms	total degree 2			total degree 3		total degree 4	
			OLS	CHS	Median	OLS	CHS	OLS	CHS
1	2365	296	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2	1270	147	0.004	0.043	0.046	0.008	0.032	0.012	0.033
3	2168	278	0.000	0.000	0.002	0.000	0.000	0.000	0.000
4	1411	165	0.000	0.004	0.005	0.016	0.006	0.035	0.004
5	1505	197	0.000	0.004	0.012	0.000	0.000	0.000	0.000
6	1206	148	0.000	0.000	0.018	0.001	0.000	0.001	0.000
7	2455	304	0.000	0.000	0.000	0.000	0.000	0.000	0.000
8	2368	317	0.000	0.000	0.000	0.000	0.000	0.000	0.000
9	1445	191	0.001	0.000	0.001	0.017	0.000	0.009	0.000
10	1414	174	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 4: p -value for test of $H_0 : \beta = 0$ with $x_{it} = (k_{it}, l_{it}, m_{it}, p_{it}, w_{it}, p_{it}^M)$. ESEE data.

industry	#obs	#firms	total degree 2			total degree 3		total degree 4	
			OLS	CHS	Median	OLS	CHS	OLS	CHS
1	2365	296	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2	1270	147	0.000	0.001	0.110	0.001	0.000	0.070	0.001
3	2168	278	0.000	0.000	0.000	0.000	0.000	0.000	0.000
4	1411	165	0.001	0.000	0.107	0.062	0.000	0.208	0.000
5	1505	197	0.000	0.000	0.004	0.000	0.000	0.000	0.000
6	1206	148	0.000	0.000	0.017	0.000	0.000	0.015	0.000
7	2455	304	0.000	0.000	0.000	0.000	0.000	0.000	0.000
8	2368	317	0.000	0.000	0.000	0.000	0.000	0.000	0.000
9	1445	191	0.000	0.000	0.009	0.000	0.000	0.007	0.000
10	1414	174	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 5: p -value for test of $H_0 : \beta = 0$ with $x_{it} = (k_{it}, l_{it}, m_{it}, p_{it}, w_{it}, p_{it}^M, mdy_{it})$. ESEE data.

ESEE data has a market dynamism variable mdy_{it} that indicates shifts in the demand the firm faces (slump, stability, or expansion). Investment is non-zero in the NBER-CES data, so that the investment variable inv_{it} (in logs) is well-defined.

Results. Tables 4 and 5 present results for the 10 industries in the ESEE data. We show the p -value for the test of $H_0 : \beta = 0$ for different estimators (OLS, CHS, and median regression) and for different specifications of $\psi(x_{it})$ (total degree 2, 3, and 4).²¹ In Table 4, we use $x_{it} = (k_{it}, l_{it}, m_{it}, p_{it}, w_{it}, p_{it}^M)$. We reject invertibility by a wide margin in all industries and for all combinations of estimators and specifications of $\psi(x_{it})$. In Table 5, we add the market dynamism

²¹Results for median regression are limited to total degree 2 because the Stata package `qreg2` fails to converge for total degree 3 and 4.

variable mdy_{it} to x_{it} . While we continue to reject invertibility in 8 industries, we cannot reject invertibility at a 5% significance level in industries 2 and 4 for some combinations of estimators and specifications of $\psi(x_{it})$. This reinforces that the ability to control for demand is a precondition for invertibility.

We note that all estimated models in Tables 4 and 5 have an R^2 of at least 0.990 and some have an R^2 of 0.998. Insufficient flexibility in modeling $\psi(x_{it})$ is thus unlikely to compromise our results.

Turning to the NBER-CES data, we alternatively use $x_{it} = (k_{it}, l_{it}, m_{it}, p_{it}, w_{it}, p_{it}^M)$ and $x_{it} = (k_{it}, l_{it}, m_{it}, p_{it}, w_{it}, p_{it}^M, inv_{it})$. The latter follows OP to invert the firm's demand for investment. Throughout the p -value for the test of $H_0 : \beta = 0$ is 0.000 and the R^2 exceeds 0.987.

B Failure of invertibility and validity of OP/LP/ACF moment condition

Omitted proofs.

Proof of Theorem 2. Using $\zeta_{it-1} = q_{it-1}^* - E[q_{it-1}|x_{it-1}]$, we have

$$\begin{aligned}
& E [g' (q_{it-1}^* - \zeta_{it-1} - f(k_{it-1}, v_{it-1})) \zeta_{it-1} | z_{it}] \\
&= E [g' (E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \zeta_{it-1} | z_{it}] \\
&= E [E[g' (E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \zeta_{it-1} | x_{it-1}] | z_{it}] \\
&= E [g' (E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) E[\zeta_{it-1}|x_{it-1}] | z_{it}] \\
&= 0,
\end{aligned} \tag{22}$$

where the second equality uses $x_{it-1} \supseteq z_{it}$ and the law of iterated expectations, the third equality uses that $g' (E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}))$ is a function of x_{it-1} because $x_{it-1} = (k_{it-1}, v_{it-1}, \dots)$, and the last equality uses $E[\zeta_{it-1}|x_{it-1}] = 0$. Finally, using equation (22) and $\zeta_{it-1} = q_{it-1}^* - E[q_{it-1}|x_{it-1}]$ establishes equation (10). \square

Proof of Theorem 3. Because $g^0 \in \mathcal{G}$, we have

$$\begin{aligned}
& \inf_{\tilde{g} \in \mathcal{G}} \|E [g^0(\omega_{it-1})|z_{it}] - E [\tilde{g} (E [\omega_{it-1}|x_{it-1}]) | z_{it}]\|_{L,1} \\
&= \inf_{\tilde{g} \in \mathcal{G}} E \left[\left| E [g^0(\omega_{it-1})|z_{it}] - E [\tilde{g} (E [\omega_{it-1}|x_{it-1}]) | z_{it}] \right| \right] \\
&\leq E \left[\left| E [g^0(\omega_{it-1})|z_{it}] - E [g^0(E[\omega_{it-1}|x_{it-1}])|z_{it}] \right| \right].
\end{aligned}$$

Using a second-order Taylor expansion of g^0 around $E[\omega_{it-1}|x_{it-1}]$, we have

$$\begin{aligned}
& \left| E[g^0(\omega_{it-1})|z_{it}] - E[g^0(E[\omega_{it-1}|x_{it-1}])|z_{it}] \right| \\
&= \left| E \left[g^{0'}(E[\omega_{it-1}|x_{it-1}])\zeta_{it-1} \middle| z_{it} \right] + E \left[g^{0''}(\tilde{\omega}_{it-1})\zeta_{it-1}^2 \middle| z_{it} \right] \right| \\
&\leq \left| E \left[g^{0'}(E[\omega_{it-1}|x_{it-1}])\zeta_{it-1} \middle| z_{it} \right] \right| + \left| E \left[g^{0''}(\tilde{\omega}_{it-1})\zeta_{it-1}^2 \middle| z_{it} \right] \right| \\
&\leq \left| E \left[g^{0'}(E[\omega_{it-1}|x_{it-1}])\zeta_{it-1} \middle| z_{it} \right] \right| + \tau E \left[\zeta_{it-1}^2 \middle| z_{it} \right],
\end{aligned}$$

where $\tilde{\omega}_{it-1}$ is some value between ω_{it-1} and $E[\omega_{it-1}|x_{it-1}]$, and the last inequality follows from the definition of τ . As in the proof of Theorem 2, using $x_{it-1} \supseteq z_{it}$ and the law of iterated expectations yields $E \left[g^{0'}(E[\omega_{it-1}|x_{it-1}])\zeta_{it-1} \middle| z_{it} \right] = 0$. Thus, we have

$$\inf_{\tilde{g} \in \mathcal{G}} \left\| E \left[g^0(\omega_{it-1})|z_{it} \right] - E \left[\tilde{g}(E[\omega_{it-1}|x_{it-1}])|z_{it} \right] \right\|_{L,1} \leq E \left[\tau E \left[\zeta_{it-1}^2 | z_{it} \right] \right] = \tau \text{Var}(\zeta_{it-1}).$$

□

C Monte Carlo exercise

Data generating process. Marginal revenue is

$$mr_{it} = -\ln(1 + \exp(\delta_{2i})) + p_{it}$$

and marginal cost is

$$mc_{it} = p_i^V + v_{it} - f(k_{it}, v_{it}) - \ln \frac{\partial f(k_{it}, v_{it})}{\partial v_{it}} - \omega_{it}.$$

We rewrite $mr_{it} = mc_{it}$ as

$$\frac{\delta_{1i} + \exp(-\delta_{2i})(f(k_{it}, v_{it}) + \omega_{it})}{1 + \exp(-\delta_{2i})} - \ln(1 + \exp(\delta_{2i})) + \ln \frac{\partial f(k_{it}, v_{it})}{\partial v_{it}} - p_i^V - v_{it} = 0$$

and solve numerically for v_{it} . With v_{it} in hand, we determine $q_{it}^* = f(k_{it}, v_{it}) + \omega_{it}$, $q_{it} = q_{it}^* + \varepsilon_{it}$, and $p_{it} = \frac{\delta_{1i} - q_{it}^*}{1 + \exp(-\delta_{2i})}$.

As capital fully depreciates between periods and the firm has static expectations, it chooses k_{it+1} in period t given ω_{it} , δ_i , p_i^K , and p_i^V whilst anticipating that it optimally chooses \tilde{v}_{it+1} in

period $t + 1$. To maximize its expected profit in period $t + 1$, in period t the firm solves

$$\begin{aligned} & \max_{k_{it+1}} \left(\max_{\tilde{v}_{it+1}} \exp(\tilde{p}_{it+1} + \tilde{q}_{it+1}^*) - \exp(p_i^K + k_{it+1}) - \exp(p_i^V + \tilde{v}_{it+1}) \right) \\ & = \max_{k_{it+1}, \tilde{v}_{it+1}} \exp(\tilde{p}_{it+1} + \tilde{q}_{it+1}^*) - \exp(p_i^K + k_{it+1}) - \exp(p_i^V + \tilde{v}_{it+1}) \end{aligned}$$

subject to

$$\begin{aligned} \tilde{q}_{it+1}^* &= f(k_{it+1}, \tilde{v}_{it+1}) + \omega_{it}, \\ \tilde{q}_{it+1}^* &= \delta_{1i} - (1 + \exp(-\delta_{2i}))\tilde{p}_{it+1}. \end{aligned}$$

The solution is

$$\begin{aligned} k_{it+1} &= \frac{1}{1-\rho} (\ln \alpha - p_i^K) + \frac{\exp(\delta_{2i}) + 1}{\exp(\delta_{2i}) + 1 - \nu} \left(\ln \nu \right. \\ &+ \left(\frac{\nu}{(\exp(\delta_{2i}) + 1)\rho} - 1 \right) \ln \left(\alpha \left(\frac{\alpha}{\exp(p_i^K)} \right)^{\frac{\rho}{1-\rho}} + (1-\alpha) \left(\frac{1-\alpha}{\exp(p_i^V)} \right)^{\frac{\rho}{1-\rho}} \right) \\ &\left. - \ln(\exp(\delta_{2i}) + 1) + \frac{1}{\exp(\delta_{2i}) + 1} \omega_{it} + \frac{\delta_{1i}}{1 + \exp(-\delta_{2i})} \right). \end{aligned}$$

Neural network estimator. The neural network uses two hidden layers of 128 neurons and ReLU activation. We standardize the input data. We train the network using the stochastic gradient decent (SDG) method on 80% of firms, setting aside 20% of firms for validation, to minimize the mean squared error. We stop the SDG iterations if the mean squared error for the validation sample does not improve after 10 epoch iterations. We use the `trainnet` command in Matlab and set the learning rate to 0.01 and the mini-batch size to 500 observations.

GMM estimator. Corresponding to moment condition (5) in step 2, define the moment function

$$m_{it}(\theta) = q_{it} - f(k_{it}, v_{it}; \theta_f) - g(e(x_{it-1}) - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g), \quad (23)$$

where we replace $e(x_{it-1})$ by the estimate from step 1. In step 2, we solve the GMM problem

$$\min_{\theta} \left(\frac{1}{NT} \sum_{i,t} h(z_{it}) m_{it}(\theta) \right)^{\top} W \left(\frac{1}{NT} \sum_{i,t} h(z_{it}) m_{it}(\theta) \right), \quad (24)$$

where the superscript \top denotes the transpose, $h(z_{it})$ is the complete set of Hermite polynomials of total degree 4 in the variables in z_{it} , and

$$W = \left(\frac{1}{NT-1} \sum_{i,t} (h(z_{it})m_{it}(\theta^0) - \hat{\mu})^\top (h(z_{it})m_{it}(\theta^0) - \hat{\mu}) \right)^{-1} \text{ with } \hat{\mu} = \frac{1}{NT} \sum_{i,t} h(z_{it})m_{it}(\theta^0),$$

is a weighting matrix evaluated at the true parameters values θ^0 .

Corresponding to moment condition (13), our modification of the OP/LP/ACF procedure redefines the moment function as

$$m_{it}(\theta) = q_{it} - f(k_{it}, v_{it}; \theta_f) - g(e(x_{it-1}) - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g) - g'(e(x_{it-1}) - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)(q_{it-1} - e(x_{it-1})), \quad (25)$$

where we replace $e(x_{it-1})$ by the estimate from step 1. We hold fixed W to facilitate the comparison between our modification and the original procedure.

Lagrange multiplier (LM) test. We first derive the LM test for the original moment function (23). To account for the plug-in nature of the GMM estimator in step 2, we have to explicitly specify how $e(x_{it-1})$ is modelled and estimated in step 1. In our Monte Carlo exercise, we model $e(x_{it-1}) = r^\top(x_{it-1})\tau$, where the superscript \top denotes the transpose, $r(x_{it-1})$ is the complete set of Hermite polynomials of total degree 4 in the variables in x_{it-1} , and τ is a parameter vector. Consequently, we define

$$m_{it}(\theta, \tau) = q_{it} - f(k_{it}, v_{it}; \theta_f) - g(r^\top(x_{it-1})\tau - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g), \quad (26)$$

where we make explicit the dependence of the moment function on τ . We estimate τ in step 1 by OLS as $\hat{\tau} = [\frac{1}{NT} \sum_{i,t} r(x_{it-1})r^\top(x_{it-1})]^{-1} \frac{1}{NT} \sum_{i,t} q_{it-1}r(x_{it-1})$.²²

Following standard asymptotic arguments, we can show that

$$\hat{\Sigma}^{-\frac{1}{2}} \frac{1}{\sqrt{NT}} \sum_{i,t} h(z_{it})m_{it}(\theta^0, \hat{\tau}) \xrightarrow{d} N(0, I),$$

where $\hat{\Sigma} = \hat{\Lambda}\hat{\Omega}\hat{\Lambda}^\top$,

$$\hat{\Lambda} = \left(I \quad \sum_{i,t} h(z_{it}) \left[\frac{\partial m_{it}(\theta^0, \hat{\tau})}{\partial \tau^\top} \right] \left[\sum_{i,t} r(x_{it-1})r^\top(x_{it-1}) \right]^{-1} \right),$$

²²We can alternatively estimate τ as $\hat{\tau} = [\frac{1}{NT} \sum_{i,t} r(x_{it})r^\top(x_{it})]^{-1} \frac{1}{NT} \sum_{i,t} q_{it}r(x_{it})$. In our simulated datasets, both current and lagged values are available for all T periods and N firms.

with I being the $\dim(h(z_{it})) \times \dim(h(z_{it}))$ identity matrix, and

$$\begin{aligned}\widehat{\Omega} &= \begin{pmatrix} \widehat{\Omega}_{11} & \widehat{\Omega}_{12} \\ \widehat{\Omega}_{12}^\top & \widehat{\Omega}_{22} \end{pmatrix}, \\ \widehat{\Omega}_{11} &= \frac{1}{NT} \sum_{t=1}^T \sum_{t'=1}^T \sum_{i=1}^N [m_{it}(\theta^0, \hat{\tau}) m_{it'}(\theta^0, \hat{\tau}) h(z_{it}) h^\top(z_{it'})], \\ \widehat{\Omega}_{12} &= \frac{1}{NT} \sum_{t=1}^T \sum_{t'=1}^T \sum_{i=1}^N [m_{it}(\theta^0, \hat{\tau}) (q_{it'-1} - r^\top(x_{it'-1}) \hat{\tau}) h(z_{it}) r^\top(x_{it'-1})], \\ \widehat{\Omega}_{22} &= \frac{1}{NT} \sum_{t=1}^T \sum_{t'=1}^T \sum_{i=1}^N [(q_{it-1} - r^\top(x_{it-1}) \hat{\tau}) (q_{it'-1} - r^\top(x_{it'-1}) \hat{\tau}) r(x_{it-1}) r^\top(x_{it'-1})].\end{aligned}$$

Note that $\widehat{\Sigma}$ incorporates firm-level clustering and thus allows for arbitrary correlation within firm across time in the data. As $N \rightarrow \infty$, the LM statistic

$$S_N = \frac{1}{NT} \left(\sum_{i,t} h(z_{it}) m_{it}(\theta^0, \hat{\tau}) \right)^\top \widehat{\Sigma}^{-1} \left(\sum_{i,t} h(z_{it}) m_{it}(\theta^0, \hat{\tau}) \right)$$

converges to a χ^2 distribution with $\dim(h(z_{it}))$ degrees of freedom. Let $c_{1-\alpha}$ be the $1-\alpha$ quantile of this distribution. Then rejecting $H_0 : E[h(z_{it}) m_{it}(\theta^0, \tau^0)] = 0$ if $S_N > c_{1-\alpha}$ is a valid LM inference procedure with asymptotic size α .

Turning to the modified moment function (25), Neyman orthogonality renders correcting for the plug-in nature of the GMM estimator in step 2 superfluous. Consequently, we use a standard LM test with firm-level clustering. As $N \rightarrow \infty$, the LM statistic

$$S_N = \frac{1}{NT} \left(\sum_{i,t} h(z_{it}) m_{it}(\theta^0, \hat{\tau}) \right)^\top \widehat{\Omega}_{11}^{-1} \left(\sum_{i,t} h(z_{it}) m_{it}(\theta^0, \hat{\tau}) \right)$$

converges to a χ^2 distribution with $\dim(h(z_{it}))$ degrees of freedom. Rejecting $H_0 : E[h(z_{it}) m_{it}(\theta^0, \tau^0)] = 0$ if $S_N > c_{1-\alpha}$ is a valid LM inference procedure with asymptotic size α .

D Modification of OP/LP/ACF moment condition

Omitted proofs.

Proof of Theorem 4. By the definition of the derivative, it suffices to show that for any non-zero

sequence λ_n with $\lim_{n \rightarrow \infty} \lambda_n = 0$, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{\lambda_n} & \left[\left(\mathbb{E} [q_{it} - f(k_{it}, v_{it}) - g(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) \right. \right. \\ & \left. \left. - g'(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}] - \lambda_n \delta(x_{it-1})) | z_{it} \right] \right. \\ & \left. - \left(\mathbb{E} [q_{it} - f(k_{it}, v_{it}) - g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \right. \right. \\ & \left. \left. - g'(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}]) | z_{it} \right] \right) \Big] = 0. \quad (27) \end{aligned}$$

Similar to the idea used in the proof of Theorem 2, the law of iterated expectations implies

$$\mathbb{E} [g'(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}]) | z_{it}] = 0, \quad (28)$$

$$\mathbb{E} [g'(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) (q_{it-1} - \mathbb{E}[q_{it-1}|x_{it-1}]) | z_{it}] = 0. \quad (29)$$

Using equations (28) and (29), equation (27) is equivalent to

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{\lambda_n} \mathbb{E} \left[g(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) - g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \mid z_{it} \right] \\ & = \lim_{n \rightarrow \infty} \mathbb{E} \left[g'(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) \mid z_{it} \right]. \quad (30) \end{aligned}$$

It therefore remains to establish equation (30). We do this by showing that both sides of equation (30) are equal to $\mathbb{E} \left[g'(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) \mid z_{it} \right]$.

Because g' is bounded and δ is integrable, there exists some $\eta > 0$ such that

$$\mathbb{E} \left[\sup_{\lambda: |\lambda| < \eta} |g'(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) \delta(x_{it-1})| \mid z_{it} \right] < \infty \quad (31)$$

almost surely.²³ Because g' is continuous, the dominated convergence theorem therefore implies that

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{E} \left[g'(\mathbb{E}[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) \mid z_{it} \right] \\ & = \mathbb{E} \left[g'(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) \mid z_{it} \right]. \quad (32) \end{aligned}$$

²³Theorem 4 continues to hold if the assumption that g' is bounded is relaxed to this condition.

Turning to the other side of equation (30), the mean value theorem implies that

$$\begin{aligned} & \frac{1}{\lambda_n} (g(E[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) - g(E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}))) \\ &= g' \left(E[q_{it-1}|x_{it-1}] + \tilde{\lambda}_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1}) \right) \delta(x_{it-1}), \end{aligned} \quad (33)$$

where $\tilde{\lambda}_n$ is some value between 0 and λ_n that implicitly depends on x_{it-1} . Because $|\tilde{\lambda}_n| \leq |\lambda_n|$ and g' is continuous, we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} g' \left(E[q_{it-1}|x_{it-1}] + \tilde{\lambda}_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1}) \right) \delta(x_{it-1}) \\ &= g' (E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) \end{aligned} \quad (34)$$

for all x_{it-1} . Given inequality (31) and equation (34), the dominated convergence theorem implies that

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{\lambda_n} E \left[g(E[q_{it-1}|x_{it-1}] + \lambda_n \delta(x_{it-1}) - f(k_{it-1}, v_{it-1})) - g(E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \middle| z_{it} \right] \\ &= E \left[g' (E[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1})) \delta(x_{it-1}) \middle| z_{it} \right]. \end{aligned}$$

This establishes equation (30). \square

E Sensitivity analysis

Assume $f(k_{it}, v_{it}; \theta_f)$ is twice continuously differentiable in θ_f and $g(\omega_{it-1}; \theta_g)$ is twice continuously differentiable in $(\theta_g, \omega_{it-1})$, so that $m_{it}(\theta, \lambda)$ is twice continuously differentiable in (θ, λ) . Assume the product of h and the first-order derivatives of $m_{it}(\theta, \lambda)$ is always integrable. Then the dominated convergence theorem implies

$$\frac{\partial E[h(z_{it})m_{it}(\theta, \lambda)]}{\partial \theta^\top} = E \left[h(z_{it}) \frac{\partial m_{it}(\theta, \lambda)}{\partial \theta^\top} \right], \quad \frac{\partial E[h(z_{it})m_{it}(\theta, \lambda)]}{\partial \lambda} = E \left[h(z_{it}) \frac{\partial m_{it}(\theta, \lambda)}{\partial \lambda} \right].$$

Similarly, assume the product of h and the second-order derivatives of $m_{it}(\theta, \lambda)$ is always integrable. Then the dominated convergence theorem implies

$$\frac{\partial^2 E[h(z_{it})m_{it}(\theta, \lambda)]}{\partial \theta \partial \theta^\top} = E \left[h(z_{it}) \frac{\partial^2 m_{it}(\theta, \lambda)}{\partial \theta \partial \theta^\top} \right], \quad \frac{\partial^2 E[h(z_{it})m_{it}(\theta, \lambda)]}{\partial \theta \partial \lambda} = E \left[h(z_{it}) \frac{\partial^2 m_{it}(\theta, \lambda)}{\partial \theta \partial \lambda} \right].$$

Under these regularity conditions, the pseudo-true value $\theta(\lambda)$ satisfies the first-order condition (16).

Assume the first-order condition (16) identifies $\theta(\lambda)$ for λ in a neighborhood of 1. Then the

implicit function theorem implies

$$\begin{aligned} \mathbb{E} \left[\left(\frac{\partial^2 m_{it}}{\partial \theta \partial \theta^\top} \frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1} + \frac{\partial^2 m_{it}}{\partial \theta \partial \lambda} \right) h^\top(z_{it}) \right] \text{WE} [h(z_{it}) m_{it}(\theta(\lambda), \lambda)] \\ + \mathbb{E} \left[\frac{\partial m_{it}}{\partial \theta} h^\top(z_{it}) \right] \text{WE} \left[h(z_{it}) \left(\frac{\partial m_{it}}{\partial \theta^\top} \frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1} + \frac{\partial m_{it}}{\partial \lambda} \right) \right] = 0, \quad (35) \end{aligned}$$

where we use $\frac{\partial m_{it}}{\partial \theta}$, $\frac{\partial m_{it}}{\partial \lambda}$, $\frac{\partial^2 m_{it}}{\partial \theta \partial \theta^\top}$, and $\frac{\partial^2 m_{it}}{\partial \theta \partial \lambda}$ to abbreviate $\frac{\partial m_{it}(\theta(\lambda), \lambda)}{\partial \theta} \Big|_{\lambda=1}$, $\frac{\partial m_{it}(\theta(\lambda), \lambda)}{\partial \lambda} \Big|_{\lambda=1}$, $\frac{\partial^2 m_{it}(\theta(\lambda), \lambda)}{\partial \theta \partial \theta^\top} \Big|_{\lambda=1}$, and $\frac{\partial^2 m_{it}(\theta(\lambda), \lambda)}{\partial \theta \partial \lambda} \Big|_{\lambda=1}$, respectively. Equation (35) can be rewritten as

$$\Gamma \frac{d\theta(\lambda)}{d\lambda} \Big|_{\lambda=1} + \gamma = 0,$$

where Γ and γ are defined in equations (18) and (19), respectively. Assume Γ is invertible. Then equation (17) follows.

Consistently estimating Γ is straightforward because

$$\frac{\partial m_{it}}{\partial \theta} = \begin{pmatrix} \frac{\partial m_{it}}{\partial \theta_f} \\ \frac{\partial m_{it}}{\partial \theta_g} \end{pmatrix} = \begin{pmatrix} -\frac{\partial f(k_{it}, v_{it}; \theta_f)}{\partial \theta_f} + \frac{\partial g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}} \frac{\partial f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f} \\ -\frac{\partial g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \theta_g} \end{pmatrix}$$

and

$$\frac{\partial^2 m_{it}}{\partial \theta \partial \theta^\top} = \begin{pmatrix} \frac{\partial^2 m_{it}}{\partial \theta_f \partial \theta_f^\top} & \frac{\partial^2 m_{it}}{\partial \theta_f \partial \theta_g^\top} \\ \left(\frac{\partial^2 m_{it}}{\partial \theta_f \partial \theta_g^\top} \right)^\top & \frac{\partial^2 m_{it}}{\partial \theta_g \partial \theta_g^\top} \end{pmatrix},$$

where

$$\begin{aligned} \frac{\partial^2 m_{it}}{\partial \theta_f \partial \theta_f^\top} &= -\frac{\partial^2 f(k_{it}, v_{it}; \theta_f)}{\partial \theta_f \partial \theta_f^\top} + \frac{\partial g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}} \frac{\partial^2 f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f \partial \theta_f^\top} \\ &\quad - \frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1}^2} \left(\frac{\partial f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f} \right) \left(\frac{\partial f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f} \right)^\top, \\ \frac{\partial^2 m_{it}}{\partial \theta_f \partial \theta_g^\top} &= \frac{\partial f(k_{it-1}, v_{it-1}; \theta_f)}{\partial \theta_f} \frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \omega_{it-1} \partial \theta_g^\top}, \\ \frac{\partial^2 m_{it}}{\partial \theta_g \partial \theta_g^\top} &= \frac{\partial^2 g(\mathbb{E}[q_{it-1}|x_{it-1}] - f(k_{it-1}, v_{it-1}; \theta_f); \theta_g)}{\partial \theta_g \partial \theta_g^\top}, \end{aligned}$$

do not depend on the lagged prediction error ζ_{it-1} . Following standard asymptotic arguments, Γ can be consistently estimated by the finite-sample analog to equation (18) using the estimate of $\mathbb{E}[q_{it-1}|x_{it-1}]$ from step 1 and the estimate of θ from step 2.

References

- Akerberg, D., Caves, K. & Frazer, G. (2015), ‘Identification properties of recent production function estimators’, *Econometrica* **83**(6), 2411–2451.
- Akerberg, D. & De Loecker, J. (2024), Production function identification under imperfect competition, Working paper, University of Texas, Austin.
- Andrews, I., Gentzkow, M. & Shapiro, J. (2017), ‘Measuring the sensitivity of parameter estimates to estimation moments’, *Quarterly Journal of Economics* **132**(4), 1553–1592.
- Aw, B., Roberts, M. & Xu, D. (2011), ‘R&D investment, exporting, and productivity dynamics’, *American Economic Review* **101**(4), 1312–1344.
- Becker, R., Gray, W. & Marvakov, J. (2021), ‘NBER-CES manufacturing productivity database (1958-2018, version 2021a)’, <https://www.nber.org/research/data/nber-ces-manufacturing-industry-database>.
- Berry, S., Levinsohn, J. & Pakes, A. (1995), ‘Automobile prices in market equilibrium’, *Econometrica* **63**(4), 841–890.
- Biondi, F. (2022), Firm productivity and derived factor demand under variable markups, Working paper, Katholieke Universiteit Leuven, Leuven.
- Blundell, R. & Bond, S. (2000), ‘GMM estimation with persistent panel data: An application to production functions’, *Econometric Reviews* **19**, 321–340.
- Brand, J. (2020), Estimating productivity and markups under imperfect competition, Working paper, University of Texas, Austin.
- Cai, L., Guo, X. & Zhong, W. (2024), ‘Test and measure for partial mean dependence based on machine learning methods’, *Journal of the American Statistical Association* **120**(550), 833–845.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. & Robins, J. (2018), ‘Double/debiased machine learning for treatment and structural parameters’, *Journal of Econometrics* **21**(1), C1–C68.
- Chernozhukov, V., Escanciano, J., Ichimura, H., Newey, W. & Robins, J. (2022), ‘Locally robust semiparametric estimation’, *Econometrica* **90**(4), 1501–1535.

- Chernozhukov, V., Hansen, C. & Spindler, M. (2015), ‘Valid post-selection and post-regularization inference: An elementary, general approach’, *Annual Review of Economics* **7**, 649–688.
- Chirinko, R. (2008), ‘ σ : The long and short of it’, *Journal of Multivariate Analysis* **30**, 671–686.
- Cobb, C. & Douglas, P. (1928), ‘A theory of production’, *American Economic Review* **18**(1), 139–165.
- De Loecker, J. (2013), ‘Detecting learning by exporting’, *American Economic Journal: Microeconomics* **5**(3), 1–21.
- De Loecker, J., Goldberg, P., Khandelwal, A. & Pavcnik, N. (2016), ‘Prices, markups and trade reform’, *Econometrica* **84**(2), 445–510.
- De Loecker, J. & Syverson, C. (2021), An industrial organization perspective on productivity, in K. Ho, A. Hortacsu & A. Lizzeri, eds, ‘Handbook of Industrial Organization’, Vol. 4, North-Holland, Amsterdam, pp. 141–223.
- De Loecker, J. & Warzynski, F. (2012), ‘Markups and firm-level export status’, *American Economic Review* **102**(6), 2437–2471.
- Delgado, M. & Manteiga, W. (2001), ‘Significance testing in nonparametric regression based on the bootstrap’, *Annals of Statistics* **29**(5), 1469–1507.
- Doraszelski, U. & Jaumandreu, J. (2013), ‘R&D and productivity: Estimating endogenous productivity’, *Review of Economic Studies* **80**(4), 1338–1383.
- Doraszelski, U. & Jaumandreu, J. (2018), ‘Measuring the bias of technological change’, *Journal of Political Economy* **126**(3), 1027–1084.
- Doraszelski, U. & Jaumandreu, J. (2024), Reexamining the De Loecker & Warzynski (2012) method for estimating markups, Working paper, University of Pennsylvania, Philadelphia.
- Foster, L., Haltiwanger, J. & Syverson, C. (2008), ‘Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?’, *American Economic Review* **98**(1), 394–425.
- Gandhi, A., Navarro, S. & Rivers, D. (2020), ‘On the identification of gross output production functions’, *Journal of Political Economy* **128**(8), 2973–3016.
- Hahn, J., Liao, Z. & Ridder, G. (2018), ‘Nonparametric two-step sieve M estimation and inference’, *Econometric Theory* **34**(6), 1281–1324.

- Hall, R. (1988), ‘The relation between price and marginal cost in U.S. industry’, *Journal of Political Economy* **96**(5), 921–947.
- Hu, Y., Huang, G. & Sasaki, Y. (2020), ‘Estimating production functions with robustness against errors in the proxy variables’, *Journal of Econometrics* **215**(2), 375–398.
- Hu, Y. & Schennach, S. (2008), ‘Instrumental variable treatment of nonclassical measurement error models’, *Econometrica* **76**(1), 195–216.
- Jaumandreu, J. (2022), The remarkable stability of us manufacturing markups, Working paper no. 17490, CEPR, London.
- Jaumandreu, J. & Lopez, R. (2024), ‘Markups in US food manufacturing accounting for non-neutral productivity’, *Journal of Applied Economics* **75**(2), 573–591.
- Levinsohn, J. & Petrin, A. (2003), ‘Estimating production functions using inputs to control for unobservables’, *Review of Economic Studies* **70**(2), 317–341.
- Marschak, J. & Andrews, W. (1944), ‘Random simultaneous equations and the theory of production’, *Econometrica* **12**(3), 143–205.
- Mundlak, Y. & Hoch, I. (1965), ‘Consequences of alternative specifications in estimation of Cobb-Douglas production functions’, *Econometrica* **33**(4), 814–828.
- Neyman, A. (1959), Optimal asymptotic tests of composite statistical hypotheses, in V. Grenander, ed., ‘Probability and Statistics’, John Wiley & Sons, New York, pp. 213–234.
- Olley, S. & Pakes, A. (1996), ‘The dynamics of productivity in the telecommunications industry’, *Econometrica* **64**(6), 1263–1297.
- Pakes, A. (1994), Dynamic structural models, problems, and prospects: mixed continuous discrete controls and market interactions, in C. Sims, ed., ‘Advances in econometrics: Sixth World Congress’, Vol. 2, Cambridge University Press, Cambridge.
- Parente, P. & Santos Silva, J. (2016), ‘Quantile regression with clustered data’, *Journal of Econometric Methods* **5**(1), 1–15.
- Ponder, M. (2021), Production functions without control functions: Estimating input elasticities when monotonicity does not hold, Working paper, University of Minnesota, Minneapolis.
- Shah, R. & Peters, J. (2020), ‘The hardness of conditional independence testing and the generalised covariance measure’, *Annals of Statistics* **48**(3), 1514–1538.

- Solow, R. (1957), ‘Technical change and the aggregate production function’, *Review of Economics and Statistics* **39**(3), 312–320.
- Syverson, C. (2001), Market structure and productivity, Ph.d. dissertation, University of Maryland, College Park.
- Van Biesebroeck, J. (2007), ‘Robustness of productivity estimates’, *Journal of Industrial Economics* **55**(3), 529–569.
- von Thünen, J. (1842), *Der isolirte Staat in Beziehung auf Landwirthschaft und Nationalökonomie*, Leopold, Rostock.
- Wicksteed, P. (1894), *An essay on the coordination of the laws of distribution*, Macmillan, London.