Estimating Factor-Biased Productivity Spillovers*

ZHEZHI HOU[1] EMIR MALIKOV[2] SHUNAN ZHAO[3]

[1]Institute of Western China Economic Research, Southwestern University of Finance and Economics, Chengdu, China
[2]Lee Business School, University of Nevada, Las Vegas, Las Vegas, USA
[3]Department of Economics, Oakland University, Rochester, USA

August 31, 2026

Abstract

Structural estimators of firm productivity predominantly assume that productivity is scalar, Hicks-neutral and independent across firms, ruling out factor-biased productivity spillovers by construction. We develop a structural framework that relaxes these restrictions jointly, treating productivity as two-dimensional, comprising a factor-neutral and a labor-augmenting component, and letting each component propagate across firms through spatiotemporal spillovers along a network of peers. Rooted in the proxy-variable tradition, our identification scheme combines input optimality conditions with productivity dynamics in a CES revenue-production framework, allowing firm-specific input prices to be correlated with productivity and output prices to be unobserved. Applying the method to Chinese primary agricultural food-processing firms over 1998–2007, with foreign direct investment as a productivity shifter, we find that productivity externalities are strongly factor-biased, operating predominantly through the labor-augmenting component, and that the direct productivity gains from foreign investment are themselves factor-biased irrespective of investor origin, albeit far more strongly so for capital from Hong Kong, Macao and Taiwan, whose labor-augmenting effect is nearly three times that of capital from other foreign sources even as the two are comparable along the factor-neutral dimension.

Keywords: productivity; biased technical change; productivity spillovers; production function estimation; foreign direct investment

JEL Classification: D24; L20; O33; C33

*Email: zhou3@binghamton.edu (Hou), emir.malikov@unlv.edu (Malikov), shunanzhao@oakland.edu (Zhao).

1 Introduction

Because firm productivity is latent, its measurement depends on assumptions about how technological heterogeneity enters production and evolves over time. Two features of that structure are increasingly difficult to defend. First, productivity is not confined to the firm that embodies it but propagates across firms through agglomeration externalities, technology transfer, supply-chain linkages and learning from peers. This cross-firm dependence is reflected not only in evidence that a firm's measured productivity moves with that of the firms surrounding it (e.g., Glaeser et al., 1992; Moretti, 2004; Greenstone et al., 2010; Ellison et al., 2010; Malikov and Zhao, 2023), but also in the repeated finding that domestic producers grow more efficient when their industry or locality attracts more inbound foreign direct investment (Javorcik, 2004; Keller and Yeaple, 2009; Agarwal et al., 2014), exports more (Alvarez and López, 2008; Zhang and Malikov, 2023) or invests more intensively in research and development (Griliches, 1979; Coe and Helpman, 1995; Bloom et al., 2013; Zacchia, 2020). Second, the technological heterogeneity across firms that productivity captures is generally not factor-neutral, in that technical change is biased, augmenting some inputs more than others and reshaping factor shares in ways that a single Hicks-neutral residual cannot represent (see Doraszelski and Jaumandreu, 2018; Raval, 2019; Dewitte et al., 2020; Demirer, 2025). The workhorse structural estimators in the manner of Olley and Pakes (1996), Levinsohn and Petrin (2003), Ackerberg et al. (2015) and Gandhi et al. (2020), however, accommodate neither regularity because they assume that latent productivity is scalar, Hicks-neutral and evolves as a Markov process that is independent across firms.

These two abstractions are not innocuous, and the literature has relaxed them only separately. One strand accommodates cross-firm dependence by letting productivity of each firm depend on the productivity of its peers through spillovers but retains a unidimensional, Hicks-neutral formulation of firm heterogeneity (Malikov and Zhao, 2023; Iyoha, 2023). A parallel strand allows productivity to be multidimensional, typically, by distinguishing labor-augmenting from factor-neutral components but maintains that each firm's productivity process evolves independently of its peers, ruling out cross-firm spillovers by construction (Doraszelski and Jaumandreu, 2018; Zhang, 2019; Malikov et al., 2024; Zhao et al., 2026). What remains absent is a framework in which cross-firm dependence operates through the multidimensional components of productivity rather than through a single neutral residual. If firms learn from one another through technique, organization, input quality, workforce composition or managerial practice, then what diffuses across firms need not be a factor-neutral shift in output but may instead alter the relative productivity of different inputs. Existing methods therefore cannot identify whether productivity spillovers are themselves factor-biased.

In this paper, we develop a structural framework for the identification of firm productivity and of the cross-firm spillovers therein when productivity is multidimensional, comprising both a factor-neutral and a labor-augmenting component. Rooted in the canonical "proxy variable" tradition, the identification methodology we develop exploits the static optimality conditions together with the evolution processes for unobservables under the information set and timing assumptions. The resulting system identifies technology parameters jointly with the own-persistence and spillover coefficients of each productivity component, and can be estimated either sequentially in three steps or as a single nonlinear system of equations.

We formalize the firm's technology as a constant elasticity of substitution (CES) production function with nonneutral technical change, in which a Hicks-neutral component scales output radially while a labor-augmenting component enters by augmenting the effective quantity of labor. The central object of our analysis is the propagation of these latent productivity components across firms.

To this end, we let each component evolve as a controlled first-order vector Markov process whose conditional mean depends on the lagged cross-sectional state of the local industry through a network of peer links governed by spatial proximity and industrial similarity across firms. A firm's productivity is therefore shaped not only by its own past and by its productivity-altering activities but also by the productivity of its peers. This formulation conceptualizes peer dependence as spatiotemporal spillovers in latent productivity, and the spillover effects we seek to identify are dynamic analogues of what Manski (1993) terms "endogenous effects." By restricting the empirical scope to a single industry, we focus on horizontal spillovers among geographically proximate producers, consistent with the agglomeration literature's emphasis on localized knowledge diffusion (see Rosenthal and Strange, 2004; Duranton and Puga, 2004) and with production research documenting strong spatial dependence in productivity/efficiency across units (e.g., Glass et al., 2016; Skevas, 2023).

We also relax the common assumption that firms within an industry face homogeneous factor prices or that cross-firm input price dispersion is exogenous. While maintaining price taking in flexible-input markets, we allow factor prices to be firm-specific and, crucially, to depend on firm characteristics, including productivity. Thus, building on Grieco et al. (2016), we do not introduce monopsony power or quantity discounts but allow the prices at which firms can hire labor or procure intermediates to reflect heterogeneity in input quality, location, supplier relationships or other firm attributes correlated with productivity. This issue is particularly acute for wages because labor-augmenting productivity scales effective labor and may partly capture worker quality, workforce composition, management practices, task organization or other labor-side differences. These same attributes are plausibly priced in firm-specific wages. Treating current wages as homogeneous would then risk folding priced labor quality into latent labor-augmenting productivity, while treating them as exogenous instruments would require an orthogonality condition with respect to productivity innovations that our framework intentionally does not impose. On the output side, because most production datasets report revenue rather than separate output quantities and output prices, we close the model with a parsimonious demand structure. Firms are allowed to have market power in the output market, which we model through constant-elasticity demand under monopolistic competition. This assumption is a measurement device that maps the physical production technology into an estimatable revenue equation when output prices are unobserved.

Our paper is most closely related to Malikov and Zhao (2023) and Doraszelski and Jaumandreu (2018), whose conceptual paradigms we bring together and extend. From Malikov and Zhao (2023), we adopt the idea that spillover-induced cross-sectional dependence should enter the productivity process itself rather than be imposed ex post on productivity estimates recovered under cross-firm independence. Our first contribution is to generalize that logic to multidimensional productivity, whereby spillovers operate separately through factor-neutral and labor-augmenting productivity rather than through a single Hicks-neutral residual. This distinction is substantive because peer learning may diffuse organizational practices, labor quality, task design or input-specific techniques that alter the relative productivity of labor, not merely the radial total factor productivity of the firm. From Doraszelski and Jaumandreu (2018), we adopt the CES formulation of nonneutral technology and the use of flexible-input choices to separate labor-augmenting from factor-neutral productivity; unlike them, however, we allow both productivity components to spill over across firms. These extensions make our framework suitable for empirical settings that neither parent model accommodates. Relative to scalar spillover models, it can ask whether productivity transmission is factor-biased. Relative to existing multidimensional productivity models, it allows the latent productivity components themselves to propagate across firms. We also complement Iyoha (2023), who jointly estimates network effects and scalar productivity over the buyer-supplier production network, by treating the analogous identification problem for horizontal rather than vertical spillovers and for multidimensional rather than scalar productivity. We likewise differ from Lee (2023), who endogenizes the R&D-collaboration network as the outcome of forward-looking strategic decisions, by treating the peer network as weakly exogenous and predetermined with respect to the current-period productivity, an assumption motivated by the high cost of firm relocation and one that lets us identify spillovers in latent productivity without modeling the network-formation game.

We assess the finite-sample performance of our proposed estimator through a series of Monte Carlo experiments. Across several data-generating regimes for input prices, ranging from prices that are exogenous to firm productivity to prices that depend on both the factor-neutral and the labor-augmenting component, the estimator recovers the underlying technology, demand and productivity-process parameters consistently well, with estimation accuracy improving as the cross-section grows. Of note, the estimator performance is essentially unaffected by the introduction of productivity-correlated input-price dispersion, supporting that our identification is robust to the departure from the homogeneous-price assumption.

We then apply the methodology to firm-level data on China's primary agricultural food processing industry over 1998–2007, an industry marked by spatial agglomeration and substantial foreign presence, using foreign direct investment, separated into equity from Hong Kong, Macao and Taiwan (HMT) versus that from other foreign countries (OFC), as the shifter for both productivity components. Our central finding is that productivity externalities are biased towards labor, with the labor-augmenting spillover elasticity roughly an order of magnitude larger than its Hicks-neutral counterpart at the baseline (0.204 versus 0.021). Thus, the cross-firm transmission of productivity in this industry operates predominantly through the labor-augmenting component while the factor-neutral component diffuses little across firms. The productivity effect of foreign investment is itself factor-biased for equity of either origin, though its strength varies by source. In its direct effect on the recipient firm's own productivity, HMT equity is the more strongly labor-biased of the two, raising labor-augmenting productivity nearly three times as much as OFC equity does, even as the two sources are comparable along the Hicks-neutral dimension. The indirect gains that a domestic firm derives from the foreign equity of its peers, compounding this direct learning with the labor-biased spillover, are likewise concentrated almost entirely in the labor-augmenting component, with their factor-neutral counterpart an order of magnitude smaller. These patterns survive alternative peer-group definitions, weighting schemes and the inclusion of group-level fixed effects, with the labor-augmenting spillover exceeding its Hicks-neutral counterpart throughout.

The rest of the paper unfolds as follows. Section 2 lays out the structural framework, stating the assumptions on technology, productivity dynamics, decision timing and the market environment. Section 3 develops the identification argument and the three-step generalized method of moments implementation. Section 4 reports Monte Carlo simulations evaluating the finite-sample performance of the proposed estimator, and Section 5 presents the empirical application. Section 6 concludes.

2 Structural framework

We consider an industry of heterogeneous firms whose production transforms physical capital K_{it}, labor L_{it} and an intermediate input such as materials M_{it} into a scalar output Y_{it}, where $i = 1,\dots,N$ indexes firms and $t = 1,\dots,T$ indexes time. Throughout, we depart from the conventional assumption that firm heterogeneity is scalar and Hicks-neutral. Instead, we let technology be nonneutral by distinguishing between a factor-neutral productivity component $\omega_{H,it}$ that scales output radially, affecting marginal products of all inputs neutrally by the same proportion, and a labor-augmenting component $\omega_{L,it}$ that enters the technology by "augmenting" the labor input.

We formalize the firm's technology by means of a popular constant elasticity of substitution (CES) production function with non-Hicks-neutral technical change and non-unitary returns to scale (RTS). Following recent work on factor-augmenting technology and its implications, notably, Doraszelski and Jaumandreu (2018) and Raval (2019), the production process of firm i in period t is given by

$$Y_{it} = \underbrace{\left(\beta_K K_{it}^{-\frac{1-\sigma}{\sigma}} + \left[e^{\omega_{L,it}} L_{it}\right]^{-\frac{1-\sigma}{\sigma}} + \beta_M M_{it}^{-\frac{1-\sigma}{\sigma}}\right)^{-\frac{\nu\sigma}{1-\sigma}}}_{F(K_{it}, e^{\omega_{L,it}} L_{it}, M_{it})} e^{\omega_{H,it} + v_{it}}, \tag{2.1}$$

where ν and σ denote the elasticities of scale and substitution, respectively; $\omega_{H,it}$ and $\omega_{L,it}$ are the persistent factor-neutral and labor-augmenting productivity components, both observed by the firm at the time of production but unobserved by the researcher; and v_{it} is a mean-zero ex-post shock. As in Doraszelski and Jaumandreu (2018), without loss of generality, in (2.1) we implicitly set the constant of proportionality to one, since it cannot be separated from the latent factor-neutral productivity $\omega_{H,it}$, and for the analogous reason we normalize to unity the distributional parameter in front of the "effective labor" $e^{\omega_{L,it}} L_{it}$.

Of note, modeling labor-augmenting productivity separately from its factor-neutral counterpart is not a refinement of the scalar firm heterogeneity paradigm but a substantive departure from it because the two components carry distinct economic content. The factor-neutral component scales output for given inputs and leaves the marginal rate of technical substitution between factors unchanged, whereas the labor-augmenting component alters the relative marginal products of inputs and hence the labor share of cost, so that the latter is the object that governs the bias of technical change. Under a Cobb-Douglas technology the two would be observationally equivalent, with any labor-augmenting heterogeneity absorbed into the neutral residual and factor shares left undisturbed (Raval, 2019), and it is the non-unitary substitution ($\sigma \neq 1$) permitted by the CES that lends the labor-augmenting component empirical content and renders factor bias a measurable quantity.

Our central object of interest is the way in which firm productivity propagates across firms within an industry/sector. We conceptualize this peer dependence in firm performance through spatiotemporal productivity spillovers by allowing latent multidimensional firm heterogeneity to exhibit cross-sectional dependence across firms. To formalize the structure of these peer interactions, we let $\mathbf{S}_t$ be an $N \times N$ matrix describing the cross-firm spillover links, where the ith row contains unit-sum peer weights $\{s_{ijt-1}\}_{j=1}^{N}$ with zero diagonal ($s_{iit-1} = 0 \ \forall i = 1, \ldots, N$) which, for any $\{x_j\}_{j=1}^{N}$, lends the weighted sum $\sum_{j \neq i} s_{ijt-1} x_j$ the convenient interpretation of the average of firm i's peers' x_j.

With the above modeling environment in mind, we now formalize the structural assumption about the firm's technology, productivity dynamics, decision timing and market environment that jointly facilitate identification. Throughout the paper, $\mathscr{I}_{it}$ denotes all information that the firm i can use to solve its period-t decision problem.

Assumption A1 (Production technology) *Nonneutral production technology $F(\cdot)$ relating the ith firm's inputs (K_{it}, L_{it}, M_{it}) to the output Y_{it}, given its factor-neutral ($\omega_{H,it}$) and labor-augmenting ($\omega_{L,it}$) productivity components, takes the CES form given in equation* (2.1)*.*

The CES specification of the firm's technology in Assumption A1 is a parsimonious and widely used choice that admits nonneutral technical change (Doraszelski and Jaumandreu, 2018; Raval, 2019; Zhang, 2019; Caselli et al., 2026). It nests the canonical Cobb-Douglas technology, itself the workhorse of the structural productivity literature, as the limiting case $\sigma \to 1$,[1] and it leaves the substitution elasticity σ to be estimated rather than fixed at unity prima facie. Of more importance, however, is suitability of the CES for modeling factor bias as it accommodates the nonneutrality of technical change by separating the factor-neutral productivity $\omega_{H,it}$ from the labor-augmenting productivity $\omega_{L,it}$. The former shifts output for given inputs, whereas the latter alters the marginal rate of technical substitution between labor and the remaining inputs and hence the labor share of income/cost. This separation is not merely notational, because under Cobb-Douglas the two components would be observationally equivalent, with any labor-augmenting heterogeneity absorbed into Hicks-neutral productivity and factor shares left unchanged (Raval, 2019). As such, it is the nonunitary substitution permitted by the CES that lends $\omega_{L,it}$ empirical bite and renders the factor bias of technical change a measurable object (Doraszelski and Jaumandreu, 2018).

Assumption A2 (Information set and timing) *(i) The firm i's period-t information set includes contemporaneous and past draws of its persistent productivity:* $\{\omega_{H,i\tau}, \omega_{L,i\tau}\}_{\tau=1}^{t} \in \mathcal{I}_{it}$. *The transitory shock* $v_{it} \notin \mathcal{I}_{it}$ *is an i.i.d. white noise process and is observed by the firm ex post after solving its period-t decision problem:* $\mathcal{P}_v(v_{it} \mid \mathcal{I}_{it}) = \mathcal{P}_v(v_{it})$. *(ii) Capital is a dynamic input, accumulated according to* $K_{it} = \kappa(K_{it-1}, I_{it-1})$, *where gross investment* I_{it-1} *is optimally chosen by the firm i in period* $t-1$. *Labor and materials* (L_{it}, M_{it}) *are flexible inputs optimally chosen statically at time t. Concurrently, adjustments to peer links* $\mathbf{S}_t$ *determining the structure of the firm i's spillover network as well as its productivity controls,* $G_{H,it}$ *and* $G_{L,it}$ *respectively aimed at shifting productivity* $\omega_{H,it}$ *and* $\omega_{L,it}$, *are chosen (possibly dynamically) at time t. All choices are profit-maximizing in expectation, conditional on the available information* $\mathcal{I}_{it}$ *that, among other things, includes input prices and the expected output demand.*

Assumption A2 fixes the firm's information set and the timing of its decisions. Unlike persistent firm productivity, the transitory shock $v_{it} \perp \mathcal{I}_{it}$ is unforecastable and realized only after all period-t production decisions are finalized. It admits the usual interpretation of an ex-post transitory production shock or, alternatively, of classical measurement error in log-output.

Capital is subject to adjustment frictions such as time-to-install, and the firm therefore optimizes it dynamically via gross investment decision I_{it-1} at time $t-1$. K_{it} is thus a non-flexible input, predetermined at time $t-1$, whereby $\{K_{i\tau}\}_{\tau=1}^{t} \in \mathcal{I}_{it}$. Labor and materials, by contrast, are freely varying inputs chosen statically at time t conditional on the already-determined capital stock, in keeping with standard practice in the proxy-variable literature (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015, and many others). From an empirical standpoint, treating labor as "static" is plausible in the developing country setting of our application: the relatively elastic supply of labor and the low cost of employment adjustment during our sample period allowed firms considerable flexibility in their period-by-period labor decisions.

Beyond input allocation, the firm is also said to optimize—or indirectly influence—the productivity shifters $G_{H,it}$ and $G_{L,it}$ affecting the evolution of its multidimensional productivity, such as research and development (R&D), exporting and foreign direct investment (FDI), as well as its peer spillover network. Albeit potentially with dynamic implications too (particularly because of the delayed manner by which they affect the productivity evolution processes), the productivity controls $G_{H,it}$ and $G_{L,it}$ are chosen by the firm contemporaneously at time t unlike the level of K_{it} which is a delayed decision made at time $t-1$. This allows for persistence in productivity-enhancing activities

[1]Together with the Leontief technology of fixed proportions as $\sigma \to 0$.

but does not force them to be subject to adjustment frictions that would also render them delayed and, hence, predetermined.[2] This distinction is important because it does not rule out a contemporaneous correlation between the firm productivities $\{\omega_{H,it},\omega_{L,it}\}$ and the shifters $\{G_{H,it},G_{L,it}\}$, a feature that would be lost were the latter forced to be predetermined. Analogously, a firm's location and other network-relevant characteristics determining its peer connections $\mathbf{S}_t$ may adjust contemporaneously in response to time-t information, including realized persistent productivity draws. Our identifying assumption is therefore not strict but weak exogeneity of the network process, whereby the lagged network exposure $\mathbf{S}_{t-1}$ is predetermined relative to the period-t innovations (see also Assumption A3). This treatment of timing follows the structural logic of Biesebroeck (2004), De Loecker (2013), Doraszelski and Jaumandreu (2013), Malikov et al. (2020) and Malikov and Zhao (2023), and it is what permits us to dispense with an explicit formalization of the firm's long-run dynamic program while still anchoring the predeterminedness of the relevant conditioning variables in economic primitives.

Assumption A3 (Controlled first-order vector Markov) *The vector process of each persistent productivity component, $\boldsymbol{\omega}_{H,t}=(\omega_{H,it})_{i=1}^{N}$ and $\boldsymbol{\omega}_{L,t}=(\omega_{L,it})_{i=1}^{N}$, is controlled first-order Markov, with a linear vector conditional expectation function depending on its lagged cross-sectional state:*

$$\boldsymbol{\omega}_{H,t}=\rho_0\mathbf{i}+\left(\rho_1\mathbf{I}+\rho_2\mathbf{S}_{t-1}\right)\boldsymbol{\omega}_{H,t-1}+\rho_3'\mathbf{G}_{H,t-1}+\boldsymbol{\zeta}_{H,t}, \tag{2.2}$$

$$\boldsymbol{\omega}_{L,t}=\left(\gamma_1\mathbf{I}+\gamma_2\mathbf{S}_{t-1}\right)\boldsymbol{\omega}_{L,t-1}+\gamma_3'\mathbf{G}_{L,t-1}+\boldsymbol{\zeta}_{L,t}, \tag{2.3}$$

where $\mathbf{S}_{t-1}$ describing a peer network structure across firms at $t-1$ is measurable w.r.t. the information set $\boldsymbol{\Upsilon}_{t-1}=\{\mathscr{I}_{it-1}\}_{i=1}^{N}$ for the industry / sector in that period; $\mathbf{G}_{H,t-1}=(G_{H,it-1})_{i=1}^{N}$ and $\mathbf{G}_{L,t-1}=(G_{L,it-1})_{i=1}^{N}$ are (possibly overlapping) vectors of productivity shifters by which each firm may influence the evolution of its productivity, with all also measurable w.r.t. $\boldsymbol{\Upsilon}_{t-1}$; and the productivity innovations are unanticipated and satisfy $\mathbb{E}\left(\boldsymbol{\zeta}_{H,t}\mid\boldsymbol{\Upsilon}_{t-1}\right)=\mathbb{E}\left(\boldsymbol{\zeta}_{H,t}\right)=\mathbf{0}$ and $\mathbb{E}\left(\boldsymbol{\zeta}_{L,t}\mid\boldsymbol{\Upsilon}_{t-1}\right)=\mathbb{E}\left(\boldsymbol{\zeta}_{L,t}\right)=\mathbf{0}$.

Assumption A3 governs the dynamics of the two persistent productivity components and is where the cross-sectional dependence central to our analysis enters. Building on Malikov and Zhao (2023), we relax the standard assumption that firm productivity follows process that is independent across i, as in Doraszelski and Jaumandreu (2018), and instead let each component evolve as a controlled first-order vector Markov process whose conditional expectation depends on the lagged cross-sectional state of the entire industry/sector. The implied scalar unit-level representations of these processes for each firm i are

$$\omega_{H,it}=\rho_0+\rho_1\omega_{H,it-1}+\rho_2\sum_{j\neq i}s_{ijt-1}\omega_{H,jt-1}+\rho_3'G_{H,it-1}+\zeta_{H,it}, \tag{2.4}$$

$$\omega_{L,it}=\gamma_1\omega_{L,it-1}+\gamma_2\sum_{j\neq i}s_{ijt-1}\omega_{L,jt-1}+\gamma_3'G_{L,it-1}+\zeta_{L,it}, \tag{2.5}$$

where $G_{H,it-1}$ and $G_{L,it-1}$ collect firm-specific factors or activities through which the firm may deliberately influence the evolution of its productivity. This controlled Markov process formulation generalizes a model of "endogenous productivity change," first introduced in Doraszelski and Jaumandreu (2013) and De Loecker (2013) and now commonly adopted in productivity studies, because it permits

[2] For clarity, if the optimal decision concerning production in period t is affected by its history, then that decision is said to be "dynamic." If, due to adjustment frictions, a decision concerning production in period t is effectively made at $t-1$, then it is also "predetermined." In this nomenclature, K_{it} is dynamic and predetermined, whereas $G_{H,it}$, $G_{L,it}$ and $\mathbf{S}_t$ are dynamic but are chosen at time t.

the firm to change its productivity over time not only through its own decisions but also by learning from peers via spillovers, accommodating the agglomeration externalities that the independent-process abstraction precludes by construction. We assume that each process is linear in its arguments, a common approximation in the literature (e.g., Biesebroeck, 2004; Ackerberg et al., 2015; Zhang, 2019; Malikov et al., 2024; Ackerberg et al., 2023, to name a few). Since $\omega_{L,it}$ is identified only up to a constant, we suppress a constant term in its evolution equation.[3]

The autoregressive coefficients in the spatiotemporal autoregressions (2.4)–(2.5) carry distinct economic content. The own-persistence parameters ρ_1 and γ_1 measure the degree to which a firm's current productivity is inherited from its own past, whereas the spillover parameters ρ_2 and γ_2 measure the strength with which the lagged productivity of a firm's peer group, channeled through the network matrix $\mathbf{S}_{t-1}$, propagates into its current productivity. The latter are the objects of interest we seek to identify and correspond to what Manski (1993) terms *endogenous* network effects, whereby a firm's outcome is directly influenced by the mean outcome of its peers. The productivity shifters $G_{H,it-1}$ and $G_{L,it-1}$, with coefficients ρ_3 and γ_3, capture internal learning whereby the firm enhances its own productivity through its deliberate activities. Of note, Assumption A3 places no restriction on the correlation between the two productivity components $\omega_{H,it}$ and $\omega_{L,it}$, which may be related through overlapping controls and through their innovations $\zeta_{H,it}$ and $\zeta_{L,it}$ that are plausibly positively correlated. The innovations are mean-independent of the lagged industry information set, being the residuals of the conditional expectation functions and therefore mean-orthogonal to all $\boldsymbol{\Upsilon}_{t-1}$-measurable variables determined prior to period t. Substantively, $\zeta_{H,it}$ collects unanticipated innovations in persistent productivity such as implementation uncertainty, random discovery and chance in the success of productivity-enhancing activities, with $\zeta_{L,it}$ admitting the analogous interpretation for labor-augmenting productivity.

The network matrix $\mathbf{S}_{t-1}$ operationalizes the peer group through which spillovers flow. Its typical element s_{ijt-1} is the (i,j)th peer weight, non-stochastic conditional on the lagged information set $\boldsymbol{\Upsilon}_{t-1}$ and constructed with a zero diagonal, $s_{iit-1}=0$ for all i. Depending on the empirical context and relevant priors, the weights may be specified as the normalized inverse geographic distance between firms or, more coarsely, as a uniform weighting applied to all peers within a firm's spatial vicinity. As a baseline, we identify firm i's peers on the basis of both spatial proximity and industrial similarity, so that, letting $\mathscr{L}(i,t)$ denote the set of spatially proximate "neighbors" of firm i in period t that also operate in the same industry/sector, the peer weights are

$$s_{ijt} = \frac{\mathbf{1}\{(j,t)\in\mathscr{L}(i,t)\}}{\sum_{k(\neq i)=1}^{n}\mathbf{1}\{(k,t)\in\mathscr{L}(i,t)\}}. \tag{2.6}$$

Restricting the scope to geographically proximate peers within the industry fits a broader narrative in regional and urban economics about the localized nature of agglomeration economies and knowledge diffusion (e.g., see Rosenthal and Strange, 2004; Duranton and Puga, 2004), and it effectively focuses the analysis on intra-industry horizontal spillovers. With the network matrix row-standardized, the weights in each row sum to unity and, thus, $\sum_{j\neq i}s_{ijt-1}\omega_{H,jt-1}$ and $\sum_{j\neq i}s_{ijt-1}\omega_{L,jt-1}$ in (2.4)–(2.5) measure average peer productivity levels. Further, the weighting scheme in (2.6) treats cross-firm spillovers symmetrically, in that all members of a peer group influence one another's productivity. This standard approach to measuring "peer effects" allows us to remain agnostic about the

[3]Because productivity is measured in relative terms, this restriction effectively normalizes $\omega_{L,it}$ to have a zero mean and does not affect firms' relative rankings in labor-augmenting productivity. Of note, we are able to allow a non-zero constant in the $\omega_{H,it}$ process because of the unit-constant-of-proportionality restriction on the production function or, otherwise, it would have also had to be suppressed.

direction of peer influence and avoids imposing the directionality of cross-firm dependence a priori.

The identification of network effects of this kind is a notoriously difficult problem, and our specification confronts its three classic obstacles deliberately (see Manski, 1993, 2000; Moffitt, 2001). The first is the so-called reflection problem, namely the difficulty of distinguishing endogenous effects, which operate through peer outcomes (here, productivity), from *contextual* effects, which operate through peer characteristics (here, productivity shifters). By postulating that cross-firm interaction occurs solely through productivity, so that a firm's productivity depends on the mean productivity of its peer group, our process is a dynamic analogue of a "pure endogenous-effects model" and thereby sets aside the contextual effects of peers' productivity-enhancing activities that would otherwise render the endogenous and contextual channels separately unidentified.[4] The second obstacle is the confounding presence of unobserved *correlated* effects, that is, common factors such as a shared regional or business environment that induce productivity correlation among firms in the same network cluster independently of any genuine spillovers. Our Assumption A3 implicitly assumes that the innovations $\zeta_{H,it}$ and $\zeta_{L,it}$ contain no such peer-group-level component, so that all cross-firm dependence is attributed to within-group spillovers rather than to common shocks. Although admittedly restrictive, this assumption is standard given the well-known challenges of accounting for group-level unobservables in network models (see Blume et al., 2011, for a comprehensive review), and it ensures comparability with the established empirical literature on R&D, FDI and export spillovers. For researchers who find it too strong, it can be relaxed by admitting time-invariant network unobservables and/or time-varying common shocks, addressed through the inclusion of group-level fixed effects and time dummies (Graham and Hahn, 2005; Bramoullé et al., 2009), as we also explore in our empirical analysis.[5]

The third obstacle is the endogeneity of group membership, which in the spatial/urban economics literature surfaces as the "endogenous sorting" problem whereby more productive firms ex ante self-select into locations populated by other highly productive firms. Were sorting left unaddressed, differences in future productivity across firms facing different levels of agglomeration could be mistakenly attributed to spillovers when they merely reflect the tendency of productive firms to cluster. Including a firm's own lagged productivity as a control partially absorbs this self-sorting, in the same spirit as the argument that De Loecker (2013) makes for the self-selection of exporters in the context of export-based learning. Beyond this, we maintain that the peer network is exogenous only weakly. The plausibility of this assumption rests on the high cost of firm relocation, which renders network structures effectively predetermined: firms neither change their location or network in anticipation of future productivity shocks nor select locations in ways systematically correlated with expected future innovations. Imposing the structural timing of Assumption A2 then makes the lagged network characterized by $\mathbf{S}_{t-1}$ predetermined with respect to the period-t productivity innovations, which is what allows us to identify the spillover effect by comparing current productivity across firms with different exposure to the lagged weighted average productivity of nearby firms, holding fixed their own past productivity and productivity-shifting activities.

Assumption A4 (Market environment) *(i) Firms are price-takers in the input markets for labor and materials, where they face prices P^M_{it} and P^L_{it}, respectively. Suppliers use linear pricing, but*

[4] For the formal statement of this indistinguishability result and the conditions under which a single channel of cross-peer dependence restores identification, see Manski (1993). Our setup may be augmented to admit contextual effects of the peers' controls $\{G_{H,jt-1}, G_{L,jt-1}\}$, at the cost of additional restrictions, although we abstract from them in this paper.

[5] Time fixed effects absorb aggregate shocks common to all units at a given point in time, which admittedly is sufficient only when such shocks affect all units homogeneously. If, however, time-varying unobserved shocks vary across peer groups or load differentially on units depending on their characteristics or network position, these components are not eliminated by time fixed effects but could be addressed through common factor structure (see Shi and Lee, 2017; Bai and Li, 2021).

the prices may vary across firms and over time, as well as depend on firm characteristics, including productivity. (ii) In the output market, firms are monopolistically competitive and face a constant-elasticity (inverse) demand function:

$$P_{it} = P_t \left(\frac{Y_{it}}{Y_t} \right)^{\frac{1}{\eta}} e^{\varrho' \chi_{it} + \xi_{it}}, \tag{2.7}$$

where P_t is the price index (the average price) for a given industry/sector at time t; Y_t is the corresponding quantity index (aggregate demand shifter); and χ_{it} captures observed demand shocks specific to firm i; and ξ_{it} is a zero-mean i.i.d. demand heterogeneity: $\mathscr{P}_{\xi}(\xi_{it} \mid \mathscr{I}_{it}) = \mathscr{P}_{\xi}(\xi_{it})$.

We assume the market environment in Assumption A4, which concerns both the input and the output sides of the firm's problem. On the input side, we maintain that firms are price-takers in the labor and materials markets but, departing from the bulk of the proxy-variable literature, we do not require that the prices P^L_{it} and P^M_{it} be homogeneous across firms or exogenous to firm characteristics. Suppliers use linear pricing, so we rule out quantity discounts that would endogenize prices with respect to purchasing decisions, but the prices are allowed to vary across firms and over time and, crucially, to depend on firm characteristics including productivity in the spirit of Grieco et al. (2016) and Malikov et al. (2025). This is a meaningful relaxation because it accommodates the empirically documented dependence of input prices on firm size (Kugler and Verhoogen, 2012; Dhawan, 2001) and exporting behavior (Manova and Zhang, 2012; Verhoogen, 2008), both of which are tightly correlated with productivity. Allowing input prices to correlate with firm productivity is especially important in our setting because labor-augmenting productivity may partly reflect differences in the quality and organization of labor that are also priced in wages. Since $\omega_{L,it}$ enters the technology by scaling the effective quantity of labor as $e^{\omega_{L,it}} L_{it}$, cross-firm heterogeneity in labor-augmenting productivity that we seek to isolate may capture worker quality, workforce composition, management practices, task organization or other differences in how labor is selected, organized and used. These same differences can generate wage dispersion across otherwise similar firms, for instance, through worker sorting, compensating differentials or efficiency-wage mechanisms. A specification that forces wages to be homogeneous would then risk attributing priced labor quality to latent labor-augmenting productivity and thus distorting inference about the direction and magnitude of factor bias, while a specification that treats (heterogeneous) current-period wages as exogenous instruments would require an orthogonality condition between wages and productivity innovations that we deliberately wish to avoid. The same logic applies, although less directly, to materials, whose firm-specific prices could also co-move with unobserved input quality and hence with both productivity components.[6] Contemporaneous correlation between input prices and productivity may also occur when less productive firms survive by accessing lower input prices or when more productive firm use higher-quality, more efficient inputs as discussed in Flynn et al. (2019). Such quality-driven correlations between prices and productivity are well documented for China, the setting of our data (Fan et al., 2018).

On the output side, we let firms be monopolistically competitive and face the Klette and Griliches (1996)-style constant-elasticity demand; also see De Loecker (2011), Gandhi et al. (2020) and Zhao et al. (2026) for recent applications. Explicit modeling of the demand side is necessary in our setting, because firm-level output quantities and prices are rarely observed and only revenue is reported, so that a demand structure is required to express the unobserved output entering (2.1) in terms of

[6]Further sources of heterogeneity in wage rates and material prices, despite firms being price-takers in the variable input markets, include transportation costs, as discussed by Grieco et al. (2016) and Zhang (2019).

observable revenue.[7] The price index P_t, the aggregate quantity index Y_t and the observed firm-specific demand shock χ_{it} are all known to the firm when it finalizes its period-t input decisions and therefore belong to $\mathscr{I}_{it}$, with the index Y_t constructible, for example, as a market-share-weighted average of firms' outputs (De Loecker, 2011). The random demand heterogeneity $\xi_{it} \perp \mathscr{I}_{it}$ plays a role symmetric to that of the transitory production shock v_{it}, and the firm must form expectations over it when solving its decision problem.

Taken together, the four assumptions A1 through A4 deliver a structural model of production in which two-dimensional, factor-neutral and labor-augmenting, productivity evolves through controlled processes that admit both internal learning and cross-firm spillovers, the firm faces heterogeneous input prices that may correlate with its productivity and exercises market power in the output market subject to constant-elasticity demand.

3 Identification and implementation methodology

Our identification scheme rests on the timing and information set restrictions coupled with the evolution process assumptions about unobservables fleshed out in A2–A4, which render the relevant conditioning variables predetermined with respect to the current-period innovations and yield the orthogonality moment restrictions that we exploit together with the static optimality conditions. We carry out our identification procedure in three steps.

Remark 1 *In Appendix A, we present our identification procedure for the special case in which firms operate in perfectly competitive output and factor markets, facing homogeneous, exogenous prices. This classical market environment assumption is a workhorse of the structural productivity literature (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2020) and is also imposed by Malikov and Zhao (2023). There, we demonstrate that identification remains feasible even in the absence of cross-sectional price variation or observable exogenous demand shifters. This theoretical benchmark may be particularly relevant for empirical work when firm-specific price data are unavailable, as oftentimes is the case.*

3.1 Identification

Substituting for P_{it} using (2.7), under Assumptions A2 and A4 the firm's static profit-maximization problem in expectation is given by

$$\max_{M_{it},L_{it}} P_t \frac{[F(K_{it}, e^{\omega_{L,it}} L_{it}, M_{it}) e^{\omega_{H,it}}]^{1+\frac{1}{\eta}}}{Y_t^{\frac{1}{\eta}}} e^{\varrho' \chi_{it}} \Lambda^* - P_{it}^M M_{it} - P_{it}^L L_{it}, \tag{3.1}$$

where $\Lambda^* \equiv \mathbb{E}(e^{v_{it}^*} \mid \mathscr{I}_{it}) = \mathbb{E}(e^{v_{it}^*})$ is the firm's expectation of a composite ex post shock $v_{it}^* \equiv \left(1+\frac{1}{\eta}\right) v_{it} + \xi_{it}$. The corresponding first-order conditions for materials and labor are

$$\left(1+\frac{1}{\eta}\right) P_t \left(\frac{F(\cdot)}{Y_t}\right)^{\frac{1}{\eta}} \frac{\partial F(\cdot)}{\partial V_{it}} e^{\left(1+\frac{1}{\eta}\right)\omega_{H,it} + \varrho' \chi_{it}} \Lambda^* = P_{it}^V \qquad \forall\ V_{it} \in \{M_{it}, L_{it}\}, \tag{3.2}$$

[7] For this reason, our output market structure is more restrictive than that in Doraszelski and Jaumandreu (2018), who treat a firm's residual demand function as nonparametric, permitting heterogeneous price elasticities.

where $\frac{\partial F(\cdot)}{\partial V_{it}}$ $\forall$ $V_{it} \in \{M_{it}, L_{it}\}$ is a respective marginal product net of Hicks-neutral $e^{\omega_{H,it}}$ and error term v_{it}:

$$\frac{\partial F(\cdot)}{\partial L_{it}} = \nu X_{it}^{-\frac{\nu\sigma}{1-\sigma}-1}[e^{\omega_{Lit}}L_{it}]^{-\frac{1}{\sigma}}e^{\omega_{Lit}} \qquad \text{and} \qquad \frac{\partial F(\cdot)}{\partial M_{it}} = \nu\beta_M X_{it}^{-\frac{\nu\sigma}{1-\sigma}-1}M_{it}^{-\frac{1}{\sigma}}, \tag{3.3}$$

with the CES aggregator kernel inside the production function $X_{it} \equiv \beta_K K_{it}^{-\frac{1-\sigma}{\sigma}} + [e^{\omega_{Lit}}L_{it}]^{-\frac{1-\sigma}{\sigma}} + \beta_M M_{it}^{-\frac{1-\sigma}{\sigma}}$.

Let us also recognize that, using (3.2), the optimal X_{it} can be expressed as

$$X_{it} = \beta_K K_{it}^{-\frac{1-\sigma}{\sigma}} + \beta_M M_{it}^{-\frac{1-\sigma}{\sigma}}\left[S_{M,it}\right]^{-1}, \tag{3.4}$$

where $S_{M,it} = \frac{P_{it}^M M_{it}}{P_{it}^L L_{it} + P_{it}^M M_{it}}$ denotes the share of materials in the cost of flexible inputs. The above expression is useful because it is rid of the unobserved $\omega_{L,it}$ which is implicitly being proxied using the optimal materials and labor demand.

Step One. We begin by identifying the elasticity of substitution σ, the CES distribution parameters (β_K, β_M), and the ω_L Markov process parameters $(\gamma_1, \gamma_2, \gamma_3)$.

Taking the ratio of the first-order conditions in (3.2) and exploiting the parametric form of $F(\cdot)$, the log ratio of materials to labor at the optimum can be expressed as

$$m_{it} - l_{it} = \sigma \ln \beta_M - \sigma\left(p_{it}^M - p_{it}^L\right) + (1-\sigma)\omega_{L,it}, \tag{3.5}$$

where, throughout the paper, we adopt the convention that uppercase Latin letters denote levels and lowercase denote their logs. The optimal material-to-labor ratio does not depend on factor-neutral $\omega_{H,it}$ since the latter influences the marginal productivity of both static inputs by the same proportion, leaving their optimal ratio unaffected. The mix of these inputs is, however, related to $\omega_{L,it}$ as, by augmenting the labor input, this labor-saving productivity affects the relative marginal products of materials and labor.

At its core, equation (3.5) is a log-linear regression model of the optimal material-to-labor ratio $m_{it} - l_{it}$ on the relative price $p_{it}^M - p_{it}^L$ with an "error term" $(1-\sigma)\omega_{L,it}$. With enough variation in heterogeneous input prices that are exogenous with respect to $\omega_{L,it}$, it in theory could identify σ and β_M from the regression parameters. However, the former exogeneity condition—particularly, between p_{it}^L and $\omega_{L,it}$—is implausible as discussed at length in Section 2 and why we intentionally relax it in Assumption A4. As such, rather than using (3.5) as an estimating equation "as is" we, instead, use it to form a proxy expression for $\omega_{L,it}$ as it provides a unique mapping from the unobservable labor-augmenting productivity to observable variation in the material-to-labor ratio and the corresponding relative price.

Specifically, we invert (3.5) for $\omega_{L,it}$ as

$$\omega_{L,it} = \tfrac{-\sigma}{1-\sigma}\ln\beta_M + \tfrac{1}{1-\sigma}(m_{it} - l_{it}) + \tfrac{\sigma}{1-\sigma}\left(p_{it}^M - p_{it}^L\right). \tag{3.6}$$

Then, under the Markov property of $\omega_{L,it}$ per Assumption A3, using the inversion in (3.6) to substitute for $\omega_{L,it}$, $\omega_{L,it-1}$ and $\{\omega_{L,jt-1}\}_{j=1}^N$ in (2.5), we arrive at the autoregressive evolution process of labor-augmenting productivity in terms of its proxies, with a single unobservable being the innovation term $\zeta_{L,it}$, rewriting which we obtain our first estimating equation:

$$m_{it} - l_{it} = \left(1-(\gamma_1+\gamma_2)\right)\sigma \ln \beta_M - \sigma\left(p^M_{it} - p^L_{it}\right) + g_L(\mathbb{X}_{L,it-1}) + \underbrace{\zeta^*_{L,it}}_{(1-\sigma)\zeta_{L,it}}, \tag{3.7}$$

where we have used $\sum_{j\neq i} s_{ijt-1} = 1$ and

$$\begin{aligned} g_L(\mathbb{X}_{L,it-1}) \equiv\ & \gamma_1(m_{it-1} - l_{it-1}) + \sigma\gamma_1\left(p^M_{it-1} - p^L_{it-1}\right) + \\ & \gamma_2 \sum_{j\neq i} s_{ijt-1}\left(m_{jt-1} - l_{jt-1}\right) + \sigma\gamma_2 \sum_{j\neq i} s_{ijt-1}\left(p^M_{jt-1} - p^L_{jt-1}\right) + (1-\sigma)\gamma_3' G_{L,it-1}, \end{aligned} \tag{3.8}$$

with $\mathbb{X}_{L,it-1} = \left(m_{it-1} - l_{it-1}, p^M_{it-1} - p^L_{it-1}, \sum_{j\neq i} s_{ijt-1}(m_{jt-1} - l_{jt-1}), \sum_{j\neq i} s_{ijt-1}(p^M_{jt-1} - p^L_{jt-1}), G_{L,it-1}\right)$ collecting all covariates on the right-hand side.

Evidently, equation (3.7) is a restatement of (3.5) with the contemporaneous $\omega_{L,it}$ replaced using its Markov process in which all lagged productivity terms have been substituted away using the proxy function. Per Assumptions A2–A3, unanticipated productivity innovations $\{\zeta_{L,it}\}_{i=1}^N$ are mean-orthogonal to the lagged information set $\mathbf{\Upsilon}_{t-1}$, and thus

$$\mathbb{E}\left(\zeta^*_{L,it} \mid \mathbb{X}_{L,it-1}\right) = \mathbb{E}\left(\zeta^*_{L,it}\right) = 0, \tag{3.9}$$

implying that the $\mathbb{X}_{L,it-1}$ regressors in equation (3.7) are all self-instrumenting. This, however, still leaves the period-t relative price $p^M_{it} - p^L_{it}$ that is susceptible to contemporaneous correlation with $\zeta^*_{L,it}$ per our Assumption A4, and at first glance, it might appear as if no gain has been had by moving from (3.5) to (3.7). But it does *not* jeopardize the identification of (3.7) as $p^M_{it} - p^L_{it}$ is not a "free" regressor. Rather, it enters the equation subject to a parameter restriction and, as such, does not require an additional moment condition. Concretely, the unknown parameters in (3.7) are $(\sigma, \beta_M, \gamma_1, \gamma_2, \gamma_3)$: γ_1 is identified off the variation in own lagged $m_{it-1} - l_{it-1}$; γ_2 off the variation in lagged peer average $\sum_{j\neq i} s_{ijt-1}(m_{jt-1} - l_{jt-1})$; γ_3 off the variation in own lagged shifters $G_{L,it-1}$; σ is identified using the lagged information about both $p^M_{it-1} - p^L_{it-1}$ and $\sum_{j\neq i} s_{ijt-1}(p^M_{jt-1} - p^L_{jt-1})$; and the unconditional moment that the innovations are mean-zero helps pin-point β_M in the intercept. Thus, the moment condition system in (3.9) that we use to identify these parameters is already over-identified without an additional instrument for the contemporaneous $p^M_{it} - p^L_{it}$.[8]

As in Doraszelski and Jaumandreu (2018), without loss of generality we set $\beta_K + \beta_M = 1$, which implies immediate identification of β_K now that we have β_M. Following this step, we will have also recovered the firm's labor-augmenting productivity $\omega_{L,it}$ and the optimal X_{it}, leaving ν as the only remaining unknown production-function parameter.

Step Two. This is an intermediate step in which we identify some auxiliary parameters. Define $\widetilde{r}_{it} \equiv \ln \frac{P_{it}Y_{it}}{P_t}$. Under Assumptions A1 and A4, from equations (2.1) and (2.7) together with the expression for the optimal CES aggregator kernel X_{it} in (3.4), we obtain the log-revenue production function in a partially reduced form from which $\omega_{L,it}$ has been concentrated out:

$$\widetilde{r}_{it} = \left(1 + \tfrac{1}{\eta}\right)\left(-\tfrac{\nu\sigma}{1-\sigma}\right) x_{it} + \left(1 + \tfrac{1}{\eta}\right)\omega_{H,it} - \tfrac{1}{\eta} y_t + \varrho' \chi_{it} + v^*_{it}, \tag{3.10}$$

where, recall, x_{it} and σ have already been identified and, thus, can be treated as known.

[8]Doraszelski and Jaumandreu (2018) analogously exclude current relative prices from the instruments.

This revenue production function contains a confounding unobservable, $\omega_{H,it}$. To construct a proxy function for it, we invert the first-order condition for materials in (3.2), yielding the following expression in logs:

$$\left(1+\tfrac{1}{\eta}\right)\omega_{H,it} = \gamma_H + \tfrac{1}{\sigma}m_{it} - \varrho'\chi_{it} + p^M_{it} - p_t + \tfrac{1}{\eta}y_t + \left(1+\tfrac{\eta+1}{\eta}\tfrac{\nu\sigma}{1-\sigma}\right)x_{it}, \tag{3.11}$$

where $\gamma_H = -\ln\left(\nu\beta_M\Lambda^*\left(1+\tfrac{1}{\eta}\right)\right)$. Substituting for $\left(1+\tfrac{1}{\eta}\right)\omega_{H,it}$ in (3.10) using the proxy (3.11), we obtain a reduced form of our model that is rid of all persistent productivity terms:

$$\widetilde{r}_{it} = \gamma_H + x_{it} + p^M_{it} - p_t + \tfrac{1}{\sigma}m_{it} + v^*_{it}. \tag{3.12}$$

It is useful because it facilitates identification of the composite noise v^*_{it}. Namely, under Assumptions A2 and A4, $\mathbb{E}\left(v^*_{it}\right) = 0$ and thus, substituting for v^*_{it} using (3.12), we have

$$\mathbb{E}\left(\widetilde{r}_{it} - x_{it} - p^M_{it} + p_t - \tfrac{1}{\sigma}m_{it}\right) = \gamma_H, \tag{3.13}$$

and $\widetilde{r}_{it} - x_{it} - p^M_{it} + p_t - \frac{1}{\sigma}m_{it}$ is known/observed following the first step. While γ_H identified from (3.13) is of little interest in itself, knowing it allows us to identify v^*_{it} as a residual in (3.12).

Step Three. Having identified σ, γ_1, γ_2, γ_3, β_M and β_K in step one, as well as v^*_{it} in step two, we next complete the recovery of the proposed production system by identifying the RTS parameter ν, the parameters governing the evolution of Hicks-neutral productivity $(\rho_0, \rho_1, \rho_2, \rho_3)$, and the demand-function parameters η and ϱ.

Define $\widetilde{r}^*_{it} = \widetilde{r}_{it} - v^*_{it}$, a known quantity following step two. From (3.10), it follows that

$$\omega_{H,it} = \tfrac{\eta}{\eta+1}\widetilde{r}^*_{it} + \tfrac{\nu\sigma}{1-\sigma}x_{it} + \tfrac{1}{\eta+1}y_t - \tfrac{\eta}{\eta+1}\varrho'\chi_{it}. \tag{3.14}$$

Using the Markov process of $\omega_{H,it}$ in (2.4) under Assumption A3 and substituting for $\omega_{H,it}$, $\omega_{H,it-1}$ and $\{\omega_{H,jt-1}\}_{j=1}^N$ using expression (3.14), we obtain

$$\widetilde{r}^*_{it} = \tfrac{\eta+1}{\eta}\rho_0 - \tfrac{\nu\sigma}{1-\sigma}\tfrac{\eta+1}{\eta}x_{it} - \tfrac{1}{\eta}y_t + \varrho'\chi_{it} + g_H(\mathbb{X}_{H,it-1}) + \underbrace{\zeta^*_{H,it}}_{\frac{\eta+1}{\eta}\zeta_{H,it}}, \tag{3.15}$$

where, like before, we have used $\sum_{j\neq i} s_{ijt-1} = 1$ and

$$\begin{aligned} g_H(\mathbb{X}_{H,it-1}) \equiv\; & \rho_1\widetilde{r}^*_{it-1} + \tfrac{\nu\sigma}{1-\sigma}\tfrac{\eta+1}{\eta}\rho_1 x_{it-1} + \tfrac{1}{\eta}(\rho_1+\rho_2)y_{t-1} - \varrho'\rho_1\chi_{it-1} + \\ & \rho_2\sum_{j\neq i} s_{ijt-1}\widetilde{r}^*_{jt-1} + \rho_2\tfrac{\eta+1}{\eta}\tfrac{\nu\sigma}{1-\sigma}\sum_{j\neq i} s_{ijt-1}x_{jt-1} - \rho_2\varrho'\sum_{j\neq i} s_{ijt-1}\chi_{jt-1} + \tfrac{\eta+1}{\eta}\rho_3' G_{H,it-1}, \end{aligned} \tag{3.16}$$

with $\mathbb{X}_{H,it-1} = \left(\widetilde{r}^*_{it-1}, x_{it-1}, y_{t-1}, \chi_{it-1}, \sum_{j\neq i} s_{ijt-1}\widetilde{r}^*_{jt-1}, \sum_{j\neq i} s_{ijt-1}x_{jt-1}, \sum_{j\neq i} s_{ijt-1}\chi_{jt-1}, G_{H,it-1}\right)$ collecting all lagged covariates. The construction of the estimating equation (3.15)–(3.16) above is analogous to that of (3.7)–(3.8) in step one.

Per Assumptions A2–A3, productivity innovations $\{\zeta_{H,it}\}_{i=1}^N$ are unanticipated and zero-mean conditional on period-$t-1$ information set, and we therefore have that

$$\mathbb{E}\left(\zeta^*_{H,it} \mid \mathbb{X}_{H,it-1}\right) = \mathbb{E}\left(\zeta^*_{H,it}\right) = 0, \tag{3.17}$$

whereby $\mathbb{X}_{H,it-1}$ in equation (3.15) is weakly exogenous with respect to the error term. But the contemporaneous covariates, (x_{it}, y_t, χ_{it}) are endogenous. x_{it} is a function of static inputs (M_{it}, L_{it}) which, per Assumption A2, are chosen after the current-period realization of $\omega_{H,it}$, which is inclusive of the productivity update $\zeta_{H,it}$. In addition, both the aggregate demand shifter y_t and the firm-specific demand shocks χ_{it} also belong to the firm's period-t information set (also by Assumption A2) and, as such, may correlate with the contemporaneous $\zeta_{H,it}$, reflecting common macroeconomic conditions or firm-level factors, such as quality improvements or managerial innovations, that jointly affect demand and productivity. Having said that, neither one of (x_{it}, y_t, χ_{it}) is a "free" regressor as is obvious from (3.15)–(3.16), so they require no external instrumentation and the model is identified based on the moments (3.17) alone. This argument is akin to our step-one discussion, but it might be helpful to see why the contemporaneous regressors need no excluded instruments from a different angle. To this end, rewrite (3.15)–(3.16) as

$$\begin{aligned}\widetilde{r}^*_{it} = {} & \tfrac{\eta+1}{\eta}\rho_0 + \rho_1 \widetilde{r}^*_{it-1} - \tfrac{\nu\sigma}{1-\sigma}\tfrac{\eta+1}{\eta}(x_{it} - \rho_1 x_{it-1}) - \tfrac{1}{\eta}\left(y_t - (\rho_1+\rho_2)y_{t-1}\right) + \varrho'(\chi_{it} - \rho_1\chi_{it-1}) + \\ & \rho_2 \sum_{j\neq i} s_{ijt-1}\widetilde{r}^*_{jt-1} + \rho_2 \tfrac{\eta+1}{\eta}\tfrac{\nu\sigma}{1-\sigma}\sum_{j\neq i} s_{ijt-1}x_{jt-1} - \rho_2\varrho'\sum_{j\neq i} s_{ijt-1}\chi_{jt-1} + \tfrac{\eta+1}{\eta}\rho_3' G_{H,it-1} + \zeta^*_{H,it}.\end{aligned} \tag{3.18}$$

The above makes it clearer that the equation we seek to identify is just a quasi-difference dynamic panel-data model of the log-revenue production function with the labor-augmenting productivity proxied out and the factor-neutral productivity removed via Cochrane-Orcutt-like transformation. Thus, with (3.17) we effectively instrument the endogenous quasi-difference $x_{it} - \rho_1 x_{it-1}$ with the own first lag in level x_{it-1}. Analogously, y_{t-1} serves as a valid instrument for $y_t - (\rho_1+\rho_2)y_{t-1}$, and χ_{it-1} for $\chi_{it} - \rho_1\chi_{it-1}$. Concluding this step, we therefore have identified ν, η, ϱ and $(\rho_0, \rho_1, \rho_2, \rho_3)$, using which we also recover the firm's Hicks-neutral productivity $\omega_{H,it}$.

3.2 Estimation

Our identification strategy can be empirically implemented via generalized method of moments (GMM) using the moment restrictions in (3.9), (3.13) and (3.17) either in three sequential steps as outlined in Section 3.1[9] or, if desired, all in one step as a multi-equation system estimator. The latter option is convenient because it makes it straightforward to establish consistency and asymptotic normality of our three-step estimator by appealing to the standard limit results for a nonlinear system GMM estimator. Incidentally, it also facilitates the derivation of an asymptotic covariance for all parameter estimators that accounts for a sequential nature of our methodology.

Concretely, letting $\theta = (\sigma, \nu, \beta_M, \beta_K, \gamma_H, \gamma_1, \gamma_2, \gamma_3, \rho_0, \rho_1, \rho_2, \rho_3, \eta, \varrho)$ be a collection of all parameters across the three steps, the unconditional version of the recursive moments in (3.9), (3.13) and

[9] This is how we implement it due to its simplicity and better numerical stability.

(3.17) stacked together for the entire system of equations is given by

$$\mathbb{E}(\mathbf{h}_{it}(\theta)) \equiv \mathbb{E}\begin{pmatrix} \zeta^*_{L,it}\left(\sigma,\beta_M,\gamma_1,\gamma_2,\gamma_3\right)\begin{bmatrix} 1 \\ \mathbb{X}_{L,it-1}\end{bmatrix} \\ 1-\beta_M-\beta_K \\ v^*_{it}\left(\gamma_H \mid \sigma,\beta_M,\beta_K\right) \\ \zeta^*_{H,it}\left(\nu,\eta,\varrho,\rho_0,\rho_1,\rho_2,\rho_3 \mid \sigma,\beta_M,\beta_K,\gamma_H\right)\begin{bmatrix} 1 \\ \mathbb{X}_{H,it-1}\left(\sigma,\beta_M,\beta_K,\gamma_H\right)\end{bmatrix}\end{pmatrix} = \mathbf{0}, \tag{3.19}$$

with $\zeta^*_{L,it}(\cdot)$, $v^*_{it}(\cdot)$ and $\zeta^*_{h,it}(\cdot)$ being the residual functions corresponding to (3.7), (3.12) and (3.18), respectively, and where we have also included the $\beta_M + \beta_K = 1$ parameter normalization. Then, the familiar system GMM estimator of θ is a sample analogue of

$$\widehat{\theta} = \arg\min_{\theta}\ \mathbb{E}(\mathbf{h}_{it}(\theta))'\,\mathbf{W}\mathbb{E}(\mathbf{h}_{it}(\theta)), \tag{3.20}$$

with $\mathbf{W}$ being a conformable symmetric positive-definite moment-weighting matrix. The corresponding optimal asymptotic covariance of $\widehat{\theta}$ is $\left[\mathbb{E}\left(\frac{\partial \mathbf{h}_{it}(\theta)}{\partial \theta'}\right)\right]^{-1}\mathbb{E}\left(\mathbf{h}_{it}(\theta)\mathbf{h}_{it}(\theta)'\right)\left[\mathbb{E}\left(\frac{\partial \mathbf{h}_{it}(\theta)}{\partial \theta}\right)\right]^{-1}$ which can be robustified using usual off-the-shelf methods.

4 Simulations

We conduct simulation exercises to evaluate the finite-sample performance of our proposed methodology in recovering parameters of the structural production system. We consider a balanced panel of $N = \{49, 245, 490\}$ firms over $T = 10$ periods, and simulate each panel 500 times. The true production technology follows a CES form with two-dimensional firm productivity, as specified in equation (2.1). We set $\sigma = 2.5$, $\nu = 0.7$, $\beta_K = 0.2$ and $\beta_M = 0.8$.

For simplicity, we abstract away from the productivity shifters $G_{H,it-1}$ and $G_{L,it-1}$ and assume that the peer spillover network $\{s_{ij}\}_{i,j=1}^{N}$ is time-invariant and strictly exogenous. Thus, consistent with (2.4)–(2.5), the Hicks-neutral and labor-augmenting productivity components follow linear first-order spatiotemporal autoregressive processes:

$$\omega_{H,it} = \rho_0 + \rho_1\omega_{H,it-1} + \rho_2\sum_{j\neq i} s_{ij}\omega_{H,jt-1} + \zeta_{H,it}, \tag{4.1}$$

$$\omega_{L,it} = \gamma_1\omega_{L,it-1} + \gamma_2\sum_{j\neq i} s_{ij}\omega_{L,jt-1} + \zeta_{L,it}, \tag{4.2}$$

where $\rho_0 = 0.2$, $\rho_1 = 0.45$ and $\rho_2 = 0.35$ govern the evolution of $\omega_{H,it}$, while $\gamma_1 = 0.6$ and $\gamma_2 = 0.2$ govern $\omega_{L,it}$. Initial productivity levels $\omega_{H,i0}$ and $\omega_{L,i0}$ are drawn independently across firms from $\mathscr{U}(-1,1)$. The innovations are independently drawn as $\zeta_{H,it} \sim \mathcal{N}(0,\sigma_H^2)$ and $\zeta_{L,it} \sim \mathcal{N}(0,\sigma_L^2)$ across i and t with $\sigma_H = \sigma_L = 0.2$. The ex post production noise satisfies $v_{it} \sim$ i.i.d. $\mathcal{N}(0,\sigma_v^2)$ with $\sigma_v = \sqrt{0.07}$.

To define a cross-firm peer network, we use a fixed, non-stochastic spatial weighting matrix W_{49} from Anselin (1988) based on 49 districts in Columbus, Ohio, constructed using a contiguity-based "queen" definition of neighbors. For $N = 245$ and $N = 490$, we construct the corresponding peer weights as $W_{245} = I_5 \otimes W_{49}$ and $W_{490} = I_{10} \otimes W_{49}$. In line with our implicit restriction of the common spillover network definition for both factor-neutral and labor-augmenting productivity components in Assumption A3, the same peer weighting matrix is used for both $\omega_{H,it}$ and $\omega_{L,it}$, although they could be easily allowed to differ.

Table 1: Finite Sample Performance of the Proposed Estimator

	Case 1			Case 2			Case 3		
	$N=49$	$N=245$	$N=490$	$N=49$	$N=245$	$N=490$	$N=49$	$N=245$	$N=490$
σ									
RMSE	0.0880	0.0362	0.0268	0.1135	0.0522	0.0385	0.0883	0.0364	0.0270
MAE	0.0688	0.0291	0.0217	0.0899	0.0409	0.0310	0.0690	0.0293	0.0217
β_M									
RMSE	0.0429	0.0111	0.0078	0.0483	0.0207	0.0137	0.0499	0.0112	0.0078
MAE	0.0271	0.0089	0.0062	0.0280	0.0099	0.0068	0.0304	0.0089	0.0062
γ_1									
RMSE	0.0387	0.0157	0.0114	0.0368	0.0172	0.0114	0.0418	0.0164	0.0121
MAE	0.0309	0.0124	0.0092	0.0294	0.0135	0.0091	0.0330	0.0130	0.0097
γ_2									
RMSE	0.1154	0.0470	0.0323	0.1037	0.0480	0.0316	0.1364	0.0578	0.0397
MAE	0.0943	0.0369	0.0255	0.0832	0.0378	0.0248	0.1126	0.0451	0.0311
ν									
RMSE	0.0562	0.0248	0.0177	0.0539	0.0274	0.0186	0.0586	0.0255	0.0170
MAE	0.0419	0.0182	0.0141	0.0397	0.0210	0.0145	0.0429	0.0180	0.0134
η									
RMSE	0.7159	0.4332	0.3644	0.7279	0.4823	0.3205	0.7766	0.4283	0.3129
MAE	0.4020	0.2766	0.2386	0.4089	0.2871	0.2259	0.4493	0.2613	0.2191
ρ_0									
RMSE	0.0699	0.0238	0.0171	0.0677	0.0267	0.0193	0.0817	0.0243	0.0170
MAE	0.0489	0.0179	0.0136	0.0482	0.0194	0.0137	0.0547	0.0180	0.0135
ρ_1									
RMSE	0.0420	0.0175	0.0125	0.0420	0.0175	0.0126	0.0421	0.0175	0.0125
MAE	0.0335	0.0140	0.0098	0.0333	0.0139	0.0099	0.0336	0.0140	0.0098
ρ_2									
RMSE	0.0592	0.0262	0.0169	0.0579	0.0258	0.0169	0.0600	0.0260	0.0172
MAE	0.0468	0.0207	0.0137	0.0459	0.0202	0.0135	0.0470	0.0204	0.0138

Per Assumption A2, each firm's capital stock evolves according to $K_{it} = I_{it-1} + (1-d_i)K_{it-1}$, where the depreciation rate d_i is drawn uniformly from $\{0.05, 0.075, 0.10, 0.125, 0.15\}$. Initial capital satisfies $K_{i0} \sim$ i.i.d. $\mathscr{U}(10, 200)$, and gross investment is given by $I_{it-1} = K_{it-1}^{\beta_1} e^{\beta_2 \omega_{H,it-1} + \beta_3 \omega_{L,it-1}}$, with $\beta_1 = 0.8$ and $\beta_2 = \beta_3 = 0.1$.

The material and labor input series are obtained for each firm and time period by numerically solving the firm's static first-order conditions in equation (3.2) after having generated $(K_{it}, \omega_{H,it}, \omega_{L,it})$ and prices, whose data-generating processes (DGPs) are described below.

We generate firm-level output prices using the inverse demand function in equation (2.7), setting $\eta = -5$. For simplicity, we dispense with observable demand shock χ_{it} and generate the unobserved ex post demand shock as $\xi_{it} \sim$ i.i.d. $\mathcal{N}(0, \sigma_\xi^2)$ with $\sigma_\xi = 0.05$. The aggregate demand shifter is defined as $Y_t = N^{-1} \sum_i Y_{it}$. We fix $P_t = 1$ for all t and normalize input prices by $\frac{P_t}{Y_t^{1/\eta}} \Lambda^*$, whereby they become relative to this industry-wide demand/price scaling factor.

We consider three alternative DGPs for input prices:

$$\text{Case 1:}\quad P^M_{it} = \frac{P_t}{Y_t^{1/\eta}}\Lambda^* e^{\epsilon^M_{it}}, \quad P^L_{it} = \frac{P_t}{Y_t^{1/\eta}}\Lambda^* e^{\epsilon^L_{it}}, \tag{4.3}$$

$$\text{Case 2:}\quad P^M_{t} = \frac{P_t}{Y_t^{1/\eta}}\Lambda^*, \quad P^L_{it} = \frac{P_t}{Y_t^{1/\eta}}\Lambda^* e^{\theta_L\omega_{L,it}+\epsilon^L_{it}}, \tag{4.4}$$

$$\text{Case 3:}\quad P^M_{it} = \frac{P_t}{Y_t^{1/\eta}}\Lambda^* e^{\epsilon^M_{it}}, \quad P^L_{it} = \frac{P_t}{Y_t^{1/\eta}}\Lambda^* e^{\theta_L\omega_{L,it}+\theta_H\omega_{H,it}+\epsilon^L_{it}}, \tag{4.5}$$

where $\theta_L = 0.2$ and $\theta_H = 0.1$. The price heterogeneity is independently drawn as $\epsilon^L_{it} \sim \mathcal{N}(0,\sigma^2_L)$ and $\epsilon^M_{it} \sim \mathcal{N}(0,\sigma^2_M)$ across i and t with $\sigma_{PL} = \sigma_{PM} = 0.2$. Case 1 represents a setting in which both input prices vary across firm-years solely due to exogenous unobservables and are therefore uncorrelated with firms' productivity. This represents a rather implausible scenario of purely exogenous cross-firm variation in input prices[10] that we aim to relax in our Assumption A4(i). Case 2 describes a setting in which firms face heterogeneous, firm-specific labor prices. As discussed in Section 2, this captures a labor market in which more productive firms attract better workers, pay higher wages or hire higher-quality labor that commands a premium, leading to a potential contemporaneous correlation between wages and labor-augmenting productivity. In contrast, the materials market remains perfectly competitive, with a single homogeneous, exogenous materials price across all firms at any given time. Case 3 allows for firm-specific input prices in both the materials and labor markets, with labor prices correlated with the firm's current latent productivity along both its dimensions. Materials prices remain heterogeneous through the idiosyncratic shock but are not directly linked to productivity.[11]

We implement the estimation procedure as described in Section 3 and obtain estimates of the key parameters of the production function and the productivity evolution process for each simulation replication. The results are summarized in Table 1, which reports the root mean squared error (RMSE) and mean absolute error (MAE) across these simulations. Both the RMSE and MAE decline rapidly as the sample size increases, consistent with the expected root-N convergence of the estimator. Overall, the simulation results indicate that our estimator performs well in finite samples and accurately recovers the true parameters, supporting the validity of our identification strategy for the proposed production system.

To illustrate the consequences of misspecifying the dimensionality and cross-firm dependence of productivity, we conduct two additional simulation exercises. In both exercises, the data continue to be generated from our proposed model, but estimation is carried out using alternative, more restrictive methodologies that are inconsistent with the true DGP. Because the purpose of this exercise is to isolate the consequences of productivity-process misspecification, we report results for Cases 1 and 3, which represent the two endpoints of the input-price designs considered above.

In the spirit of Doraszelski and Jaumandreu (2018), the first alternative model rules out cross-firm productivity transmission by imposing $\rho_2 = \gamma_2 = 0$. It therefore retains separate factor-neutral and labor-augmenting productivity components but assumes that each component evolves independently across firms. The second alternative model—in the tradition of Malikov and Zhao's (2023) conceptualization of firm heterogeneity—does away with heterogeneous labor-augmenting productivity by imposing $\omega_{L,it} = \overline{\omega}_L$ $\forall(i,t)$ for some constant $\overline{\omega}_L$, which is observationally equivalent to reinstating a constant CES distribution parameter on labor, $\beta_L \equiv e^{-\frac{1-\sigma}{\sigma}\overline{\omega}_L}$, in place of the unit nor-

[10]Of note, such a DGP, in theory, makes input prices valid instruments for static inputs correlated with unobserved persistent productivity, obviating a more structural "proxy variable" approach to the production function estimation.

[11]We also considered the intermediate case in which $P^M_{it} = P_tY_t^{-1/\eta}\Lambda^* e^{\theta_L\omega_{L,it}+\epsilon^M_{it}}$ and $P^L_{it} = P_tY_t^{-1/\eta}\Lambda^* e^{\theta_L\omega_{L,it}+\epsilon^L_{it}}$ and found the results to be similar to those in Case 3.

malization on effective labor adopted in (2.1), subject to the usual $\beta_K + \beta_L + \beta_M = 1$ normalization.[12] This model retains cross-firm dependence only in the remaining scalar, factor-neutral productivity process. Of note, this is not the normalization maintained in the true DGP, where the weight on *effective* labor is set to unity and $\beta_K + \beta_M = 1$, so that the three distribution parameters sum to two rather than to one. Because the CES distribution parameters are identified only up to a common scale, with any rescaling of the aggregator kernel absorbed by the constant of the ω_H process, the pseudo-true value against which $\widehat{\beta}_M$ is to be benchmarked under this alternative model is $\beta_M/(\beta_K + e^{-\frac{1-\sigma}{\sigma}\overline{\omega}_L} + \beta_M)$ rather than $\beta_M = 0.8$ itself. For our DGP, the former is approximately 0.4, since $\mathbb{E}(\omega_{L,it}) = 0$ implies a labor weight of approximately unity. The same rescaling propagates to the intercept of the factor-neutral productivity process, which is what absorbs it. Since scaling the CES aggregator kernel by a shifts $\omega_{H,it}$ by $\frac{\nu\sigma}{1-\sigma}\ln a$ while $\mathbb{E}(\omega_{H,it}) = \rho_0/(1-\rho_1-\rho_2)$, the benchmark for the intercept is correspondingly $\rho_0 - (1-\rho_1-\rho_2)\frac{\nu\sigma}{1-\sigma}\ln\big(\beta_K + e^{-\frac{1-\sigma}{\sigma}\overline{\omega}_L} + \beta_M\big)$, which for our DGP is approximately 0.36 rather than $\rho_0 = 0.2$ once (ρ_1, ρ_2) are held at their DGP values. These restrictions also imply that the two alternative models contain different sets of parameters to be estimated; Table 2 reports the RMSE and MAE for these parameters.

The results show that imposing an incorrect productivity structure can have substantial consequences. Under Alternative Model 1, the estimates of σ and β_M remain relatively accurate, with their RMSEs and MAEs declining markedly as N increases. Conceivably, this is because the lagged terms enter (3.7) only through the composite index $(m_{it-1} - l_{it-1}) + \sigma\big(p^M_{it-1} - p^L_{it-1}\big)$, which by (3.6) equals $(1-\sigma)\omega_{L,it-1} + \sigma\ln\beta_M$, entering both in own-firm form and as a peer average, so that imposing $\gamma_2 = 0$ deletes the latter and retains the former. The deletion is absorbed by $\widehat{\gamma}_1$, whose pseudo-true value differs from γ_1 by approximately γ_2 times the coefficient from the projection of lagged peer-average labor-augmenting productivity on its own-firm counterpart, whereas σ, identified by the ratio of the two coefficients within the index, and β_M, pinned down by the unconditional-mean moment, are left undisturbed. Panel A demonstrates this, whereby γ_1 is the one first-step parameter whose errors cease to decline with the sample size, stalling at roughly 0.03 with its RMSE and MAE nearly coinciding, which is the signature of an estimator concentrating on a pseudo-true value rather than on the truth. The same signature appears, at far greater magnitude, among the parameters of the revenue-production equation and the factor-neutral productivity process, whose errors show essentially no tendency to vanish as the cross-section grows. In particular, the RMSEs and MAEs of ν, η and ρ_0 remain large and nearly constant across all three sample sizes, while those of ρ_1 likewise fail to decline. This pattern contrasts sharply with Table 1, where the corresponding errors under the correctly specified model generally fall substantially as N increases. Thus, when the true productivity processes contain peer spillovers, imposing $\rho_2 = \gamma_2 = 0$ does more than preclude recovery of the spillover effects themselves: the omitted cross-firm dependence is absorbed by other technology, demand and productivity-process parameters, generating distortions that persist even in substantially larger cross-sections.

The consequences of misspecification take a different form under Alternative Model 2, which collapses heterogeneous labor-augmenting productivity into a constant CES distribution parameter. Unlike Alternative Model 1, the errors for σ, ν, ρ_1 and ρ_2 decline at approximately the same rate as under the correctly specified estimator in Table 1, so the omitted productivity dimension appears to operate on these parameters by inflating sampling variability rather than by relocating the estimator onto a pseudo-true value. The inflation is nonetheless severe, e.g., at $N = 490$ the RMSE of the spillover coefficient ρ_2 is seven to nine times its counterpart in Table 1. Systematic distor-

[12]We adapt the Malikov and Zhao (2023) two-step estimation procedure based on variable input share equations to the CES specification.

Table 2: Finite-Sample Performance under Misspecified Alternative Models

		Case 1			Case 3		
		$N = 49$	$N = 245$	$N = 490$	$N = 49$	$N = 245$	$N = 490$
Panel A: Alternative Model 1—No Cross-Firm Productivity Spillovers							
σ	RMSE	0.0858	0.0355	0.0260	0.0857	0.0357	0.0263
	MAE	0.0666	0.0283	0.0209	0.0663	0.0287	0.0210
β_M	RMSE	0.0221	0.0094	0.0068	0.0215	0.0092	0.0066
	MAE	0.0173	0.0076	0.0054	0.0168	0.0074	0.0052
γ_1	RMSE	0.0428	0.0328	0.0309	0.0428	0.0328	0.0309
	MAE	0.0338	0.0288	0.0285	0.0339	0.0288	0.0285
ν	RMSE	0.2484	0.2507	0.2506	0.2480	0.2512	0.2506
	MAE	0.2450	0.2497	0.2502	0.2444	0.2504	0.2501
η	RMSE	2.5794	2.5721	2.5764	2.5773	2.5735	2.5755
	MAE	2.5592	2.5680	2.5741	2.5574	2.5699	2.5733
ρ_0	RMSE	0.3048	0.2969	0.2967	0.3006	0.2946	0.2935
	MAE	0.2871	0.2929	0.2947	0.2826	0.2910	0.2914
ρ_1	RMSE	0.0777	0.0800	0.0784	0.0780	0.0794	0.0786
	MAE	0.0655	0.0768	0.0771	0.0660	0.0765	0.0773
Panel B: Alternative Model 2—Scalar Hicks-Neutral Productivity							
σ	RMSE	0.2232	0.1033	0.0774	0.2174	0.1002	0.0734
	MAE	0.1657	0.0820	0.0600	0.1632	0.0799	0.0572
β_M	RMSE	0.2393	0.1620	0.1290	0.2327	0.1472	0.1130
	MAE	0.1962	0.1276	0.1044	0.1902	0.1147	0.0899
ν	RMSE	0.2355	0.1360	0.0829	0.2270	0.1173	0.0673
	MAE	0.1691	0.0735	0.0402	0.1587	0.0608	0.0314
η	RMSE	0.5110	0.3224	0.3191	0.4851	0.3297	0.3126
	MAE	0.3445	0.2508	0.2463	0.3390	0.2469	0.2415
ρ_0	RMSE	0.5719	0.4680	0.4190	0.5831	0.4651	0.3918
	MAE	0.4534	0.3850	0.3284	0.4648	0.3775	0.3039
ρ_1	RMSE	0.1398	0.0642	0.0421	0.1327	0.0607	0.0369
	MAE	0.1092	0.0469	0.0300	0.1018	0.0432	0.0269
ρ_2	RMSE	0.4393	0.2652	0.1458	0.4385	0.2295	0.1167
	MAE	0.2992	0.1294	0.0590	0.2919	0.1062	0.0462

Notes: The data are generated from the proposed model, whereas the two alternative models impose restrictions that are inconsistent with the true DGP. Alternative Model 1 suppresses both productivity-spillover coefficients by imposing $\rho_2 = \gamma_2 = 0$. Alternative Model 2 suppresses firm-specific labor-augmenting productivity by imposing $\omega_{L,it} = \beta_L$ and $\beta_K + \beta_L + \beta_M = 1$. Each design uses $T = 10$ and 500 Monte Carlo replications.

tion is instead concentrated in β_M and ρ_0, whose errors flatten rather than continue to decline and which must be assessed against the rescaled benchmarks implied by the alternative normalization as discussed earlier. So measured, the RMSE of $\widehat{\beta}_M$ at $N = 490$ is about thirty percent of the parameter and that of $\widehat{\rho}_0$ exceeds the parameter outright, against relative errors of roughly one and nine percent under our proposed estimator. Suppressing heterogeneous labor-augmenting productivity therefore leaves the omitted component to act as a persistent, spatially correlated error that biases the CES weights and the intercept of the factor-neutral process while also degrading the precision with which the surviving spillover parameter is recovered, even though cross-firm dependence is itself permitted.

Taken together with Table 1, the simulations show that our proposed estimator accurately recovers the structural parameters with errors that generally shrink as the cross section grows, whereas restricting either the cross-firm dependence or the dimensionality of productivity generates distortions in parameters beyond those directly restricted by the misspecified model.

5 Empirical application

We apply our methodology to firm-level data from China's primary agricultural food processing industry (officially "Agricultural and Sideline Food Processing," SIC two-digit code 13) over 1998–2007.[13] This industry exhibits pronounced spatial clustering and industrial agglomeration, primarily near coastal regions; see Figure 1(a). During the sample period, it was one of the largest—by employment and gross output—manufacturing sectors in China. The dense spatial concentration and a labor-intensive production process make it especially well-suited for studying factor-biased productivity spillovers.

We use FDI to define the productivity shifter $G_{H,it}$ and $G_{L,it}$ for both Hicks-neutral and labor-augmenting components of persistent firm productivity. FDI is a natural candidate given the important role foreign investment has played in China's economic growth and the government's long-standing preferential policies toward foreign investors. A large body of literature documents that foreign ownership can enhance firm productivity through multiple channels, including technology transfer, organizational restructuring, improved management practices and integration into global markets (e.g., Arnold and Javorcik, 2009; Keller and Yeaple, 2009; Guadalupe et al., 2012; Javorcik and Poelhekke, 2017). Yet, this literature typically treats productivity as a factor-neutral scalar, leaving the effects of FDI on distinct productivity components understudied. Our framework allows foreign ownership to differentially affect factor-neutral and labor-augmenting productivity, while separating the direct effect of a firm's own foreign equity from the indirect spillover effect transmitted through peer networks. The distinction is economically substantive because foreign ownership need not shift the two components together or in equal measure. Its productivity gains may be predominantly factor-neutral, predominantly labor-augmenting or some mix of the two, and a scalar-productivity framework cannot reveal which pattern prevails. While our framework cannot attribute productivity changes to particular mechanisms such as technology transfer or workforce training, it is able to identify the factor bias of FDI's overall effect on productivity.

In what follows, we first describe the data and variable construction, and then present the baseline estimates of the production function, productivity evolution processes and the demand function,

[13] It covers processing activities that directly use agriculture, forestry, animal-husbandry and fishery products as raw materials, including grain milling, feed processing, vegetable-oil and sugar processing, slaughtering and meat processing, aquatic-product processing, and processing of vegetables, fruits and nuts.

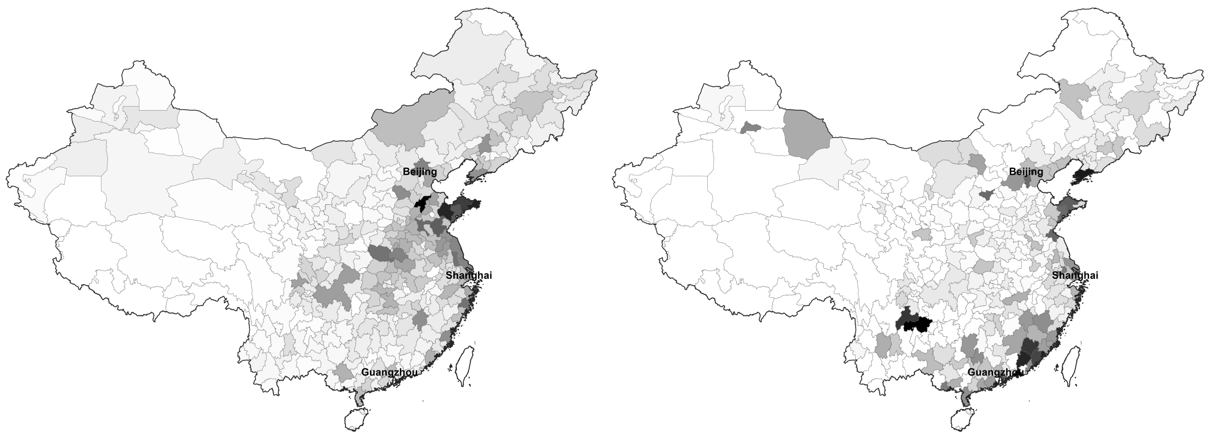

(a) (b)

Figure 1: Geographic Distribution of Firms and Foreign Equity
Notes: Panel (a) shows the number of sample firms by county; panel (b) shows the county-level mean foreign equity share (HMT & OFC). Darker shading indicates higher values.

followed by analysis of the dynamics of both productivity components in the industry. Next, we examine the sensitivity of the productivity spillover estimates—the object of main interest—to alternative peer-group definitions, weighting schemes and the inclusion of group-level fixed effects.

5.1 Data and variable construction

Our data are drawn from the Chinese Industrial Enterprises Database, a survey conducted by China's National Bureau of Statistics (NBS) that covers all manufacturing, mining, and public utility sectors. The survey includes all state-owned enterprises (SOEs) and non-SOEs with annual sales exceeding 5 million RMB (approximately 0.7 million USD).

Output, factors and prices. We construct input and output variables following methodologies established in prior studies using Chinese firm-level data (e.g., Brandt et al., 2017; Zhao et al., 2020; Malikov et al., 2020; Malikov and Zhao, 2023). Revenue R_{it} is defined as the gross industrial output value, and its real value is obtained by deflating it using the industry-level producer price index. Labor input L_{it} is measured by the number of employees. Materials input M_{it} is defined as total expenditure on intermediate goods, including raw materials and other production-related inputs, deflated by the purchasing price index for industrial inputs. Capital stock K_{it} is constructed via the perpetual inventory method following Brandt et al. (2012). All price indices are sourced from the NBS and the World Bank. Except for labor, all input and output variables are expressed in thousands of real RMB.

Firm-specific labor prices P^L_{it} are calculated by dividing the total wage bill (including benefits) by the number of employees, which furnishes the heterogeneous and potentially productivity-correlated wage that Assumption A4(i) is meant to accommodate. The materials price is less tractable, however. Because the data report only total expenditure on intermediate goods rather than the underlying quantities, as is typical of firm-level manufacturing datasets, a firm-specific P^M_{it} cannot be recovered, and we instead proxy it with the purchasing price index for industrial inputs, which varies over time but not across firms. This standard treatment (see, e.g., Levinsohn and Petrin, 2003; Zhao et al., 2020; Malikov et al., 2020; Malikov and Zhao, 2023) effectively replaces the firm-specific P^M_{it}

permitted by our model with a homogeneous P_t^M, thereby imposing for materials the special case of price homogeneity that our model otherwise relaxes. This concession is largely innocuous for our purposes[14] because the price-productivity correlation that is economically first-order, and that our framework is built to accommodate, arguably operates through the wage: labor-augmenting productivity scales effective labor and is the component most plausibly reflected in firm-specific input costs.

Following De Loecker (2011), the log aggregate demand shifter y_t is calculated as the market-share-weighted average of log real revenues within each of the 15 four-digit subindustries of the two-digit sector we study. For the firm-specific demand shock χ_{it} we consider the log of total assets, which captures observable differences in firm scale that may shift the selling price through distribution reach, cost advantage or brand recognition.

Productivity shifters. We follow the standard practice of measuring FDI exposure by the share of foreign equity in a firm's ownership structure. We further distinguish between equity investments from different foreign sources, namely, the share of equity provided by investors from Hong Kong, Macao and Taiwan (HMT) versus that from other foreign countries (OFC), since these sources have been shown to have distinct effects on firm performance in China (see Huang et al., 2013; Lee et al., 2023). Investments from the former played a particularly prominent role in China's manufacturing transformation during this period: HMT capital was among the earliest sources of foreign investment, flowed disproportionately into labor-intensive coastal industries and often arrived bundled with managerial and workforce practices transferred through ethnic Chinese business networks (Huang et al., 2013). OFC investment, in contrast, more often originated from Western or Japanese multinationals with weaker ties to local business communities. This distinction motivates our decision to estimate the productivity effects of HMT and OFC foreign equity separately rather than pooling them into a single foreign-ownership measure.

Figure 1(b) maps the county-level mean foreign equity share (HMT and OFC combined) for the industry. Comparing panels (a) and (b), both the firm count and foreign equity are concentrated along the eastern seaboard, but their geographic distributions do not coincide closely. Firm density is broadly concentrated in the major coastal manufacturing regions, with particularly dense clusters around the Bohai Rim, the Yangtze River Delta, and the southeastern coast, whereas high foreign-equity shares are more unevenly distributed across localities and are especially prominent in parts of the southeastern coastal region. Thus, even within the broadly more industrialized eastern part of the country, areas with the greatest concentration of firms are not necessarily those with the highest foreign-equity shares. This partial decoupling of agglomeration and foreign presence provides useful variation for separately identifying the spillover and direct learning channels of FDI.

Figures B1(a) and B1(b) in the appendix show the distribution of foreign equity share across HMT- and OFC-invested firms (i.e., $\text{HMT}_{it} \mid \text{HMT}_{it} > 0$ and $\text{OFC}_{it} \mid \text{OFC}_{it} > 0$). Both conditional distributions are slightly right-skewed and dispersed throughout the interior of the unit interval, giving the foreign-equity shares the continuous, graded variation that our estimation exploits in treating FDI as an intensity of foreign ownership rather than a discrete indicator of its presence.

To address potential reporting errors and outliers, we apply the following data-cleaning procedures: (i) exclude observations with missing values for key variables or incomplete postcode information (i.e., postal codes not exactly six digits); (ii) trim observations outside the 0.5th and 99.5th percentiles for production variables, total assets, labor prices and the labor-to-material ratio; (iii)

[14] Because p_{it}^M enters the first-step equation (3.7)–(3.8) only through the material-to-labor price $(p_{it}^M - p_{it}^L)$, treating it as the common index p_t^M leaves the firm-specific wage to carry the cross-firm variation in this relative price.

Table 3: Data Summary Statistics

	Mean	Q1	Median	Q3
Revenue ($\widetilde{R}$)	42,628	9,819	20,940	48,030
Capital (K)	9,118	1,628	3,813	9,307
Labor (L)	122	38	70	140
Materials (M)	29,240	6,644	14,239	32,559
Total Assets (TA)	23,275	4,690	10,000	24,162
Labor Price (P_L)	11.121	6.000	8.900	13.232
$\mathbb{1}\{\text{HMT} > 0\}$	0.0230			
HMT | HMT > 0	0.4162	0.1442	0.3420	0.6521
$\mathbb{1}\{\text{OFC} > 0\}$	0.0450			
OFC | OFC > 0	0.4422	0.1827	0.3983	0.6843

Notes: Q1 and Q3 denote the first and third quartiles, respectively; $\widetilde{R}$ is the deflated revenue; HMT and OFC denote equity shares held by investors from Hong Kong, Macao, and Taiwan and from other foreign countries, respectively. $\mathbb{1}\{\text{HMT} > 0\}$ and $\mathbb{1}\{\text{OFC} > 0\}$ are indicators for positive HMT and OFC equity shares, respectively; and HMT | HMT > 0 and OFC | OFC > 0 are the corresponding equity shares conditional on being positive. Monetary variables are expressed in thousands of real RMB.

exclude observations with labor-to-output or material-to-output ratios, or foreign equity shares, outside the unit interval; and (iv) exclude firms with fewer than eight employees, consistent with the dataset's focus on "above-scale" enterprises. We further restrict the sample to non-state-owned firms excluding SOEs, whose distinct objectives and institutional constraints make their production decisions difficult to reconcile with our profit-maximization framework. In addition, following Lee et al. (2023), we exclude firms that switch ownership status more than once during the sample period.[15] The resulting sample comprises an unbalanced panel of 19,803 firms and 71,844 firm-year observations.

Table 3 reports summary statistics for the key variables. The distributions of firm size and production variables are noticeably right-skewed, with mean revenue, capital, labor, materials, and total assets all exceeding their corresponding medians. The median firm employs 70 workers and faces a labor price of about 8.9 thousand real RMB per worker. In the sample period, 2.3% of firm-year observations have positive HMT equity, while 4.5% have positive OFC equity. Conditional on positive foreign ownership, however, the foreign equity stakes are substantial, averaging 41.6% for HMT and 44.2% for OFC, with median shares of 34.2% and 39.8%, respectively.

5.2 Baseline results

Before proceeding to the results, we first define the productivity measures of direct and indirect learning from FDI. The parametric structure of the productivity processes (2.4) and (2.5) yields closed-form measures of these effects. Throughout, $G_{H,it} = G_{L,it}$ is a 2×1 vector collecting the HMT

[15] The dataset reports firms' equity shares by source: state, HMT, OFC and private. We use the state, HMT and OFC shares to classify firms into four mutually exclusive ownership types: HMT-owned, OFC-owned, state-owned and domestic-private. A firm is classified as HMT-owned, OFC-owned or state-owned if at least one of these three shares exceeds 25%, with the firm assigned to the category corresponding to the largest share; all remaining firms are classified as domestic-private.

and OFC equity shares, so the related derivative-based measures defined below are likewise 2×1 vectors, with each element corresponding to a distinct source of foreign equity.

First, the *direct learning* (DL) effect measures the marginal impact of a firm's own foreign equity on its next-period productivity:

$$DL_{H,it} = \frac{\partial \mathbb{E}[\omega_{H,it} \mid \cdot]}{\partial G_{H,it-1}} = \rho_3, \qquad DL_{L,it} = \frac{\partial \mathbb{E}[\omega_{L,it} \mid \cdot]}{\partial G_{L,it-1}} = \gamma_3. \tag{5.1}$$

This captures the channel through which foreign ownership enhances productivity via technology transfer, knowledge acquisition or organizational improvements originating from the firm's own foreign investors. Second, the productivity *spillover* (SP) effect measures the strength of the network externality, that is, the marginal effect of peer average productivity on the firm's next-period productivity:

$$SP_{H,it} = \frac{\partial \mathbb{E}[\omega_{H,it} \mid \cdot]}{\partial \sum_{j(\neq i)} s_{ijt-1}\omega_{H,jt-1}} = \rho_2, \qquad SP_{L,it} = \frac{\partial \mathbb{E}[\omega_{L,it} \mid \cdot]}{\partial \sum_{j(\neq i)} s_{ijt-1}\omega_{L,jt-1}} = \gamma_2. \tag{5.2}$$

Combining the two channels via the chain rule, we define the *indirect learning* (IL) of firm i from firm j's foreign equity as

$$IL_{H,ijt} = \frac{\partial \mathbb{E}[\omega_{H,it} \mid \cdot]}{\partial G_{H,jt-2}} = SP_{H,it} \times s_{ijt-1} \times DL_{H,jt-1}, \tag{5.3}$$

$$IL_{L,ijt} = \frac{\partial \mathbb{E}[\omega_{L,it} \mid \cdot]}{\partial G_{L,jt-2}} = SP_{L,it} \times s_{ijt-1} \times DL_{L,jt-1}. \tag{5.4}$$

This captures external spillovers from FDI, whereby a firm's productivity-altering foreign investment generates secondary effects on other firms within its peer network. Aggregating over all peers and using the fact that DL is constant across firms (so that $\sum_{j\neq i} s_{ijt-1}DL_{jt-1} = DL_{jt-1}$ under row-standardized weights), we obtain the *total indirect learning* (TIL) effect for each firm:

$$TIL_{H,it} = \sum_{j(\neq i)} IL_{H,ijt} = \rho_2\rho_3, \qquad TIL_{L,it} = \sum_{j(\neq i)} IL_{L,ijt} = \gamma_2\gamma_3. \tag{5.5}$$

Because the spatiotemporally autoregressive productivity processes are linear, DL, SP and TIL effects are constant and common across firms.

5.2.1 Point estimates

Table 4 reports the baseline parameter estimates. In this specification, a firm's peer group $\mathscr{L}(i,t)$ is defined to consist of all firms located in the same city and operating in the same four-digit subindustry, with uniform row-standardized weights as defined in (2.6). To absorb common secular technical change, we include year fixed effects in both productivity evolution equations.

The estimated elasticity of substitution statistically exceeds one, suggesting that production inputs are gross substitutes in the studied industry. Since the CES collapses to Cobb-Douglas at $\sigma = 1$, where labor-augmenting heterogeneity is subsumed into the Hicks-neutral residual and the two components are observationally equivalent, the rejection of unit substitution elasticity indicates that the model does not degenerate to this indistinguishable case, so the data are consistent with identifying the two productivity components separately rather than forcing their collapse. The estimate of the RTS parameter, $\widehat{\nu} = 1.454$, implies increasing returns, while the materials distribution parameter

Table 4: Parameter Estimates of the Production Function, Productivity, and Demand Function

Parameter	Estimate	Std. Error
Production Function		
Elasticity of substitution (σ)	2.772	0.035
Elasticity of scale (ν)	1.454	0.056
Distribution parameter (β_M)	0.889	0.031
Demand Function		
Price elasticity of demand (η)	–4.815	0.750
Total assets (ϱ)	–0.203	0.013
Hicks-Neutral Productivity Process		
Intercept (ρ_0)	–0.501	0.131
Own-persistence coefficient (ρ_1)	0.748	0.008
Cross-peer spillover SP_H coefficient (ρ_2)	0.021	0.006
$DL_{H,HMT}$ ($\rho_{3,1}$)	0.182	0.015
$DL_{H,OFC}$ ($\rho_{3,2}$)	0.208	0.012
$TIL_{H,HMT}$ ($\rho_2 \times \rho_{3,1}$)	0.004	0.001
$TIL_{H,OFC}$ ($\rho_2 \times \rho_{3,2}$)	0.004	0.001
Labor-Augmenting Productivity Process		
Own-persistence coefficient (γ_1)	0.499	0.006
Cross-peer spillover SP_L coefficient (γ_2)	0.204	0.007
$DL_{L,HMT}$ ($\gamma_{3,1}$)	0.640	0.047
$DL_{L,OFC}$ ($\gamma_{3,2}$)	0.238	0.034
$TIL_{L,HMT}$ ($\gamma_2 \times \gamma_{3,1}$)	0.130	0.011
$TIL_{L,OFC}$ ($\gamma_2 \times \gamma_{3,2}$)	0.048	0.007

Notes: Reported are the results based on our baseline specification of the productivity processes, in which each firm's peer group is restricted to firms in the same city and the same 4-digit subindustry and the secular technical change is modeled using year effects. Asymptotic cluster-robust standard errors are in parentheses. The standard error of TIL is calculated using delta method.

$\widehat{\beta}_M = 0.889$ indicates that materials enter the CES aggregator with a weight comparable to that of labor, whose distribution parameter is normalized to unity. On the demand side, the price elasticity $\widehat{\eta} = -4.815$ implies a markup of roughly $\frac{\widehat{\eta}}{1+\widehat{\eta}} = 1.26$, and the negative coefficient on the demand-shifting total assets ($\widehat{\varrho} = -0.203$) suggests that at a given supply quantity larger firms command lower prices, consistent with their occupying a lower-price, mass-market position.

Turning to the productivity evolution processes, the own-persistence coefficients are $\widehat{\rho}_1 = 0.748$ and $\widehat{\gamma}_1 = 0.499$, indicating that firm productivity is highly persistent over time and that the factor-neutral productivity is more so than the labor-augmenting one. The most striking contrast between the two productivity dimensions emerges in the spillover parameters, both of which are positive and statistically significant. The Hicks-neutral spillover is economically modest ($\widehat{SP}_H = \widehat{\rho}_2 = 0.021$), whereas the labor-augmenting spillover is an order of magnitude larger ($\widehat{SP}_L = \widehat{\gamma}_2 = 0.204$).[16]

[16] Of note, the temporal and spatial autoregressive lag parameter estimates satisfy the stability conditions for both components ($\widehat{\rho}_1 + \widehat{\rho}_2 = 0.769$ and $\widehat{\gamma}_1 + \widehat{\gamma}_2 = 0.703$, each below unity), ensuring that the implied spatiotemporal processes are non-explosive.

This asymmetry suggests that the productivity externalities flowing through local peer networks in China's primary agricultural food processing industry are biased toward labor, operating predominantly through the labor-augmenting component of firm heterogeneity and plausibly reflecting the localized diffusion of workforce practices, training methods and process know-how rather than factor-neutral technological shifts. A scalar, Hicks-neutral specification would have folded these two dimensions into a single spillover estimate, averaging the strong labor-augmenting transmission against the negligible neutral one and thereby concealing the uncovered factor bias that motivates our framework. To interpret the magnitudes, a 1% increase in the average labor-augmenting productivity of a firm's peers is expected to raise its labor-augmenting productivity next period by roughly 0.2%, while the analogous spillover effect in the Hicks-neutral productivity is only 0.02%.

Read together, the own-persistence and spillover asymmetries point to a common mechanism. The component a firm retains most strongly over time, factor-neutral productivity, is the one that diffuses least to its peers, whereas the less persistent labor-augmenting component is the one that propagates most strongly across firms. A plausible interpretation is that both margins are governed by the transferability of the underlying carrier of each kind of heterogeneity. Labor-augmenting productivity is carried by workers, training and managerial practices that move across firms through hiring and the circulation of personnel, so that the very mobility preventing any single firm from retaining this advantage ($\widehat{\gamma}_1 = 0.499$) is what transmits it to nearby peers ($\widehat{\gamma}_2 = 0.204$). Factor-neutral productivity, by contrast, reflects more appropriable and firm-level advantages rooted in proprietary technology and organizational capital, which are retained internally better ($\widehat{\rho}_1 = 0.748$) precisely because they do not travel easily to competitors ($\widehat{\rho}_2 = 0.021$). On this reading, internal persistence and external diffusion are two sides of the same trade-off, jointly shaped by how mobile the carrier of each component is, and the localized labor-biased pattern that we find comports with the agglomeration literature's emphasis on the local circulation of tacit, person-embodied knowledge (Marshall, 1890; Jaffe et al., 1993; Gertler, 2003; Storper and Venables, 2004; Almeida and Kogut, 1999; Combes and Duranton, 2006).

Consider next the productivity effects of FDI, beginning with the within-firm direct learning. The estimates in Table 4 indicate that foreign equity of either origin enhances both components of firm productivity, with all four direct-learning coefficients positive and significant. Interpreting the factor bias of these effects requires care, however, because the two coefficients associated with each source are not output-equivalent. As formalized in Section 2, a shift in $\omega_{H,it}$ scales output radially and leaves relative marginal products unchanged, so the component of learning captured by ρ_3 is Hicks-neutral irrespective of its magnitude, whereas *any* nonzero effect on $\omega_{L,it}$ alters the marginal rate of technical substitution between labor and the other inputs (so long as $\sigma \neq 1$) and thereby constitutes factor-biased technical change. Factor neutrality of FDI's productivity effect therefore corresponds to the restriction $\gamma_3 = 0$ rather than to the closeness of γ_3 to ρ_3. Moreover, the comparison of raw coefficients would mix objects that are not output-equivalent, since equal proportional shifts in $e^{\omega_{H,it}}$ and $e^{\omega_{L,it}}$ raise output one-for-one and in proportion to the output elasticity of labor, respectively (a point we formalize below). In gauging the factor bias of direct learning, the meaningful magnitude comparisons accordingly run across investor origins within each productivity margin rather than across the two margins within an FDI source.

With this distinction in mind, our estimates imply that the direct productivity gains from foreign equity are factor-biased for both investor origins, with the strength of the bias differing sharply by source. Along the factor-neutral margin, the two origins are comparable, with $\widehat{DL}_{H,HMT} = \widehat{\rho}_{3,1} = 0.182$ and $\widehat{DL}_{H,OFC} = \widehat{\rho}_{3,2} = 0.208$ implying that a one-percentage-point increase in the foreign equity share raises Hicks-neutral productivity by roughly 0.2% regardless of the capital's origin, whereas along the factor-biased margin the labor-augmenting effect of HMT equity ($\widehat{DL}_{L,HMT} = \widehat{\gamma}_{3,1} = 0.640$) is nearly three times that of OFC equity ($\widehat{DL}_{L,OFC} = \widehat{\gamma}_{3,2} = 0.238$), or about 0.64% against 0.24% per percentage point of equity. The difference between the two sources of foreign capital is therefore concentrated almost entirely in the non-neutral component of their productivity effects. Because the estimated elasticity of substitution exceeds unity, these labor-augmenting gains raise the marginal product of labor relative to those of capital and materials at given input quantities and are thus labor-biased in the Hicks sense. The implied tilt in the marginal rate of technical substitution between labor and either non-labor input is $\frac{\widehat{\sigma}-1}{\widehat{\sigma}}\widehat{\gamma}_3$ log points per unit increase in the equity share, or roughly 0.41% and 0.15% per percentage point of HMT and OFC equity, respectively.[17] This source asymmetry aligns with the institutional differences noted earlier, in that the managerial and workforce practices bundled with HMT capital plausibly load more heavily on the labor-augmenting margin than do the formal technology licensing and capital-embodied transfer through which arm's-length OFC investment operates.

The total indirect learning estimates combine the spillover and direct-learning channels. Consistent with the dominance of labor-augmenting productivity spillovers, $\widehat{TIL}_{L,HMT} = 0.130$ and $\widehat{TIL}_{L,OFC} = 0.048$ are substantially larger than their Hicks-neutral counterparts ($\widehat{TIL}_{H,HMT} = \widehat{TIL}_{H,OFC} = 0.004$). That is, the indirect productivity gains that domestic firms derive from nearby foreign-invested peers are concentrated almost entirely in the labor-augmenting component. Quantitatively, a uniform one-percentage-point increase in the HMT (OFC) equity share among a firm's peers raises its own labor-augmenting productivity by roughly 0.13% (0.05%) through the indirect learning channel, which is about one-fifth the size of the corresponding direct learning effect. Indirect gains in factor-neutral productivity are an order of magnitude smaller, reflecting the combination of a weak Hicks-neutral spillover and more moderate direct learning effects.

Of note, both DL and TIL capture only first-round effects of FDI: DL measures how a firm's own foreign equity shifts its productivity in the following period, whereas TIL measures how its peers' foreign equity shifts the firm's productivity with a two-period lag through the spillover channel. They abstract from the iterated propagation—over time and across space—that follows, whereby once a firm's productivity rises, that increase feeds back into its own next-period productivity via persistence and into its peers' productivity via the spillover diffusion, and so on indefinitely. Analogous to "long-run" multipliers in spatiotemporal autoregressive models, the full steady-state effect of a permanent, uniform increase in G_H and G_L across all firms can be obtained from the reduced form of the vector productivity process in Assumption A3. Exploiting the row-normalization of peer weights matrix and holding the network fixed, this multiplier simplifies to $1/(1-\gamma_1-\gamma_2)$ for the labor-augmenting productivity and $1/(1-\rho_1-\rho_2)$ for the factor-neutral productivity, with stability requiring $\gamma_1+\gamma_2 < 1$ and $\rho_1+\rho_2 < 1$. Evaluating at our baseline estimates yields long-run steady-state multipliers of approximately 3.4 and 4.3, respectively, indicating that the total productivity gains from FDI substantially exceed the first-round direct and indirect learning effects reported in Table 4. For each productivity component, the long-run amplification is governed by the sum of own-persistence and spillover, though its composition differs across components. Specifically, the factor-neutral multiplier ($\widehat{\rho}_1+\widehat{\rho}_2 = 0.748+0.021$) is driven almost entirely by internal persistence and is thus essentially a within-firm phenomenon, whereas the labor-augmenting multiplier ($\widehat{\gamma}_1+\widehat{\gamma}_2 = 0.499+0.204$) draws a substantially larger share of its amplification from the cross-firm spillover channel.

[17] Formally, the reported quantity is $\partial \ln\left(\frac{\partial F(\cdot)}{\partial L_{it}} \Big/ \frac{\partial F(\cdot)}{\partial Z_{it}}\right)/\partial G_{L,it-1} = \frac{\sigma-1}{\sigma}\gamma_3$ for either non-labor input $Z_{it} \in \{K_{it}, M_{it}\}$, evaluated at fixed input quantities. It is therefore the gap between the proportional responses of the two marginal products to foreign equity. From (3.3), $\omega_{L,it}$ enters the marginal products of capital and materials only through the aggregator X_{it}, which cancels in the ratio and renders the tilt the same against either non-labor input. The marginal products of all three inputs in fact rise, since $\widehat{\nu} > (\widehat{\sigma}-1)/\widehat{\sigma}$, so that the factor bias resides in the gap between two positive effects.

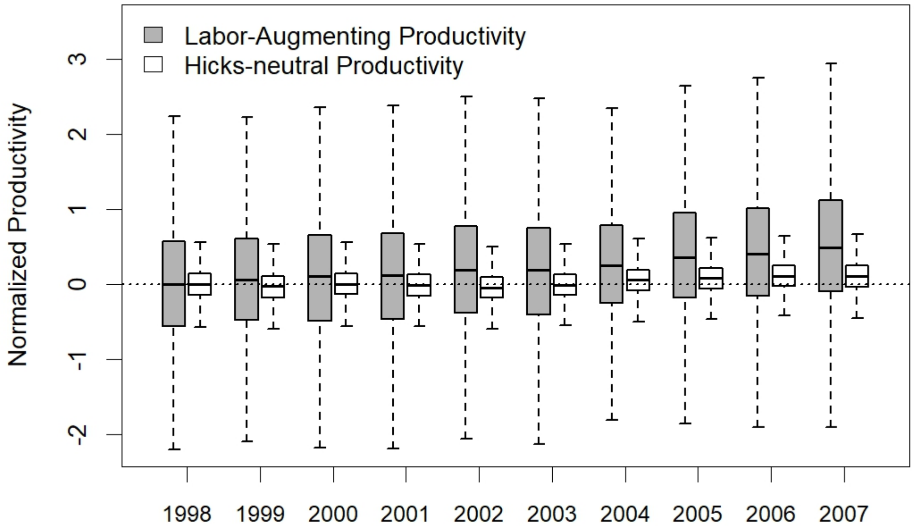

Figure 2: Cross-Firm Distributions of Labor-Augmenting and Hicks-Neutral Productivity
Notes: Each pair of boxplots shows the cross-sectional distribution of the two productivity components in a given year. Both measures are normalized to a median of zero in 1998.

5.2.2 Productivity dynamics

Figure 2 compares the cross-sectional distributions of the two productivity components over the sample period using side-by-side boxplots. For ease of comparison, both measures are normalized to a median of zero in 1998. Two features stand out. First, the two components exhibit markedly different dynamics. The distribution of labor-augmenting productivity drifts upward over the sample period, suggesting a sustained shift toward more labor-efficient production techniques in the industry. Hicks-neutral productivity, by contrast, remains comparatively stable, with its median hovering near zero and only a modest upward drift in more recent years. The raw magnitudes of the two components are not directly comparable, however, as ω_H and ω_L enter the production function differently (we revisit this point later). Second, labor-augmenting productivity exhibits substantially greater cross-firm dispersion than its Hicks-neutral counterpart throughout the sample period, indicating that firm productivity heterogeneity in this industry is concentrated primarily in the labor-augmenting dimension.

To better characterize aggregate productivity dynamics in the industry, we apply the popular Olley and Pakes (1996) decomposition, which separates industry-level aggregate productivity—a revenue-weighted average of firm-level productivity—into the simple (unweighted) cross-firm average and a covariance term capturing the extent to which more productive firms command larger market shares. The latter component reflects the efficiency of factor allocation across firms. We implement this decomposition for both productivity measures.[18] Figure 3 plots the weighted and unweighted average industry productivity series for both components over time, with the vertical

[18]Concretely, aggregate productivity levels are given by $\Omega_t = \sum_{i=1}^{N_t} \mu_{it}\Omega_{it} = \bar{\Omega}_t + \sum_{i=1}^{N_t}(\mu_{it} - \bar{\mu}_t)(\Omega_{it} - \bar{\Omega}_t)$ and $\Phi_t = \sum_{i=1}^{N_t} \mu_{it}\Phi_{it} = \bar{\Phi}_t + \sum_{i=1}^{N_t}(\mu_{it} - \bar{\mu}_t)(\Phi_{it} - \bar{\Phi}_t)$, where N_t denotes the number of firms in year t, μ_{it} is firm i's revenue share in period t, $\Omega_{it} = \exp(\omega_{H,it})$, $\Phi_{it} = \exp(\omega_{L,it})$, and $\bar{a}_t \equiv \frac{1}{N_t}\sum_{i=1}^{N_t} a_{it}$ for $a \in \{\Omega, \Phi, \mu\}$.

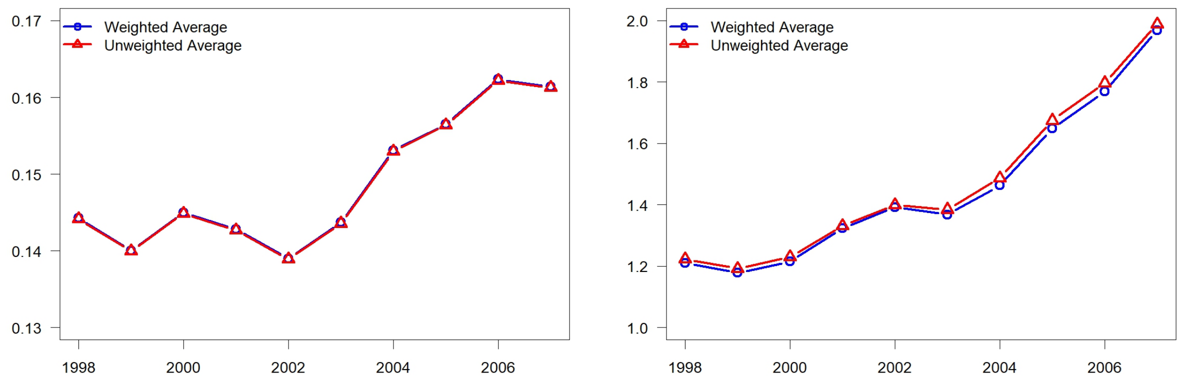

(a) Hicks-Neutral Productivity

(b) Labor-Augmenting Productivity

Figure 3: Aggregate Productivity Trends

Notes: Summarized productivity estimates are from our main model. Vertical distance between two plotted averages represents factor allocation across firms.

gap between the two corresponding to the allocation term.

Consistent with the findings above, growth trajectories of the two productivity components are not the same. Aggregate Hicks-neutral productivity rose modestly over the sample period, from approximately 0.14 to 0.16, an increase of roughly 14%. Aggregate labor-augmenting productivity, by contrast, grew substantially over the same period, from around 1.2 to nearly 2.0, an increase of roughly 67%. For both components, most of the acceleration occurred after 2003. However, this does not by itself imply that labor-augmenting progress is the dominant driver of aggregate output growth, since the contribution of ω_L to output depends on the output elasticity of labor in the CES aggregator.

Further, the weighted and unweighted series are nearly indistinguishable in both subfigures, indicating that the allocative-efficiency term is quantitatively small for both productivity components. This finding is consistent with prior evidence on Chinese manufacturing (see Malikov et al., 2020; Huang et al., 2022) documenting that the covariance between firm productivity and market share is modest in this period, suggesting that factor allocation across firms does not strongly favor the most productive producers. For Hicks-neutral productivity, the two series essentially coincide throughout the sample period, suggesting that firm-level Hicks-neutral productivity is largely uncorrelated with market share. For labor-augmenting productivity, the weighted average sits slightly below the unweighted average, implying a mildly negative covariance between ω_L and market share. This is suggestive of modest misallocation along the labor-augmenting dimension, whereby firms with higher labor-augmenting productivity command somewhat smaller revenue shares. The magnitude of this misallocation is small, however, and the dominant feature of the data is the strong upward trend in the simple average of labor-augmenting productivity itself.

A note on comparability is in order. Because ω_H enters the production function as a multiplicative shifter on output while ω_L augments only the labor input, the two components are not directly comparable in levels. From the production function, the (semi-)elasticity of output with respect to ω_L is

$$\frac{\partial \ln F(K_{it}, e^{\omega_{L,it}} L_{it}, M_{it})}{\partial \omega_{L,it}} = \frac{\partial \ln F(K_{it}, e^{\omega_{L,it}} L_{it}, M_{it})}{\partial \ln L_{it}} = \left(\frac{1}{1+1/\eta}\right) \frac{P^L_{it} L_{it}}{P_{it} Y_{it}}. \tag{5.6}$$

The first equality is a primal, purely technological identity. In the CES production function (2.1), $\omega_{L,it}$ enters through $\widetilde{L}_{it}$, and since $\ln \widetilde{L}_{it} = \omega_{L,it} + \ln L_{it}$, the labor-augmenting term and the log labor input are perfect substitutes inside the log effective-labor bundle.[19] The (semi-)elasticity of output with respect to $\omega_{L,it}$ therefore coincides with the output elasticity of labor, which from the marginal product in (3.3) equals $\partial \ln F(\cdot)/\partial \ln L_{it} = \nu\big[\widetilde{L}_{it}\big]^{-(1-\sigma)/\sigma}/X_{it}$. The second equality is the *dual* counterpart of this technological object. Substituting the firm's static first-order condition for labor (3.2)—which, under the monopolistic-competition demand (2.7) of Assumption A4 and the static-input timing of Assumption A2, equates the expected marginal revenue product of labor to the wage—and multiplying by L_{it}/Y_{it} re-expresses the same elasticity through observed prices and quantities. The result is the markup-adjusted labor revenue share, in which $1/(1+1/\eta)$ is the markup and $P^L_{it}L_{it}/(P_{it}Y_{it})$ is the labor share of revenue. This is the production-approach markup identity of De Loecker and Warzynski (2012) read in reverse, i.e., whereas they recover a firm's markup as the ratio of a flexible input's output elasticity to its revenue share, we invert the relationship, using the demand-side markup $\eta/(\eta+1)$ to recover the output elasticity of labor from its observed revenue share. The two elasticity formulas are equivalent under the model, with the dual form recoverable from the observed labor revenue share and the estimated demand elasticity η alone.

Evaluating at the sample mean, the markup-adjusted labor revenue share in this industry is approximately 0.055. It means that a 1% increase in e^{ω_L} raises output by roughly 0.055% at the mean, compared with the direct one-for-one effect of e^{ω_H}. This adjustment matters for interpreting the aggregate trends in Figure 3. The 67% cumulative growth in labor-augmenting productivity translates into roughly 3.66% of output growth, whereas the corresponding 14% growth in Hicks-neutral productivity translates directly into 14% of output growth. The bias of technical change in this industry is therefore strongly toward augmenting labor productivity, but in output-equivalent terms, Hicks-neutral progress dominates aggregate productivity growth over the sample period.

To assess the extent to which the two estimated productivity components capture distinct dimensions of firm technology, we examine their contemporaneous cross-sectional correlations, reported in Appendix Table B1. The annual Pearson correlations are uniformly positive but small, ranging from 0.005 in 2007 to 0.086 in 2000, with a pooled year-demeaned correlation of 0.051. The corresponding Spearman rank correlations are somewhat larger but remain modest, ranging from 0.029 to 0.114, with a pooled year-demeaned value of 0.077. Thus, firms with relatively high labor-augmenting productivity tend to have slightly higher Hicks-neutral productivity, but the association is economically weak. This limited comovement is consistent with the two measures capturing distinct underlying components of technology rather than different manifestations of a single latent productivity factor. It therefore provides additional empirical support for modeling firm productivity as multidimensional instead of summarizing technological heterogeneity with a single Hicks-neutral index.

5.3 Sensitivity checks

Of our main interest are the productivity spillovers across firms. Here, we assess the sensitivity of our spillover estimates along three dimensions: the treatment of peer-group-level unobservables, the definition of the peer network, and the construction of peer weights. The results are summarized in Table 5, with the baseline estimates from Table 4 reproduced in the first column for ease of comparison.

Group-level unobservables. Our baseline framework for identifying productivity spillovers, formalized in Assumption A3, rules out peer-group-level correlated effects, that is, common unobserv-

[19] Concretely, a one-percent increase in labor efficiency $e^{\omega_{L,it}}$ scales effective labor $\widetilde{L}_{it}$ by the same proportion as a one-percent increase in L_{it}.

Table 5: Estimates of the Productivity Effects: Sensitivity Analysis

Estimand	Baseline	F1	F2	W1	P1	P2	P3
				—Hicks-Neutral Productivity—			
SP_H	0.021	0.039	0.045	0.083	0.099	0.056	0.099
	(0.009, 0.032)	(0.025, 0.053)	(0.029, 0.061)	(0.054, 0.112)	(0.077, 0.120)	(0.035, 0.077)	(0.069, 0.129)
$DL_{H,HMT}$	0.182	0.032	0.062	0.173	0.154	0.232	–0.027
	(0.152, 0.212)	(0.009, 0.055)	(0.031, 0.092)	(0.124, 0.222)	(0.128, 0.180)	(0.196, 0.267)	(–0.047, –0.007)
$DL_{H,OFC}$	0.208	0.071	0.040	0.430	–0.042	0.798	–0.006
	(0.186, 0.231)	(0.056, 0.086)	(0.020, 0.061)	(0.382, 0.477)	(–0.055, –0.029)	(0.716, 0.879)	(–0.023, 0.011)
$TIL_{H,HMT}$	0.004	0.001	0.003	0.014	0.015	0.013	–0.003
	(0.002, 0.006)	(0.000, 0.002)	(0.001, 0.005)	(0.008, 0.021)	(0.011, 0.020)	(0.008, 0.018)	(–0.004, –0.001)
$TIL_{H,OFC}$	0.004	0.003	0.002	0.036	–0.004	0.045	–0.001
	(0.002, 0.007)	(0.002, 0.004)	(0.001, 0.003)	(0.023, 0.048)	(–0.006, –0.003)	(0.027, 0.062)	(–0.002, 0.001)
				—Labor-Augmenting Productivity—			
SP_L	0.204	0.186	0.196	0.205	0.307	0.128	0.332
	(0.190, 0.217)	(0.172, 0.199)	(0.183, 0.210)	(0.192, 0.218)	(0.290, 0.325)	(0.114, 0.142)	(0.301, 0.362)
$DL_{L,HMT}$	0.640	0.435	0.399	0.337	0.249	0.350	0.371
	(0.547, 0.732)	(0.345, 0.525)	(0.311, 0.487)	(0.253, 0.421)	(0.164, 0.334)	(0.267, 0.433)	(0.282, 0.461)
$DL_{L,OFC}$	0.238	0.395	0.407	0.321	0.411	0.672	0.504
	(0.171, 0.304)	(0.325, 0.464)	(0.338, 0.475)	(0.256, 0.386)	(0.342, 0.479)	(0.601, 0.744)	(0.421, 0.587)
$TIL_{L,HMT}$	0.130	0.081	0.078	0.069	0.077	0.045	0.123
	(0.110, 0.151)	(0.063, 0.099)	(0.060, 0.097)	(0.051, 0.087)	(0.050, 0.103)	(0.033, 0.056)	(0.094, 0.152)
$TIL_{L,OFC}$	0.048	0.073	0.080	0.066	0.126	0.086	0.167
	(0.035, 0.062)	(0.059, 0.087)	(0.065, 0.095)	(0.052, 0.080)	(0.104, 0.148)	(0.073, 0.099)	(0.141, 0.193)

Notes: The table reports estimates for the productivity evolution processes in (2.4) and (2.5). Two-sided 95% asymptotic cluster-robust percentile confidence intervals are reported in parentheses. In the baseline specification, (i) each firm's peers are restricted to firms located in the same city and operating in the same four-digit sub-industry, and (ii) technical change is flexibly controlled for using a series of year effects. Alternative specifications modify the baseline as follows. F1: adds location fixed effects at the peer-group level (province); F2: adds location-by-industry fixed effects at the peer-subgroup level (province and four-digit sub-industry); W1: peers are size-weighted, $s_{ijt} = L_{jt}\mathbf{1}\{(j,t)\in\mathscr{L}(i,t)\}/\sum_k L_{kt}\mathbf{1}\{(k,t)\in\mathscr{L}(i,t)\}$; P1: peers in $\mathscr{L}(i,t)$ are restricted to the same four-digit sub-industry within the same province; P2: peers are restricted to the same two-digit industry within the same city; P3: peers are restricted to the same two-digit industry within the same province.

ables such as shared regional or industrial environments that could induce productivity correlation among firms in the same network independently of any genuine spillovers. As discussed in Section 2, this assumption can be relaxed by allowing for time-invariant network unobservables and controlling for them via group-level fixed effects. Columns F1 and F2 of Table 5 implement this approach at progressively finer levels of aggregation: F1 includes province fixed effects, absorbing location-specific common factors, while F2 adds province-by-four-digit-industry fixed effects, absorbing common factors that vary jointly across location and subindustry. The Hicks-neutral spillover SP_H, while still small in absolute terms, rises when group-level fixed effects are introduced (0.039 and 0.045 versus 0.021 in the baseline). The labor-augmenting spillover SP_L is essentially unchanged across these specifications (0.186 and 0.196 versus 0.204), indicating that the estimated peer effect on ω_L is not driven by regional or subindustry-level common factors. Both direct and indirect learning estimates for Hicks-neutral productivity decline once group fixed effects are absorbed. For labor-augmenting productivity, the estimates remain economically large but the baseline asymmetry between HMT and OFC narrows considerably, with $DL_{L,HMT}$ falling from 0.640 and $DL_{L,OFC}$ rising from 0.238 to around 0.4.

Taken together, these results suggest that part of the baseline HMT–OFC source asymmetry may reflect correlation between foreign-investor presence and unobserved regional or subindustry characteristics rather than intrinsic differences between these two types of foreign capital. The factor bias of the network effect, however, is unaffected. Namely, the labor-augmenting productivity spillover and indirect learning from foreign investment continue to dominate their Hicks-neutral counterparts, so the concentration of FDI's network effects in the labor-augmenting dimension does not depend on controlling for group-level common factors.

Alternative weighting scheme. The baseline specification assigns uniform weights to all peers within $\mathscr{L}(i,t)$, treating each member of the reference group as equally influential. Column W1 of Table 5 replaces this with a size-weighted scheme,

$$s_{ijt} = \frac{L_{jt}\mathbf{1}\{(j,t)\in\mathscr{L}(i,t)\}}{\sum_k L_{kt}\mathbf{1}\{(k,t)\in\mathscr{L}(i,t)\}},$$

in which larger peers exert proportionally greater influence on a given firm's productivity. This reflects the intuition that larger firms are more visible and that knowledge or practices are more likely to diffuse from them.

The Hicks-neutral spillover SP_H rises noticeably (from 0.021 to 0.083), suggesting that the small baseline estimate partly reflects averaging across peers of heterogeneous size and that the relevant Hicks-neutral knowledge originates disproportionately from larger firms.[20] The labor-augmenting spillover SP_L, by contrast, remains virtually unchanged (0.205 versus 0.204 in the baseline), again confirming that our central peer-effect finding does not hinge on the weighting choice. The direct learning estimates respond asymmetrically across productivity components. For Hicks-neutral productivity, $DL_{H,HMT}$ is essentially unchanged (0.173 versus 0.182) but $DL_{H,OFC}$ more than doubles (0.430 versus 0.208). Being a partial effect, the coefficient on own foreign equity rests, in the Frisch-Waugh sense, on the equity variation left after projecting out the included peer-exposure terms, and re-weighting the network redefines this projection. The estimate of $DL_{H,OFC}$ therefore shifts to the

[20] To see the intuition, under the baseline equal-weight scheme, each peer contributes equally to the peer-productivity index regardless of firm size, whereas under W1, a peer's contribution is proportional to its employment. The latter therefore shifts the peer-productivity measure toward the productivity of larger firms. If productivity emanating from larger peers is more strongly associated with a firm's subsequent Hicks-neutral productivity, this reweighting raises the estimated spillover coefficient, as observed in the data.

extent that OFC equity covaries with peer exposure differently under the two weighting schemes, whereas no correlation under either scheme would have left the direct-learning estimates invariant to the re-weighting altogether. The contrast with the stable $DL_{H,HMT}$ thus points to OFC presence specifically being correlated with the size profile and productivity of nearby peers, rather than to a change in the learning channel itself. For labor-augmenting productivity, the baseline asymmetry between HMT and OFC again narrows: $DL_{L,HMT}$ falls from 0.640 to 0.337 while $DL_{L,OFC}$ rises moderately from 0.238 to 0.321, bringing the two sources nearly to parity. The indirect learning estimates follow these patterns, with TIL_H rising appreciably under size-weighting because both the spillover term and the OFC direct effect strengthen simultaneously. Under size-weighting, the Hicks-neutral channels strengthen (SP_H nearly quadruples and $DL_{H,OFC}$ more than doubles) and the labor-augmenting margin narrows. The ordering is nonetheless preserved, whereby SP_L continues to exceed SP_H and TIL_L still exceeds TIL_H for both sources. As such, the labor-augmenting concentration of FDI's network effects is preserved in direction, if not in degree, under this alternative weighting.

Alternative peer-set definitions. Columns P1–P3 broaden the scope of local interactions relative to the baseline specification of city-by-four-digit-subindustry peer groups. P1 retains the four-digit subindustry restriction but expands the geographic scope to the entire province; P2 retains the city restriction but expands the industrial scope to the full two-digit industry; P3 broadens both, defining peers as all firms in the same two-digit industry within the same province. These specifications progressively dilute the peer network with firms that are likely less relevant for direct knowledge diffusion.

The factor-neutral spillover SP_H remains positive and significant, with estimates between 0.056 and 0.099, which are noticeably larger than the baseline value of 0.021, suggesting that broader peer definitions absorb additional sources of cross-firm productivity correlation into the spillover term. The labor-augmenting spillover SP_L also remains positive and statistically significant across all three specifications, with point estimates ranging from 0.128 to 0.332. The direct and indirect learning estimates show more variation across the P specifications. For labor-augmenting productivity, the estimates remain positive and economically meaningful throughout, although the relative magnitudes of $DL_{L,HMT}$ and $DL_{L,OFC}$ shift across columns. For Hicks-neutral productivity, both DL_H and TIL_H become unstable: $DL_{H,OFC}$ turns negative under P1 and P3, while several TIL_H estimates flip sign under the broadest peer definitions. This instability suggests that the Hicks-neutral learning channels are sensitive to how the peer set is drawn. By contrast, the labor-augmenting spillover stays positive and larger than its Hicks-neutral counterpart in every column, and labor-augmenting direct and indirect learning remain positive and economically meaningful across all peer-set choices. This channel thus captures a more robust feature of cross-firm productivity transmission in the industry.

Taken together, the sensitivity exercise buttresses the central finding of our baseline analysis that productivity externalities are biased towards labor. The labor-augmenting spillover SP_L is sizable, positive and statistically significant in every specification and exceeds the Hicks-neutral spillover SP_H throughout, though the margin between them ranges from roughly two- to ten-fold as the peer set and weighting are varied. The same labor-augmenting concentration characterizes the indirect (through peer network diffusion) learning from FDI, with TIL_L exceeding TIL_H in all specifications. The direct learning estimates are less stable. Those for labor-augmenting productivity remain positive and economically meaningful throughout, but the baseline HMT-OFC source asymmetry does not survive broader peer definitions.

6 Conclusion

We develop a structural framework for estimating productivity spillovers when firm productivity is multidimensional. Existing methods typically impose two restrictions: (i) productivity is scalar and Hicks-neutral, and (ii) its evolution is independent across firms. These assumptions rule out factor-biased productivity spillovers by construction. We relax them simultaneously by modeling productivity as the combination of factor-neutral and labor-augmenting components, each evolving through a controlled spatiotemporal Markov process that depends on the firm's own history, productivity-shifting activities, and the lagged productivity of its peers. The resulting framework allows cross-firm productivity transmission to operate differently across the two components of technological heterogeneity.

In the proxy-variable tradition, our identification strategy combines the static input optimality conditions with the productivity evolution processes under the maintained timing and information-set assumptions, and the resulting moment system pins down the technology parameters jointly with the own-persistence and spillover coefficients of each productivity component. We allow factor prices to be firm-specific and to depend on productivity rather than imposing homogeneity or exogeneity. This matters most for wages, since labor-augmenting productivity scales the effective quantity of labor and is thus plausibly reflected in the firm-specific wage that a homogeneous-price specification would otherwise absorb into the very productivity it is meant to measure. Because most datasets, including ours, report only revenue rather than separate output prices and quantities, we close the model with a constant-elasticity demand under monopolistic competition that serves as a parsimonious measurement device, mapping the physical technology into a revenue equation that can be estimated.

Applied to China's primary agricultural food processing industry, our framework uncovers a pronounced factor bias in the way productivity spills across firms within the sector. Productivity externalities operate predominantly through the labor-augmenting component, with the labor-augmenting spillover exceeding its Hicks-neutral counterpart by roughly an order of magnitude at the baseline. This bias has a counterpart in the internal dynamics of the two components, whereby the factor-neutral component that firms retain most strongly over time is the one that diffuses least to peers, while the less persistent labor-augmenting component propagates most readily across firms. A natural reading is that what spreads locally is carried by workers, managerial practices and process know-how that circulate through hiring and the circulation of personnel, rather than the more appropriable, firm-level advantages rooted in proprietary technology and organizational capital that travel less easily across firm boundaries. The direct productivity gains from foreign investment display a parallel bias, present for equity of either origin but far stronger for HMT capital, whose labor-augmenting effect is nearly three times that of OFC capital even as the two contribute comparably along the factor-neutral dimension. The labor-biased nature of these productivity externalities persists across alternative peer-group definitions, weighting schemes and the inclusion of group-level fixed effects.

Although we study foreign investment, the productivity shifters enter our framework generically, so the same methodology recovers the factor bias of other productivity-enhancing activities, such as R&D, exporting, technology adoption or worker training, requiring only that the shifter be observed and predetermined. The spillover network itself may likewise be redefined along supply-chain or ownership ties rather than spatial proximity. Directional spillovers along such networks, contemporaneous productivity spillovers and an endogenously formed network are more substantive extensions that we leave to future research.

Acknowledgement

This paper is supported by the National Natural Science Foundation of China Young Scholar Project (grant number: 72301218)

References

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

Ackerberg, D. A., G. Frazer, K. i. Kim, Y. Luo, and Y. Su (2023). Under-identification of structural models based on timing and information set assumptions. Journal of Econometrics 237(1), 105463.

Agarwal, N., C. Milner, and A. Riaño (2014). Credit constraints and spillovers from foreign firms in china. Journal of Banking & Finance 48, 261–275.

Almeida, P. and B. Kogut (1999). Localization of knowledge and the mobility of engineers in regional networks. Management Science 45(7), 905–917.

Alvarez, R. and R. A. López (2008). Is exporting a source of productivity spillovers? Review of World Economics 144(4), 723–749.

Anselin, L. (1988). Lagrange multiplier test diagnostics for spatial dependence and spatial heterogeneity. Geographical analysis 20(1), 1–17.

Arnold, J. M. and B. S. Javorcik (2009). Gifted kids or pushy parents? foreign direct investment and plant productivity in indonesia. Journal of International Economics 79(1), 42–53.

Bai, J. and K. Li (2021). Dynamic spatial panel data models with common shocks. Journal of Econometrics 224(1), 134–160.

Biesebroeck, J. V. (2004). Exporting raises productivity in sub-saharan african manufacturing plants. Journal of International Economics.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying technology spillovers and product market rivalry. Econometrica 81(4), 1347–1393.

Blume, L. E., W. A. Brock, S. N. Durlauf, and Y. M. Ioannides (2011). Identification of social interactions. In Handbook of social economics, Volume 1, pp. 853–964. Elsevier.

Bramoullé, Y., H. Djebbari, and B. Fortin (2009). Identification of peer effects through social networks. Journal of econometrics 150(1), 41–55.

Brandt, L., J. Van Biesebroeck, L. Wang, and Y. Zhang (2017). Wto accession and performance of chinese manufacturing firms. American Economic Review 107(9), 2784–2820.

Brandt, L., J. Van Biesebroeck, and Y. Zhang (2012). Creative accounting or creative destruction? firm-level productivity growth in chinese manufacturing. Journal of development economics 97(2), 339–351.

Caselli, M., A. Chatterjee, and S. Li (2026). Productivity and quality of multi-product firms. International Finance Discussion Papers 1430, Board of Governors of the Federal Reserve System, Washington, DC.

Coe, D. T. and E. Helpman (1995). International r&d spillovers. European Economic Review 39(5), 859–887.

Combes, P.-P. and G. Duranton (2006). Labour pooling, labour poaching, and spatial clustering. Regional Science and Urban Economics 36(1), 1–28.

De Loecker, J. (2011). Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. Econometrica 79(5), 1407–1451.

De Loecker, J. (2013). Detecting learning by exporting. American Economic Journal: Microeconomics 5(3), 1–21.

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

Demirer, M. (2025). Production function estimation with factor-augmenting technology: An application to markups. Econometrica.

Dewitte, R., M. Dumont, B. Merlevede, G. Rayp, and M. Verschelde (2020). Firm-heterogeneous biased technological change: A nonparametric approach under endogeneity. European Journal of Operational Research 283(3), 1172–1182.

Dhawan, R. (2001). Firm size and productivity differential: theory and evidence from a panel of us firms. Journal of economic behavior & organization 44(3), 269–293.

Doraszelski, U. and J. Jaumandreu (2013). R&d and productivity: Estimating endogenous productivity. Review of economic studies 80(4), 1338–1383.

Doraszelski, U. and J. Jaumandreu (2018). Measuring the bias of technological change. Journal of Political Economy 126(3), 1027–1084.

Duranton, G. and D. Puga (2004). Micro-foundations of urban agglomeration economies. In J. V. Henderson and J.-F. Thisse (Eds.), Handbook of Regional and Urban Economics, Volume 4, Chapter 48, pp. 2063–2117. Amsterdam: Elsevier.

Ellison, G., E. L. Glaeser, and W. R. Kerr (2010). What causes industry agglomeration? evidence from coagglomeration patterns. American Economic Review 100(3), 1195–1213.

Fan, H., Y. A. Li, and S. R. Yeaple (2018). On the relationship between quality and productivity: Evidence from china's accession to the wto. Journal of International Economics 110, 28–49.

Flynn, Z., A. Gandhi, and J. Traina (2019). Measuring markups with production data. Working paper, University of Pennsylvania.

Gandhi, A., S. Navarro, and D. A. Rivers (2020). On the identification of gross output production functions. Journal of Political Economy 128(8), 2973–3016.

Gertler, M. S. (2003). Tacit knowledge and the economic geography of context, or the undefinable tacitness of being (there). Journal of Economic Geography 3(1), 75–99.

Glaeser, E. L., H. D. Kallal, J. A. Scheinkman, and A. Shleifer (1992). Growth in cities. Journal of Political Economy 100(6), 1126–1152.

Glass, A. J., K. Kenjegalieva, and R. C. Sickles (2016). A spatial autoregressive stochastic frontier model for panel data with asymmetric efficiency spillovers. Journal of Econometrics 190(2), 289–300.

Graham, B. S. and J. Hahn (2005). Identification and estimation of the linear-in-means model of social interactions. Economics Letters 88(1), 1–6.

Greenstone, M., R. Hornbeck, and E. Moretti (2010). Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. Journal of Political Economy 118(3), 536–598.

Grieco, P. L., S. Li, and H. Zhang (2016). Production function estimation with unobserved input price dispersion. International Economic Review 57(2), 665–690.

Griliches, Z. (1979). Issues in assessing the contribution of research and development to productivity growth. Bell Journal of Economics 10(1), 92–116.

Guadalupe, M., O. Kuzmina, and C. Thomas (2012). Innovation and foreign ownership. American economic review 102(7), 3594–3627.

Huang, M., S. Zhao, and S. C. Kumbhakar (2022). Decomposition of output, productivity and market structure changes. European Journal of Operational Research 303(1), 422–437.

Huang, Y., L. Jin, and Y. Qian (2013). Does ethnicity pay? evidence from overseas chinese fdi in china. Review of Economics and Statistics 95(3), 868–883.

Iyoha, E. (2023). Estimating productivity in the presence of spillovers: Firm-level evidence from the US production network. Working Paper 24-033, Harvard Business School, Boston, MA.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. Quarterly Journal of Economics 108(3), 577–598.

Javorcik, B. and S. Poelhekke (2017). Former foreign affiliates: Cast out and outperformed? Journal of the European Economic Association 15(3), 501–539.

Javorcik, B. S. (2004). Does foreign direct investment increase the productivity of domestic firms? in search of spillovers through backward linkages. American Economic Review 94(3), 605–627.

Keller, W. and S. R. Yeaple (2009). Multinational enterprises, international trade, and productivity growth: firm-level evidence from the united states. The review of economics and statistics 91(4), 821–831.

Klette, T. J. and Z. Griliches (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. Journal of applied econometrics 11(4), 343–361.

Kugler, M. and E. Verhoogen (2012). Prices, plant size, and product quality. The Review of Economic Studies 79(1), 307–339.

Lee, K. (2023). Strategic network decisions and knowledge spillovers: Evidence from R&D collaborations of the U.S. firms. Working paper, Korea Institute of Public Finance.

Lee, Y., M. E. Lovely, and H. Pham (2023). Dynamic and non-neutral productivity effects of foreign ownership: A nonparametric approach. Journal of Applied Econometrics 38(1), 24–48.

Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. The review of economic studies 70(2), 317–341.

Malikov, E. and S. Zhao (2023). On the estimation of cross-firm productivity spillovers with an application to fdi. Review of Economics and Statistics 105(5), 1207–1223.

Malikov, E., S. Zhao, and S. C. Kumbhakar (2020). Estimation of firm-level productivity in the presence of exports: Evidence from china's manufacturing. Journal of Applied Econometrics 35(4), 457–480.

Malikov, E., S. Zhao, S. C. Kumbhakar, and D. H. Bernstein (2025). Estimating production functions using costs when outputs are restricted. Econometric Reviews 44(4), 433–461.

Malikov, E., S. Zhao, and J. Zhang (2024). A system approach to structural identification of production functions with multi-dimensional productivity. Advances in Econometrics 46.

Manova, K. and Z. Zhang (2012). Export prices across firms and destinations. The Quarterly Journal of Economics 127(1), 379–436.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. The review of economic studies 60(3), 531–542.

Manski, C. F. (2000). Economic analysis of social interactions. Journal of Economic Perspectives 14(3), 115–136.

Marshall, A. (1890). Principles of Economics. London: Macmillan.

Moffitt, R. A. (2001). Policy interventions, low-level equilibria, and social interactions. In S. N. Durlauf and H. P. Young (Eds.), Social Dynamics, pp. 45–82. Cambridge, MA: MIT Press.

Moretti, E. (2004). Workers' education, spillovers, and productivity: Evidence from plant-level production functions. American Economic Review 94(3), 656–690.

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

Raval, D. R. (2019). The micro elasticity of substitution and non-neutral technology. The RAND Journal of Economics 50(1), 147–167.

Rosenthal, S. S. and W. C. Strange (2004). Evidence on the nature and sources of agglomeration economies. In J. V. Henderson and J.-F. Thisse (Eds.), Handbook of Regional and Urban Economics, Volume 4, Chapter 49, pp. 2119–2171. Amsterdam: Elsevier.

Shi, W. and L.-f. Lee (2017). Spatial dynamic panel data models with interactive fixed effects. Journal of Econometrics 197(2), 323–347.

Skevas, I. (2023). A novel modeling framework for quantifying spatial spillovers on total factor productivity growth and its components. American Journal of Agricultural Economics 105(4), 1221–1247.

Storper, M. and A. J. Venables (2004). Buzz: Face-to-face contact and the urban economy. Journal of Economic Geography 4(4), 351–370.

Verhoogen, E. A. (2008). Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. The quarterly journal of economics 123(2), 489–530.

Zacchia, P. (2020). Knowledge spillovers through networks of scientists. The Review of Economic Studies 87(4), 1989–2018.

Zhang, H. (2019). Non-neutral technology, firm heterogeneity, and labor demand. Journal of Development Economics 140, 145–168.

Zhang, J. and E. Malikov (2023). Detecting learning by exporting and from exporters. Journal of Productivity Analysis 60(1), 1–19.

Zhao, S., E. Malikov, and S. C. Kumbhakar (2026). On the estimation of nonneutral production technologies without homothetic separability. Working Paper.

Zhao, S., B. Qian, and S. C. Kumbhakar (2020). Estimation of productivity and markups with price dispersion: Evidence from chinese manufacturing during economic transition. Southern Economic Journal 87(2), 666–699.

Appendix

A Special Case of Perfect Competition in the Factors and Output Markets with Homogeneous Prices

The assumption of homogeneous and exogenous prices of variable inputs and the output is commonly made, either explicitly or implicitly, in many productivity studies. Here we show that our identification strategy still applies in such a classical perfectly competitive setting.

Concretely, in what follows, firms are price-takers in the input markets for labor and materials as well as in the output market, facing $P_{it} = P_t$, $P^M_{it} = P^M_t$ and $P^L_{it} = P^L_t$ for all i, and these prices are mean-orthogonal to period-t productivity innovations in firms.

Step One. The firm's static profit maximization problem becomes

$$\max_{M_{it},L_{it}} P_t F(K_{it}, e^{\omega_{Lit}} L_{it}, M_{it}) e^{\omega_{it}} \Lambda - P^M_t M_{it} - P^L_t L_{it}, \tag{A.1}$$

where $\Lambda \equiv \mathbb{E}(e^{v_{it}})$. Following the same steps used to derive equation (3.5), the log-ratio of materials to labor under homogeneous input prices is

$$m_{it} - l_{it} = \sigma \ln \beta_M - \sigma(p^M_t - p^L_t) + (1-\sigma)\omega_{L,it}. \tag{A.2}$$

Since, under this market structure, the homogeneous p^M_t and p^L_t are exogenously determined and therefore uncorrelated with firm characteristics, including productivity, we can identify $\sigma \neq 1$ and β_M directly from (A.2) using the moment condition

$$\mathbb{E}\Big((1-\sigma)\omega_{L,it} \mid p^M_t - p^L_t\Big) = \mathbb{E}\big(\omega_{L,it}\big) = 0, \tag{A.3}$$

and then recover β_K via the normalizing restriction $\beta_K + \beta_M = 1$. This step also delivers X_{it} via the definition in (3.4), which remains unchanged, as well as allows us to recover the firm's labor productivity $\omega_{L,it}$ from residuals, leaving ν as the only remaining unknown production-function parameter.

Step Two. Inverting the first-order condition for materials derived from (A.1) for $\omega_{H,it}$, we get

$$\omega_{H,it} = \gamma_H + \tfrac{1}{\sigma} m_{it} + p^M_t - p^Y_t + \left(1 + \tfrac{\nu\sigma}{1-\sigma}\right) x_{it}, \tag{A.4}$$

where $\gamma_H = -\ln(\nu \beta_M \Lambda)$, substituting which for Hicks-neutral productivity in the production function in logarithmic form, we arrive at

$$\begin{aligned} y_{it} &= -\tfrac{\nu\sigma}{1-\sigma} x_{it} + \omega_{H,it} + v_{it} & \text{(A.5)} \\ &= \gamma_H + x_{it} + \tfrac{1}{\sigma} m_{it} + p^M_t - p_t + v_{it}, & \text{(A.6)} \end{aligned}$$

which, under our structural assumptions, can be identified from the zero-mean moment condition for v_{it}:

$$\mathbb{E}\Big(\Big[y_{it} - x_{it} - \tfrac{1}{\sigma} m_{it} - p^M_t + p_t\Big] - \gamma_H\Big) = 0, \tag{A.7}$$

where the bracketed term is "observed" following the first step. With γ_H identified, v_{it} can be obtained directly from (A.6).

Step Three. We now identify the remaining unknown production-function parameter ν together with the parameters governing the ω_H process, $(\rho_0, \rho_1, \rho_2, \rho_3)$. To do so, we first express Hicks-neutral productivity as a residual in the production function, in parallel with the identification strategy employed in Step Three of Section 3.1. From (A.5), we have

$$\omega_{H,it} = y_{it} + \frac{\nu\sigma}{1-\sigma} x_{it} - v_{it}. \tag{A.8}$$

Since v_{it} has already been recovered in Step Two, define

$$\widetilde{y}_{it} \equiv y_{it} - v_{it}, \tag{A.9}$$

so that (A.8) can be written compactly as

$$\omega_{H,it} = \widetilde{y}_{it} + \frac{\nu\sigma}{1-\sigma} x_{it}. \tag{A.10}$$

Thus, conditional on ν, the production function provides a direct proxy for the persistent Hicks-neutral productivity component.

We then substitute (A.10) for $\omega_{H,it}$, $\omega_{H,it-1}$, and $\{\omega_{H,jt-1}\}_{j=1}^N$ in the Markov process (2.4). Rearranging yields

$$\begin{aligned} \widetilde{y}_{it} = &-\frac{\nu\sigma}{1-\sigma} x_{it} + \rho_0 + \rho_1 \left(\widetilde{y}_{it-1} + \frac{\nu\sigma}{1-\sigma} x_{it-1} \right) \\ &+ \rho_2 \sum_{j\neq i} s_{ijt-1} \left(\widetilde{y}_{jt-1} + \frac{\nu\sigma}{1-\sigma} x_{jt-1} \right) + \rho_3' G_{H,it-1} + \zeta_{H,it}. \end{aligned} \tag{A.11}$$

Equivalently, defining

$$\begin{aligned} g_H(\mathbb{X}_{H,it-1}) \equiv\ & \rho_0 + \rho_1 \widetilde{y}_{it-1} + \rho_2 \sum_{j\neq i} s_{ijt-1} \widetilde{y}_{jt-1} + \rho_3' G_{H,it-1} \\ &+ \frac{\nu\sigma}{1-\sigma} \left(\rho_1 x_{it-1} + \rho_2 \sum_{j\neq i} s_{ijt-1} x_{jt-1} \right), \end{aligned} \tag{A.12}$$

we can write the estimating equation compactly as

$$\widetilde{y}_{it} = -\frac{\nu\sigma}{1-\sigma} x_{it} + g_H(\mathbb{X}_{H,it-1}) + \zeta_{H,it}, \tag{A.13}$$

where

$$\mathbb{X}_{H,it-1} \equiv \left(\widetilde{y}_{it-1}, x_{it-1}, \sum_{j\neq i} s_{ijt-1} \widetilde{y}_{jt-1}, \sum_{j\neq i} s_{ijt-1} x_{jt-1}, G_{H,it-1} \right).$$

Analogous to Step Three in Section 3.1, the zero-mean innovations $\{\zeta_{H,it}\}_{i=1}^N$ are mean-orthogonal to the period-$t-1$ information set, whereby

$$\mathbb{E}\left(\zeta_{H,it} \mid \mathbb{X}_{H,it-1}\right) = \mathbb{E}\left(\zeta_{H,it}\right) = 0. \tag{A.14}$$

Hence, all lagged variables in $\mathbb{X}_{H,it-1}$ are valid instruments. The contemporaneous x_{it}, by contrast, is endogenous because it is a function of the static inputs (M_{it}, L_{it}), which the firm chooses after observing the current-period persistent productivity innovations. Importantly, however, x_{it} is not a "free" regressor: its coefficient is restricted to $-\nu\sigma/(1-\sigma)$, where σ has already been identified in Step One. Consequently, no separate instrument for contemporaneous x_{it} is required.

The moment restrictions in (A.14), together with (A.13), identify ν jointly with $(\rho_0, \rho_1, \rho_2, \rho_3)$. In particular, ρ_1 is identified from variation in own lagged productivity proxy $\widetilde{y}_{it-1}$, ρ_2 from variation in the lagged peer-average productivity proxy $\sum_{j\neq i} s_{ijt-1}\widetilde{y}_{jt-1}$, and ρ_3 from variation in $G_{H,it-1}$, while the zero-unconditional-mean restriction helps pin down ρ_0. The remaining parameter ν is identified through the cross-equation restrictions it imposes on the coefficients associated with contemporaneous x_{it}, own lagged x_{it-1}, and the lagged peer average $\sum_{j\neq i} s_{ijt-1}x_{jt-1}$. Concluding this step, all production-function parameters as well as the parameters governing the ω_H process have been identified.

The remaining unknown parameters are those governing the ω_L process, viz., $(\gamma_1, \gamma_2, \gamma_3)$. Recall that, upon Step One, we have already identified $\omega_{L,it}$ for all i and t. Since the labor-augmenting productivity component is now known, we can estimate (2.5) using $\omega_{L,it}$, $\omega_{L,it-1}$, and $\{\omega_{L,jt-1}\}_{j=1}^{N}$ as observed variables:

$$\omega_{L,it} = \gamma_1 \omega_{L,it-1} + \gamma_2 \sum_{j\neq i} s_{ijt-1}\omega_{L,jt-1} + \gamma_3' G_{L,it-1} + \zeta_{L,it}, \tag{A.15}$$

which is immediately identified based on

$$\mathbb{E}\left(\zeta_{L,it} \mid \omega_{L,it-1}, \sum_{j\neq i} s_{ijt-1}\omega_{L,jt-1}, G_{L,it-1}\right) = 0. \tag{A.16}$$

B Additional Results

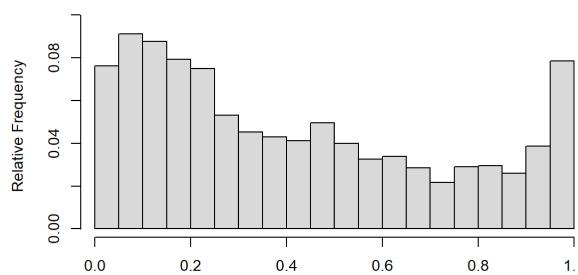

(a) HMT Equity Share | HMT > 0

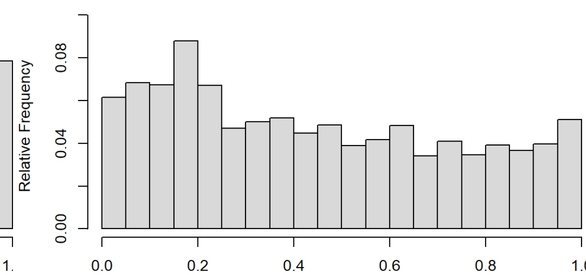

(b) OFC Equity Share | OFC > 0

Figure B1: Empirical Distribution of the Foreign Equity Share

Table B1: Correlation Between the Estimated Productivity Components

Year	No. of Firms	Pearson Correlation	Spearman Correlation
1998	3,171	0.054	0.083
1999	4,558	0.072	0.099
2000	4,685	0.086	0.114
2001	5,085	0.082	0.109
2002	5,693	0.071	0.097
2003	6,147	0.063	0.097
2004	8,761	0.073	0.107
2005	10,601	0.034	0.055
2006	12,303	0.039	0.063
2007	10,839	0.005	0.029
Pooled Sample	19,803	0.051	0.077

Notes: The table reports contemporaneous cross-sectional correlations between the estimated labor-augmenting productivity component, $\widehat{\omega}_{L,it}$, and the estimated Hicks-neutral productivity component, $\widehat{\omega}_{H,it}$. Annual correlations are computed separately for each year. The pooled correlations are calculated after demeaning both productivity components by their respective annual sample means. The pooled sample contains 71,843 firm-year observations from 19,803 distinct firms.