



Structural Identification of Productivity under Biased Technological Change

Laurens Cherchye
Department of Economics, KULeuven

Thomas Demuyck
ECARES, Université libre de Bruxelles

Bram De Rock
ECARES, Université libre de Bruxelles
Department of Economics, KU Leuven

Cédric Duprez
National Bank of Belgium

Glenn Magerman
ECARES, Université libre de Bruxelles

Marijn Verschelde
Department of Economics and Quantitative Methods, ESEG School of Management
Lille Economie Management, Department of Economics, KU Leuven

December 2021

ECARES working paper 2021-28

Structural identification of productivity under biased technological change*

Laurens Cherchye[†] Thomas Demuyne[‡] Bram De Rock[§]
Cédric Duprez[¶] Glenn Magerman^{||} Marijn Vershelde^{**}

December 10, 2021

Abstract

We propose a novel method for structural production analysis in the presence of unobserved heterogeneity in productivity. Our approach is intrinsically nonparametric and does not require the stringent assumption of Hicks neutrality. We assume cost minimization as the firms' behavioral objective, and we model productivity on which

*This is a substantial revision of an earlier paper that circulated under the title “Nonparametric production analysis with unobserved heterogeneity in productivity”. We are grateful to the National Bank of Belgium for supporting this project. We also thank Jan De Loecker, Knox Lovell, Johannes Van Biesebroeck, Ruben Dewitte, Klaas Mulier, Veerle Hennebel, Catherine Fuss, the seminar participants at Queensland University and the participants of EWEPA 2017, EEA 2018, EARIE 2018, the 2019 Workshop on Firm Heterogeneity in Technical Change at Ghent University and AFSE 2021 for useful comments. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium or any other institution to which the authors are affiliated. This work was supported by the Fonds de la Recherche Scientifique - FNRS under grant number 40003508.

[†]Department of Economics, University of Leuven. E-mail: laurens.cherchye@kuleuven.be.

[‡]ECARES, Université Libre de Bruxelles. E-mail: thomas.demuyne@ulb.be.

[§]ECARES, Université Libre de Bruxelles, and Department of Economics, University of Leuven (KU Leuven). E-mail: bram.de.rock@ulb.be.

[¶]National Bank of Belgium. E-mail: Cedric.Duprez@nbb.be

^{||}ECARES, Université Libre de Bruxelles, National Bank of Belgium, and CEPR. E-mail: glenn.magerman@ulb.be

^{**}Department of Economics and Quantitative Methods, IÉSEG School of Management, Lille Économie Management (UMR-CNRS 9221) and Department of Economics, KU Leuven. E-mail: m.vershelde@ieseg.fr.

firms condition the input demand. Our model can equivalently be represented in terms of endogenously chosen latent input costs, which avoids an endogeneity bias in a natural way. Our empirical application to unique and detailed Belgian manufacturing data shows that our method allows for drawing strong and robust conclusions, despite its nonparametric orientation. For example, we confirm the well-documented productivity slowdown, and we highlight a potential bias when using a common-scale intermediate inputs price deflator in the estimation of productivity. In addition, we provide robust empirical evidence against the assumption of Hicks neutrality for the setting at hand.

Keywords: productivity, unobserved heterogeneity, simultaneity bias, nonparametric production analysis, cost minimisation, manufacturing

JEL: C14, D21, D22, D24.

1 Introduction

Existing empirical analyses of the heterogeneity in firms' production behavior usually rely on the debatable assumption of Hicks neutral technical change and impose stringent parametric structure on the production processes under study. In this paper, we propose a novel method for structural production analysis that is intrinsically nonparametric and does not impose Hicks neutrality. Our method is operationalized through linear programming, which makes it easy to use in practice. Moreover, it is flexible and can be tailored to address specific empirical challenges for the production setting at hand. Our empirical findings for Belgian manufacturing firms generate informative conclusions, despite our nonparametric set-up. For example, we document a productivity slowdown since the early 2000s, and we provide strong empirical evidence against the assumption of Hicks neutrality.

Research question. The increasing prevalence of firm heterogeneity in global sourcing (Antras and Helpman (2004)) and input cost shares changes (Autor et al. (1998), Autor et al. (2003), Autor et al. (2020) and Kehrig and Vincent (2021)) lies at the heart of the industrial policy debate. Paradoxically, these phenomena are excluded by construction under the assumption of Hicks neutral technical change. The few empirical production studies that do relax this assumption typically rely on a specific parametrization of the production

technology or impose a common structure on the factor bias across firms.¹ However, as one may expect, both empirical evidence and economic theory show that there can be firm heterogeneity in factor biased technical change (see, e.g., Acemoglu et al. (2015) and references therein), which –evidently– makes an a priori parametrization difficult.

In a series of seminal papers, Afriat (1972), Hanoch and Rothschild (1972), Diewert and Parkan (1983) and Varian (1984) proposed an intrinsically nonparametric approach to address the identification of production functions.² It recovers the production possibilities directly from the data and avoids functional specification bias by not imposing any (non-verifiable) parametric structure on the production technology. Its identifying power comes from a structural specification of the firms’ objectives that underlie the observed production behavior.

Despite this conceptually appealing starting point, the more recent literature on the identification and estimation of production functions has largely ignored this nonparametric alternative. We interpret this lack of attention as principally originating from the fact that the existing nonparametric methods are unable to deal with heterogeneity in unobserved productivity. The importance of effectively dealing with unobserved productivity is by now well-established in the literature (see, e.g., the reviews of Syverson (2011) and, more recently, De Loecker and Syverson (2021)). Basically, incorporating unobservables in the empirical analysis is a prerequisite to account for endogeneity between input choice and unobserved productivity. This endogeneity issue was first pointed out by Marschak and Andrews (1944), and originates from the fact that a firm’s productivity transmits to its optimal input choices. It implies that standard OLS-type estimation techniques will suffer from a simultaneity bias (see also Olley and Pakes (1996) and Griliches and Mairesse (1998)).³

¹See, e.g., the recent study of Doraszelski and Jaumandreu (2018), which parametrizes the labor augmenting technological change next to the Hicks neutral technological change in a constant elasticity of substitution (CES) framework.

²We refer to Grifell-Tatjé et al. (2018) (and references therein) for a recent review of alternative approaches of productivity measurement that have been proposed in the Economics and OR/MS literature.

³The literature on the estimation and identification of production functions has paid considerable attention to developing techniques that address this endogeneity problem. Notable examples include Olley and Pakes (1996); Levinsohn and Petrin (2003); Wooldridge (2009); Akerberg et al. (2015); Gandhi et al. (2020). A main difference with our nonparametric approach is that the empirical implementation of these existing approaches requires a (semi)parametric specification of the production technology.

Methodological contribution. The principal aim of the current paper is to re-establish the nonparametric approach as a full-fledged alternative for empirical production analysis. To this end, we present a methodology that uses minimal assumptions to address identification of unobserved productivity differences across firms and of productivity and cost share changes over time. Specifically, we assume cost minimization as the firms’ behavioral objective and we model unobserved heterogeneity as an unobserved productivity factor, on which we condition the demand for the observed inputs. This avoids the endogeneity bias in a natural way, by explicitly accounting for the simultaneity between productivity and input decisions in our structural specification of the firm’s optimization problem. We also provide a novel and intuitive way to quantify unobserved heterogeneity in terms of endogenously chosen latent inputs.

An attractive feature of our method is its empirical applicability. It is operationalized through linear programming, which makes it easy to apply in practice. Moreover, in contrast to most production function estimators, it is based on gross output rather than on value added, which makes that our methodology follows closely the theory of the firm.⁴ Furthermore, as our methodology is not based on structuring the timing of (input) decisions, it can be applied to both panel and cross-sectional data.

Importantly, our methodology is also flexible. It can be tailored to the specific empirical setting under study to effectively address empirical challenges that received substantial attention in the literature. We will illustrate this versatility in a number of ways. First, our approach only requires that the empirical analyst observes a single input price, while imperfectly observed price heterogeneity for the remaining inputs can be accounted for through shadow pricing. See, e.g., Klette and Griliches (1996), Foster et al. (2008), De Loecker (2011), Collard-Wexler and De Loecker (2016), De Loecker et al. (2016) and Grieco et al. (2016) for discussions on the importance of price heterogeneity. Next, we customize our approach to the standardly used behavioral model that assumes dynamically optimizing firms characterized by capital accumulation over time subject to (non-linear) adjustment costs (see, e.g., Akerberg et al. (2015)). Specifically, we can do so by assuming short term cost minimization conditional on fixed capital levels. Finally, we address the issue of unobserved quality in some output and/or input dimension by restricting the set of comparison firm

⁴See, e.g., Gandhi et al. (2020) for a comparison of gross output and value added production function estimates.

observations when evaluating cost minimizing behavior. For example, firms with both lower market shares and lower output prices are removed from the set of comparison partners for the firm observation that is evaluated, as economic theory predicts that these firms produce a product of lower (subjective or objective) quality (see, e.g., De Loecker et al. (2016) and references therein).

Application to Belgian manufacturing data. We apply our nonparametric methodology a uniquely rich data set on highly internationalized Belgian manufacturing firms that covers the period 2002-2014. Using detailed quantity and price information on both the input and output side of production, we study the micro-dynamics of quantity-based productivity and input usage (including latent inputs) while addressing the empirical challenges indicated above. Our output equals gross output, and we distinguish between foreign sourcing, domestic sourcing, service sourcing and wholesale, retail and energy expenditure as inputs. As discussed in Duprez and Magerman (2018) and Bernard et al. (Forthcoming), a distinguishing feature of our data set is that no common-scale intermediate inputs price deflator has to be used.

Our application shows that our method does allow for drawing strong and robust conclusions, despite its nonparametric orientation. For the period under study, productivity growth and the cost share of latent input costs is non-increasing over time, which is in accordance with the well-documented productivity slowdown in manufacturing since the early 2000s (see, e.g., Syverson (2017)). In addition, we document technological change in favor of the use of intermediate inputs. Even when taking latent inputs into account, we find that the share of (particularly foreign) intermediates is increasing, in line with the recent trend towards outsourcing and offshoring. As price changes cannot explain these identified cost changes, we see this as strong empirical evidence against the assumption of Hicks neutrality. Lastly, we highlight a potential bias when using a common-scale intermediate inputs price deflator in the estimation of productivity, especially when the focus is on the linkage between production and internationalization. The difference between productivity change estimates, with and without decomposed firm-year intermediate input price and quantity information, correlates significantly with the changing intensity of capital usage and foreign sourcing.

Outline. The remainder of this paper unfolds as follows. Section 2 presents our novel methodology for nonparametric production analysis with unobserved heterogeneity. Section 3 motivates our application to Belgian manufacturing firms, and discusses the input and output data that we use. Section 4 presents our main empirical findings. Section 5 concludes and discusses possible avenues for follow-up research. Appendix A provides additional discussion (including proofs) of our main theoretical results, and Appendix B presents a Monte Carlo simulation exercise. The Online Appendix presents extra details regarding our empirical application.

2 Methodology

We begin this section by considering a basic specification of the firm’s optimization problem with variable inputs and complete input price information, to subsequently establish the associated nonparametric characterization of optimizing firm behavior under productivity heterogeneity. This will pave the way for introducing our concept of a productivity factor to empirically identify differences in productivity between firms. In most cases, the empirical analyst is confronted with dynamically optimizing firms that accumulate capital over time under adjustment costs. Therefore, in a next step we show how we can modify our basic identification strategy to apply to the behavioral model of short run cost minimization with fixed capital usage. We then proceed by indicating how to impose Hicks neutrality in our characterization of optimizing behavior. This will allow us to investigate the empirical validity of this assumption in our following application. We conclude this section by discussing a number of specific issues that pertain to applying our nonparametric methodology to empirical data. Particularly, we explain how we can account for (small) deviations from “exactly” optimizing behavior in empirical applications by using a nonparametric measure of goodness-of-fit, and we indicate how we can sharpen identification by imposing additional structure on the unknown input prices. Appendix B contains a Monte Carlo simulation that shows the finite sample properties of our novel methodology. In this simulation, a particular focus is on evaluating the empirical performance of our methodology with and without the assumption of Hicks neutrality.

2.1 Production with heterogeneity in productivity

In our basic specification, firms' production levels depend on observed inputs, as well as on unobserved productivity. Formally, we assume a production function F that defines:

$$Q = F(\mathbf{V}, \Omega),$$

with $Q \in \mathbb{R}_+$ the output, $\mathbf{V} \in \mathbb{R}_+^M$ a M -dimensional vector of flexible and freely adjustable observed inputs, and $\Omega \in \mathbb{R}_+$ a single-dimensional measure of the unobserved productivity heterogeneity in the production process across firms. The assumption that unobserved productivity differences are one-dimensional follows the standard practice in the literature (see, e.g., Olley and Pakes (1996), Levinsohn and Petrin (2003), Wooldridge (2009), Akerberg et al. (2015) and Gandhi et al. (2020)). A useful implication is that it allows for a transparent empirical analysis of heterogeneity patterns, as we will demonstrate in our empirical application.

Generally, we can interpret the unobserved Ω in two ways.⁵ In the first interpretation, Ω falls beyond the firms' control and stands for external drivers of productivity and random productivity shocks. For instance, firms with higher Ω have access to better technologies, thereby increasing their output $F(\mathbf{V}, \Omega)$ for the same level of observed inputs \mathbf{V} . Alternatively, we can also interpret Ω as a latent input, which implies that it is optimally chosen by the firm. This interpretation includes all factors under the control of the firm that influence productivity, such as managerial input and information technology. Importantly, while the two interpretations are clearly distinct, we will show that the associated models of optimizing firm behavior are empirically equivalent in terms of their nonparametric testable implications. As a result, our following characterization of optimizing behavior does not depend on the specific meaning that is attached to Ω .

Throughout, we assume that the function F is strictly increasing, continuous and quasi-concave in (\mathbf{V}, Ω) . In addition, we postulate that the production technology is characterized by a parameter γ that defines the returns-to-scale (RTS) associated with observed inputs

⁵See Syverson (2011) for a general discussion on alternative interpretations of productivity differences that appeared in the literature.

(for the given unobserved productivity factor Ω), which means that, for all $t > 0$:

$$F(t\mathbf{V}, \Omega) = t^\gamma F(\mathbf{V}, \Omega).$$

To focus our discussion, we will impose a constant returns-to-scale (CRS) assumption in what follows, which corresponds to setting γ equal to one. However, it is worth emphasizing that our theoretical arguments are easily generalized for alternative RTS assumptions.⁶ The CRS assumption is usually motivated by a replication argument: if one doubles all the observed inputs, one can always double the output given the latent input. In our empirical analysis, we only impose CRS within a specific firm size category, as productivity heterogeneity is analyzed for each firm size category separately. This effectively implies that we (only) assume CRS to hold “locally” (i.e., for the given firm size), so avoiding the “global” CRS postulate, which –admittedly– may seem overly strong in many practical settings.

Importantly, imposing some RTS assumption (by specifying γ) is crucial for obtaining nonparametric identifying restrictions in terms of observed input and output. It can be verified that our assumption of cost minimizing behavior with latent input would define vacuous conditions for rationalizable production behavior (as specified in the following Definition 1) if we put no structure on the RTS of the production technology. Intuitively speaking, we obtain this “negative” conclusion by setting the cost of the latent input sufficiently high, so that the observed input does not generate meaningfully testable implications associated with cost minimization for the observed output.⁷ Essentially, our RTS assumption disciplines the latent cost structure in a way that excludes such a trivial rationalization of the observed production behavior.

2.2 Cost minimizing production behavior

Throughout, we assume that firms are price takers in the input market and we impose no structure on the form of the output market. As shown by Carvajal et al. (2013, 2014), it is possible to impose alternative (e.g., Cournot or Bertrand) structures on the output market in our advocated nonparametric framework. In our following analysis, we purposely do not

⁶In Online Appendix D, we show that our main results are robust for moderate changes of the RTS assumption.

⁷Compare with Varian (1988), who formalized a similar argument in a consumption context.

impose any such assumption, so showing that our identification results are independent of the output market form.

Let $\mathbf{W} \in \mathbb{R}_{++}^M$ be the price vector for the observed inputs. Our above two interpretations of the heterogeneity factor Ω yield two different models of optimizing firm behavior. First, if we assume that Ω is beyond the firm's control, then the firm solves the optimization problem:

$$\min_{\mathbf{V}} \mathbf{W}'\mathbf{V} \text{ s.t. } F(\mathbf{V}, \Omega) \geq Q \quad (OP.I).$$

That is, the firm's input choice \mathbf{V} is conditional on the unobserved factor Ω . Second, if Ω is a latent input factor that is chosen by the optimizing firm, then this firm solves

$$\min_{\mathbf{V}, \Omega} \{\mathbf{W}'\mathbf{V} + \Gamma\Omega\} \text{ s.t. } F(\mathbf{V}, \Omega) \geq Q, \quad (OP.II)$$

for $\Gamma \in \mathbb{R}_{++}$ the unobserved price of Ω . In both scenarios, the simultaneity bias is absent by construction, because either the (observed) inputs \mathbf{V} are optimally chosen conditionally on the unobserved Ω or, alternatively, these inputs are defined simultaneously with the latent input Ω .

We demonstrate the empirical equivalence of optimizing behavior in terms of (OP.I) and (OP.II) by establishing the associated testable implications. To this end, we assume to observe a data set

$$S = \{\mathbf{W}_i, \mathbf{V}_i, Q_i\}_{i \in N},$$

with $\mathbf{W}_i \in \mathbb{R}_{++}^M$ the observed input prices, $\mathbf{V}_i \in \mathbb{R}_{++}^M$ the observed inputs, and $Q_i > 0$ the observed outputs for a set of $|N|$ firm observations. The data set can be a cross-section, a time-series or, as in our own empirical application, a panel with firm observations specified at the firm-year level. The set S contains all information on observed production behavior that is used by the empirical analyst. In principle, it is possible to integrate in our set-up extra information on indicators that are (assumed to be) correlated with the unobserved technological heterogeneity (e.g., R&D investments). Again, we intentionally restrict to our minimalistic setting to show the generality of our identification results.

The functional form of the production function F is unknown to the empirical analyst. Our nonparametric method basically checks whether there exists at least one specification of

F that represents the observed firm behavior in terms of the optimization problems (OP.I) and (OP.II). If such a function exists, we say that the data set S is rationalizable in terms of (OP.I) and (OP.II).

Definition 1. Let $S = \{\mathbf{W}_i, \mathbf{V}_i, Q_i\}_{i \in N}$ be a given data set. S is (OP.I)-rationalizable if there exist numbers $\Omega_i \in \mathbb{R}_{++}$ and a strictly increasing, continuous and quasi-concave production function $F : \mathbb{R}_+^{M+1} \rightarrow \mathbb{R}_+$ which is CRS in \mathbf{V} and such that, for all firm observations $i \in N$:

$$\mathbf{V}_i \in \arg \min_{\mathbf{V}} \mathbf{W}'_i \mathbf{V} \text{ s.t. } F(\mathbf{V}, \Omega_i) \geq Q_i.$$

The data set S is (OP.II)-rationalizable if, in addition, there exist prices $\Gamma_i \in \mathbb{R}_+$ such that, for all firm observations $i \in N$:

$$(\mathbf{V}_i, \Omega_i) \in \arg \min_{\mathbf{V}, \Omega} \{\mathbf{W}'_i \mathbf{V} + \Gamma_i \Omega\} \text{ s.t. } F(\mathbf{V}, \Omega) \geq Q_i.$$

The following proposition states that (OP.I)-rationalizability and (OP.II)-rationalizability generate exactly the same nonparametric testable implications for a given data set S .

Proposition 1. Let $S = \{\mathbf{W}_i, \mathbf{V}_i, Q_i\}_{i \in N}$ be a given data set. The following statements are equivalent:

- (i) The data set S is (OP.I)-rationalizable;
- (ii) The data set S is (OP.II)-rationalizable;
- (iii) There exist $\Omega_i \in \mathbb{R}_{++}$ and $\Lambda_i \in \mathbb{R}_{++}$ that satisfy, for all $i, j \in N$, the inequalities:

$$\Lambda_j \mathbf{W}'_j \mathbf{V}_j + \Omega_j \leq \Lambda_j \mathbf{W}'_j \left(\frac{Q_j}{Q_i} \mathbf{V}_i \right) + \Omega_i.$$

Appendix A contains the proof of Proposition 1. To sharpen the intuition behind condition (iii), we start from the observation that the input bundle (\mathbf{V}_i, Ω_i) can produce the output Q_i . Then, it follows from our RTS assumption that, given Ω_i , the rescaled input bundle $\frac{Q_j}{Q_i} \times \mathbf{V}_i$ must be able to produce the output level Q_j . Under (OP.II)-rationalizability,

the cost of using this last input combination at the prices (W_j, Γ_j) equals:

$$\mathbf{W}'_j \left(\frac{Q_j}{Q_i} \mathbf{V}_i \right) + \Gamma_j \Omega_i.$$

On the other hand, cost minimizing production behavior (as specified in (OP.II)) also requires that, at the prices (\mathbf{W}_j, Γ_j) , the input bundle (\mathbf{V}_j, Ω_j) produces the output Q_j at a lower or equal cost. Thus, we must have:

$$\mathbf{W}'_j \mathbf{V}_j + \Gamma_j \Omega_j \leq \mathbf{W}'_j \left(\frac{Q_j}{Q_i} \mathbf{V}_i \right) + \Gamma_j \Omega_i.$$

Using $\Lambda_j = 1/\Gamma_j$, then effectively obtains condition (iii) in Proposition 1. Interestingly, the inequalities in condition (iii) are linear in unknowns (Ω_i and Λ_i), which means that they can be checked by standard linear programming techniques.

In Section 2.5 we show how the inequalities in condition (iii) of Proposition 1 can be used in practice to specify values of Ω_i that rationalize the data set S . Moreover, when interpreting these numbers Ω_i as representing latent input quantities, the associated numbers Λ_i give the inverse of the corresponding shadow input prices ($1/\Gamma_i$). Interestingly, we can use this to nonparametrically quantify productivity heterogeneity in terms of the latent input, which we refer to as our nonparametric estimate of productivity. It readily follows from our above discussion that Ω has a direct interpretation as capturing productivity differences. All else equal, higher Ω values indicate that the same output can be produced with less observed costs, which effectively reveals a higher (unobserved) productivity level.

As a concluding note, Appendix A provides a numerical example that illustrates the testable implications in Proposition 1. It shows that our empirical conditions for cost minimization with unobserved productivity differences can be falsified (i.e., have empirical content) even in a minimalistic setting with only two firm observations and two observed inputs. Generally, the empirical bite of the conditions will increase with the number of observations and observed inputs.

2.3 Short run cost minimization

So far, we limited our discussion to statically optimizing firms choosing only flexible inputs that are freely adjustable. In our empirical application, however, we will consider firms facing a dynamic optimization problem in which capital is a dynamic input that is accumulated via investments and characterized by adjustment costs. We next show how to adapt our previous reasoning to this modified setting. For the sake of brevity, and because we believe the required extensions are straightforward, we will not explicitly state the formal analogues of Definition 1 and Proposition 1.

The relevant production function becomes:

$$Q = F(\mathbf{V}, K, \Omega),$$

where \mathbf{V} is again a vector of variable inputs that are freely adjustable by the firm (such as labor and intermediate inputs) and K denotes capital stock which may face adjustment costs and is fixed in the short run. Once more, we assume CRS on the observed inputs, meaning that, for all $t > 0$,

$$F(t\mathbf{V}, tK, \Omega) = tF(\mathbf{V}, K, \Omega).$$

Like before, Ω and F are unobserved by the empirical analyst. Thus, for a set N of firm observations, we now have a data set:

$$S = \{\mathbf{W}_i, \mathbf{V}_i, K_i, Q_i\}_{i \in N},$$

with $K_i > 0$ the observed capital stock for each firm observation $i \in N$.

For the given data set S , we can define the analogues of the rationalizability requirements in condition (iii) of Proposition 1 in a similar way as before. From the observation that the input bundle $(\mathbf{V}_i, K_i, \Omega_i)$ can produce the output Q_i , our CRS assumption implies that the rescaled input bundle $\left(\frac{Q_j}{Q_i} \times \mathbf{V}_i, \frac{Q_j}{Q_i} \times K_i, \Omega_i\right)$ can produce the output level Q_j . Then, short run cost minimizing production behavior requires that, at the prices (\mathbf{W}_j, Γ_j) , the variable input bundle (\mathbf{V}_j, Ω_j) , under fixed K_j , produces the output Q_j at a lower or equal variable cost. Fixed input K is together with \mathbf{V} subject to a CRS assumption, limiting the

comparison partners of observation j to the rescaled input bundle of observations i with $\frac{K_j}{K_i} = \frac{Q_j}{Q_i}$. For these comparison observations i , we obtain the rationalizability restriction:

$$\mathbf{W}'_j \mathbf{V}_j + \Gamma_j \Omega_j \leq \frac{Q_j}{Q_i} (\mathbf{W}'_j \mathbf{V}_i) + \Gamma_j \Omega_i.$$

Again using $\Lambda_j = 1/\Gamma_j$, we get that short run cost minimizing behavior is characterized by feasibility of the inequalities:

$$\Lambda_j \mathbf{W}'_j \mathbf{V}_j + \Omega_j \leq \frac{Q_j}{Q_i} (\Lambda_j \mathbf{W}'_j \mathbf{V}_i) + \Omega_i, \quad (1)$$

for all $i, j \in N$ such that $\frac{K_j}{K_i} = \frac{Q_j}{Q_i}$. Once more, these inequalities are linear in unknowns, making them easy to use in practice.

Obviously, limiting the comparison partners of observation j to the rescaled input bundle of observations i with $\frac{K_j}{K_i} = \frac{Q_j}{Q_i}$ is too data demanding. Therefore, in practice some approximation is needed such that observations i with $\frac{K_j}{K_i} \approx \frac{Q_j}{Q_i}$ are considered. More specifically, our approach to select comparison observations via kernels with compact support, together with estimating the appropriate bandwidth size via leave-one-out least squares cross-validation, is in the spirit of conditional full-frontier measurement. See Daraio and Simar (2007) for a methodological overview, and Dewitte et al. (2020) for an application of kernel weighting in the context of nonparametric technical change measurement under moderate noise.⁸

2.4 Imposing Hicks neutrality

In our following empirical application, we will also assess the empirical validity of the Hicks neutrality assumption for the production setting at hand. To this end, we next present the additional rationalizability restrictions that need to be satisfied under this assumption. Particularly, we will introduce the restrictions that allow us to interpret Ω as equalling Hicks neutral total factor productivity.

⁸In Online Appendix A we provide more discussion of kernel-based conditioning.

As a first step, we use that imposing a Hicks neutral technology implies that Ω is multiplicatively separable. Formally, for any $t > 0$ we have:

$$F(\mathbf{V}, K, t\Omega) = t^{\gamma^\Omega} F(\mathbf{V}, K, \Omega).$$

Then, when setting $\gamma^\Omega = 1$, we obtain that Ω has the attractive feature that it equals total factor productivity (i.e., the ratio of output over input):

$$\Omega = \frac{F(\mathbf{V}, K, \Omega)}{F(\mathbf{V}, K, 1)}.$$

Thus, a readily similar argument as above yields that Ω represents Hicks neutral total factor productivity (by using $\gamma^\Omega = 1$) only if the data set S meets the following rationalizability constraints:

$$\Lambda_j \mathbf{W}'_j \mathbf{V}_j + \Omega_j \leq \Lambda_j \mathbf{W}'_j \mathbf{V}_i + \frac{Q_j}{Q_i} \Omega_i, \text{ for all } i, j \in N \text{ with } K_j = K_i,$$

where we treat \mathbf{V} as variable input and K as fixed input, like before.

2.5 Bringing our methodology to data

We conclude this section by discussing three practical issues that pertain to the empirical application of our nonparametric methodology. First, we show how we can account for small deviations from exactly optimizing behavior by making use of a nonparametric goodness-of-fit parameter. Next, we explain how we can use the testable conditions outlined above to identify the unobserved productivity factor Ω . Finally, we indicate how to address incomplete information on input prices.

Exactly versus nearly optimizing behavior. The rationalizability conditions in Proposition 1 are strict: either the data set S satisfies them “exactly” or it does not. In practice, it is often useful to allow for small deviations from exactly rationalizable behavior. Such deviations may be due to (small) unanticipated shocks experienced by the firms or, alternatively, data imperfections (e.g., ill-measured input/output quantities or deviations from

optimizing behavior).⁹ To include these possibilities, we define a nonparametric goodness-of-fit parameter that has an intuitive economic interpretation in terms of departures from the cost minimization hypothesis that we maintain as our core identifying assumption (see Afriat (1972) and Varian (1990) for more discussion). By fixing our goodness-of-fit parameter at a value close to (but different from) one, we can take account of observed behavior that is close to (but not exactly) rationalizable in the sense of Definition 1.

More precisely, we increase the right hand sides of the inequalities (1) by using the goodness-of-fit parameter θ (with $\theta \geq 1$), to obtain:

$$\Lambda_j \mathbf{W}'_j \mathbf{V}_j + \Omega_j \leq \theta \left(\frac{Q_j}{Q_i} (\Lambda_j \mathbf{W}'_j \mathbf{V}_i) + \Omega_i \right). \quad (2)$$

Obviously $\theta = 1$ yields the exact conditions in (1), while higher values for θ weaken the rationalizability requirements. For a fixed value of θ , this defines restrictions that are linear in the unknowns Λ_i and Ω_i . We can use simple linear programming tools to check if there exists a solution and, thus, to conclude if the data set is exactly rationalizable (when using $\theta = 1$) or nearly rationalizable (when using $\theta > 1$).

Identifying productivity. The linear restrictions (2) will generally define a multitude of feasible specifications of the unknowns Λ_i and Ω_i . However, since our rationalizability conditions are falsifiable (see Appendix A), we can use these conditions to recover lower and upper bounds on these unknowns (i.e., set identification). To empirically evaluate the importance of unobserved heterogeneity in firm productivity, a natural choice is to use the specification that minimizes the average cost share of the latent inputs that are required for rationalizability (as characterized in (2)) or, equivalently, that maximizes the role played by

⁹In fact, it is also possible to explicitly account for measurement errors in prices and quantities in our nonparametric analysis. e.g., we may use the procedure suggested by Varian (1985) and Epstein and Yatchew (1985), which is fairly easily adjusted to our setting. This complies with the more standard econometric use of a minimum distance criterion. To facilitate our exposition, we will not consider this extension in this paper. In the current context, measurement error in the output quantity can also be interpreted as reflecting productivity shocks that are not anticipated by the firm. Our simulation exercise in Appendix B includes output error and shows how our goodness-of-fit parameter θ incorporates this error.

the observed inputs. This corresponds to minimizing the objective:

$$\frac{1}{|N|} \sum_{i \in N} \frac{\Gamma_i \Omega_i}{\mathbf{W}'_i \mathbf{V}_i + \Gamma_i \Omega_i} = \frac{1}{|N|} \sum_{i \in N} \frac{(1/\Lambda_i) \Omega_i}{\mathbf{W}'_i \mathbf{V}_i + (1/\Lambda_i) \Omega_i}, \quad (3)$$

subject to the linear restrictions (2). In Online Appendix A, we present a binary search routine to deal with the technical issue that the objective function in (3) is not linear in unknowns. This routine allows us to use standard linear programming techniques in our empirical analysis.

Incomplete information on input prices. The recent literature on the identification of heterogeneity in firm productivity acknowledges that the empirical analyst's information on the flexible input prices is mostly partial (Klette and Griliches (1996), Foster et al. (2008) and Collard-Wexler and De Loecker (2016)). Specifically, while wages are usually observed, there is often incomplete information on intermediate input prices. To formalize this setting, let $\mathbf{V} = [V^1, \dots, V^M]$ denote the flexible input quantities and $\mathbf{W} = [W^1, \dots, W^M]$ the corresponding prices. We use that W^1 captures the observed wages and W^2 to W^M the unobserved intermediate input prices. We can then approximate the unobserved intermediate input prices W^2 to W^M by shadow prices that are endogenously defined.

Specifically, in (2) we use $\alpha_j^m = \Lambda_j W_j^m$ for $m = 2, \dots, M$. This obtains the restrictions:

$$\Lambda_j W_j^1 V_j^1 + \sum_{m=2}^M \alpha_j^m V_j^m + \Omega_j \leq \theta \left(\frac{Q_j}{Q_i} \left(\Lambda_j W_j^1 V_i^1 + \sum_{m=2}^M \alpha_j^m V_i^m \right) + \Omega_i \right),$$

which remain linear in unknowns (Λ_j , Ω_j and α_j^m) for a given value of the goodness-of-fit parameter θ .

In our following application, we will strengthen our identification analysis by imposing observation-specific lower and upper bounds on the unobserved prices W^2, \dots, W^M . We will ensure that the shadow prices are situated reasonably close to observed proxies $\widehat{W}^2, \dots, \widehat{W}^M$ by using the parameter β (with $0 \leq \beta \leq 1$), to specify the constraints (for $m = 2, \dots, M$):

$$(1 - \beta) \widehat{W}_i^m \leq \frac{\alpha_i^m}{\Lambda_i} \leq (1 + \beta) \widehat{W}_i^m.$$

As a further strengthening, we will impose additional structure on the (shadow) price of the latent productivity factor Γ . We will partition our sample N into subsets G capturing observations of the same firm, and we assume that the price Γ is constant for all firm observations $i \in G$, which gives the restrictions $\Lambda_j = \Lambda_i$ for all $i, j \in G$.

3 Application set-up and data

Our empirical analysis makes use of a rich micro-level data structure on production networks of Belgian manufacturing firms in the period 2002-2014. The data structure brings together information drawn from four comprehensive panel-level data sets that all cover the respective period: (i) the National Bank of Belgium’s Central Balance Sheet Office (CBSO) (ii) the Belgian Prodcom Survey, (iii) the National Bank of Belgium (NBB) B2B Transactions data set, and (iv) the International Trade data at the NBB. As described in detail in Dhyne et al. (2015), Duprez and Magerman (2018), the resulting data set provides unique quantity and price information on firms’ inputs and outputs. See Online Appendix C for more details.

On the output-side, our data set is based on the Belgian Prodcom Survey. As such, we can measure productivity in terms of output quantities, yielding productivity estimates with a technological interpretation that are free from demand structuring assumptions.¹⁰ We use produced quantity of manufacturing goods as our output Q . By using output quantity data, we avoid the issues that relate to revenue-based total factor productivity estimation. Particularly, revenue-based estimates risk to not only include the pure technological features of the firm (e.g., innovation, intangibles and managerial quality) and to be confounded by potential influences from firm-level price setting behavior on the output market. See Klette and Griliches (1996), Foster et al. (2008), De Loecker (2011) and De Loecker and Warzynski (2012) for discussions.

On the input side, we use labor (in full time equivalents) and intermediates as flexible inputs (with quantities \mathbf{V} and prices \mathbf{W}). We define labor use V^{labor} in terms of FTE’s employed and wages W^{labor} as total labor costs divided by labor use. For intermediate inputs, we only have imperfect proxies of intermediate input prices. In this respect, an important consideration is that using industry-wide deflators for intermediate inputs, which is common

¹⁰Recent studies that make use of the Belgian Prodcom database include De Loecker et al. (2014), Dhyne et al. (2014), Forlani et al. (2016), Bernard et al. (2019) and Bernard et al. (Forthcoming).

practice in the empirical productivity literature, implies vulnerability of productivity estimates for confounding (imported) input price heterogeneity; see, e.g., Goldberg et al. (2010), Halpern et al. (2015) and Duprez and Magerman (2018). Our detailed data structure therefore uses firm-to-firm linkages as included in the NBB B2B Transactions data set and the International Trade data at the NBB, to provide a unique granularity on intermediate input quantities and prices. In particular, we make use of a Theil-type Paasche price index to disaggregate intermediate input information to the level of i) imported intermediates, ii) domestically produced intermediates, iii) purchased services, and iv) wholesalers, retailers and energy supply. As such, we obtain \widehat{W}^{imp} , \widehat{W}^{dom} , \widehat{W}^{serv} , $\widehat{W}^{wholesale}$, \widehat{V}^{imp} , \widehat{V}^{dom} , \widehat{V}^{serv} and $\widehat{W}^{wholesale}$.¹¹

For all four intermediate inputs, we account for potential errors in the input prices by using the method that we introduced in Section 2.5. We estimate intermediate input prices as shadow prices that are endogenously defined, and we ensure that these shadow prices are sufficiently close to the constructed intermediate input price indices by setting the parameter β equal to 0.1.¹²

Next, we measure our fixed input capital K as tangible fixed assets, which we deflate with a sector-year specific gross fixed capital formation (gfcf) deflator. Given the empirical difficulties related to the measurement of capital and the associated implications for productivity estimation (see, e.g., Collard-Wexler and De Loecker (2016) and references therein), we consider capital as a difficult-to-measure dynamic input, facing potentially non-linear adjustment costs. As explained in the previous section, our empirical analysis builds on restrictions for short term cost minimization that are conditional on (a ratio of) capital stock. In our empirical implementation of these restrictions, we make use of kernel-based conditioning to select comparison firm observations with similar (instead of exactly equal) capital stock to output ratio. This kernel weighting approach also allows for limited mismeasurement of capital stock.

To account for unobserved heterogeneity related to the size of production, we separately apply our identification methodology to three firm size groups: small, medium and large

¹¹Imposing parametric structure within the subcomponents of the intermediate input indices does not imply parametric structure on the production function $Q = F(\mathbf{V}, K, \Omega)$ under consideration. See Online Appendix B for a detailed discussion.

¹²Sensitivity analysis shows that our main results also hold when we assume that the intermediate input prices are perfectly observed (i.e., $\beta = 0$). See Online Appendix D.

firms. We follow the European Union SME definition to characterize these groups.¹³ Within each firm size group, we conduct our estimation for sectors with at least 250 firm observations. Sectors are specified according to the NACE aggregates in the National Accounts system (A64); see Table 1. Our sample contains 10,995 small firms, 7,601 medium firms and 1,180 large firms in total. Tables 1 and 2 provide summary information on the firms under study; see also Online Appendix C for additional details. In line with the smaller sample size of the raw data over time, Table 2 shows a decreasing sample size of the cleaned data over time. Therefore, to control for the changing sample size, we will also provide the results of our main regression exercises (reported in Tables 3 and 6 below) for a balanced panel in Online Appendix D.

Table 1: Firm observations and firms per sector and firm size group

Sector	Small		Medium		Large	
	Obs.	Firms	Obs.	Firms	Obs.	Firms
16 - Wood and products of wood and cork, except furniture; articles of straw and plaiting materials	655	149	309	58	–	–
17 - Paper and paper products	441	88	315	56	–	–
20 - Chemicals and chemical products	775	150	785	139	519	85
22 - Rubber and plastics products	765	151	597	114	–	–
23 - Other non-metallic mineral products	806	203	647	117	–	–
24 - Basic metals	323	85	372	75	–	–
25 - Fabricated metal products, except machinery and equipment	1855	441	1021	211	–	–
26 - Computer, electronic and optical products	285	63	267	61	–	–
27 - Electrical equipment	314	76	–	–	–	–
28 - Machinery and equipment n.e.c.	745	166	425	92	–	–
10-12 - Food products; beverages; tobacco products	1911	399	1667	294	661	120
13-15 - Textiles; wearing apparel; leather and related products	1432	358	819	167	–	–
31-32 - Furniture; other manufactured goods	688	157	377	72	–	–
Total	10995	2486	7601	1456	1180	205

¹³We define small firms to have labor quantity below 50, and either sales or total assets lower than or equal to 10mln euro. Medium firms are defined to have labor quantity between 50 and 250, and either sales lower than or equal to 50mln euro or total assets lower than or equal to 43mln euro. See ec.europa.eu/growth/smes/sme-definition_en.

Table 2: Firm observations per year and firm size group

Year	Small	Medium	Large	Total
2002	1276	738	94	2108
2003	1159	643	84	1886
2004	1145	641	88	1874
2005	1069	611	85	1765
2006	1045	592	100	1737
2007	953	592	100	1645
2008	718	576	93	1387
2009	746	554	79	1379
2010	658	569	95	1322
2011	573	565	96	1234
2012	610	547	100	1257
2013	522	470	80	1072
2014	521	503	86	1110
Total	10995	7601	1180	19776

Figure 1 summarizes the observed cost shares of the flexible inputs for small, medium and large firms. For small firms, we document an increase in the cost share of foreign intermediates of 35 percent, from 28.68 in 2002 to 35.84 in 2014.¹⁴ Domestic sourcing reduced with 29 percent, while service sourcing and wholesale sourcing remained largely constant, as we observe a modest decline of a bit more than 2 percent over the time span under study. Further, the labor cost share experienced a decline of slightly less than 11 percent, from 23.82 in 2002 to 21.29 in 2014. For medium firms, the cost share changes show a similar pattern, with a 18 percent increase in foreign sourcing, a 31 percent decline in domestic sourcing, largely constant cost shares for servicing and wholesale expenditures, and a decline of 9 percent for labor. Finally, while our small sample size makes us cautious to interpret the cost share changes for large firms, we may safely conclude that the observed cost shares remained fairly stable over the considered time span.

Cost share changes can be the result of within-firm changes or reallocation (see, e.g., the recent micro-analysis of labor revenue shares by Kehrig and Vincent (2021) for U.S. firms). For the purpose of our paper and given the limitations of our sample in terms


¹⁴Our sensitivity analysis in Online Appendix D shows that the variable cost share changes are not driven by differences in capital intensity.

of representativeness for the population of manufacturing firms, we will mainly focus on within-firm changes. We report regression results on the relationship within-firm between variable cost shares and year of observation in Table 3.¹⁵ We notice a sharp within-firm technological change in favor of foreign intermediates and against domestic intermediate inputs in small/medium firms. In the period before the 08-09 financial crisis, the changes in observed cost shares are most pronounced. In the pre-crisis period, input substitution was also significantly against labor. For large firms, the subcomponents of intermediate inputs show no significant change over time (at the 5 percent significance level) for the 2002-2014 period. For these firms, we do notice a significant within-firm substitution in favor of intermediate inputs and against labor.

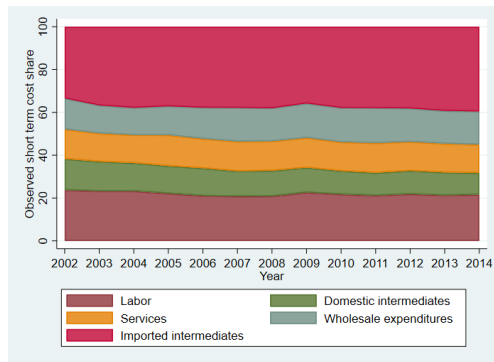
All in all, our descriptive statistics show that variable input cost shares are changing over time. Arguably, these changes cannot be explained by input price changes, which highlights the importance of accounting for the possibility of biased technological change when studying firm heterogeneity in productivity. This observation directly motivates the empirical usefulness of our methodology, which is nonparametric and does not require the assumption of Hicks neutrality.

As a final note, a main shortcoming of our data structure is that it does not include information on quality. On the input side, we resolve this issue –at least partially– by using firm-year disaggregate price information on intermediate inputs. On the output side, we will account for product quality heterogeneity in our following analysis by excluding firm observations with both lower output prices and lower market shares when evaluating cost minimization of a given firm. This follows the argument of De Loecker et al. (2016) that it are especially these firm observations that are expected to produce lower quality products (corresponding to less costly production).


¹⁵The reported patterns remain unaffected when using endogenously defined shadow prices instead of observed prices for the intermediate inputs. These regression results are available upon request.



(a) Small



(b) Medium



(c) Large

Figure 1: Observed cost shares, totalling to 100 each year, over the time span 2002-2014

Table 3: Within-firm temporal variation of the observed variable cost shares

	All years (2002-2014)			Pre-crisis (2002-2007)			Post-crisis (2010-2014)		
	Small	Medium	Large	Small	Medium	Large	Small	Medium	Large
	Dependent variable: Cost share labor								
Year	-0.0427*	-0.00125	-0.241***	-0.412***	-0.408***	-0.391***	0.0768	0.0919**	0.00795
	(0.0232)	(0.0242)	(0.0435)	(0.0391)	(0.0414)	(0.0689)	(0.0515)	(0.0428)	(0.0639)
Constant	108.2**	24.52	498.0***	848.7***	840.0***	800.0***	-132.8	-163.4*	-2.256
	(46.56)	(48.60)	(87.35)	(78.45)	(82.94)	(138.0)	(103.7)	(86.10)	(128.6)
	Dependent variable: Cost share imported intermediates								
Year	0.333***	0.467***	0.268*	0.699***	1.199***	-0.296	0.555***	0.0682	-0.0847
	(0.0594)	(0.0593)	(0.136)	(0.0992)	(0.115)	(0.321)	(0.147)	(0.106)	(0.223)
Constant	-635.4***	-900.0***	-495.2*	-1,369***	-2,367***	634.2	-1,083***	-98.72	212.8
	(119.2)	(119.0)	(274.0)	(198.8)	(230.4)	(644.1)	(296.0)	(213.8)	(449.7)
	Dependent variable: Cost share domestic intermediates								
Year	-0.371***	-0.413***	-0.141	-0.401***	-0.752***	-0.331*	-0.147*	-0.110	-0.208
	(0.0402)	(0.0428)	(0.0904)	(0.0688)	(0.0889)	(0.194)	(0.0864)	(0.0712)	(0.298)
Constant	755.3***	840.4***	296.9	816.5***	1,520***	676.2*	305.9*	231.5	430.9
	(80.69)	(85.83)	(181.5)	(137.9)	(178.3)	(389.5)	(173.7)	(143.2)	(599.4)
	Dependent variable: Cost share services inputs								
Year	0.0217	-0.115***	-0.108	0.0478	-0.0440	-0.0843	-0.121	-0.0282	0.159
	(0.0378)	(0.0418)	(0.0848)	(0.0718)	(0.0882)	(0.171)	(0.0840)	(0.0675)	(0.310)
Constant	-27.86	244.0***	229.3	-80.27	101.9	181.8	259.5	70.30	-308.1
	(75.77)	(83.87)	(170.4)	(143.9)	(176.8)	(342.8)	(169.0)	(135.8)	(624.2)
	Dependent variable: Cost share wholesale expenditures								
Year	0.0592	0.0617	0.222*	0.0670	0.00414	1.102***	-0.363***	-0.0221	0.125
	(0.0391)	(0.0381)	(0.123)	(0.0711)	(0.0789)	(0.313)	(0.0949)	(0.0809)	(0.197)
Constant	-100.3	-109.0	-429.0*	-116.1	5.682	-2,192***	750.2***	60.26	-233.4
	(78.52)	(76.55)	(246.4)	(142.5)	(158.1)	(627.8)	(191.0)	(162.7)	(396.1)
	Dependent variable: Cost share intermediate inputs								
Year	0.0427*	0.00125	0.241***	0.412***	0.408***	0.391***	-0.0768	-0.0919**	-0.00795
	(0.0232)	(0.0242)	(0.0435)	(0.0391)	(0.0414)	(0.0689)	(0.0515)	(0.0428)	(0.0639)
Constant	-8.180	75.48	-398.0***	-748.7***	-740.0***	-700.0***	232.8**	263.4***	102.3
	(46.56)	(48.60)	(87.35)	(78.45)	(82.94)	(138.0)	(103.7)	(86.10)	(128.6)
Firm observations	10,995	7,601	1,180	6,647	3,817	551	2,884	2,654	457
Firms	2,486	1,456	205	2,099	1,110	146	1,096	896	154

Notes: The dependent variable is the observed variable cost share (times 100) of respectively labor, imported intermediates, domestic intermediates, purchased services, wholesale expenditures, intermediates. All regressions use firm-level fixed effects. Robust standard errors in parentheses.*** p < 0.01, ** p < 0.05, * p < 0.1.

4 Empirical results

A main distinguishing feature of our novel methodology is that it can relax the frequently made –but often debatable– assumption of Hicks neutrality. We start our following empirical analysis by comparing the goodness-of-fit of our rationalizability conditions with and without imposing Hicks neutrality.¹⁶ This will provide nonparametric evidence against Hicks neutral technical change, so complementing our descriptive findings in Figure 1 and Table 3. In a second step, we then present a descriptive analysis of our nonparametric estimates of unobserved firm productivity while accounting for biased technological change. We specifically consider the evolution of productivity over time and investigate patterns of substitution between latent productivity and the observed variable inputs. In a final exercise, we exploit

¹⁶Our linear programming problems are operationalized in MATLAB, with links to the IBM ILOG CPLEX Optimization Studio. All our codes are available upon request.

the granular nature of our data set to assess the potential bias that originates from using a common-scale intermediate inputs deflator. We will focus on our main results over all considered sectors in the current section, and we refer to Online Appendix D for sector-specific estimates. This appendix also contains a sensitivity analysis that shows robustness of our principal findings for small changes of the RTS assumption (by setting γ equal to respectively 0.9 and 1.1), for assuming perfect knowledge about intermediate input prices (by setting $\beta = 0$), and for only considering firm observations with at least 20 comparison observations.

Empirical evidence against Hicks neutrality. As a first step of our analysis, we assess the empirical validity of the Hicks neutrality assumption by using our goodness-of-fit parameter θ . Particularly, we compute the value of θ that allows us to rationalize the observed firm behavior with and without this assumption of Hicks neutrality. Throughout our analysis, we assume a constant (shadow) price of the latent productivity factor (I) for all observations of the same firm.

When allowing for biased technological change, we find that setting $\theta = 1.05$ can rationalize the observed production behavior (as “nearly optimizing”) of any firm group (listed in Table 1) that we study. We can compare this outcome with the results in Table 4, which contains descriptive statistics on the smallest θ -values needed to rationalize the observed firm behavior as short-term cost minimizing under Hicks neutrality. For the different sectors under consideration, θ^{HN} and $\theta^{DMAT,HN}$ represent these goodness-of-fit values when using, respectively, granular intermediate input price data and a common-scale intermediate inputs deflator. If the Hicks neutrality assumption effectively holds (i.e., does not imply additionally binding rationalizability constraints), these goodness-of-fit values should not exceed 1.05.

When using this criterion, Table 4 provides clear evidence against Hicks neutrality for the production setting under study. In particular, no firm group can be rationalized as nearly cost minimizing under Hicks neutrality when using a goodness-of-fit value that is as small as 1.05. The smallest value of θ^{HN} amounts to 1.08; and its average for any firm size group is even above 1.16. Interestingly, Table 4 also reveals stronger evidence against Hicks neutrality when using more granular input quantity and price data than when using a common-scale intermediate input price deflator.

Table 4: Goodness-of-fit when assuming Hicks neutrality

	Mean	St.Dev.	Min.	Med.	Max.
Small firms					
θ^{HN}	1.195	0.051	1.130	1.130	1.450
$\theta^{DMAT,HN}$	1.132	0.058	1.080	1.080	1.430
Medium firms					
θ^{HN}	1.165	0.045	1.080	1.080	1.230
$\theta^{DMAT,HN}$	1.098	0.021	1.070	1.070	1.130
Large firms					
θ^{HN}	1.162	0.025	1.140	1.140	1.190
$\theta^{DMAT,HN}$	1.083	0.015	1.070	1.070	1.100

Notes: To construct the descriptive statistics of the sector-specific goodness-of-fit measures, we weight according to the number of firm observations.

Productivity growth and changes in input usage. From our first step analysis, we conclude that we can rationalize firm behavior under biased technological change when setting $\theta = 1.05$. Under these conditions we next focus on identifying unobserved heterogeneity in firm productivity.

From Table 5, we learn that productivity is highly persistent over time, which is indeed what we could expect a priori. When studying within-firm growth via a fixed effects regression, we find non-increasing to decreasing productivity. This is shown in Table 6, which uses as dependent variables both $\log(\Omega)$ and the cost share of the latent input (defined as $\Gamma\Omega/(\Gamma\Omega + \mathbf{W}'\mathbf{V})$). These results provide a nonparametric confirmation of the well-documented productivity slowdown in manufacturing since the early 2000s. Relatedly, when interpreting productivity as a latent input factor, Figure 2 shows that the cost share of this latent input (indicating the importance of unobserved productivity) gradually decreases over time, while the share of (particularly foreign) intermediates increases. This suggests biased technological change in favor of intermediate inputs, in line with the well-documented trend towards outsourcing and offshoring. Finally, when taking into account latent inputs, our empirical results suggest no robust evidence that technological change is biased against labor for the sample of firms that we study.

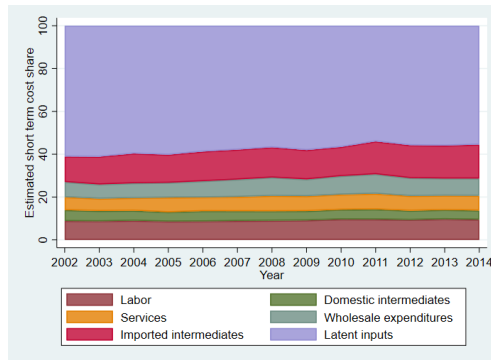
Table 5: Persistence: Spearman correlation with lagged variable

	Small	Medium	Large
$Log(\Omega)$	0.914	0.921	0.891
$Log(\Gamma\Omega)$	0.920	0.934	0.944


Table 6: Within-firm temporal variation of log productivity and the total variable cost share of latent inputs

	All years (2002-2014)			Pre-crisis (2002-2007)			Post-crisis (2010-2014)		
	Small	Medium	Large	Small	Medium	Large	Small	Medium	Large
	Dependent variable: $log(\Omega)$								
Year	-0.027*** (0.005)	-0.024*** (0.005)	-0.032** (0.014)	-0.039*** (0.008)	-0.004 (0.009)	-0.024 (0.023)	-0.023 (0.015)	-0.058*** (0.013)	0.005 (0.031)
Constant	57.672*** (9.543)	53.489*** (10.261)	70.917** (28.191)	80.398*** (15.879)	12.960 (18.605)	56.108 (45.898)	49.211 (30.386)	120.345*** (26.421)	-3.721 (61.558)
	Dependent variable: total variable cost share of latent inputs ($\times 100$)								
Year	-0.980*** (0.084)	-0.933*** (0.089)	-1.251*** (0.225)	-1.391*** (0.136)	-1.020*** (0.154)	-1.282*** (0.418)	-0.422* (0.245)	-1.208*** (0.207)	-0.489 (0.435)
Constant	2,024.392*** (168.619)	1,928.532*** (178.565)	2,575.685*** (451.723)	2,847.237*** (271.685)	2,102.504*** (309.543)	2,638.289*** (838.881)	905.431* (493.005)	2,484.619*** (417.401)	1,043.358 (875.843)
Firm obs.	10,995	7,601	1,180	6,647	3,817	551	2,884	2,654	457
Firms	2,486	1,456	205	2,099	1,110	146	1,096	896	154


Notes: The dependent variable is respectively log productivity and the total variable cost share of latent inputs (times 100). All regressions use firm-level fixed effects. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.



(a) Small



(b) Medium



(c) Large

Figure 2: Estimated cost shares for given firm size

Using a common-scale intermediate inputs deflator. From above, we conclude that using granular price and quantity data for intermediate inputs does not affect the well-

established finding of a productivity slowdown since the early 2000s. In this respect, an interesting feature of our detailed data is that they allow us to investigate the value-added of using granular data. We can do so by assessing the bias that would be introduced when using less granular data for intermediate inputs, so assessing the need to use decomposed intermediate input price and quantity indices.

To address this question, we compare the results for the latent productivity factor Ω that we reported above with those obtained for a one-dimensional sector-year specific deflator as price index and deflated intermediate expenditures as quantity index. We let Ω^{DMAT} denote productivity identified in this new exercise. Table 7 reveals that the log growth of Ω^{DMAT} correlates only moderately with the log growth of Ω . Moreover and importantly, as shown in Table 8, the absolute difference between the log growth of Ω and Ω^{DMAT} is non-random; it significantly relates to capital usage and foreign sourcing. These results highlight the risk of generating biased empirical results when using a common-scale deflator as price index for intermediate inputs. This conclusion complements the findings of Foster et al. (2008), who showed the potential bias implied by the use of a common-scale output deflator. In our opinion, these findings provide once more a strong empirical argument pro using of granular price and quantity data in empirical production analysis.

Table 7: Spearman correlation between the log growth of Ω and Ω^{DMAT}

Firm size	Log growth Ω^{DMAT}
Small	0.540
Medium	0.572
Large	0.450

Table 8: Dependent variable: Absolute difference between the log growth of Ω and Ω^{DMAT}

	All years (2002-2014)			Pre-crisis (2002-2007)			Post-crisis (2010-2014)		
	Small	Medium	Large	Small	Medium	Large	Small	Medium	Large
Year	-0.005*** (0.002)	-0.008*** (0.002)	0.000 (0.004)	-0.005 (0.005)	-0.020*** (0.006)	0.016 (0.013)	-0.005 (0.007)	-0.012** (0.006)	-0.016 (0.016)
Log growth Ω	-0.034* (0.018)	0.061*** (0.017)	0.028 (0.047)	-0.066*** (0.025)	0.059** (0.023)	0.034 (0.096)	0.009 (0.030)	0.063** (0.027)	0.067 (0.058)
Log Ω	-0.033*** (0.009)	-0.089*** (0.010)	-0.047** (0.023)	-0.014 (0.012)	-0.071*** (0.014)	-0.024 (0.044)	-0.083*** (0.016)	-0.104*** (0.016)	-0.095*** (0.035)
Log \hat{V}^{imp}	0.017*** (0.003)	0.019*** (0.005)	0.060*** (0.012)	0.018*** (0.005)	0.028*** (0.007)	0.070*** (0.018)	0.016*** (0.005)	0.024*** (0.007)	0.054*** (0.020)
Log \hat{V}^{dom}	-0.010** (0.004)	0.003 (0.005)	0.017 (0.014)	-0.006 (0.006)	0.002 (0.007)	0.036* (0.019)	-0.020** (0.008)	0.000 (0.008)	-0.001 (0.023)
Log \hat{V}^{serv}	0.017** (0.008)	-0.010 (0.009)	-0.007 (0.026)	0.019* (0.010)	-0.001 (0.013)	0.020 (0.027)	0.020 (0.015)	-0.004 (0.014)	-0.091** (0.045)
Log \hat{V}^{whole}	0.013* (0.008)	0.002 (0.008)	0.022 (0.017)	0.001 (0.010)	-0.008 (0.010)	0.001 (0.028)	0.020 (0.015)	-0.002 (0.014)	0.038 (0.025)
log K	-0.099*** (0.009)	-0.042*** (0.010)	-0.078*** (0.024)	-0.118*** (0.012)	-0.053*** (0.015)	-0.096** (0.041)	-0.051*** (0.017)	-0.032** (0.015)	-0.001 (0.038)
log V^{labor}	0.017 (0.014)	-0.025** (0.012)	-0.011 (0.027)	0.031 (0.019)	-0.047*** (0.017)	-0.016 (0.042)	-0.000 (0.026)	-0.013 (0.019)	-0.003 (0.038)
Constant	9.725*** (3.327)	16.137*** (3.278)	0.298 (7.259)	9.995 (10.200)	39.891*** (11.827)	-31.406 (26.348)	10.321 (13.532)	24.662** (12.125)	32.530 (32.023)
Sector f.e.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm char.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.097	0.105	0.115	0.115	0.099	0.100	0.117	0.132	0.156
Firm obs.	7,561	5,663	913	4,252	2,612	405	2,180	2,159	358

Notes: The dependent variable is the absolute difference between the log growth of Ω and Ω^{DMAT} . The included firm characteristics are exporter, entry and firm age. All regressions use firm-level fixed effects. Robust standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

5 Conclusion

We introduced a novel structural method for production analysis that recovers unobserved productivity in a fully nonparametric fashion. We model unobserved heterogeneity as an unobserved productivity factor on which we condition the demand of the observed inputs. Our method deals with the simultaneity bias in a natural way, and it empirically quantifies productivity differences across firms in terms of differences in latent input. Our nonparametric methodology is easy to implement as it merely requires the use of linear programming techniques. It allows for a powerful identification analysis, while avoiding (nonverifiable and often debatable) assumptions of functional form regarding the relationship between inputs and outputs (including the hypothesis of Hicks neutral technical change).

Our empirical application has shown that the method does allow for drawing strong empirical conclusions, despite its nonparametric nature. For a set of Belgian manufacturing firms, we have recovered productivity differences at the firm-year level over the period 2002-2014 for broad industry categories. We use this to provide robust empirical evidence against the assumption of Hicks neutrality for the setting at hand. Next, in line with the

existing literature, our results suggest a productivity slowdown in manufacturing since the early 2000s. Lastly, we highlight a potential bias when using a common-scale intermediate input price deflator in the estimation of productivity. In our opinion, all this convincingly demonstrates the empirical attractiveness of our methodology.

Finally, from a methodological point of view, we emphasize that we see the current paper primarily as providing a fruitful starting ground, rather than a complete toolkit for nonparametric production analysis with unobserved productivity differences. Most notably, we have focused on a single-output setting throughout. As discussed in De Loecker et al. (2016), a multiproduct framework (also involving the identification of input allocations across products) is warranted to obtain a more detailed insight into influences of exogenous trade or cost shocks. To develop this multi-output version of our methodology, a useful starting point is the study of Cherchye et al. (2014), who presented a nonparametric framework (abstracting from the interdependence between observed input choice dependency and unobserved productivity) for the analysis of firms producing multiple products. A closely related issue concerns dealing with non-competitive output markets. In this respect, Carvajal et al. (2013, 2014) show how to analyze alternative (e.g., Cournot or Bertrand) structures on output markets in the advocated nonparametric framework. In our opinion, integrating these authors' insights with our newly developed methodology may constitute another fertile avenue for follow-up research.

References

- Acemoglu, D., Gancia, G., Zilibotti, F., 2015. Offshoring and directed technical change. *American Economic Journal: Macroeconomics* 7, 84–122.
- Akerberg, D., Caves, K., Frazer, G., 2015. Identification properties of recent production function estimators. *Econometrica* 83, 2411–2451.
- Afriat, S., 1972. Efficiency estimation of production functions. *International Economic Review* 13, 568–598.
- Antras, P., Helpman, E., 2004. Global sourcing. *Journal of Political Economy* 112, 552–580.

- Autor, D., Dorn, D., Katz, L.F., Patterson, C., Van Reenen, J., 2020. The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics* 135, 645–709.
- Autor, D.H., Katz, L., Krueger, A., 1998. Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics* 113, 1169–1213.
- Autor, D.H., Levy, F., Murnane, R.J., 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118, 1279–1333.
- Bernard, A., Blanchard, E., Van Beveren, I., Vandebussche, H., 2019. Carry-along trade. *The Review of Economic Studies* 86, 526–563.
- Bernard, A.B., Dhyne, E., Magerman, G., Manova, K., Moxnes, A., Forthcoming. The origins of firm heterogeneity: A production network approach. *Journal of Political Economy* .
- Carvajal, A., Deb, R., Fenske, J., Quah, K.H., 2013. Revealed preference tests of the cournot model. *Econometrica* 81, 2351–2379.
- Carvajal, A., Deb, R., Fenske, J., Quah, K.H., 2014. A nonparametric analysis of multi-product oligopolies. *Economic Theory* 57, 253–277.
- Cherchye, L., Demuynck, T., De Rock, B., De Witte, K., 2014. Nonparametric analysis of multioutput production with joint inputs. *Economic Journal* 124, 735–775.
- Collard-Wexler, A., De Loecker, J., 2016. Production Function Estimation and Capital Measurement Error. Working paper 22437. National Bureau of Economic Research.
- Daraio, C., Simar, L., 2007. Advanced robust and nonparametric methods in efficiency analysis: Methodology and applications. Springer Science & Business Media.
- De Loecker, J., 2011. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica* 79, 1407–1451.
- De Loecker, J., Fuss, C., Van Biesebroeck, J., 2014. International competition and firm performance : Evidence from Belgium. Working Paper 269. National Bank of Belgium.
- De Loecker, J., Goldberg, P., Khandelwal, A., Pavcnik, N., 2016. Prices, markups, and trade reform. *Econometrica* 84, 445–510.

- De Loecker, J., Syverson, C., 2021. An Industrial Organization Perspective on Productivity. Working paper 29229. National Bureau of Economic Research.
- De Loecker, J., Warzynski, F., 2012. Markups and firm-level export status. *American Economic Review* 102, 2437–2471.
- Dewitte, R., Dumont, M., Merlevede, B., Rayp, G., Vershelde, M., 2020. Firm-heterogeneous biased technological change: A nonparametric approach under endogeneity. *European Journal of Operational Research* 283, 1172–1182.
- Dhyne, E., Magerman, G., Rubínová, S., 2015. The Belgian production network 2002-2012. Working Paper 288. National Bank of Belgium.
- Dhyne, E., Petrin, A., Smeets, V., Warzynski, F., et al., 2014. Import competition, productivity and multi-product firms. Working Paper 268. National Bank of Belgium.
- Diewert, W., Parkan, C., 1983. Linear programming tests of regularity conditions for production frontiers, in: Eichhorn, W., Henn, R., Neumann, K., Shephard, R. (Eds.), *Quantitative Studies on Production and Prices*. Physica-Verlag, Würzburg.
- Doraszelski, U., Jaumandreu, J., 2018. Measuring the bias of technological change. *Journal of Political Economy* 126, 1027–1084.
- Duprez, C., Magerman, G., 2018. Price updating in production networks. Working Paper 352. National Bank of Belgium.
- Epstein, L., Yatchew, A., 1985. Non-parametric hypothesis testing procedures and applications to demand analysis. *Journal of Econometrics* 30, 149–169.
- Forlani, E., Martin, R., Mion, G., Muûls, M., 2016. Unraveling firms: Demand, productivity and markups heterogeneity. Working Paper 293. National Bank of Belgium.
- Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review* 98, 394–425.
- Gandhi, A., Navarro, S., Rivers, D.A., 2020. On the identification of gross output production functions. *Journal of Political Economy* 128, 2973–3016.

- Goldberg, P., Khandelwal, A., Pavcnik, N., Topalova, P., 2010. Imported intermediate inputs and domestic product growth: evidence from india. *The Quarterly Journal of Economics* 125, 1727–1767.
- Grieco, P.L., Li, S., Zhang, H., 2016. Production function estimation with unobserved input price dispersion. *International Economic Review* 57, 665–690.
- Grifell-Tatjé, E., Lovell, C., Sickles, R. (Eds.), 2018. *The Oxford Handbook of Productivity Analysis*. Oxford University Press.
- Griliches, Z., Mairesse, J., 1998. Production functions: The search for identification, in: *Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium*. New York: Cambridge University Press, pp. 169–203.
- Halpern, L., Koren, M., Szeidl, A., 2015. Imported inputs and productivity. *American Economic Review* 105, 3660–3703.
- Hanoch, G., Rothschild, M., 1972. Testing assumptions of production theory: A nonparametric approach. *Journal of Political Economy* 80, 256–275.
- Hayfield, T., Racine, J.S., 2008. Nonparametric econometrics: The np package. *Journal of Statistical Software* 27, 1–32.
- Kehrig, M., Vincent, N., 2021. The micro-level anatomy of the labor share decline. *The Quarterly Journal of Economics* 136, 1031–1087.
- Klette, T.J., Griliches, Z., 1996. The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics* 11, 343–361.
- Levinsohn, J., Petrin, A., 2003. Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies* 70, 317–341.
- Li, Q., Racine, J.S., 2007. *Nonparametric econometrics: theory and practice*. Princeton University Press.
- Marschak, J., Andrews, W., 1944. Random simultaneous equations and the theory of production. *Econometrica* 12, 143–205.

- Olley, G.S., Pakes, A., 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64, 1263–1297.
- Simar, L., Zelenyuk, V., 2007. Statistical inference for aggregates of Farrell-type efficiencies. *Journal of Applied Econometrics* 22, 1367–1394.
- Syverson, C., 2011. What determines productivity? *Journal of Economic Literature* 49, 326–365.
- Syverson, C., 2017. Challenges to mismeasurement explanations for the US productivity slowdown. *Journal of Economic Perspectives* 31, 165–186.
- Varian, H., 1984. The nonparametric approach to production analysis. *Econometrica* 52, 579–598.
- Varian, H., 1985. Non-parametric analysis of optimizing behavior with measurement error. *Journal of Econometrics* 30, 445–458.
- Varian, H., 1988. Revealed preference with a subset of goods. *Journal of Economic Theory* 46, 179–185.
- Varian, H., 1990. Goodness-of-fit in optimizing models. *Journal of Econometrics* 46, 125–140.
- Wooldridge, J.M., 2009. On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters* 104, 112–114.

Appendix A: Theoretical results

Appendix A.1: Proof of Proposition 1

(i) or (ii) \Rightarrow (iii). Since F is quasi-concave, we have that $\mathcal{U}_j = \{(\mathbf{V}, \Omega) : F(\mathbf{V}, \Omega) \geq Q_j\}$ is convex. Consider a supporting hyperplane through (\mathbf{V}_j, Ω_j) at \mathcal{U}_j . Under OP.II, by cost minimization, we can assume that the slope of this hyperplane equals (\mathbf{W}_j, Γ_j) . Under OP.I, by cost minimization, we can assume that the \mathbf{V} component equals \mathbf{W}_j and we denote the Ω -component by $\Gamma_j \in \mathbb{R}_{++}$.

For every $(\mathbf{V}, \Omega) \in \mathcal{U}_j$, i.e. with $F(\mathbf{V}, \Omega) \geq Q_j$, we then have that (by definition of a supporting hyperplane):

$$\mathbf{W}'_j \mathbf{V}_j + \Gamma_j \Omega_j \leq \mathbf{W}'_j \mathbf{V} + \Gamma_j \Omega.$$

Next, by assumption, the input (\mathbf{V}_i, Ω_i) can produce Q_i (i.e., $F(\mathbf{V}_i, \Omega_i) \geq Q_i$). So rescaling the vector \mathbf{V}_i by Q_j/Q_i , constant returns to scale in \mathbf{V} implies that we have:

$$F\left(\frac{Q_j}{Q_i} \mathbf{V}_i, \Omega_i\right) = \frac{Q_j}{Q_i} F(\mathbf{V}_i, \Omega_i) \geq Q_j.$$

As such, $\left(\frac{Q_j}{Q_i} \mathbf{V}_i, \Omega_i\right) \in \mathcal{U}_j$ and:

$$\mathbf{W}'_j \mathbf{V}_j + \Gamma_j \Omega_j \leq \frac{Q_j}{Q_i} \mathbf{W}'_j \mathbf{V}_i + \Gamma_j \Omega_i.$$

If we divide both sides by Γ_j and define $\Lambda_j = \frac{1}{\Gamma_j} \in \mathbb{R}_{++}$, we obtain the desired linear inequalities.

(iii) \Rightarrow (ii) and (i). For all observations $i \in N$, let us define:

$$\begin{aligned} D_i &= \{\Omega \geq 0 : \mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega) > 0\}, \\ &= \left[0, \frac{\mathbf{W}'_i \mathbf{V}_i}{\Gamma_i} + \Omega_i\right). \end{aligned}$$

Let us also define $D_0 = \mathbb{R}_+ = [0, +\infty)$.

Consider a vector $\mathbf{A} \in \mathbb{R}_{++}^M$ large enough such that for all $i \in N$: $\Omega_i \mathbf{A}' \mathbf{V}_i > \max_{j \in N} Q_j$ and define:

$$F_0(\mathbf{V}, \Omega) = \Omega \mathbf{A}' \mathbf{V}.$$

Next, for all $i \in N$ we define $F_i : \mathbb{R}_+^M \times \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{+\infty\}$:

$$F_i(\mathbf{V}, \Omega) = \begin{cases} \frac{Q_i \mathbf{W}'_i \mathbf{V}}{\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega)}, & \text{if } \Omega_i \in D_i, \\ +\infty & \text{else.} \end{cases}$$

Finally, define the production function $F : \mathbb{R}_+^M \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$:

$$F(\mathbf{V}, \Omega) = \min_{i \in N \cup \{0\}} F_i(\mathbf{V}, \Omega).$$

Note that, as $D_0 = \mathbb{R}_+$ and $F(\mathbf{V}, \Omega) \leq \Omega \mathbf{A}' \mathbf{V} < \infty$, we have that F indeed maps to finite numbers. This also implies that below we can restrict our attention to the domains $\mathbb{R}^N \times D_i$ for $i \in N$.

It is clear that F is homogeneous of degree 1 in \mathbf{V} , as all functions F_i ($i \in N \cup \{0\}$) are homogeneous of degree one on the domain $\mathbb{R}^N \times D_i$. Also, F is strictly increasing in (\mathbf{V}, Ω) by the same reasoning.

Let us then show that F is continuous, quasi-concave and that the data is (OP.I)- and (OP.II)-rationalizable for this production function:

F is continuous. Let $(\mathbf{V}^n, \Omega^n) \rightarrow^n (\mathbf{V}, \Omega)$. To show continuity of F , it suffices to show that there is a subsequence $(\mathbf{V}^{n_j}, \Omega^{n_j})_{j \in \mathbb{N}}$ of $(\mathbf{V}^n, \Omega^n)_{n \in \mathbb{N}}$ such that $\lim_{j \rightarrow \infty} F(\mathbf{V}^{n_j}, \Omega^{n_j}) \rightarrow F(\mathbf{V}, \Omega)$.

For all $n \in \mathbb{N}$, by definition, $F(\mathbf{V}^n, \Omega^n) = F_i(\mathbf{V}, \Omega)$ for some $i \in N \cup \{0\}$. As $N \cup \{0\}$ is finite, we have that there is an element $i^* \in N \cup \{0\}$ and a subsequence $(\mathbf{V}^{n_j}, \Omega^{n_j})_{j \in \mathbb{N}}$ such that for all $(\mathbf{V}_j^n, \Omega_j^n)$ in this subsequence:

$$F(\mathbf{V}^{n_j}, \Omega^{n_j}) = F_{i^*}(\mathbf{V}^{n_j}, \Omega^{n_j}).$$

We consider two cases:

- If $\Omega \in D_{i^*}$, then as F_{i^*} is continuous on $\mathbb{R}_+^N \times D_{i^*}$, we have that:

$$\lim_{j \rightarrow \infty} F(\mathbf{V}^{n_j}, \Omega^{n_j}) = F_{i^*}(\mathbf{V}, \Omega),$$

so it suffices to show that $F(\mathbf{V}, \Omega) = F_{i^*}(\mathbf{V}, \Omega)$.

If not, there is a $k \neq i^*$ such that:

$$F(\mathbf{V}, \Omega) = F_k(\mathbf{V}, \Omega) < F_{i^*}(\mathbf{V}, \Omega).$$

Now, given that $\mathbb{R}_+^M \times D_k$ is open (relative to $\mathbb{R}_+^M \times \mathbb{R}_+$), F_k is continuous on $\mathbb{R}_+^M \times D_k$ and $(\mathbf{V}^{n_j}, \Omega^{n_j}) \rightarrow^j (\mathbf{V}, \Omega)$, there should be an n_j large enough such that:

$$F_k(\mathbf{V}^{n_j}, \Omega^{n_j}) < F_{i^*}(\mathbf{V}^{n_j}, \Omega^{n_j}) = F(\mathbf{V}^{n_j}, \Omega^{n_j}),$$

a contradiction.

- Next, consider the case where $\Omega \notin D_{i^*}$. Then as $\Omega^{n_j} \rightarrow \Omega$ and $\Omega^{n_j} \in D_{i^*}$ for all n_j , it follows that $\Omega = \frac{\mathbf{W}'_i V_i}{\Gamma_i} + \Omega_i$.

Then:

$$F_{i^*}(\mathbf{V}^{n_j}, \Omega^{n_j}) = \frac{Q_i \mathbf{W}'_i \mathbf{V}^{n_j}}{\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega^{n_j})} \rightarrow^j +\infty.$$

This also means that we can find an n_j large enough such that:

$$F(\mathbf{V}^{n_j}, \Omega^{n_j}) = F_{i^*}(\mathbf{V}^{n_j}, \Omega^{n_j}) > \Omega \mathbf{A}' \mathbf{V} + 1 > \Omega^{n_j} \mathbf{A}' \mathbf{V}^{n_j},$$

a contradiction with the definition of F .

F is quasi-concave. As F is the minimum of a finite number of functions F_i ($i \in N \cup \{0\}$) it suffices to show that every function F_i is quasi-concave.

As $F_0(\mathbf{V}, \Omega) = \Omega \mathbf{A}' \mathbf{V}$ is the product of two non-negative concave (linear) functions (i.e., Ω and $\mathbf{A}' \mathbf{V}$) it is quasi-concave.

To show that F_i ($i \in N$) is quasi-concave, let $\beta \in \mathbb{R}$ and assume that:

$$\begin{aligned} F_i(\mathbf{V}, \Omega) &\geq \beta, \\ F_i(\mathbf{V}', \Omega') &\geq \beta. \end{aligned}$$

For $\alpha \in [0, 1]$, define:

$$\begin{aligned} V_\alpha &= \alpha \mathbf{V} + (1 - \alpha) \mathbf{V}', \\ \Omega_\alpha &= \alpha \Omega + (1 - \alpha) \Omega'. \end{aligned}$$

We need to show that $F(V_\alpha, \Omega_\alpha) \geq \beta$.

We need to consider three cases.

- If $\Omega_\alpha \notin D_i$, then $F_i(V_\alpha, \Omega_\alpha) = +\infty \geq \beta$, so there is nothing to prove.
- If $\Omega_\alpha \in D_i$ and both $\Omega, \Omega' \in D_i$, then using the definition of F_i , we have:

$$\begin{aligned} Q_i \mathbf{W}'_i \mathbf{V} &\geq \beta(\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega)), \\ Q_i \mathbf{W}'_i \mathbf{V}' &\geq \beta(\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega')). \end{aligned}$$

Taking the convex combination of these two inequalities, gives:

$$Q_i \mathbf{W}'_i \mathbf{V}_\alpha \geq \beta(\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega_\alpha)),$$

which proves that, indeed, $F(\mathbf{V}_\alpha, \Omega_\alpha) \geq \beta$.

- As Ω_α is a convex combination of Ω and Ω' , the case where $\Omega_\alpha \in D_i$ and $\Omega, \Omega' \notin D_i$ can not occur. So, for the final case we can assume that $\Omega_\alpha \in D_i$ but one of Ω or Ω' is not an element of D_i . Without loss of generality, assume that $\Omega' \in D_i$ and $\Omega \notin D_i$ (i.e. $\Omega > \Omega'$).

Define $\tilde{\alpha}$ such that:

$$\Omega_{\tilde{\alpha}} = \tilde{\alpha}\Omega + (1 - \tilde{\alpha})\Omega' = \frac{\mathbf{W}'_i \mathbf{V}_i}{\Gamma_i} + \Omega_i.$$

Notice that, as $\Omega_\alpha \in D_i$, we have that $\tilde{\alpha} > \alpha$.

If we then define $(\mathbf{V}_\gamma, \Omega_\gamma) = \gamma(\mathbf{V}, \Omega) + (1 - \gamma)(\mathbf{V}', \Omega')$, we obtain

$$\lim_{\substack{\gamma \rightarrow \tilde{\alpha} \\ \gamma < \tilde{\alpha}}} F_i(\mathbf{V}_\gamma, \Omega_\gamma) = \lim_{\substack{\gamma \rightarrow \tilde{\alpha} \\ \gamma < \tilde{\alpha}}} \frac{\mathbf{W}'_i \mathbf{V}_\gamma}{\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega_\gamma)} \rightarrow +\infty.$$

So we can find a $\alpha < \gamma < \tilde{\alpha}$, close enough to $\tilde{\alpha}$, such that:

$$F_i(\mathbf{V}_\gamma, \Omega_\gamma) > \beta.$$

Notice that $(\mathbf{V}_\alpha, \Omega_\alpha) = \delta(\mathbf{V}_\gamma, \Omega_\gamma) + (1 - \delta)(\mathbf{V}', \Omega')$ with $0 \leq \delta = \frac{\alpha}{\gamma} \leq 1$. As $\Omega', \Omega_\gamma \in D_i$

it follows from the previous part that:

$$F_i(\mathbf{V}_\alpha, \Omega_\alpha) \geq \min\{F_i(\mathbf{V}', \Omega'), F_i(\mathbf{V}_\gamma, \Omega_\gamma)\} \geq \beta.$$

F (OP.I)- and (OP.II)-rationalizes the data. Let us first show that for all observations $i \in N$, $F(\mathbf{V}_i, \Omega_i) = Q_i$. Since $\Omega_i \in D_i$, we have

$$F(\mathbf{V}_i, \Omega_i) \leq F_i(\mathbf{V}_i, \Omega_i) = Q_i.$$

If the inequality would be strict, then there should be a $j \in N \cup \{0\}$ such that

$$F_j(\mathbf{V}_i, \Omega_i) < F_i(\mathbf{V}_i, \Omega_i).$$

If $j = 0$ then $\Omega_j \mathbf{A}' \mathbf{V}_j < F_i(\mathbf{V}_i, \Omega_i) \leq Q_i$, a contradiction given the definition of \mathbf{A} . If $j \in N$ then we have

$$\frac{Q_j \mathbf{W}'_j \mathbf{V}_i}{\mathbf{W}'_j \mathbf{V}_j + \Gamma_j(\Omega_j - \Omega_i)} < Q_i,$$

and thus

$$\frac{Q_j}{Q_i} \mathbf{W}'_j \mathbf{V}_i + \Gamma_j \Omega_i < \mathbf{W}'_j \mathbf{V}_j + \Gamma_j \Omega_j.$$

This contradicts with the linear inequalities in (iii).

For (OP.II)-rationalizability, towards a contradiction, assume that there exists an observation $i \in N$ and an input combination (\mathbf{V}, Ω) such that $F(\mathbf{V}, \Omega) \geq Q_i$ and $\mathbf{W}'_i \mathbf{V} + \Gamma_i \Omega < \mathbf{W}'_i \mathbf{V}_i + \Gamma_i \Omega_i$. Then, $\mathbf{W}'_i \mathbf{V} < \mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega)$, which means that the right hand side is strictly positive (i.e., $\Omega \in D_i$). As such we obtain

$$Q_i \leq F(\mathbf{V}, \Omega) \leq F_i(\mathbf{V}, \Omega) = \frac{Q_i \mathbf{W}'_i \mathbf{V}}{\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega)},$$

which in turn implies

$$\mathbf{W}'_i \mathbf{V}_i + \Gamma_i \Omega_i \leq \mathbf{W}'_i \mathbf{V} + \Gamma_i \Omega,$$

a contradiction.

Finally, for (OP.I)-rationalizability, towards a contradiction, assume that there is an input

\mathbf{V} such that $F(\mathbf{V}, \Omega_i) \geq Q_i$ and $\mathbf{W}'_i \mathbf{V} < \mathbf{W}'_i \mathbf{V}_i$. Then we obtain

$$Q_i \leq F(\mathbf{V}, \Omega_i) \leq F_i(\mathbf{V}, \Omega_i) = \frac{Q_i \mathbf{W}'_i \mathbf{V}}{\mathbf{W}'_i \mathbf{V}_i + \Gamma_i(\Omega_i - \Omega_i)},$$

which leads to

$$\mathbf{W}'_i \mathbf{V}_i \leq \mathbf{W}'_i \mathbf{V},$$

again a contradiction.

Appendix A.2: Testability

The following example illustrates the testable implications in Proposition 1. It shows that these implications can be rejected even in a minimalistic setting with only two firm observations and two observed inputs.

Consider a data set S capturing two firm observations that produce the same output quantity with input prices $\mathbf{W}_1 = (1, 2)$ and $\mathbf{W}_2 = (2, 1)$ and input quantities $\mathbf{V}_1 = (1, 2)$ and $\mathbf{V}_2 = (2, 1)$. The rationalizability conditions in Proposition 1 require:

$$\begin{aligned} 5\Lambda_1 + \Omega_1 &\leq 4\Lambda_1 \frac{Q_1}{Q_2} + \Omega_2 \\ 5\Lambda_2 + \Omega_2 &\leq 4\Lambda_2 \frac{Q_2}{Q_1} + \Omega_1. \end{aligned}$$

Reformulating these inequalities obtains:

$$\begin{aligned} (\Omega_2 - \Omega_1) &\leq \left(4 \frac{Q_2}{Q_1} - 5\right) \Lambda_2 \text{ and} \\ (\Omega_2 - \Omega_1) &\geq \left(5 - 4 \frac{Q_1}{Q_2}\right) \Lambda_1, \end{aligned}$$

which implies that $\left(5 - 4 \frac{Q_1}{Q_2}\right) \Lambda_1 \leq \left(4 \frac{Q_2}{Q_1} - 5\right) \Lambda_2$. If we then assume that the (observed) output levels Q_1 and Q_2 are such that $\frac{4}{5} < \frac{Q_1}{Q_2} < \frac{5}{4}$, we obtain that there can never exist strict positive Λ_1 and Λ_2 that satisfy this inequality restriction (since $4 \frac{Q_2}{Q_1} - 5 < 0$ and

$$5 - 4\frac{Q_1}{Q_2} > 0).$$

Appendix B: Monte Carlo simulation

We demonstrate the usefulness of our advocated nonparametric methodology for our general setting of cost minimization by means of a Monte Carlo simulation analysis. This will show the proper working of our methodology for noisy production data and settings with biased technological change. We will focus on the setting with only flexible inputs that are freely adjustable. This setting also provides insight into the working of our methodology under short term cost minimization in the presence of fixed inputs with adjustment costs, as the latter setting is a special case of the former setting, with additional constraints on the set of comparison firm observations. In addition, we will consider the possibility to impose additional structure to sharpen identification if a priori information is available on the nature of technological change. Particularly, in our following simulation, we will show and compare estimates without and with a Hicks neutrality assumption (as discussed in Section 2.4).

Simulation set-up. We generate data that resemble actual production data as used in empirical firm-level productivity studies. Following Doraszelski and Jaumandreu (2018), we start from a production function that is characterized by Constant Elasticity of Substitution (CES):

$$Q = \left(\Omega [0.5 (\Omega^\delta V_1)^\rho + 0.5 V_2^\rho]^{\frac{\gamma}{\rho}} \right) e^\varepsilon,$$

where $Q \in \mathbb{R}_+$ represents output and $\mathbf{V} \in \mathbb{R}_+^2 = [V_1, V_2]$ is a 2-dimensional vector of observed inputs. W_1 is considered to be known by both the firm and the empirical analyst while W_2 is considered to be known by the firm but unobserved by the empirical analyst. ε represents mean-zero normally distributed noise (e.g., unanticipated productivity shocks) with standard deviation σ_ε . Ω and V_2 are drawn randomly with mean 10 and covariance matrix $R = \begin{pmatrix} 4 & 1 \\ 1 & 4 \end{pmatrix}$. By introducing correlation between Ω and V_2 , an endogeneity issue arises, rendering Ordinary Least Squares (OLS) a biased estimator of output elasticities for this data generating process (DGP). In addition, we set $\gamma = 1$ (i.e., CRS) and $\rho = 0.75$. W_2 is uniformly drawn from the interval $[0.9, 1.1]$ and Γ is uniformly drawn for sets of 10

observations from the interval [3.9, 4.1]. For the setting with biased technical change, we set $\delta = 1$. To impose Hicks neutrality, we set $\delta = 0$. We consider different levels of noise by setting σ_ε equal to 0 and 0.05: $\sigma_\varepsilon = 0.05$ reflects a moderately noisy data set, while $\sigma_\varepsilon = 0$ corresponds to a deterministic setting. Using the generated values for $(\Omega, \Gamma, \gamma, \rho, V_2, \delta, W_2)$, we can construct observed input price W_1 and input usage V_1 that correspond with cost minimizing behavior. Lastly, we obtain Q from introducing the generated values into the production function. Based on this DGP, we generated B samples of size $|N|$ for four cases, reflecting two levels of noise and whether or not Hicks neutrality holds.

Simulation results. For each case, we set $B = 1000$. We focus on the small sample properties of our methodology by setting $|N| = 250$, which corresponds to the lowest sample size considered in our own empirical application. Next, $|G| = 10$ approximates the number of observations of one and the same firm in our empirical setting. We presume that W_2 is unobserved for the empirical analyst, but that the empirical analyst knows that equality of Γ can be imposed for observations of the same firm.

Table 9 reports on whether or not the simulated data satisfy our nonparametric rationalizability conditions (i.e., there exists a feasible solution or not). The first (last) four columns give these results when the true DGP is (not) Hicks neutral (i.e., $\delta = 0$ ($\delta = 1$)); and columns 1, 2, 5 and 6 pertain to deterministic settings (i.e., no noise; $\sigma_\varepsilon = 0$) while columns 3, 4, 7 and 8 pertain to settings with moderate stochastic noise ($\sigma_\varepsilon = 0.05$). Columns 1, 3, 5 and 7 report on feasibility of our rationalizability conditions when additionally imposing the restrictions associated with the assumption of Hicks neutrality; the other columns provide feasibility results when not using these Hicks neutrality restrictions (thus allowing for biased technological change). Further, we have computed identification results by both minimizing the objective function (3) subject to our linear rationalizability conditions. The results of this identification exercise are summarized in Table 10.

Table 9 reveals that, when the true DGP satisfies Hicks neutrality (i.e., $\delta = 0$) and there is no noise (i.e., $\sigma_\varepsilon = 0$), the simulated data pass our rationalizability conditions with and without the additional Hicks neutrality restrictions (i.e., the average goodness of fit parameter $\theta = 1$). Evidently, this is unsurprising as this DGP is perfectly consistent with the structural models at hand, and the models are nested. By contrast, in the deterministic setting with biased technological change (i.e., $\delta = 1$), the data only pass the strict ratio-

nalizability conditions that do not impose Hicks neutrality. In this case, falsely imposing Hicks neutrality implies infeasibility for above 60 percent of our simulated data sets. This clearly demonstrates that wrongly imposing Hicks neutrality can hugely impact the empirical findings.

Next, adding moderate noise to the data (i.e., $\sigma_\epsilon = 0.05$) implies that some firm observations are no longer perfectly compatible with our structural model specifications. As before, the models with and without Hicks neutrality imposed perform equally well when Hicks neutrality effectively holds (i.e., $\delta = 0$), while we observe divergent performance patterns under biased technological change (i.e., $\delta = 1$). Moreover, in both cases, we find that our model without the Hicks neutrality conditions is showing close-to rational behavior, which is exactly in line with the GDP specification. Particularly, our nonparametric rationalizability conditions are met for an average goodness-of-fit parameter θ that amounts to 1.050 (for $\delta = 0$) and 1.034 (for $\delta = 1$).

Table 9: Feasibility, $|N|=250, B=1000$

$ N = 250, B=1000, G=10$	$\sigma_\epsilon = 0, \delta = 0$		$\sigma_\epsilon = 0.05, \delta = 0$		$\sigma_\epsilon = 0, \delta = 1$		$\sigma_\epsilon = 0.05, \delta = 1$	
Hicks neutrality imposed?	Yes	No	Yes	No	Yes	No	Yes	No
Perc. infeasible	0.000	0.000	0.000	0.000	0.642	0.000	0.683	0.000
Av. θ when feasible	1.000	1.000	1.050	1.050	1.198	1.000	1.201	1.034

Notes: We denote a replication as *infeasible* if no feasible solution can be reached with $\theta \leq 1.3$.

Table 10 shows the empirical performance of our nonparametric identification methods. Specifically, we compare the (simulated) true and estimated values of the productivity factor Ω , the associated cost $\Gamma\Omega$ and the cost share of V_1 (equalling its output elasticity). The subscript *HN* denotes that Hicks neutrality is imposed in the linear rationalizability restrictions. We focus on both the Pearson rank correlation and the mean absolute percentage error (MAPE) between the true and estimated values. As absolute levels of Ω and $\Gamma\Omega$ cannot be identified, we normalize the true and estimated values for these variables by taking their ratios over the minimum levels in our simulated samples.

We learn that, when the data generation process effectively satisfies Hicks neutrality and exhibits no noise (i.e., $\delta = 0$ and $\sigma_\epsilon = 0$), the correlations are very high (i.e., exceeding 0.95) and the average errors are low (i.e., below 5 percent) for all three variables under study. Correctly imposing Hicks neutrality overall sharpens identification in terms of ranks as well as levels. This is particularly the case for Ω , for which the correct assumption of

Hicks neutrality lowers the MAPE from 4.6 percent to 1 percent. Under moderate noise (i.e., $\sigma_\epsilon = 0.05$), there is only a minor drop of the Pearson correlations, and the MAPE comfortably stays below 3 percent for CS^{V_1} . However, the MAPE for Ω and $\Gamma\Omega$ now exceeds 20 percent, indicating that the absolute levels of these variables are not well identified under moderate noise.

Further, as we discussed above, when the DGP exhibits biased technological change (i.e., $\delta = 1$), the rationalizability restrictions with Hicks neutrality (wrongly) imposed are no longer feasible for reasonably small θ -values. In turn, this prevents an (informative) identification analysis under the assumption of Hicks neutrality when this assumption is truly violated. Interestingly, this negative conclusion does not carry over to the rationalizability conditions that do not impose Hicks neutrality. When using this more flexible specification, we again obtain informatively precise identification of both the levels and ranks of CS^{V_1} and of the ranks of Ω and $\Gamma\Omega$, even for moderately noisy data.

Table 10: Identification, $|N|=250, B=1000$

	Ω_{HN}	Ω	$\Gamma_{HN}\Omega_{HN}$	$\Gamma\Omega$	$CS_{HN}^{V_1}$	CS^{V_1}
$\sigma_\epsilon = 0, \delta = 0$						
ρ	0.999	0.997	0.999	0.995	0.986	0.954
MAPE	0.010	0.046	0.013	0.049	0.030	0.028
$\sigma_\epsilon = 0.05, \delta = 0$						
ρ	0.957	0.958	0.946	0.928	0.960	0.945
MAPE	0.232	0.259	0.226	0.259	0.028	0.026
$\sigma_\epsilon = 0, \delta = 1$						
ρ	–	0.998	–	0.996	–	0.964
MAPE	–	0.042	–	0.044	–	0.036
$\sigma_\epsilon = 0.05, \delta = 1$						
ρ	–	0.984	–	0.961	–	0.951
MAPE	–	0.182	–	0.191	–	0.044

Notes: We denote a replication as *infeasible* if no feasible solution can be reached with $\theta \leq 1.3$. No output is shown if more than 50 percent of replications are denoted as infeasible. ρ stands for the average Pearson correlation over the replications with feasible results. *MAPE* stands for the mean absolute percentage error over the replications with feasible results.