INSTITUTE OF ECONOMICS

Scuola Superiore Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM
WORKING PAPER SERIES

Exporting and productivity as part of the growth process: Causal evidence from a data-driven structural VAR

Tommaso Ciarli [a]
Alex Coad [b]
Alessio Moneta [c]

[a] SPRU, University of Sussex, Brighton, UK.
[b] CENTRUM Católica Graduate Business School (CCGBS) and
Pontificia Universidad Católica del Perú (PUCP), Lima, Perú.
[c] Institute of Economics and Department EMbeDS, Scuola Superiore Sant'Anna, Pisa, Italy.

2019/39 **December 2019**
ISSN(ONLINE) 2284-0400

Does Exporting Cause Productivity Growth? Evidence from a Structural VAR Analysis*

Tommaso Ciarli[†] Alex Coad[‡] Alessio Moneta[§]

June 2022

Abstract

Does exporting more increase firm productivity? Or is it only firms that manage to increase productivity that increase their sales in the export market? This paper provides new empirical evidence on the causal relation between trade and productivity adopting a structural vector autoregressive analysis combined with identification algorithms in the machine learning literature. We focus on a well-studied country (Chile) and on already-exporting firms (intensive margin). We identify the contemporaneous and lagged causal structure between firm productivity and export growth using two different machine learning algorithms based on Independent Components Analysis (ICA), which exploit the non-Gaussian distribution of the data to recover the independent structural shocks that drive the observed variables. Our findings show that, for Chilean firms, productivity growth causes export growth in the same year, but not the other way around. Export growth also has no causal effect on TFP growth in subsequent years. To increase sales in the foreign market, firms must first increase productivity. The increased presence in the foreign market does not contribute to such productivity growth.

JEL Classification: L21; D24; F14; C22
Keywords: Productivity; Exporting; Learning-by-exporting; Causality; Structural VAR; Independent Component Analysis

---

*We are grateful to Sebastian Vergara for sharing with us his data and variable construction pre-2003, and for helpful initial discussions on the construction of the 2001-2007 variables. We thank Matté Hartog and Cristina Chaminade for insightful discussions that helped improving the paper. We thank conference/seminar participants at Carnegie Mellon University, University of Jena, Queen Mary University of London, Free University of Berlin, University of Kassel, University of Rovira i Virgili, EMAEE 13 (Nice), and Globelics 15 (Havana) for useful comments. Thanks also to Doris Entner for sharing the R code of the LiNG algorithm with us. The usual caveat applies.

[†]UNU-MERIT, Maastricht & SPRU, University of Sussex, ciarli@merit.unu.edu

[‡]Waseda University, Shinjuku, Tokyo, Japan, alex.coad@waseda.jp

[§]*Corresponding Author*. Institute of Economics and EMbeDS, Sant'Anna School of Advanced Studies, Piazza Martiri della Libertà 33, 56127 Pisa, Italy. Email: a.moneta@santannapisa.it

1 Introduction

This paper investigates the still-unresolved causal relationship between exporting and productivity. Do firms learn from trade? Or is exporting a game that only the best firms can play, and once they export there is little scope for productivity growth? Since the 1990s economists have used micro data to study how firms take advantage from exporting, in both low income and high income countries (see e.g. Cirera et al., 2012). The seminal work of Bernard and Jensen (1995) started a prolific field of enquiry, using firm and plant survey data from a large number of different countries. The results from this literature relatively unambiguously indicate that exporting firms, on average, do better than non-exporting firms on different performance measures.[1] However, Bernard and Jensen (1999) subsequently highlighted that, in order to enter the global market, firms need to be more productive than average. On the one hand, firms need to increase productivity before entering the export market – because of trade costs, stronger competition, and investments required to increase the scale. On the other hand, firms may increase productivity while exporting – because of learning from foreign buyers, use of excess capacity, or stronger competition.[2] Which comes first, the chicken (growth of exports) or the egg (productivity growth)?

We focus on Chile, a small country, with a small domestic market, an open economy, a well-tested firm survey, and a large literature on the relation between exporting and productivity (e.g. Alvarez and Crespi, 2007; Alvarez and López, 2005, 2008; López, 2009; Pavcnik, 2002).

Two contributions help to distinguish our work from previous attempts to disentangle the causal relations between firm productivity and exporting. First, we apply a class of methods for causal inference – linear non-Gaussian vector autoregression models – which bring together recent results from both the time-series econometrics and machine learning literatures. Our methodology exploits the non-Gaussian feature of the data to recover the set of independent structural shocks that drive the movements of the observed variables. The identification of the shocks, and the manner in which these are linearly mixed to form the observed variables, allow us to arrange the variables in a causal ordering. The causal structure is identified through two different ICA-based[3] algorithms, which are

[1] See for example surveys by Greenaway and Kneller (2007); International Study Group on Exports and Productivity (2008); Wagner (2007).

[2] Increases in exports may have implications for firm productivity through a number of mechanisms discussed in the literature: larger number of clients and/or markets from which the firm could learn; increased technology/knowledge transfer from buyers; greater incentives to increase productivity and compete in markets with higher quality; learning about new market opportunities; less vulnerability and dependence on a single market; having a larger scale of output; and improved utilisation of existing production capacity.

[3] Independent Component Analysis.

based on different assumptions. The first is the LiNGAM (Linear Non-Gaussian Acyclic Model) algorithm, which assumes acyclicity – i.e. that there are no causal feedback loops that take place within the period of observation (e.g. within the year), although lagged causal effects can go in both directions between variables. The LiNGAM estimator has already been featured in previous econometric work (e.g. Moneta et al., 2013; Brenner et al., 2018; Brancaccio et al., 2020). The second one is the LiNG algorithm (Lacerda et al., 2012), which – to our knowledge – has not yet been applied to economics research questions. The LiNG estimator relaxes LiNGAM's acyclicity condition, and allows for simultaneous causal feedback loops between variables even within the period of observation.

The second contribution is that we focus on the small number of firms (Bernard et al., 2007) that have already accessed the export market (intensive margin), and we analyse changes in export and productivity, rather than levels. Focusing on export growth rather than on entry, we can ignore the problem of self-selection into exporting,[4] and focus on the identification of whether (i) firms that export more improve their productivity or (ii) the other way round: firms must improve their productivity in order to be able to export more.

Our main results show that, within a time period, export growth follows growth in productivity, and not the other way round. There is no evidence of export growth causing productivity growth. Export growth seems to be driven by an increase in employment, which has a contemporaneous negative effect on productivity. This means that Chilean firms in the foreign market need to first improve their productivity (while also growing in size), in order to increase exports.

Over time, there is a weak positive impact of productivity growth on export growth with one year lag, but the result is not robust, and is not significant when using the LiNG estimator that relaxes the acyclicity condition and allows for simultaneous causal feedback loops between variables.

Our result about the direction of the contemporaneous causal relationship between productivity growth and export growth is robust to a number of checks in which we vary indicators of productivity, sub-samples (looking at different sectors, following various criteria), and methods of estimation of the VAR model.

Our results are relevant for industrial and trade policy. They seem to suggest that exporting firms in a small economy, which has been open for a few decades, do not benefit a lot from increasing sales in the foreign market, at least within two years. This might be because most of these firms operate in non-technology intensive industries (e.g. Wang et al., 2021), where most competition is based on price (Garcia-Marin and Voigtländer, 2019). Small changes in relative labour costs can affect productivity (and firm size) and exporting. Whether these changes

[4] As summarised by Park et al. (2010, p. 822): "conceptually, the fundamental problem is that nonexporters are an inappropriate counterfactual for exporters."

bring about a longer term effect on technological learning and upgrading is left for future research.

Whether results depend or not on industry specialisation, they suggest that policies that support firm productivity, are needed, because firms cannot rely solely on learning from the foreign market to increase productivity and therefore increase exports.

Our paper contributes to the literature attempting to identify the effect of firm changes in exporting on changes in productivity. This is a large literature. For example, Park et al. (2010) find that firm specific export shocks have a positive effect on productivity growth in China. To identify the effect of exports, they instrument exporting growth with shocks in the currency exchange rate of the destination country. Berman and Rebeyrol (2010), using data on French firms, find that entry and persistence in the export market has no significant effect on productivity, although they find positive effects of export *growth* on subsequent productivity growth. More generally, we also contribute to the literature that has focused on export intensity, rather than entry in the export markets (e.g. Castellani, 2002; Fernandes and Isgut, 2005; Antolín et al., 2013; De Loecker, 2013; Dalgiç et al., 2021).

In the rest of the paper we first discuss the heterogeneous evidence on exporting and productivity (Section 2). Next, Section 3 discusses the methodological contribution. Section 4 presents the dataset and discusses the measurement of productivity in this paper. We then present and discuss the results in Section 5. In Section 6 we discuss implications for policy and future research.

2 The mixed and heterogeneous relation between exporting and firm productivity

A large applied literature in trade has attempted to identify the direction of causality between exports and growth at the firm level. Wagner (2007) conducts a systematic literature review and shows that: (i) exporting firms are always more productive than non-exporters; (ii) exporters are often more productive even before entering the export market; (iii) results on learning-by-exporting (LBE) are mixed, and when matching estimators are used no significant effect of exporting emerges; and (iv) firms that exit the export market tend to reduce productivity. In a parallel review of empirical literature, Greenaway and Kneller (2007) also report that results on LBE are not conclusive. Similar inconclusive results on LBE are found when analysing the service sector (Wagner, 2012).

Mixed results are confirmed also by studies that make an attempt to better control for causal relations using matching estimators. On the one hand, Girma et al. (2004) introduce matching estimators to this strand of literature and find a significant positive effect of exporting on productivity for UK manufacturing firms. Tsou et al. (2008) analyze a census of Taiwanese firms repeated for three different periods, and find that firms staying in the export market experience a larger increase in productivity than non-exporters. Similarly, Baldwin and Yan (2012) find that, following changes in the real exchange rate, firms that are already in the export market experience a relatively larger gain in productivity than new entrants. Manjón et al. (2013) find evidence of LBE for Spanish manufacturing firms.

On the other hand, several studies do not find a significant positive effect of exporting on productivity. No sustained effect of learning from exporting is observed by Mukim (2011) on Indian firms, by Arnold and Hussinger (2005) on German firms, by Damijan and Kostevc (2006) on Slovenian firms, and by Eliasson et al. (2012) for small and medium firms. The International Study Group on Exports and Productivity (2008) use panel data from 14 different countries, and find no evidence of LBE.

Only a handful of studies investigate the exporting-productivity nexus using growth rates rather than levels. Park et al. (2010) use exogenous shocks on the demand for exporting firms (exchange rates shocks), and find evidence of LBE for Chinese firms, especially when the destination is a high income country. Berman and Rebeyrol (2010), using data on French firms, find that entry and persistence in the export market has no significant effect on productivity. In contrast, they find positive effects of export growth on subsequent productivity growth. They explain this effect by referring to how exporting growth strengthens the incentives for firms to innovate as well as by enhancing access to finance for investment. Fernandes and Isgut (2005) focus on the level of exports ("export experience"), rather than on export participation, finding a positive effect of LBE for Colombian firms exporting to high income markets.

To explain such heterogeneity in findings, the literature has moved on to study the heterogeneous conditions under which a firm's exporting may positively affect its productivity.

One explanation focuses on export destination: firms may learn more from exporting to more advanced countries. Wagner (2012) reports that the relation between exporting and productivity is influenced by export destination for a number of papers included in his review. Similar results were documented by Fernandes and Isgut (2005) for Colombian exporting firms and by Dalgiç et al. (2021) for Turkish firms. More relevant to our paper, Martins and Yang's (2009) meta-analysis on LBE finds that, relative to industrialised countries, firms in developing countries enjoy a stronger impact of exporting on productivity. Focusing on Chile, a middle income country in the period of our analysis, we partially complement existing research with valuable new evidence (we do not have information on export destination).

A second explanation focuses on firms' ability to learn, by building complementary intangible assets. This literature suggests that specific investments may be a necessary condition to benefit from exporting. Among others, decisions to export may complement decisions to innovate (Ito and Lechevalier, 2010; Aw et al., 2011). For instance, Dai and Yu (2013) find that firms with higher R&D investment enjoy significant improvements in productivity after exporting, contrary to firms that do not invest before in R&D. Aw and Song (2013) finds that South Korean firms that only export, or only invest in R&D, do not experience the same increase in productivity as firms that do both. Serrano and Myro (2019) documents a different mechanism for Spanish firms, more relevant to our paper: they tend to increase in firm size (in terms of sales and employees), and increase R&D expenditure after they enter the export markets. An implication of these papers for our results is that investments (e.g. in intangible complementary assets) may be simultaneously driving productivity and exporting.

A third explanation, particularly relevant to our study, focuses on firm size.[5] (Damijan et al., 2010) find that larger firms benefit more from exporting. Firms with a higher export intensity (with respect to domestic sales) also tend to profit more from exporting (Girma et al., 2004).

A fourth explanation focuses on measure and methods. For instance, De Loecker (2013) proposes a method for computing productivity that includes exporting in the firm decision (and therefore as a determinant of productivity) and finds significant LBE for Slovenian firms. Garcia-Marin and Voigtländer (2019) observe a lack of growth of revenue productivity in the years after plants start to export, but show that this can be decomposed into a decrease in marginal costs which occurs alongside a commensurate decrease in prices. Hence, efficiency gains appear to be transmitted to consumers via lower prices, rather than leading to higher markups. This reconciles the two suggestions that plants enjoy efficiency gains after starting to export, but that these efficiency gains do not translate into productivity growth when this latter is measured in terms of revenue productivity.

The first three explanations consider time to be an important factor, as learning may take time, which may differ across sectors (Wang et al., 2021) and firms. Using measures of learning, Crespi et al. (2008) find evidence of LBE with a lag. However, Segarra-Blasco et al. (2020) find that only in more innovative countries firms benefit from a longer learning time, whereas firms in less innovative countries learn more from their local peers.

To summarize, therefore, we highlight three points. First, firms tend to invest in order to improve productivity before they increase their sales in a more com-

[5] The literature has also investigated the role of firm age (Alvarez and López, 2005; Girma et al., 2004; Fernandes and Isgut, 2005), although this dimension is not explored in our paper.

petitive foreign market. In this case, we should observe an increase in investment and productivity that precedes growth in exports (or export intensity). Second, while some mechanisms – such as market size, firm size, vulnerability, and use of existing capacity – may have an immediate effect on firm productivity, other mechanisms – such as learning from buyers, and from markets – may take one or more years to have a visible effect on productivity growth. Third, considering that most existing studies use yearly data, we should acknowledge that within the same year a firm takes many decisions. Firms may implement measures to increase productivity and export simultaneously, as complementary activities.

Combining contemporaneous and lagged firm decisions in relation to employment, domestic sales, and exporting, our paper makes an important contribution in determining the causal relation between growing exports and productivity in a middle income country.

3 Econometric method

3.1 VAR and SVAR models

Consider the general structural vector autoregressive (SVAR) model

$$y_{it} = By_{it} + \sum_{\ell=1}^{p} \Gamma_\ell y_{i,t-\ell} + \varepsilon_{it}, \tag{1}$$

in which y_{it} is a vector of k variables, observed for firm i ($i = 1, \ldots, N$) in year t ($t = 1, \ldots, T$); B and Γ_ℓ ($\ell = 1$,) are $k \times k$ matrices of structural coefficients; and ε_{it} is a vector of white-noise structural disturbances (shocks). The model in equation (1) generalises the standard SVAR model because it can be applied to panel data. However, differently from other formulations of panel VAR models (see, e.g., Breitung, 2005), in which both T and N are large, the matrices of coefficients do not contain the (firm) index i. In fact, since we deal with T relatively small with respect to N, we assume that structural-causal relations are stable both cross-sectionally (across firms) and over time.

A reduced-form VAR model is obtained by rearranging (1):

$$y_{it} = \sum_{\ell=1}^{p} (I - B)^{-1} \Gamma_\ell y_{i,t-\ell} + (I - B)^{-1} \varepsilon_{it}, \tag{2}$$

which can be rewritten as

$$y_{it} = \sum_{\ell=1}^{p} A_\ell y_{i,t-\ell} + u_{it}, \tag{3}$$

where $A_\ell = (I - B)^{-1}\Gamma_\ell$ (for $\ell = 1, \dots, p$) and $u_{it} = (I - B)^{-1}\varepsilon_{it}$. Differently from equation (1), equation (3) can be estimated without involving potential problems of endogeneity. Notice, however, that the reduced-form VAR model in (3) does not allow estimating the matrix of instantaneous causal effects B, and — importantly — nor does it allow to properly estimate the matrices of lagged causal effects $\Gamma_1, \dots, \Gamma_p$. But knowledge of the matrix B is sufficient to recover Γ_ℓ, having estimated A_ℓ (for $\ell = 1, \dots, p$) from equation (3) .

Let us call $\Gamma_0 = (I - B)$ and we assume it to be invertible. The matrix Γ_0 relates the k-dimensional vector of shocks ε_{it} to the k-dimensional vector of reduced-form residuals u_{it}. This why Γ_0 is also known as the *unmixing* matrix and Γ_0^{-1} as the *mixing* matrix. We have

$$u_{it} = \Gamma_0^{-1}\varepsilon_{it}. \tag{4}$$

Our approach to identify the structural model (1) is based on independent component analysis (ICA), which is a probabilistic method for finding linear combinations of the data that are maximally independent (Hyvärinen, 2013).

In our empirical analysis, y_{it} will comprise four variables, and we will estimate both one-lag and two-lag models. In keeping with previous applications of VARs and SVARs on panel data of firm growth rates (e.g., Coad et al., 2011; Moneta et al., 2013), we pool together observations under the assumption that there are no significant firm-specific time-invariant component ('fixed effects') in the growth rates of these series. This assumption that firms undergo similar structural patterns in their growth process seems plausible, because we are focusing on growth rates rather than levels, and any firm-specific components affecting levels will already have been differenced out. As a consequence, we can simplify notation by omitting henceforth the subscript i.

3.2 Identification strategy

As in Moneta et al. (2013), our identification strategy is based on a method in which we first estimate the reduced-form VAR model

$$y_t = A_1 y_{t-1} + \ldots + A_p y_{t-p} + u_t \tag{5}$$

and then we search for the 'unmixing' matrix Γ_0 such that

$$\Gamma_0 y_t = \Gamma_0 A_1 y_{t-1} + \ldots + \Gamma_0 A_p y_{t-p} + \Gamma_0 u_t \tag{6}$$

In more compact form, we have:

$$\Gamma_0 y_t = \Gamma_1 y_{t-1} + \ldots + \Gamma_p y_{t-p} + \varepsilon_t, \tag{7}$$

where $\Gamma_\ell = \Gamma_0 A_\ell$ for $\ell = 1, \ldots, p$. Assuming that the k elements of ε_t are mutually independent and (at least $k-1$ of them) non-Gaussian, the method is able to identify Γ_0 and, consequently, (having estimated all coefficient matrices of equation 5) all the coefficient matrices of equation (7). The underlying idea is to search for a mixture of the elements of u_t such that the resulting components are minimally dependent and maximally non-Gaussian (cfr. Hyvärinen and Oja, 2000; Hyvärinen et al., 2001). Since there are different measures of statistical dependence and non-Gaussianity, and different optimization methods, there are correspondingly different ICA algorithms. In our application we use FastICA, which is a fixed-point algorithm for maximum likelihood estimation and measures non-Gaussianity with an approximation of negentropy (Hyvärinen and Oja, 2000).

Regardless which algorithm is used, ICA leaves undetermined the scale, sign, and order of the latent sources or structural shocks. In other words, Γ_0^{-1} is identified up to the post multiplication by CD, where C is a permutation matrix[6] and D is a diagonal matrix with non-zero diagonal elements (Eriksson and Koivunen, 2004; Lanne et al., 2017; Gouriéroux et al., 2017). Further steps are needed to fully identify Γ_0 and ε_t. We adopt here two different ICA-based search methods to identify the shocks and more generally the structural VAR model. The first was proposed by Shimizu et al. (2006) and named LiNGAM (for linear, non-Gaussian, acyclic model), and when applied to VAR models, it is known as VAR-LiNGAM (Hyvärinen et al., 2008; Moneta et al., 2013; Coad et al., 2017). The second was proposed by Lacerda et al. (2012) and was named LiNG (for linear, non-Gaussian model). To our knowledge, this is the first time that the LiNG algorithm has been applied either in a VAR context (i.e. "VAR-LiNG") or in the discipline of economics. The algorithms VAR-LiNGAM and VAR-LiNG are described in the frames below.

Both algorithms, after having estimated the reduced-form VAR (step 1), run an ICA algorithm (e.g., FastICA) on the estimated residuals obtaining a mixing matrix P ($\equiv \Gamma_{ICA}^{-1}$) which is able to generate a vector of independent components (step 2). But the order and scaling of these independent components is arbitrary.

Algorithm 1 (VAR-LiNGAM) solves the order indeterminacy by assuming that the underlying causal structure among the contemporaneous variables contains no cycle (in other words can be represented by a directed acyclic graph). This assumption, jointly with the fact that the diagonal elements of Γ_0 must be nonzero (and should be normalised to one), ensures that if we find an ordering of the components $\hat{\varepsilon}_{1t}, \ldots, \hat{\varepsilon}_{kt}$ (output of the ICA algorithm) that produces a correspondence with the data $\hat{\varepsilon}_t = \tilde{\Gamma}_0 \hat{u}_t$ such that $\tilde{\Gamma}_0$ has non-zero elements in its

[6]A permutation matrix is a square matrix in which exactly one entry in each row and column is equal to 1 and all other entries are 0 (see, e.g., Horn and Johnson, 2012).

main diagonal, this ordering must be the correct one.[7] Exploiting this fact, step 3 is devoted to find the permutation of the matrix Γ_{ICA} generating the independent components from $\hat{u}_t$ which produces a correct matching between structural and reduced-form shocks. Step 4 solves the scale indeterminacy. This is simply done by normalising the rows of $\tilde{\Gamma}_0$ (the correctly row-permuted version of Γ_{ICA}), so that all diagonal elements equal unity. Let $\hat{\tilde{\Gamma}}_0$ denote this row normalised matrix and $\hat{B} = I - \hat{\Gamma}_0$ (step 5). Since it is assumed that there are no causal loops or feedback, there is a permutation (applied equally to columns and rows) of $\hat{\Gamma}_0$ which should be lower triangular. The same can be said for $\hat{\Gamma}_0^{-1}$ and $\hat{B}$. In practice, however, even under the correct assumptions, these matrices are not exactly lower triangular, because the ICA algorithm applied to finite data sets yields estimates with errors. Therefore step 6 searches for an approximate lower triangularity. This step is not essential for the sake of estimation of the structural model and is run only for identifying the contemporaneous causal order. Step 7 estimates the matrices of the lagged coefficients of the structural model.

Algorithm 2 (VAR-LiNG) solves the order indeterminacy by simply exploiting the assumption that Γ_0 has a zeroless diagonal, which is valid in the structural VAR model by construction. Step 3 tests which entries of the Γ_{ICA} are significantly different from zero. This is done through the following bootstrap procedure: resampling $\hat{u}_t$ and applying ICA to each $s = 1, \ldots, S$ bootstrap samples, one obtains S unmixing matrices $\Gamma_{ICA}^{(s)}$. Each $\Gamma_{ICA}^{(s)}$ is row-permuted (with a possible change of sign in each row) in a $\widetilde{\Gamma}_{ICA}^{(s)}$ to match Γ_{ICA} (the row-signed permutation is chosen so that the Frobenius norm of $(\widetilde{\Gamma}_{ICA}^{(s)} - \Gamma_{ICA})$ is minimised). Hence, one obtains an empirical distribution for each entry of Γ_{ICA}, and, therefore, a nonparametric quantile test to decide whether 0 is an outlier for each entry. Step 4 finds the permutation of the matrix Γ_{ICA} which produces a matrix $\tilde{\Gamma}_{0,j}$ which has a zeroless diagonal. There might be several such matrices: we index each of them with $h = 1, \ldots, m$. Thus, the algorithm will output h possible causal structures. However, some of them can be excluded a priori by excluding unstable contemporaneous causal structures, i.e. $\tilde{\Gamma}_{0,h}$ such that $\tilde{\Gamma}_{0,h}^{-1}$ has eigenvalues whose modulus is greater than one, for some h between 1 and m. Step 5 and 6 solve the indeterminacy of scaling in the same way as algorithm 1. Step 7 is also analogous to step 7 in algorithm 1.

[7] In other words, under acyclicity Γ_0 and Γ_0^{-1} are *essentially* triangular (i.e. $Z\Gamma_0 Z'$ is triangular for some permutation matrix Z). ICA identifies $\Gamma_0^{-1}DC$, where D is a diagonal matrix and C is an arbitrary permutation matrix. Since Γ_0^{-1} is essentially triangular, any permutation C (different from I) will yield a matrix $\Gamma_0^{-1}DC$ with some zeros on the main diagonal. To find out the correct permutation matrix C is sufficient to search for a permutation C' such that $\Gamma_0^{-1}DCC'$ has no zeros on the main diagonal. Notice that row-permuting Γ_0 through C is equivalent to column-permuting (in the same way) Γ_0^{-1} or row-permuting (in the inverse way) the rows of ε_t, since from $C\Gamma_0 u_t = \varepsilon_t$ it follows that $u_t = \Gamma_0^{-1}C'\varepsilon_t$.

To recapitulate, both algorithms are able to identify the structural model (or a class of possible structural models) from the estimated reduced form model. The assumptions which permit such an inference are, for both algorithms, non-Gaussianity and independence of the structural shocks. As regards the first algorithm, a further assumption is acyclicity, i.e. the assumption that there are no feedbacks or loops. The second algorithm relaxes this assumption, but the class of admissible models is now broader, which leads us to assume stability to restrict the number of causal structures. It should also be noted that an implicit assumption of both algorithms is causal sufficiency, i.e. the assumption that all the causally relevant variables have been modelled.

Algorithm 1: VAR-LiNGAM

1. Estimate the reduced form VAR model of equation (5), obtaining estimates $\hat{A}_\ell$ of the matrices A_ℓ for $\ell = 1, \dots, p$. Denote by $\hat{U}$ the $k \times T$ matrix of the corresponding estimated VAR residuals (T is the number of observations), that is each column of U is $\hat{u}_t \equiv (\hat{u}_{1t}, \dots, \hat{u}_{kt})'$, $(t = 1, \dots, T)$. Check whether u_{jt} (for each row $j = 1, \dots, k$ of U) is indeed non-Gaussian, and proceed only if this is the case.

2. Use FastICA or any other applicable ICA algorithm (Hyvärinen et al., 2001) to obtain a decomposition $\hat{U} = P\hat{E}$, where P is $k \times k$ and $\hat{E}$ is $k \times T$, such that the rows of $\hat{E}$ are the estimated independent components of $\hat{U}$. Then validate non-Gaussianity and (at least approximate) statistical independence of the estimated components before proceeding.

3. Let $\Gamma_{ICA} = P^{-1}$. Find $\tilde{\Gamma}_0$, the row-permuted version of Γ_{ICA} which minimizes $\sum_{j=1}^{k} 1/|\tilde{\Gamma}_{0_{jj}}|$ with respect to the permutation. Note that this is a *linear matching problem* which can be easily solved even for high k (Shimizu et al., 2006).

4. Divide each row of $\tilde{\Gamma}_0$ by its diagonal element, to obtain a matrix $\hat{\Gamma}_0$ with all ones on the diagonal.

5. Let $\tilde{B} = I - \hat{\Gamma}_0$.

6. Find the permutation matrix Z which yields a matrix $\hat{B} = Z\tilde{B}Z'$ which is as close as possible to strictly lower triangular. This can be formalized as minimizing the sum of squares of the permuted upper-triangular elements, and minimized using a heuristic procedure (Shimizu et al., 2006). Set the upper elements of $\hat{B}$ to zero.

7. Calculate estimates of $\hat{\Gamma}_i$ for lagged effects using $\hat{\Gamma}_\ell = (I - \hat{B})\hat{A}_\ell$, for $\ell = 1, \dots, p$.

Algorithm 2: VAR-LiNG

1. Same as step 1 in algorithm 1.
2. Same as step 2 in algorithm 1.
3. Let $\Gamma_{ICA} = P^{-1}$. Test which entries of Γ_{ICA} are zero. This is done using a bootstrap procedure (see main text).
4. Find all admissible row-permuted matrices $\tilde{\Gamma}_{0,1}, \ldots, \tilde{\Gamma}_{0,m}$ of Γ_{ICA} such that each $\tilde{\Gamma}_{0,h}$ has zeroless diagonal for $h = 1, \ldots, m$.
5. Divide each row of $\tilde{\Gamma}_{0,h}$ by its diagonal element, to obtain a matrix $\hat{\Gamma}_{0,h}$ with all ones on the diagonal, for each $h = 1, \ldots, m$.
6. Let $\tilde{B}_h = I - \hat{\Gamma}_{0,h}$, for each $h = 1, \ldots, m$.
7. Calculate estimates of $\hat{\Gamma}_{\ell,h}$ for lagged effects using $\hat{\Gamma}_{\ell,h} = (I - \hat{B}_h)\hat{A}_\ell$, for $\ell = 1, \ldots, p$, for $h = 1, \ldots, m$.

4 Data

We use the annual survey of manufacturing plants (Encuesta Nacional Industrial Manufacturera – ENIA) collected by the Chilean Statistical Institute (Instituto Nacional de Estadísticas – INE). The ENIA covers the universe of Chilean plants in the manufacturing sector and has been widely used by researchers (see e.g. Alvarez et al., 2016; Crespi et al., 2019). We use the database that covers the period from 2001 to 2007.[8] The database includes all firms with more than 10 employees that have registered some activity for at least one semester during an year, divided by manufacturing sector (ISIC version 3, at the 4-digit level). For more information on the database see INE (2006, 2009a).

[8]Data are available since 1979, but the INE changed the data collection and in particular the registration of firms in 2001, which, at the time of our analysis, does not allow to correctly track plants/firms across the pre- and post-2000 periods. Attempts to match the two periods and build a longer panel are part of future work.

After some preliminary data cleaning,[9] we create our SVAR variables. The variables used for the SVAR are size, proxied by employment ($empl$); output, which is proxied by total sales ($output$), and can be sub-divided into domestic sales ($domsales$) and exports (exp); and also productivity. Sales, exports, and employment are easily derived from the ENIA database, while the estimation of productivity requires further discussion.

All variables in the ENIA are in nominal values. We thus deflate the variables used in this paper to real values before computing the productivity. For output and material inputs we use the deflators computed by the INE for each of the 4-digit (ISIC) sectors (INE, 2009b). Unfortunately the report includes deflators only until 2006. Although we could use deflators from other sources for 2007, we prefer to drop the year 2007 from the data instead of having constant price variables computed from different sources. Also, INE (2009b) does not include deflators for a number of 4-digit sectors. We attempted some aggregations to avoid losing firms in those sectors, but the differences among sectors were too large, leading to an increase in the error of the computation of constant price variables, which seems less desirable than dropping a few observations across the years.

The INE computes different deflators for the gross value of production, used for total sales ($output$) and exports (exp), for overall input costs, used for variable inputs ($Material$), i.e. excluding capital, and for material inputs not completely transformed in the production process, used to compute beginning of the year and end of the year raw and input materials (respectively $Privap$, $Privaf$, $Matvap$ and $Matvaf$). To compute value added at constant prices (Va) we use the generally preferred method of double deflation, and we remove initial inputs and add left overs at the end of the year: $Va = output - Material - (Privap + Matvap) + (Privaf + Matvaf)$.

To compute the value of capital at constant prices we follow, in part, Crespi (2004) and use the implicit deflator for gross fixed capital formation released by the Central Bank (Banco Central de Chile, 2004, 2006, 2009). For our purposes, we did not consider estimating different deflators for different types of capital (machinery, buildings land and vehicles), because we could not find accounting information available for vehicles and land.

Finally, we deflate the input variables used to compute productivity with the gross value of production ($output$): primary inputs, input materials purchased, primary and material inputs from other plants (of the same firm), office material – deflator for non-completely transformed inputs – and fuel – deflator for com-

[9] We first check for inconsistencies in the data (Benavente and Ferrada, 2004) i.e. plants that report 0 days in operation, a negative gross value of production, 0 or negative number of employees, labour cost equal or less than 0, sales lower than exports, value added larger than sales and an ISIC code lower than 1500. A non significant number of observations need to be dropped across the 7 years.

pletely transformed inputs.

We then proceed to estimate total factor productivity (TFP) employing the Levinsohn and Petrin (2003) method (see also Petrin et al., 2004), and using the quantity of consumed electricity as an intermediate input. It is worth nothing that estimations of the TFP using value added and the whole sample of firms is highly correlated with labour productivity with a Spearman's correlation index of 0.96. However, for the sake of comparability with most other studies on the relation between export and productivity we use TFP estimations.

Although differences are again quite small, we choose to estimate TFP using output rather than value added. The main advantage of using output is that there is a non-negligible number of firms that in some years have negative value added (at constant prices), requiring a further drop of observations.

Arguably, plants may differ quite substantially in their production technology. It follows that using one single production function with labour and capital (and one intermediate input) may produce biased estimates. To overcome this problem we attempt a large number of estimations, taking into account different combinations of the following dimensions: size, labour, and sector.

Using the ISIC Rev3 2-digit classification we create the following relatively homogeneous sectors: (1) Manufacture of Food, Beverages and Tobacco; (2) Textile, Wearing Apparel and Leather Industries, traditional industries; (3) Manufacture of Wood and Wood Products, Including Furniture ; (4) Manufacture of Chemicals and Chemical, Petroleum, Coal, Rubber and Plastic Products; (5) Manufacture of other non-metallic mineral products and basic metals; (6) Manufacture of fabricated metal products, except machinery and equipment; (7) Manufacture of machinery and equipment, office, accounting and computing machinery, electrical machinery and apparatus, radio, television and communication equipment and apparatus, medical, precision and optical instruments, watches and clocks, motor vehicles, trailers and semi-trailers, and other transport equipment; (8) Publishing, printing and reproduction of recorded media; (9) Manufacture of paper and paper products; and (10) Other manufacturing sectors.

We create sub-samples for different size categories, based on number of employees: small (< 50), medium ($50 \leq empl < 250$) and large (≥ 250) firms. Furthermore, we attempt different measures of labour skills as variable inputs in the production function.

As expected, TFP estimations, as well as returns to scale, differ significantly when computed for different sectors and plant sizes. The distinction between different types of workers also significantly affects TFP and returns to scale. We leave the discussion on these significant differences for a different paper. For this paper it suffices to say that we consider as our most reliable estimates those obtained separating the different sectors and including in the production function 'blue collars', 'white collars', material inputs, and capital (tfp). However, in this paper we also attempt some robustness checks, using a TFP estimated with no distinction between different types of employment ($tfp2$), leading to no significant differences in the relation between exporting and productivity.

Finally, we remove firms that we consider outliers. For each of the VAR series – growth of sales, employment, exports and productivity – we impose a threshold for outliers corresponding to tenfold growth/decline in the space of one year. Observations beyond this threshold are dropped.

Table 1 summarises the variables used for the analysis.

Table 1: Summary statistics

Variable	**Description**	**Obs**	**Mean**	**Std. Dev.**	**Min**	**Max**
gr_empl	Employees	2303	0.027	0.284	-2.223	2.137
gr_exp	Export sales	2303	-0.011	0.606	-2.256	2.254
gr_tfp	TFP^b	2303	-0.011	0.266	-1.507	2.030
gr_tfp2	TFP^c	2303	-0.015	0.270	-1.793	2.243
gr_domsales	Domestic market sales	2303	0.003	0.543	-4.196	5.059

Notes: [b] Estimated for different sectors, and differentiating between blue white collars. [c] Estimated for different sectors, without differentiating between blue and white collars

Table 2: Correlation matrix.

	gr_domsales	gr_empl	gr_exp	gr_tfp
gr_domsales	1	*0.1343*	*-0.0132*	*0.4283*
gr_empl	0.0924	1	*0.1436*	*-0.1881*
gr_exp	-0.0717	0.0915	1	*0.1605*
gr_tfp	0.3327	-0.2599	0.0836	1

Notes: Lower triangle: Pearson correlation coefficients; upper triangle (and italics): Spearman's rank correlation coefficients. 4021 observations. All correlations significant at the 1% level, except for the Spearman rank correlation between gr_domsales and gr_exp (ρ=-0.0132, p-value=0.4029)

5 Results

5.1 Preliminary evidence

In this section, we present the findings of our SVAR-based causal analysis. As preliminary evidence, Table 2 shows the correlations (both Pearson and Spearman's coefficients) between the variables of interest. Our four main variables are significantly correlated between them — although the magnitudes of the correlations are far below the typical values of around 70% that are associated with problems of multicollinearity (e.g. Hair et al., 1998). The relationship between growth of exporting and growth of TFP, which is the main object of our analysis, displays a significantly positive correlation. The problem with correlation analysis, of course, is that it cuts off many possible channels of interactions among variables, especially between lagged values, and that is silent about causal directions.

These aspects are accounted for in our SVAR analysis. As first step (corresponding to step 1 in both algorithm 1 and 2), we estimate a reduced-form VAR model (equation 5). The coefficient matrices are estimated with median regression, also called least absolute deviation (LAD) regression. This is in line with suggestions in Moneta et al. (2013) in the context of non-Gaussian data with the motivation of improving the robustness to outliers. The VAR model is estimated both with one lag ($p = 1$) and two lags ($p = 2$). We will refer to these specifications as the 1-lag and the 2-lag model, respectively.[10]

We then investigate whether the assumption of non-Gaussian shocks is plausible (this task is also part of step 1 in both algorithms). Previous research has shown that the distribution of firm growth rates is heavy-tailed and non-Gaussian (Bottazzi et al., 2002; Capasso et al., 2013). Although the assumption of non-Gaussian shocks is not directly testable, it is possible to assess the departure from normality of the estimated reduced-form residuals, which are assumed to be linear combinations of the independent shocks. In Figure 1, we graphically compare the empirical distribution of the estimated reduced-form residuals with a theoretical normal distribution with the same mean and variance. We do this by comparing the histograms with the theoretical normal density function (top panel). The relative peakedness of the four residuals suggest the presence of leptokurtic (i.e. supergaussian) distributions. We also plot (bottom panel) the quantiles of the empirical distribution of the reduced-form residuals against the quantiles of the theoretical normal distributions. These q-q plots, being far from straight lines, suggest heavy departure from normality. This is also confirmed by a battery of tests, in which the common null hypothesis is normality of the reduced-form residuals: the Shapiro-Wilk, the Shapiro-Francia, and the Jarque-Bera tests reject in all cases the null hypothesis at the 0.01 level of significance.

The second step of both algorithms (VAR-LiNGAM and VAR-LiNG) aims at delivering the mixing matrix linking structural shocks to reduce-form residuals. Table 3 presents the mixing matrix (point estimate coefficients) P for both the

[10]According to different information criteria (Akaike, Hannan-Quinn, Schwarz), the 2-lag model is slightly preferable to the 1-lag model. We do not explore further lags since this would reduce considerably our sample size.

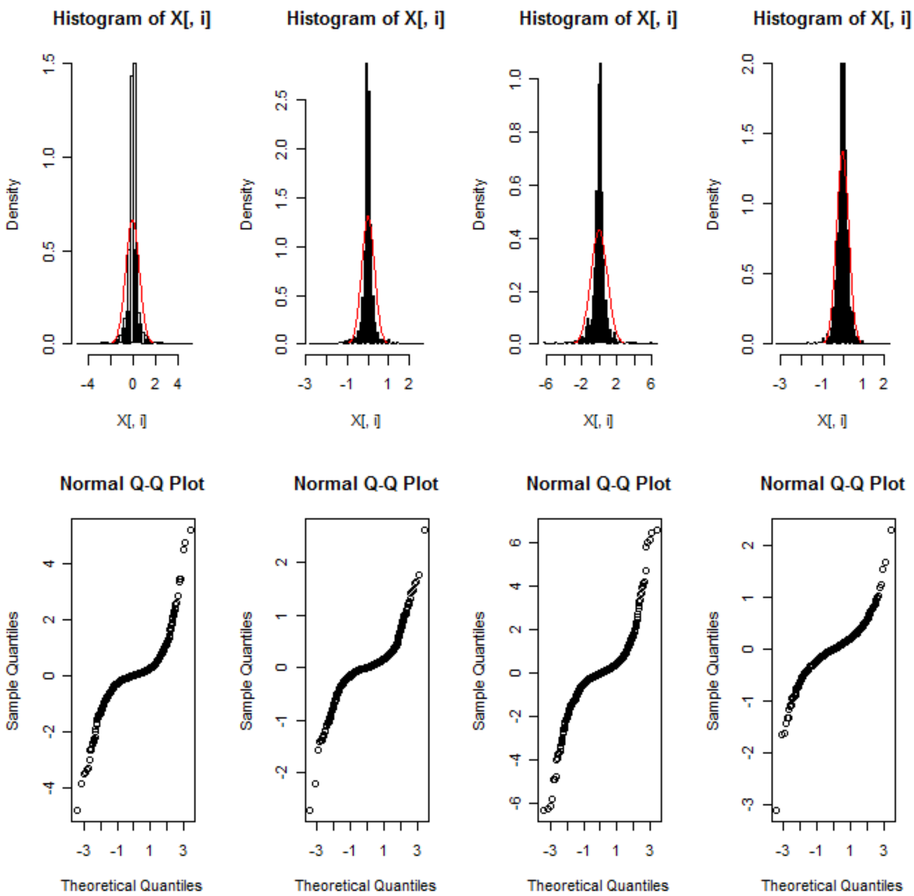

Figure 1: Empirical distributions (top panel, red lines show the corresponding theoretical normal distribution) and quantile-quantile plots (bottom panel; quantiles of the empirical distribution vs. quantiles of the theoretical normal distribution) of the reduced-form residuals (corresponding to the four variables in the following order: growth of domestic sales, growth of employment, growth of exports, and growth of TFP), 2-lag model.

1-lag and the 2-lag model. This matrix is the output of the FastICA algorithm (Hyvärinen et al., 2001) applied to the reduced-form residuals. The matrix P is, in our sample, a (4×4) matrix such that $\hat{U} = P\hat{E}$, where $\hat{U}$ is the $(4 \times T)$ matrix of the estimated reduced-form residuals and $\hat{E}$ is a $(4 \times T)$ matrix of independent components.[11] We recall that the scale (sign) and order of these components, and therefore of the columns of the matrix P, is undetermined. The columns of P displayed in Table 3 are rescaled in order to produce components with unit variance. Although the order of the columns is completely arbitrary, the order of the rows is determined by the order of the variables entering in y_t.[12] We note that some entries of this matrix are close to zero. For each column we can easily identify a coefficient (in absolute value) that is maximally loading on a particular variable. For example, looking at the fourth column of P for the 1-lag model (left part of the table) we see that there is an entry which has the highest value (among the column-entries, in absolute value) for exporting growth (0.9466). Looking at the third row, we also see that this is the maximum value (among the row-entries). This means that the shock labelled as e_4 (in the 1-lag model) is mostly loading on exporting growth. The same shock has a minimal impact on productivity growth (0.0017), which is the smallest entry both in the fourth column and in the fourth row. If we look at the matrix P of the 2-lag model (right part of table 3), the impact of the shock labelled as e_3 has very similar characteristics (and almost equal values) to the shock labelled as e_4 in the 1-lag model. Taken together, this suggests that the exporting growth shock does not transmit (within the one or two years periods) to productivity growth. In other words, results suggest that there is no contemporaneous causal relationship from exporting to productivity growth.

Table 3: Mixing matrix

	1-lag model				2-lag model			
	e_1	e_2	e_3	e_4	e_1	e_2	e_3	e_4
gr_domsales	-0.1341	0.5872	0.0457	-0.0581	0.0357	-0.5783	-0.0518	-0.1561
gr_empl	-0.0058	0.0171	0.3061	0.0098	0.3036	-0.0164	0.0099	-0.0161
gr_exp	-0.1063	-0.0124	0.0692	0.9466	0.0452	0.0101	0.9168	-0.1219
gr_tfp	-0.2968	0.0472	-0.0887	0.0017	-0.0950	-0.0436	0.0023	-0.2723

Notes: Point estimates coefficients of the mixing matrix P as output of fastICA algorithm. This corresponds to step 2 of both Algorithm 1 and 2. Left part: estimation from residuals of a 1-lag VAR. Right part: estimation from residuals of a 2-lag VAR.

Steps 3-6 of the algorithms we use (VAR-LiNGAM and LiNG) aim at infer-

[11] $T = 2658$ with one lag and is equal to 1667 with two lags.

[12] In other words, each time we run fastICA we get a (randomly) column-permuted version of P with, in addition, random changes of sign for each column. A part from column permutation and changes of sign, all the entries of the output matrices from multiple (1000) realisations of fastICA are identical, which confirms a stable convergence of the fastICA algorithm.

ring the causal relationships, such as the one just elicited, in a more formal and rigorous way. As mentioned, VAR-LiNGAM assumes that there is a recursive causal-structure. This means that the mixing matrix P contains at least $k(k-1)/2$ entries (in our case 6 zeros since $k = 4$) that are (statistically close to) zero. Recursiveness also implies that if any entry (i, j) of P is (significantly) different from zero, then the entry (j, i) must be (statistically close to) zero. In the VAR-LiNGAM setting this property is not tested empirically, but assumed a priori. In order to overcome this limitation of the VAR-LiNGAM method, and improve the empirical reliability of our causal inference, we first identify the model through VAR-LiNGAM, and then use a bootstrap procedure to check whether the causal directions found are robust under resampling. It turns out that most of the causal directions are robust, but some of them are reversed in artificial samples. We finally apply VAR-LiNG, which does not assume recursiveness, to see whether causal loops emerge. The causal relationship we are interested in, namely between productivity and export growth, emerges as robustly identified.

5.2 VAR-LiNGAM results

Table 4 shows the coefficients of structural VAR matrices (see equation 7) estimated through VAR-LiNGAM (algorithm 1 in section 3.2). The upper (lower) part of the table refers to the 1-lag (2-lag) model. The first block of 4 columns corresponds to the estimated coefficients of matrix B (contemporaneous effects) (recall $B = I - \Gamma_0$), the second block (columns 5-8) refers to the coefficients of matrix Γ_1 (one-period-lag causal effects), while the third block (columns 9-12) presents the estimated coefficients of Γ_2 (two-period-lag causal effects).

As the literature on algorithmic causal inference has demonstrated (Spirtes et al., 2000; Pearl, 2009; Peters et al., 2017), structural models can be represented as directed graphs, and directed acyclic graphs (DAGs) in case of recursive structure. We thus represent the SVAR model, output of the VAR-LiNGAM algorithm, as a DAG, in order to improve the causal interpretation of the model. The DAG is built on the criterion that a non-zero entry in the (i, j) position of B corresponds to a directed edge (i.e. arrow) from the j^{th} to the i^{th} variable in the sub-graph referring to the contemporaneous values. Analogously, a (statistically significant) non-zero entry in the (i, j) position of the Γ_1 corresponds to a directed edge (i.e. arrow) from the j^{th} variable at time $t-1$ to the i^{th} variable at time t.[13] Figure 2 shows the resulting DAG for the 1-lag model.

[13]Since the asymptotic distribution of the VAR-LiNGAM-estimated coefficients is unknown, we cannot rely on a formal significance test. As rule of thumb, we do not represent a causal arrow if the corresponding coefficient is significantly close to zero according to a standard t statistic, where the standard errors are calculated following a bootstrap procedure. In Table 4 coefficients significantly different from zero are represented in bold.

Table 4: VAR-LiNGAM estimates

					first lag				second lag			
	gr_domsales	gr_empl	gr_exp	gr_tfp	l1_gr_domsales	l1_gr_empl	l1_gr_exp	l1_gr_tfp	l2_gr_domsales	l2_gr_empl	l2_gr_exp	l2_gr_tfp
gr_domsales	0	**0.4787**	**-0.0871**	**0.8301**	**-0.2248**	**0.1159**	-0.0148	**0.2906**				
	0	0.0619	0.0127	0.0773	0.0384	0.033	0.008	0.0472				
gr_empl	0	0	0	0	0.0073	-0.0229	0.0087	0.0162				
	0	0	0	0	0.0068	0.0234	0.0041	0.0122				
gr_exp	0	**0.4369**	0	**0.4023**	-0.0144	-0.0216	**-0.1442**	0.1088				
	0	0.0889	0	0.0707	0.0157	0.0344	0.0392	0.0442				
gr_tfp	0	**-0.2712**	0	0	0	**-0.0579**	**0.0121**	**-0.2699**				
	0	0.0288	0	0	0.0096	0.0178	0.0044	0.02				
gr_domsales	0	**0.5323**	**-0.0897**	**0.9468**	**-0.2182**	0.0655	-0.0047	**0.3025**	-0.0653	0.0355	0.0139	**0.1317**
	0	0.0839	0.0145	0.102	0.0437	0.038	0.0085	0.0673	0.0542	0.0341	0.01	0.0371
gr_empl	0	0	0	0	0.0145	-0.0196	0.0065	0.0321	0.0073	0.0122	0.0045	**0.0498**
	0	0	0	0	0.0085	0.0309	0.0073	0.0174	0.0087	0.0261	0.0064	0.0184
gr_exp	0	**0.4345**	0	**0.4743**	0.0237	-0.0421	**-0.1472**	0.0886	0.0208	0.1027	**-0.0944**	0.0743
	0	0.082	0	0.0812	0.021	0.0477	0.0272	0.067	0.0205	0.0491	0.03	0.051
gr_tfp	0	**-0.2466**	0	0	-0.0016	-0.0359	0.0109	**-0.261**	-0.0232	0.0124	0.0001	**-0.0836**
	0	0.0307	0	0	0.0119	0.0211	0.0051	0.0375	0.0108	0.0274	0.006	0.0238

Notes: Structural coefficient matrices, output of the VAR-LiNGAM algorithm. Top part of the table: 1-lag model. Bottom part of the table: 2-lag model. Left block: contemporaneous effects matrices (B). Middle block: lagged effects matrix Γ_1. Right block: lagged effects matrix Γ_2. Standard errors are obtained through bootstrap (1000 iterations). Bold coefficients refer to significant entries (95% confidence intervals).

Figure 2: Causal graph resulting from VAR-LiNGAM, 1-lag model

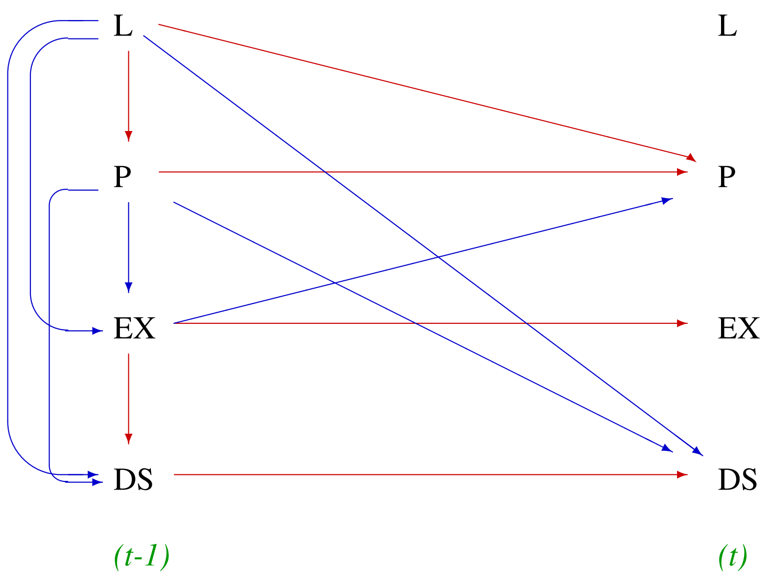

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

Both the 1-lag and 2-lag models in Table 4 show that the *primus motor* is employment growth, which has large positive effects on growth of domestic sales and growth of exports. These can both be interpreted as firm strategic decisions: when the firm intends to increases output, they first need to increase the scale. Note that the sum of these two coefficients is close to unity (0.4787 + 0.4369 = 0.9156 in the 1-lag model; and 0.5323 + 0.4345 = 0.9668 in the 2-lag model), which implies that the elasticity of employment growth to combined growth of outputs (i.e. domestic sales + exports) is close to unity when considering instantaneous effects.

Another central result is that employment growth has a negative effect on contemporaneous growth of TFP, presumably because efficiency increases when fewer inputs (i.e. employees) are used to produce the same amount of output. Downsizing firms achieve higher productivity in the same time year than firms that invest in recruiting and training new employees.

Growth of TFP has positive impacts on growth of domestic sales and, to a lesser extent, growth of export. Firms that experience an increase in their productivity are therefore more likely to grow in terms of domestic and export sales. This might suggest that firms would be better off pursuing productivity growth as a prerequisite for subsequent sales growth, instead of vice versa.

Growth of exporting has a negative impact on growth of domestic sales. This reflects the tension between domestic vs. exporting sales strategies, that was already visible in the negative correlations between these variables in Table 2. It is interesting to observe that export seems to determine domestic sales rather than vice versa. This could be because internationalized firms have already 'conquered' their home markets and become 'outward-focused' in the sense that they pay more attention to how they fare in the more competitive export markets. But as we discuss in Section 5.3 this result is not robust.

With regards to the causal link between TFP and export, our results suggest that it is an increase in TFP that causes an increase in export in the same year, rather than the other way around. Our VAR-LiNGAM estimates provide an interesting perspective on the export-productivity debate: not only firms need to be more productive to enter the export market; once they export, they need to keep increasing TFP to remain competitive. As noted, this comes hand in hand with an increase in size, measured as an increase in labour (which has a negative effect on TFP). Note, however, that export growth has a small positive impact on TFP growth in the next year. The increased export partly contributes to increasing productivity.

Very few coefficients are significant in the 2-lags model. TFP maintains its positive impact on domestic sales, and interestingly pushes to a further increase in firm size through employment.

5.3 VAR-LiNGAM robustness analysis

We run a robustness analysis to check whether the causal links depicted in Figure 2 are stable over 1000 bootstrap samples, which were created by resampling with replacement from the original data. We focus here only on the contemporaneous causal structure. As Figure 3 shows, all the causal links found by VAR-LiNGAM are very robust across bootstrap samples except the link between growth of domestic sales (DS in figure 3) and growth of exporting sales (EX in figure 3), which is reversed almost half of the times.

5.4 VAR-LiNG results

Table 5 reports the estimates of the application of algorithm 2 (VAR-LiNG, i.e. the algorithm which allows the possibility of feedback loops in the contemporaneous structure) for the model with one lag.[14] Figure 4 depicts graphically the contemporaneous causal structure, while Figure 5 shows the lagged causal structure.

With respect to the VAR-LiNGAM algorithm, results mainly differ in relation the contemporaneous structure, for which it is most difficult to estimate causal

[14]The two-lag model produces very similar results.

Figure 3: Bootstrap robustness analysis on the contemporaneous causal structure. VAR-LiNGAM is applied to each bootstrap iteration. Numbers associated to edges indicate the percentage each causal link is inferred out of 1000 bootstrap iterations.

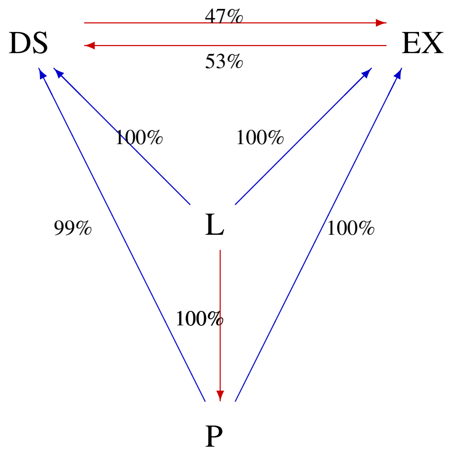

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

relations. The estimated contemporaneous causal structure shows a bi-directional negative effect between growth of domestic sales and growth of exporting sales, as suggested by the robustness results shown in Figure 3. This confirms the improvements in the causal estimations with the VAR-LiNG algorithm discussed in Section 3.2

Feedback loops (positive) emerge also between growth of domestic sales and growth of productivity and between growth of domestic sales and growth of employment (positive) in the contemporaneous causal structure.

The main finding about the causal relation between productivity and exporting growth which resulted from the application of the first algorithm is confirmed: there is no causal influence from exporting to productivity (growth) in the contemporaneous causal structure. But the viceversa holds: firms must keep increasing productivity to increase their presence in the foreign market. The coefficient which measures the instantaneous influence from productivity (growth) to exporting is very close to the coefficient obtained from the first algorithm: 0.3782 vs. 0.4023 (standard errors are also very similar: 0.0773 vs. 0.0707).

Results using the VAR-LiNG algorithm do not confirm the small positive lagged impact of export growth on productivity growth that was found with the VAR-LiNGAM algorithm (respectively Tables 5 and 4). These suggest that, not only firms must improve TFP to increase their presence in the export market, but there is also no causal evidence of "learning-by-exporting" that translates in

Figure 4: Contemporaneous causal graph resulting from VAR-LiNG 1 lag

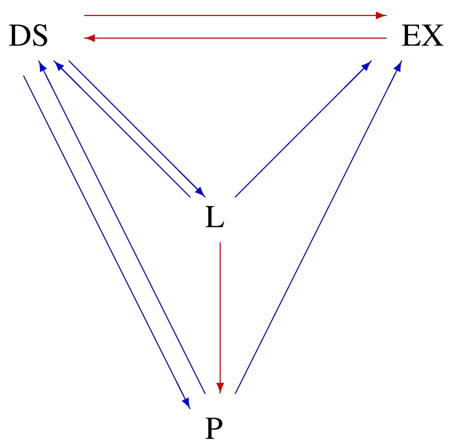

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

an increase in TFP in the year following an increase in exports. While the results from the VAR-LiNGAM algorithm showed a (weak) influence from lagged export growth to current productivity growth, the results from the VAR-LiNG algorithm show no such causal link (see Figure 5 for the one-year lagged effects).

Table 5: VAR - LiNG estimates

	gr_domsales	gr_empl	gr_exp	gr_tfp	l_gr_domsales	l_gr_empl	l_gr_exp	l_gr_tfp
gr_domsales	0	**0.2981**	**-0.0653**	**0.4620**	**-0.2240**	**0.0939**	-0.0066	**0.1925**
	0	0.0478	0.0131	0.1115	0.0062	0.0197	0.0038	0.0126
gr_empl	**0.0292**	0	0.0000	0.0000	0.0137	-0.0249	0.0085	0.0141
	0.0083	0	0.0077	0.0122	0.0062	0.0197	0.0038	0.0126
gr_exp	**-0.0615**	**0.3449**	0	**0.3782**	-0.0274	-0.0208	**-0.1426**	0.1080
	0.0199	0.0602	0	0.0773	0.0170	0.0462	0.0374	0.0530
gr_tfp	**0.0914**	**-0.3058**	0.0000	0	0.0205	**-0.0648**	0.0116	**-0.2757**
	0.0234	0.0291	0.0077	0	0.0111	0.0200	0.0045	0.0239

Notes: Structural coefficient matrices, output of the VAR-LiNG algorithm, 1-lag model. Left block: contemporaneous effects matrices (B). Right block: lagged effects matrix Γ_1. Standard errors are obtained through bootstrap (1000 iterations). Bold coefficients refer to significant entries (95% confidence intervals).

5.5 Further robustness checks

We performed several robustness studies to explore how stable our finding about a contemporaneous causal effect from productivity growth to exports growth (and a lack of causal effect in the opposite direction) is, when we look at different

Figure 5: Lagged causal structure resulting from VAR-LiNG 1 lag

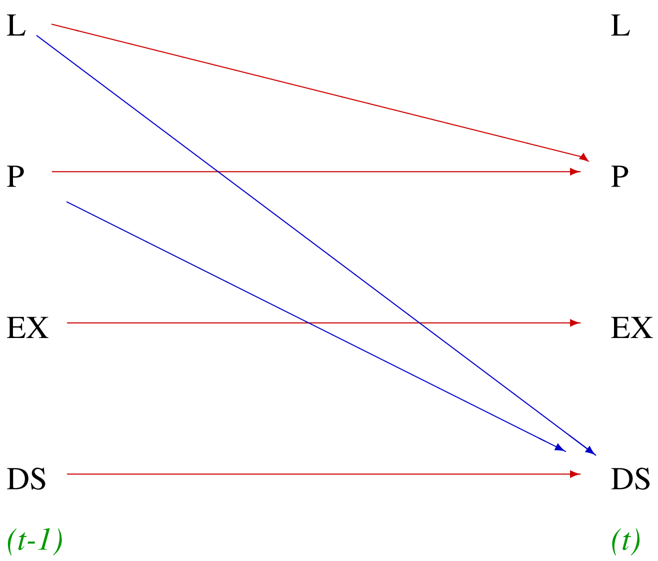

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

variable measures, sub-samples, and methods of estimation. As mentioned, TFP can be estimated without distinguishing between different types of employment. This does not bring any difference to the VAR-LiNGAM estimate with respect to the baseline results. However, the contemporaneous causal effect matrix resulting from VAR-LiNG is not sparse enough to deliver an informative and reliable output. We also looked at different sub-samples guided by different criteria: sectors (following the taxonomy by Ferraz et al., 1996), size, and productivity. The VAR-LiNGAM results is remarkable stable in delivering a causal structure in which productivity growth is always ordered before exporting growth. VAR-LiNG delivers in many cases a lack of causal influence from export growth to productivity growth (in line with the baseline results). But in many other cases the output of VAR-LING is not stable under variations of initial conditions, which we attributed to a reduction of the sample size. We also considered the results from a VAR model estimated by OLS rather than by LAD. In terms of causal structures, there are no difference with the baseline results, both as regards VAR-LiNGAM and VAR-LiNG. All these results are reported in Appendix (supplementary material).

6 Discussion

In this paper we revisit a well-known debate, nurtured by growing but contradicting empirical evidence. Does exporting activity increase firm productivity or do firms remain in the export market only if they managed to increase productivity?

To address this important question we choose a rather different strategy from previous papers, while exploring a relatively well-studied country. First, we focus on the intensive margin of exporting firms and compare firms that experience different growth rates of exports (i.e. become more or less competitive in the international market). Second, we explicitly take into account the bi-directional causal relation between productivity and export growth, including within the same year and with up to two lags.

Applying VAR-LiNGAM and VAR-LiNG, a class of SVAR models that estimates causal networks, we find that export growth does not have a direct and instantaneous causal impact on firm productivity growth within the same period. Although results from the VAR-LiNGAM model show that with one year lag export growth does have a small causal effect on TFP growth, when we apply the more robust VAR-LING model, which allows to account for feedback loops in the contemporaneous causal structure, this lagged causal effect vanishes. Instead, we observe that TFP growth has a direct large contemporaneous causal effect on export growth, which is robust to the application of both the VAR-LiNGAM and VAR-LiNG algorithms, and to resampling robustness checks.

Our results are estimates of causal effects and therefore have interesting implications for policy. In particular, it appears that firms should focus on improving their productivity in order to be able to increase their exports, because it is productivity growth that drives growth of exports, and not the other way round. While there are certainly benefits from exporting, we do not find that firms in an emerging economy such as Chile can rely on exporting as a means to increase their productivity – and there remain competitive in the export market.

For example, firms should keep improving their productivity through e.g. redesigning their production routines, upgrading their capital and IT systems, and improving management (Bloom et al., 2012) alongside appropriate organizational innovations (e.g. Cruz et al., 2018). As a result, they will be in a better position to experience growth in their participation in global trade. There is instead no strong influence of exporting on TFP growth – including with a lag. This effect is relatively small and not robust to accounting for the possibility of feedback loops in the contemporaneous structure between variables.

Our study is not without limitations. First, although we have no reason to expect that our data is unrepresentative, it is nevertheless not clear how our results can be generalized to other countries and other periods. However, we are interested and curious to check how our method would change earlier LBE results in other countries where the hypothesis has been tested, using the same data.

Second, we focus on exporting undertaken by firms that have already taken the binary decision of exporting, and succeeded. There may be differences in the exporting-productivity relationship at the time when a firm enters the export market and may have more to learn.

The application in this paper has allowed to shed new light on the long-lived controversy of whether firms benefit from exporting, or need to boost productivity to be able to increase exporting. We have show that, in the case of Chilean firms, between 2001-07, the contemporaneous causal direction runs from productivity growth to exporting. A one or two years lag does not show robust causal evidence that an increase in export will lead to increased productivity after one or two years. Future work could apply the family of techniques developed here to a broad range of controversial empirical evidence to produce new valuable evidence for academics, practitioners and policymakers.

This paper has also shown how data-driven techniques for causal inference can be introduced from the machine learning community into economics, and adapted to time-series and VAR contexts, to provide new evidence on the causal relations governing economic systems.

Appendix. Supplementary material

Supplementary material related to this article can be found online. The supplementary material contains impulse response function analysis and further robustness checks.

References

Alvarez, R., C. Bravo-Ortega, and L. Navarro (2016). Product mix changes and performance in Chilean plants. *Industrial and Corporate Change 25*(6), 1001–1017.

Alvarez, R. and G. Crespi (2007). Multinational firms and productivity catching-up: the case of Chilean manufacturing. *International Journal of Technological Learning, Innovation and Development 1*(2), 136 – 152.

Alvarez, R. and R. López (2005). Exporting and performance: evidence from Chilean plants. *Canadian Journal of Economics 38*(4), 1384–1400.

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

Antolín, M. M., J. A. Máñez, M. E. Rochina Barrachina, and J. A. Sanchis Llopis (2013). Export intensity and the productivity gains of exporting. *Applied Economics Letters 20*(8), 804–808.

Arnold, J. M. and K. Hussinger (2005, July). Export behavior and firm productivity in german manufacturing: A firm-level analysis. *Review of World Economics (Weltwirtschaftliches Archiv) 141*(2), 219–243.

Aw, B. Y., M. J. Roberts, and D. Y. Xu (2011). R&D investment, exporting, and productivity dynamics. *American Economic Review 101*(4), 1312–44.

Aw, B. Y. and B. Song (2013). Firm's choice between export and r&D and its effect on the firm's productivity and survival. *Global Economic Review 42*(3), 291–310.

Baldwin, J. and B. Yan (2012). Export market dynamics and plant-level productivity: Impact of tariff reductions and exchange-rate cycles. *The Scandinavian Journal of Economics 114*(3), 831–855.

Banco Central de Chile (2004). Anuario de cuentas nacionales de Chile 2003. Technical report, Banco Central de Chile.

Banco Central de Chile (2006). Cuentas nacionales de Chile 1996-2005. Technical report, Banco Central de Chile.

Banco Central de Chile (2009). Cuentas nacionales de Chile 2003-2008. Technical report, Banco Central de Chile.

Benavente, J. M. and C. Ferrada (2004). Probability of survival of new manufacturing plants: the case of Chile. mimeo, Economics Department, University of Chile.

Berman, N. and V. Rebeyrol (2010). Exporter dynamics and productivity growth. EUI MWP 2010/26, European University Institute.

Bernard, A. B. and J. B. Jensen (1995). Exporters, jobs, and wages in the u.s. manufacturing: 1976-1987. *Brookings Papers on Economic Activity*, 67 – 119.

Bernard, A. B. and J. B. Jensen (1999). Exceptional exporter performance: cause, effect, or both? *Journal of International Economics 47*(1), 1–25.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007, Summer). Firms in international trade. *Journal of Economic Perspectives 21*(3), 105–130.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2012, nov). Does Management Matter? Evidence from India. *The Quarterly Journal of Economics 128*(1), 1–51.

Bottazzi, G., E. Cefis, and G. Dosi (2002). Corporate growth and industrial structures: some evidence from the Italian manufacturing industry. *Industrial and Corporate Change 11*(4), 705–723.

Brancaccio, E., A. Califano, M. Lopreite, and A. Moneta (2020). Nonperforming loans and competing rules of monetary policy: A statistical identification approach. *Structural Change and Economic Dynamics 53*, 127–136.

Breitung, J. (2005). A parametric approach to the estimation of cointegration vectors in panel data. *Econometric Reviews 24*(2), 151–173.

Brenner, T., M. Capasso, M. Duschl, K. Frenken, and T. Treibich (2018). Causal relations between knowledge-intensive business services and regional employment growth. *Regional Studies 52*(2), 172–183.

Capasso, M., E. Cefis, and K. Frenken (2013). On the existence of persistently outperforming firms. *Industrial and Corporate Change 23*(4), 997–1036.

Castellani, D. (2002, December). Export behavior and productivity growth: Evidence from italian manufacturing firms. *Review of World Economics (Weltwirtschaftliches Archiv) 138*(4), 605–628.

Cirera, X., D. Lederman, J. Manez, M. Rochina, and J. Sanchis (2012). Trade status, trade policy and productivity: The brazilian case. mimeo, IDS.

Coad, A., M. Cowling, and J. Siepel (2017). Growth processes of high-growth firms as a four-dimensional chicken and egg. *Industrial and Corporate Change 26*(4), 537–554.

Coad, A., R. Rao, and F. Tamagni (2011). Growth processes of italian manufacturing firms. *Structural Change and Economic Dynamics 22*(1), 54–70.

Crespi, G., C. Criscuolo, and J. Haskel (2008, May). Productivity, exporting, and the learning-by-exporting hypothesis: direct evidence from UK firms. *Canadian Journal of Economics 41*(2), 619–638.

Crespi, G., E. Tacsir, and M. Pereira (2019). Effects of innovation on employment in Latin America. *Industrial and Corporate Change 28*(1), 139–159.

Crespi, G. A. (2004). *The Microfoundations of Aggregate Productivity in a Developing Country: The Case of Chilean Manufacturing*. Phd thesis, Science and Technology Ploicy Research (SPRU), Brighton.

Cruz, M., M. Bussolo, and L. Iacovone (2018). Organizing knowledge to compete: Impacts of capacity building programs on firm organization. *Journal of International Economics 111*, 1–20.

Dai, M. and M. Yu (2013). Firm R&D, absorptive capacity and learning by exporting: Firm-level evidence from china. *The World Economy 36*(9), 1131–1145.

Dalgiç, B., B. Fazlioğlu, and M. Gasiorek (2021). Does it matter where you export and does productivity rise with exporting? *Journal of International Trade and Economic Development 2021*(5), 766–791.

Damijan, J. P. and C. Kostevc (2006). Learning-by-exporting: Continuous productivity improvements or capacity utilization effects? evidence from Slovenian firms. *Review of World Economics (Weltwirtschaftliches Archiv) 142*(3), 599–614.

Damijan, J. P., C. Kostevc, and S. Polanec (2010). From innovation to exporting or vice versa? *World Economy 33*(3), 374–398.

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

Eliasson, K., P. Hansson, and M. Lindvert (2012). Do firms learn by exporting or learn to export? evidence from small and medium-sized enterprises. *Small Business Economics 39*(2), 453–472.

Eriksson, J. and V. Koivunen (2004). Identifiability, separability, and uniqueness of linear ica models. *IEEE signal processing letters 11*(7), 601–604.

Fernandes, A. M. and A. E. Isgut (2005, March). Learning-by-doing, learning-by-exporting, and productivity : evidence from colombia. Policy Research Working Paper Series 3544, The World Bank.

Ferraz, J. C., D. Kupfer, and L. Haguenauer (1996). El desafío competitivo para la industria brasileña. *Revista de la CEPAL 1996*(58), 145–173.

Garcia-Marin, A. and N. Voigtländer (2019). Exporting and plant-level efficiency gains: It's in the measure. *Journal of Political Economy 127*(4), 000–000.

Girma, S., D. Greenaway, and R. Kneller (2004). Does exporting increase productivity? a microeconometric analysis of matched firms. *Review of International Economics 12*(5), 855–866.

Gouriéroux, C., A. Monfort, and J.-P. Renne (2017). Statistical inference for independent component analysis: Application to structural VAR models. *Journal of Econometrics 196*(1), 111–126.

Greenaway, D. and R. Kneller (2007, 02). Firm heterogeneity, exporting and foreign direct investment. *Economic Journal 117*(517), F134–F161.

Hair, J. F., W. C. Black, B. J. Babin, and R. E. Anderson (1998). *Multivariate data analysis*, Volume 5. Prentice hall Upper Saddle River, NJ.

Horn, R. A. and C. R. Johnson (2012). *Matrix analysis*. Cambridge university press.

Hyvärinen, A. (2013). Independent component analysis: recent advances. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 371*(1984), 20110534.

Hyvärinen, A., J. Karhunen, and E. Oja (2001). *Independent Component Analysis*. Wiley.

Hyvärinen, A. and E. Oja (2000). Independent component analysis: Algorithms and applications. *Neural Networks 13*(4-5), 411–430.

Hyvärinen, A., S. Shimizu, and P. O. Hoyer (2008). Causal modelling combining instantaneous and lagged effects: an identifiable model based on non-Gaussianity. In *Proceedings of the 25th International Conference on Machine Learning (ICML-2008), pp. 424-431, Helsinki, Finland.*

INE (2006). Microdatos industriales revisados. informe metodológico 1998-2005. Technical report, Instituto Nacional de Estadísticas.

INE (2009a). Encuesta nacional industrial anual (enia). tomo i and tomo ii. Technical report, Instituto Nacional de Estadísticas.

INE (2009b). Metodología de construcción de deflactores. encuesta nacional industrial anual 1998-2006. Technical report, Instituto Nacional de Estadísticas.

International Study Group on Exports and Productivity (2008). Understanding cross-country differences in exporter premia: Comparable evidence for 14 countries. *Review of World Economics (Weltwirtschaftliches Archiv) 144*(4), 596–635.

Ito, K. and S. Lechevalier (2010). Why some firms persistently out-perform others: investigating the interactions between innovation and exporting strategies. *Industrial and Corporate Change 19*(6), 1997–2039.

Lacerda, G., P. L. Spirtes, J. Ramsey, and P. O. Hoyer (2012). Discovering cyclic causal models by independent components analysis. *arXiv preprint arXiv:1206.3273*.

Lanne, M., M. Meitz, and P. Saikkonen (2017). Identification and estimation of non-gaussian structural vector autoregressions. *Journal of Econometrics 196*(2), 288–304.

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

López, R. A. (2009, Aug). Do firms increase productivity in order to become exporters? *Oxford Bulletin of Economics and Statistics 71*(11), 621–642.

Manjón, M., J. Máñez, M. E. Rochina-Barrachina, and J. Sanchis-Llopis (2013). Reconsidering learning by exporting. *Review of World Economics 149*(1), 5–22.

Martins, P. S. and Y. Yang (2009). The impact of exporting on firm productivity: a meta-analysis of the learning-by-exporting hypothesis. *Review of World Economics 145*(3), 431–445.

Moneta, A., D. Entner, P. O. Hoyer, and A. Coad (2013). Causal inference by independent component analysis: Theory and applications. *Oxford Bulletin of Economics and Statistics 75*(5), 705–730.

Mukim, M. (2011). Does exporting increase productivity? evidence from india. Working paper, London School of Economics.

Park, A., D. Yang, X. Shi, and Y. Jiang (2010). Exporting and firm performance: Chinese exporters and the Asian financial crisis. *The Review of Economics and Statistics 92*(4), 822–842.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *The Review of Economic Studies 69*(1), 245–276.

Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

Peters, J., D. Janzing, and B. Schölkopf (2017). *Elements of causal inference: foundations and learning algorithms*. MIT press.

Petrin, A., B. P. Poi, and J. Levinsohn (2004, June). Production function estimation in stata using inputs to control for unobservables. *Stata Journal 4*(2), 113–123.

Segarra-Blasco, A., M. Teruel, and S. Cattaruzzo (2020). Innovation, productivity and learning induced by export across European manufacturing firms. *Economics of Innovation and New Technology*.

Serrano, J. and R. Myro (2019, dec). From domestic to exporter, what happens? Evidence for Spanish manufacturing firms. *Structural Change and Economic Dynamics 51*, 380–392.

Shimizu, S., P. O. Hoyer, A. Hyvärinen, and A. Kerminen (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research 7*, 2003–2030.

Spirtes, P., C. N. Glymour, and R. Scheines (2000). *Causation, prediction, and search*. MIT press.

Tsou, M.-W., J.-T. Liu, J. K. Hammitt, and K.-H. Wang (2008). Exporting and productivity growth: Evidence from the taiwan electronics plants. *Scottish Journal of Political Economy 55*(2), 190–209.

Wagner, J. (2007). Exports and productivity: A survey of the evidence from firm-level data. *The World Economy 30*(1), 60–82.

Wagner, J. (2012). International trade and firm performance: a survey of empirical studies since 2006. *Review of World Economics 148*(2), 235–267.

Wang, F., Z. Xu, and X. Dai (2021). Is learning by exporting technology specific? Evidence from Chinese firms. *Economics of Innovation and New Technology*.