

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

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

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

December 20, 2019

Abstract

This paper introduces a little known category of estimators – Linear Non-Gaussian vector autoregression models that are acyclic or cyclic – imported from the machine learning literature, to revisit a well-known debate. Does exporting increase firm productivity? Or is it only more productive firms that remain in the export market? We focus on a relatively well-studied country (Chile) and on already-exporting firms (i.e. the intensive margin of exporting). We explicitly look at the co-evolution of productivity and growth, and attempt to ascertain both contemporaneous and lagged causal relationships. Our findings suggest that exporting does not have any causal influence on the other variables. Instead, export seems to be determined by other dimensions of firm growth. With respect to learning by exporting (LBE), we find no evidence that export growth causes productivity growth within the period and very little evidence that exporting growth has a causal effect on subsequent TFP 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 VARLiNG algorithm with us. The usual caveat applies.

[†]SPRU, University of Sussex, t.ciarli@sussex.ac.uk

[‡]CENTRUM Católica Graduate Business School (CCGBS), Lima, Perú; Pontificia Universidad Católica del Perú (PUCP), Lima, Perú, acoad@pucp.edu.pe

[§]Institute of Economics and EMbeDS, Scuola Superiore Sant’Anna, a.moneta@santannapisa.it

1 Introduction

Effective economic policy requires a causal understanding of the relationships between variables. However, the toolkit of causal estimators available to econometricians is limited, and not all estimators are appropriate for certain research contexts. For example, instrumental variables analysis can be problematic if there are no suitable instruments that satisfy the exogeneity requirements. Regression discontinuity design cannot be applied if the data are not arranged in a certain way, e.g. if the quality of applications, whose causal effects are object of study, is not ranked. Randomized controlled trials can be prohibitively difficult to implement in economic contexts with large numbers of interconnected actors and over long time periods (but see the pioneering efforts by Atkin et al., 2017).

This paper introduces a little known category of estimators – Linear Non-Gaussian vector autoregression models that are acyclic or cyclic – imported from the machine learning literature, to complement the existing literature with new results on the causal relationship between exporting and productivity.

Our identification strategy applies Independent Components Analysis (ICA) to generate a set of SVAR residuals, generated from the reduced-form VAR residuals, that are maximally statistically independent. The statistical independence of these SVAR residuals allows us to arrange the variables in a causal ordering, exploiting the simple consideration that independent exogenous shocks can simultaneously affect different sets of variables precisely because of the presence or absence of causal links among the variables. The LiNGAM estimator (Linear Non-Gaussian Acyclic Model) assumes acyclicity - i.e. that there are no causal feedback loops that take place within the same 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). We also present the LiNG estimator (Lacerda et al., 2008) which – to our knowledge – has not yet been introduced to economics. The LiNG estimator relaxes LiNGAM’s acyclicity condition, and allows for simultaneous causal feedback loops between variables even within a single year. LiNG therefore relaxes the assumption of an acyclic causal structure, although it introduces a new set of computational complexities, that we discuss.

Our new estimators provide new insights to a vibrant debate, because we can complement previous results with new perspectives obtained from a new method of identification. In the context of increasing concerns about the replicability and rigour of causal estimates, our results can help triangulate between existing results.

Our empirical application focuses on the causal relationship between exporting and productivity. Do firms take advantage from trade? Or is exporting a game that only the best firms can play, and once they export there is little else they can

learn? 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 very 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 (see for example surveys by Greenaway and Kneller, 2007; International Study Group on Exports and Productivity, 2008; Wagner, 2007). 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. What comes first, the chicken (growth of exports) or the egg (productivity growth)?

Two crucial aspects distinguish our work from previous studies on the effect of export on productivity. First, we focus on the small number of firms (Bernard et al., 2007) that have already accessed the export market, and we analyse changes in export and productivity, rather than levels. In the literature, most studies have looked at the extensive margin of exporting (entering the foreign markets), whereas only few focus on the intensive margin (increasing exports). Increases in exports may have implications for firm productivity (in the short or long run) through a number of mechanisms: 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.

Focussing on export growth rather than on entry, we can ignore the problem of self-selection into exporting,¹ and focus on the identification of whether (i) firms that export more improve their performance (productivity) or (ii) the other way round: firms must improve their productivity in order to be able to export more.

Second, in order to identify the causal order of the mechanisms between exporting and productivity growth, we use Independent Component Analysis to find independent residuals from SVAR estimations. The proposed method uncovers the causal structure within the same period, providing a unique understanding of the short term relation between export and productivity growth.

¹As summarised by Park et al. (2010, p. 822): “conceptually, the fundamental problem is that nonexporters are an inappropriate counterfactual for exporters.”

The proposed method can be useful for situations in which matching estimators and instrumental variables are problematic. Indeed, in many empirical cases, finding settings that resemble the randomized control trial setting may be difficult. A crucial strength of our method is that it avoids both theoretical restrictions and it contains general assumptions of the possible causal structure at a minimum. We are able to identify whether firms simultaneously decide to invest to increase productivity and to export (Aw et al., 2011), or whether one decision precedes (and induces) the other, by investigating the causal effect in the same time period.

In order to better identify the methodological and empirical contribution, we focus on Chile, a small country, with a small domestic market, an open economy, a well-tested firm survey, and large evidence on the relation between exporting and productivity (e.g. Alvarez and Crespi, 2007; Alvarez and López, 2005, 2008; López, 2009; Pavcnik, 2002).

Our main results suggest that in the short run export growth follows growth in productivity, and not the other way round. There is no evidence of export growth causing productivity growth. The dynamic seems to be driven by employment, which has a contemporaneous negative effect on productivity, which in turn affects export. This means that, once Chilean firms enter the foreign market, changes in exporting have no detectable effect on productivity in the short run. On the contrary, Chilean firms need to improve their performance in order to increase exports. We should also note that Chilean firms, on average, choose between the domestic and the foreign market: when export increases, domestic sales decrease. This seem to confirm that for a given level of output there is an explicit choice to increase export, following other changes in the firm.

Our results are relevant for industrial and trade policy. They seem to suggest that exporting firms in a small open economy, which has been open for a few decades, are not managing to learn a lot in the short period, at least in the short run. This might be because most firm's sectors are not technology intensive, and most competition is based on price. Small changes in labour costs can boost productivity and increase exporting. Whether this dynamics brings about a longer term effect on technological learning and upgrading is left for future research.

Our paper contributes to the literature attempting to identify the effect of firm specific changes in exporting on changes in productivity. 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 the authors instrument export 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. In contrast, they find positive effects of export growth on subsequent productivity growth. More generally, we also contribute to the literature that has focussed on export intensity, rather than entry in the export markets (e.g. Castellani, 2002; Fernandes and Isgut,

2005; Antolín et al., 2013).

In the rest of the paper we first discuss the heterogeneous evidence on exporting and productivity (Section 2). Next, Section 3 discusses the core of the paper: 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 future research in this and similar topics.

2 The relation between exporting and firm productivity

A large amount of research has attempted to identify the direction of causality between exports and growth at the firm level. Wagner (2007) conducts a systematic literature review and finds that: (i) exporting firms are always more productive than non-exporters; (ii) exporters very often are more productive even before entering the export market; (iii) results on learning-by-exporting (LBE) are very 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. Wagner (2012) updates his previous review, suggesting that the relation between exporting and productivity is influenced by export destination: self-selection is stronger when exporting to high income countries, but results on LBE are still mixed, firms are more likely to increase productivity by exporting when they export to high income countries. Similar inconclusive results on LBE are found when analysing the service sector.

For instance, Girma et al. (2004), who introduce matching techniques to this strand of literature, find a significant positive effect of export on productivity for UK manufacturing firms. Tsou et al. (2008), using a census of Taiwanese firms repeated for three different periods, 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. De Loecker (2013) proposes a different method for computing productivity, which includes export in firm decision (and therefore as a determinant of productivity) and finds significant LBE for Slovenian firms.

A number of studies fail to find a significant positive effect of exporting on productivity. Using Indian data, Mukim (2011) finds that there is no sustained effect of learning from exporting. Eliasson et al. (2012) find similar results when

focussing on small and medium firms: evidence of learning to export, but no significant effect of exporting on learning. The International Study Group on Exports and Productivity (2008) use panel data from 14 different countries, finding no evidence of LBE. Arnold and Hussinger (2005) use matching techniques to investigate the LBE on German firms, but also find no significant effect. Damijan and Kostevc (2006) find similar results on Slovenian firms. Tsou et al. (2008) find mixed evidence for the LBE hypothesis in the case of Taiwanese firms, while evidence for self-selection is much stronger.

Heterogeneous effects seem to explain part of the difference in findings. We have already mentioned the ‘distance-to-frontier’ effect of the country of destination: firms may learn more from industrialised countries. Martins and Yang’s (2009) meta-analysis on LBE finds that, relative to high income countries, firms in developing countries enjoy a stronger impact of exporting on productivity. Younger firms may benefit more from exporting (Alvarez and López, 2005; Girma et al., 2004; Fernandes and Isgut, 2005), as well as larger firms (Damijan et al., 2010). Harris and Cher Li (2011) observe heterogeneous effects across industry sectors as well as within sectors. Firms with a higher export intensity (with respect to domestic sales) also tend to profit more from exporting (Girma et al., 2004).

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.

Time is crucial for productivity increases via technological learning (Fernandes and Isgut, 2005). Although younger firms benefit more, the increase in productivity may last for a few years (Blalock and Gertler, 2004; De Loecker, 2007; Hosono et al., 2015), which is consistent with the notion of diminishing returns to export experience (Fernandes and Isgut, 2015). Trade may also not have an immediate effect on productivity: Crespi et al. (2008) use learning measures to estimate the effect of export on learning, which may affect productivity only in a second stage, and find evidence of LBE.

Timing is relevant also because specific investments may be a necessary condition to benefit from exporting. Firms may not rely on LBE solely, as they would be selected out of the international market early on. Decisions to innovate and to export may thus be complementary (Ito and Lechevalier, 2010), as suggested also by Aw et al. (2011).

For instance, Dai and Yu (2013) use matching estimators to study the effect of pre-export R&D on LBE in Chinese firms and find that firms with higher R&D in-

vestment enjoy significant improvements in productivity after exporting, contrary to firms that do invest before. Aw and Song (2013) show that for Korean firms the level of productivity plays a crucial role in determining their investment in future R&D and in entering export markets. They also find evidence that the decision to improve productivity and to export are simultaneous and have additional effects on productivity.

Only a handful of studies look at changes in exports. The already cited 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 with an incentive for firms to innovate more and with more access to finance for investment. Although they focus on levels and not on changes, Fernandes and Isgut (2005) are relatively close to our work. They focus on the level of export (“export experience”), rather than on export participation, finding a positive effect of LBE for Colombian firms exporting to high income markets.

Summarising, first, it is important to distinguish between short and long term LBE. Some mechanisms – such as market size, scale, vulnerability, use of existing capacity – may have an immediate effect on firm productivity, whereas other mechanisms – such as learning from buyers, and from markets – may take longer to show in growth of productivity. In this case we should observe some effect of exporting on productivity in the long term, and some effect in the medium and long term.

Second, some firms may decide to invest in order to improve productivity before they increase their sales on a more competitive foreign market. In this case, we should observe an increase in productivity that precedes growth in export (or export intensity).

Third, considering that most existing estimates (including ours) use yearly data, we should acknowledge that within the same year a firm takes many decisions. They may decide on measures to increase productivity and export simultaneously, as complementary activities. This is where the existing empirical evidence is silent, and where we make our main contribution.

3 Econometric method

3.1 VAR and SVAR models

Consider the following structural vector autoregressive (SVAR) model that features a vector of two variables, exports growth (EXP_{it}) and productivity growth ($PROD_{it}$), for firm i at time t .² For clarity, we omit the other variables and include only one time lag.

$$\begin{aligned} EXP_{it} &= b_{12}PROD_{it} + \gamma_{11}EXP_{i,t-1} + \gamma_{12}PROD_{i,t-1} + e_{it}^{exp}, \\ PROD_{it} &= b_{21}EXP_{it} + \gamma_{13}EXP_{i,t-1} + \gamma_{14}PROD_{i,t-1} + e_{it}^{prod} \end{aligned} \quad (1)$$

Denoting by $y_{it} = (EXP_{it}, PROD_{it})'$ and $\varepsilon_{it} = (e_{it}^{exp}, e_{it}^{prod})'$, and using matrix notation, the same model can be rewritten as:

$$y_{it} = By_{it} + \Gamma_1 y_{i,t-1} + \varepsilon_{it} \quad (2)$$

As long as we do not attempt to estimate the matrix of instantaneous effects B , then we estimate a reduced-form vector autoregressive (VAR) model by rearranging (2):

$$y_{it} = (I - B)^{-1}\Gamma_1 y_{i,t-1} + (I - B)^{-1}\varepsilon_{it} \quad (3)$$

Or, equivalently,

$$y_{it} = A_1 y_{i,t-1} + u_{it}, \quad (4)$$

in which $A_1 = (I - B)^{-1}\Gamma_1$ and $u_{it} = (I - B)^{-1}\varepsilon_{it}$. Notice that the reduced-form VAR model in (4) does not allow estimating the matrix of instantaneous causal effects B , and – importantly – nor does it allow to properly estimate the matrix of lagged causal effects Γ_1 . In order to estimate this latter, we would need to isolate it from the term $(I - B)^{-1}$. However, by identifying the matrix of instantaneous causal effects B , we can also properly estimate the matrix of lagged causal effects Γ_1 .

By referring to equation (1), the matrix of instantaneous effects B can be written as follows:

$$B = \begin{pmatrix} 0 & b_{12} \\ b_{21} & 0 \end{pmatrix} \quad (5)$$

In equation (1), the contemporaneous causal effect of $PROD_{it}$ on EXP_{it} is represented by b_{12} , while the contemporaneous causal effect of EXP_{it} on $PROD_{it}$

²A similar didactic approach is in Coad and Grassano (2019).

is represented by b_{21} . Notice, however, that equation (1) cannot be consistently estimated through linear regression techniques, due to the endogeneity problem. Assuming that the model is acyclic (i.e. that there are no feedback loops), then we can impose that B is a lower-triangular matrix (or that it can be row-permuted to become lower-triangular), such that either b_{12} or b_{21} must be equal to zero. But knowing the presence of an acyclic contemporaneous structure, without knowing which coefficient between b_{12} and b_{21} is zero, is not sufficient to consistently estimate equation (1).

The standard textbook definition of endogeneity asserts that, in a regression equation of the type $y_{it} = ax_{it} + e_{it}$, the explanatory variable x_{it} is endogenous if it is correlated with the error term e_{it} (see for example Wooldridge, 2010, p. 50). If, however, x_{it} is uncorrelated with e_{it} , then x_{it} is seen to be exogenous, and therefore the causal channel goes from x to y .

Further developments in statistical theory regarding the concept of causality has put forward that, in a non-Gaussian setting (or, alternatively, in a non-linear setting), x_{it} must not only be uncorrelated with e_{it} , but fully statistically independent of e_{it} in a structural model representing a causal relationship from x to y . This is because a zero correlation is a sufficient condition for statistical independence only in a linear Gaussian setting and is a flawed indicator of statistical independence in more general contexts³ (Mooij et al., 2009; Peters et al., 2017).

A key problem affecting causal inference in social science, though, is that “everything correlates to some extent with everything else” – which has been dubbed the ‘crud factor’ by Meehl (1990, p. 204). Our approach to unravel the directions of causal influence, combines independent component analysis (ICA) with SVAR analysis. The goal is to recover both the SVAR residuals e_{it} that are statistically independent of the explanatory variables and the coefficients of the matrix that ‘mixes’ them to form the reduced-form VAR residuals u_t . ICA is a probabilistic method for finding a linear transformation of the data that are maximally-independent and non-Gaussian⁴ (Hyvärinen, 2013).

Thus, the idea here is to apply ICA to the reduced form residuals u_{it} , which can be estimated from equation (4). ICA will deliver a linear combinations of the elements of u_{it} that are maximally independent and non-Gaussian. Under some

³In general, zero correlation is only a necessary (but not sufficient) condition for statistical independence.

⁴The microphone analogy can be helpful (Stone, 2004, p.204). Consider the case of two microphones, one which records voice A, and the other which records A and B. ICA would lead to identify two independent components: the signal of voice A; as well as the independent component corresponding to voice B which is a function of the recorded message on the second microphone, adjusted to remove the signals coming from voice A. In the case of the first microphone, the recorded signal corresponds to one of the two extracted independent components. Note that the assumption of acyclicity rules out that both microphones record both voices (Coad and Grassano, 2019).

mild identifying assumptions (see section 3.2), from these mixtures we will be able to recover the terms e_{it}^{exp} and e_{it}^{prod} of equation (1) and the matrix $\Gamma_0 = (I - B)$ such that $\Gamma_0 u_{it} = \varepsilon_t$. In this manner, one is able to recover all the structural coefficients of the SVAR model in equation (2). This framework can be easily extended to the case in which the number of variables is greater than two and the number of lags is greater than one. In our empirical analysis, y_{it} will comprise four variables, and we will estimate both one-lag and two-lag models.

3.2 Identification strategy

Let us consider the general framework in which y_{it} comprises k variables. Since, in our application, firms (indexed by i) are pooled together under the assumption that different firms undergo similar structural patterns in their growth process, we omit henceforth the subscript i . The ‘unmixing’ matrix Γ_0 , which relates the k -dimensional vector of the structural residuals (shocks) ε_t to the k -dimensional vector of reduced-form residuals (errors) u_t , is an invertible $k \times k$ matrix such that

$$u_t = \Gamma_0^{-1} \varepsilon_t, \quad (6)$$

Γ_0^{-1} is called the ‘mixing’ matrix.

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

$$y_t = \mu + A_1 y_{t-1} + \dots + A_p y_{t-p} + u_t \quad (7)$$

(the model is analogous to equation 4 with the addition of a constant term vector μ), and then we search for Γ_0 such that

$$\Gamma_0 y_t = \Gamma_0 \mu + \Gamma_0 A_1 y_{t-1} + \dots + \Gamma_0 A_p y_{t-p} + \Gamma_0 u_t \quad (8)$$

or, in more compact form

$$\Gamma_0 y_t = \xi + \Gamma_1 y_{t-1} + \dots + \Gamma_p y_{t-p} + \varepsilon_t, \quad (9)$$

where $\Gamma_i = \Gamma_0 A_i$ for $i = 1, \dots, p$. Assuming that the k elements of ε_t are mutually independent and (at least $k - 1$ of them) non-Gaussian, i.e. non-normally distributed, the method is able to identify Γ_0 and, consequently (having estimated all coefficient matrices of equation 7), all the coefficient matrices of equation (9). 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).

No matter 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 matrix⁵ 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 $\hat{\varepsilon}_{1t}, \dots, \hat{\varepsilon}_{kt}$ (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 main diagonal, this ordering must be the correct one.⁶ 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

⁵A 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).

⁶In 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 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 $u_t = \Gamma_0^{-1}C'\varepsilon_t$.

by normalising the rows of $\tilde{\Gamma}_0$ (the correctly row-permuted version of Γ_{ICA}), so that all diagonal elements equal unity. Let $\hat{\Gamma}_0$ 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 estimated 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 can be done through a bootstrap procedure. 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 of such a matrix: we index each of them with $j = 1, \dots, m$. Thus, the algorithm will output m possible causal structures. However, some of them can be excluded a priori by excluding unstable contemporaneous causal structures, i.e. $\tilde{\Gamma}_{0,j}$ such that $\tilde{\Gamma}_{0,j}^{-1}$ has eigenvalues whose modulus is greater than one. 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.

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 (7), obtaining estimates \hat{A}_i of the matrices A_i for $i = 1, \dots, p$. 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 $\hat{\Gamma}_i = (I - \hat{B})\hat{A}_i$, for $i = 1, \dots, p$.

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 can be done using a bootstrap procedure.
 4. Find all admissible row-permuted matrices $\tilde{\Gamma}_{0,1}, \dots, \tilde{\Gamma}_{0,m}$ of Γ_{ICA} such that each $\tilde{\Gamma}_{0,h}$ has zeroless diagonal for $h = 1, \dots, m$.
 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 $h = 1, \dots, m$.
 6. Let $\tilde{B}_j = I - \hat{\Gamma}_{0,j}$, for each $j = 1, \dots, m$.
 7. Calculate estimates of $\hat{\Gamma}_{i,j}$ for lagged effects using $\hat{\Gamma}_{i,j} = (I - \hat{B}_h)\hat{A}_i$, for $i = 1, \dots, p$, for $h = 1, \dots, m$.
-

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 2007.⁷ 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).

After some preliminary data cleaning,⁸ 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 assumptions that are explained in what follows.

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*

⁷Data 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.

⁸We 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.

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: $Va = output - Material - (Privap + Matvap) + (Priva_f + Matva_f)$.

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 completely 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 noting 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

^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. 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* ($\rho=-0.0132$, p -value=0.4029)

	<i>gr_domsales</i>	<i>gr_empl</i>	<i>gr_exp</i>	<i>gr_tfp</i>
<i>gr_domsales</i>	1	<i>0.1343</i>	<i>-0.0132</i>	<i>0.4283</i>
<i>gr_empl</i>	0.0924	1	<i>0.1436</i>	<i>-0.1881</i>
<i>gr_exp</i>	-0.0717	0.0915	1	<i>0.1605</i>
<i>gr_tfp</i>	0.3327	-0.2599	0.0836	1

5 Results

5.1 Correlations

In this section, we present the results using our identification method. To emphasize the difference between causal and descriptive analysis, 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 negative correlation. The problem with a correlation analysis, of course, is that it cuts off many possible channels of interactions among variables, especially between lagged values.

5.2 VAR analysis

This particular aspect is, on the contrary, accounted for in the VAR analysis. Table 3 presents the reduced-form VAR results, which describe the intertemporal associations between the variables. 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. Table 3 presents both the model estimated with one lag (upper part) and the model estimated with two lags (bottom part).⁹

As noted in the previous sections, the reduced-form model delivers coefficients that, although being consistent (under stationarity of the series), and implying meaningful relations of statistical dependence, cannot be interpreted in causal

⁹We also repeated our analysis using OLS regression (see appendix B).

terms, given the fact that possible contemporaneous influences among the variables are omitted. Of interest to our analysis is that the autocorrelation coefficients (i.e. the entries of the main diagonal of each matrix) are always negative in both models, with the exception of one coefficient in the second matrix of the two-lags model. This suggests that the hypothesis of increasing returns and sustained growth is not supported by these data. We also notice that lagged growth of exports is positively associated with subsequent growth of TFP in the one-lag model (coefficient = 0.0098 with standard error 0.0047) and two lags model (coefficient = 0.0093 with standard error 0.0045).

5.3 SVAR analysis

Before applying our identification method, which allows us to recover the coefficients of the SVAR model, we investigate whether the assumption of non-Gaussian shocks is plausible. 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). In our data, the evidence from the quantile-quantile plots suggests that the VAR reduced-form residuals are far from Gaussian (see Figure 1), providing support for our econometric strategy. 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 the null hypothesis at the 0.01 level of significance.

We present now the results from our identification procedure. Table 4 presents the mixing matrix P for both the 1-lag and the 2-lag model. The matrix P is the output of the FastICA algorithm (Hyvärinen et al., 2001) applied to the reduced-form residuals. Producing the matrix P , after having estimated the reduced form VAR with its residuals, is one of the steps which is common to both the algorithms presented in section 3.2, namely step 2 of VAR-LiNGAM and VAR-LiNG. The matrix P is 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. 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 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 .¹⁰ 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,

¹⁰In 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 algorithm.

Table 3: Reduced-form VAR estimated using median (LAD) regression. Standard errors in parentheses are obtained through bootstrap (1000 iterations). Upper part of the table (first four rows): 1-lag VAR (3648 observations). Bottom part of the table (last four rows): 2-lag VAR (3123 observations).

	first lag				second lag			
	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	-0.2219 (0.0399)	0.0666 (0.029)	0.0094 (0.0095)	0.0702 (0.0396)				
gr_empl	0.0073 (0.0063)	-0.0229 (0.0207)	0.0087 (0.0039)	0.0162 (0.0108)				
gr_exp	-0.012 (0.0185)	-0.0524 (0.0402)	-0.1364 (0.0334)	0.0055 (0.0464)				
gr_tfp	-0.002 (0.0094)	-0.0517 (0.0208)	0.0098 (0.0047)	-0.2743 (0.0254)				
gr_domsales	-0.2178 (0.0467)	0.0315 (0.0312)	0.0201 (0.0093)	0.0672 (0.0428)	-0.0862 (0.0494)	0.0408 (0.0354)	0.0236 (0.0094)	0.0629 (0.0352)
gr_empl	0.0145 (0.0094)	-0.0196 (0.027)	0.0065 (0.0071)	0.0321 (0.0149)	0.0073 (0.0087)	0.0122 (0.0234)	0.0045 (0.0051)	0.0498 (0.0179)
gr_exp	0.0276 (0.0216)	-0.0654 (0.0507)	-0.1399 (0.0367)	-0.025 (0.058)	0.0121 (0.0166)	0.1124 (0.054)	-0.0929 (0.0294)	0.0504 (0.0528)
gr_tfp	-0.0052 (0.0124)	-0.0311 (0.0268)	0.0093 (0.0045)	-0.2689 (0.0391)	-0.025 (0.0113)	0.0094 (0.0254)	-0.001 (0.0061)	-0.0959 (0.0254)

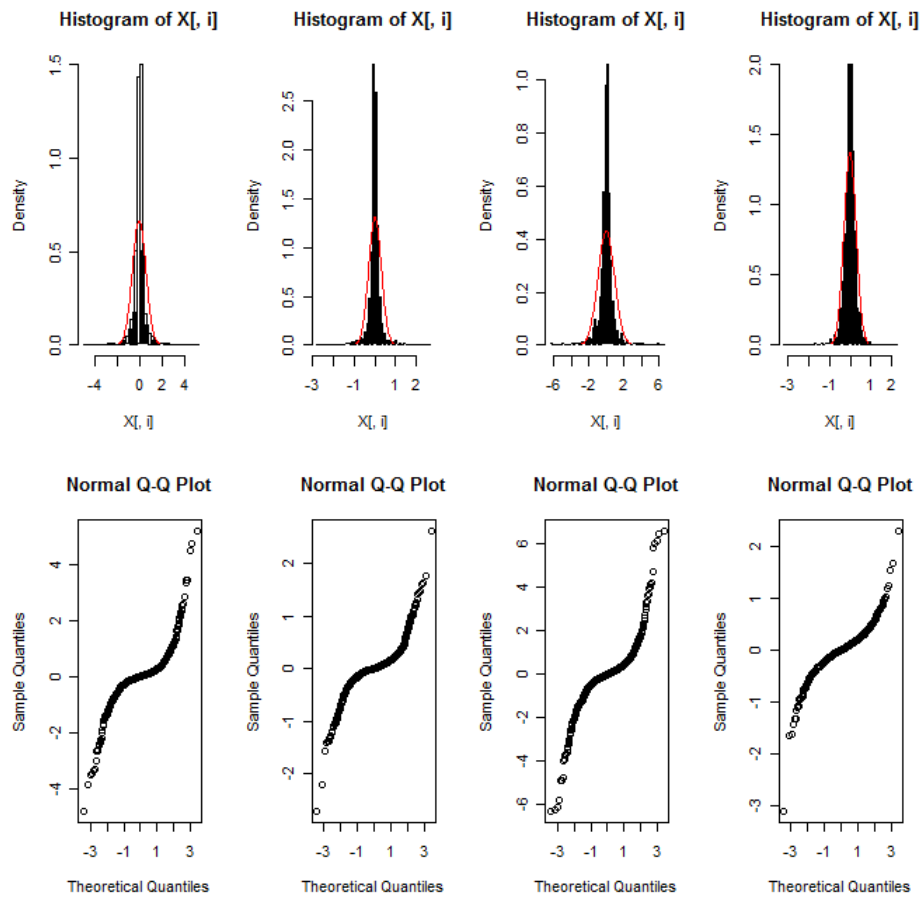


Figure 1: Quantile-quantile plots of the distributions of the four reduced-form residuals, for the 2-lag model.

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 4), 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. In both cases, this strongly suggests that the exporting growth shock does not transmit (within the one year period) to productivity growth, or, in other words, there is no contemporaneous causal relationship from exporting to productivity growth.

Table 4: Mixing matrix P , 1-lag and 2-lag model.

	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

The algorithms we use (VAR-LiNGAM and LiNG) aim at inferring 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. This is something that VAR-LiNGAM does not check empirically, rather it is simply assumed a priori and this constitutes an objective weakness of this method. In order to overcome this issue and improve the empirical reliability of our causal inference, we first identify the model through VAR-LiNGAM, and then use a bootstrap procedure in order 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 actually 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 exporting growth, emerges as robustly identified.

Table 5 shows the coefficients of structural VAR matrices (see equation 9) 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 .¹¹ Figure 2 shows the resulting DAG for the 1-lag model.

Both the 1-lag and 2-lag models in Table 5 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 sheer scale effects – employment growth leads to subsequent increases in outputs. 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 main result is that employment growth has a negative effect on contemporaneous growth of TFP, presumably because productive efficiency is attained when fewer inputs (i.e. employees) produce a given output. Downsizing firms might be better able to improve their productivity 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 exporting. 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.

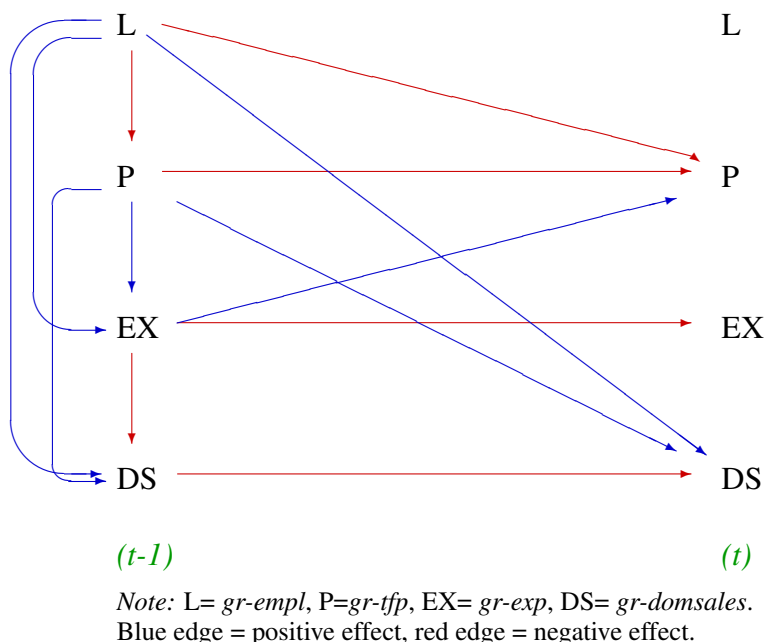
Growth of exporting has a negative impact on growth of domestic sales. This no doubt reflects the tension between domestic vs. exporting sales strategies, that

¹¹Since 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 5 coefficients significantly different from zero are represented in bold.

Table 5: VAR - LiNGAM estimates, 1- and 2-lag models. Standard errors are obtained through bootstrap (1000 iterations).

	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						
gr_empl	0	0.0619	0.0127	0.0773	0.0384	0.033	0.008	0.0472						
gr_exp	0	0	0	0	0.0073	-0.0229	0.0087	0.0162						
gr_tfp	0	0.4369	0	0.4023	0.0068	0.0234	0.0041	0.0122						
	0	0.0889	0	0.0707	-0.0144	-0.0216	-0.1442	0.1088						
	0	-0.2712	0	0	0.0157	0.0344	0.0392	0.0442						
	0	0.0288	0	0	0	-0.0579	0.0121	-0.2699						
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		
gr_empl	0	0.0839	0.0145	0.102	0.0437	0.038	0.0085	0.0673	0.0542	0.0341	0.01	0.0371		
gr_exp	0	0	0	0	0.0145	-0.0196	0.0065	0.0321	0.0073	0.0122	0.0045	0.0498		
gr_tfp	0	0.4345	0	0.4743	0.0085	0.0309	0.0073	0.0174	0.0087	0.0261	0.0064	0.0184		
	0	0.082	0	0.0812	0.0237	-0.0421	-0.1472	0.0886	0.0208	0.1027	-0.0944	0.0743		
	0	-0.2466	0	0	0.021	0.0477	0.0272	0.067	0.0205	0.0491	0.03	0.051		
	0	0.0307	0	0	-0.0016	-0.0359	0.0109	-0.261	-0.0232	0.0124	0.0001	-0.0836		
	0				0.0119	0.0211	0.0051	0.0375	0.0108	0.0274	0.006	0.0238		

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




was already visible in the negative correlations between these variables in Table 2. However, it is interesting to observe that exporting 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.

With regards to the causal link between TFP and exporting, our results suggest that it is TFP that causes exporting, rather than vice versa. Our VAR-LiNGAM estimates therefore provide an interesting perspective on the exporting-productivity debate. Note however that the first lag of exporting growth has a positive impact on subsequent TFP growth.

5.4 SVAR 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

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 iteration.




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

reversed almost half of the time.

5.5 VAR-LiNG

Table 6 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) as regards the model with one lag. We do not report here the results of the two-lag model, which are qualitatively similar, for reasons of space. Figure 4 depicts the contemporaneous causal structure, while Figure 5 shows the lagged causal influence. The estimated causal structure presents now, within the period of estimation, a bi-directional link between growth of domestic sales and growth of exporting sales. The possibility of such bi-directionality was already suggested by the robustness results shown in Figure 3. Feedback loops emerge also between growth of domestic sales and growth of productivity and between growth of domestic sales and growth of employment in the contemporaneous causal structure. Nevertheless, the main findings about the causal nexus between productivity and exporting growth which resulted from the application of the first algorithm are confirmed: there is no causal influence from exporting to productivity (growth) in the contemporaneous causal structure. 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

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.

are also very similar: 0.0773 vs. 0.0707). However, the empirical evidence resulting from the second algorithm is even more consistent with the theoretical hypothesis which denies the existence of a “learning-by-exporting” phenomenon: while the results from algorithm 1 showed some (weak) influence from lagged exporting (growth) to current productivity (growth), the results from algorithm 2 showed no such causal link (see Figure 5 for the lagged effects).


Table 6: VAR - LiNG estimates, 1-lag model.

	<i>gr_domsales</i>	<i>gr_empl</i>	<i>gr_exp</i>	<i>gr_tfp</i>	<i>l_gr_domsales</i>	<i>l_gr_empl</i>	<i>l_gr_exp</i>	<i>l_gr_tfp</i>
<i>gr_domsales</i>	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
<i>gr_empl</i>	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
<i>gr_exp</i>	-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
<i>gr_tfp</i>	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

6 Discussion

In this paper we revisit a well-known debate, which has grown exponentially in the last two decades. Does exporting activity increase firm performance, in particular productivity, as it is expected from some case study evidence? Or is it only more productive firms that enter and remain in the export market? We choose a rather different strategy from previous papers, while exploring a relatively well-

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.

studied country. First, we focus on exporting firms. Second, we do not compare them with non-exporting firms, but with other firms whose exports grow more or less (i.e. become more or less competitive in the international market). Third, we explicitly look at the co-evolution of the the two variables, productivity and growth, including the causal relation within the period and with up to two lags.

Our most interesting finding is in relation with the extensively investigated LBE hypothesis. Applying VAR-LiNGAM and VAR-LiNG, a class of SVAR models that estimates causal networks, it seems that exporting does not have any direct and instantaneous causal impact on firm performance. Instead, it seems to be determined by other dimensions of firm growth. Of interest to the inconclusive literature on LBE is that we find no evidence that it causes productivity growth within the period. However, a result from VAR-LiNGAM is that the first lag of exporting growth does have a small causal effect on subsequent TFP growth. But when we apply VAR-LiNG, the algorithm which allows the possibility of feedback loops in the contemporaneous structure, this lagged causal effect vanishes. Instead, we observe that TFP growth has a direct causal effect on exporting growth within-the-period, which is robust under the application of the different algorithms.

Our results are estimates of causal effects (rather than mere associations) and therefore have interesting implications for policy. In particular, it appears that

firms should focus on improving their productivity before attempting to increase their exports, because it is productivity growth that drives growth of exports. For example, firms should first improve their productivity through e.g. redesigning their production routines, and upgrading their capital and IT systems, alongside appropriate organizational innovations (e.g. Cruz et al., 2018), and as a result, they will be in a better position to experience growth of exports. There is negligible influence of exporting on TFP growth, however – only with a lag does exporting feed back into TFP growth, and moreover this effect is relatively small.

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.

Second, we focus on exporting undertaken by firms that have already taken the binary decision of whether to export. There may be differences in the exporting-productivity relationship at the time when a firm first decides to transform from a non-exporter to an exporter.

To conclude, this paper has 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. Our application has shed new light on the LBE controversy by showing that the causal direction runs from productivity growth to exporting in our panel of exporters. Future work could apply the family of techniques developed here to a broad range of contexts to get valuable new evidence for academics, practitioners and policymakers.

Appendix A: Impulse response analysis

In the standard SVAR analysis, dynamic causal effects are studied through impulse response functions (see e.g. Kilian and Lütkepohl, 2017). In this study, we have focused on contemporaneous and short-run causal effects, given also the relatively short time window of our data set. We report nevertheless results from impulse response analysis. Ignoring the constant term μ (which is irrelevant for the impulse response analysis) a stable (stationary) reduced-form VAR model (equation 7) can be written as

$$y_t = \sum_{j=0}^{\infty} \Phi_j u_{t-j},$$

where $\Phi_0 = I$ and $\Phi_j = \sum_{i=1}^j \Phi_{j-i} A_i$, for $j = 1, 2, \dots$. Since $\Gamma_0 u_t = \varepsilon_t$ we can also write:

$$y_t = \sum_{j=0}^{\infty} \Phi_j \Gamma_0^{-1} \Gamma_0 u_{t-j} = \sum_{j=0}^{\infty} \Phi_j \Gamma_0^{-1} \varepsilon_{t-j} = \sum_{j=0}^{\infty} \Psi_j \varepsilon_{t-j}.$$

The matrix $\Psi_j = \Phi_j \Gamma_0^{-1}$ contains the marginal effects of the structural shocks ε_t on the variables over time. In particular the (m, n) element of Ψ_j contains the reaction of the variable $y_{m,t}$ to a shock $\varepsilon_{n,t}$, after j periods. Since we have 4 variables, we obtain 16 impulse response functions. Figure 6 reports impulse response functions, while Figure 7 show cumulated impulse response functions. We can see that in both cases the productivity shock ε_P has a greater impact on exporting (growth and levels) than the other way around ($\varepsilon_{EX} \rightarrow P$), which is in tune with the other results so far.

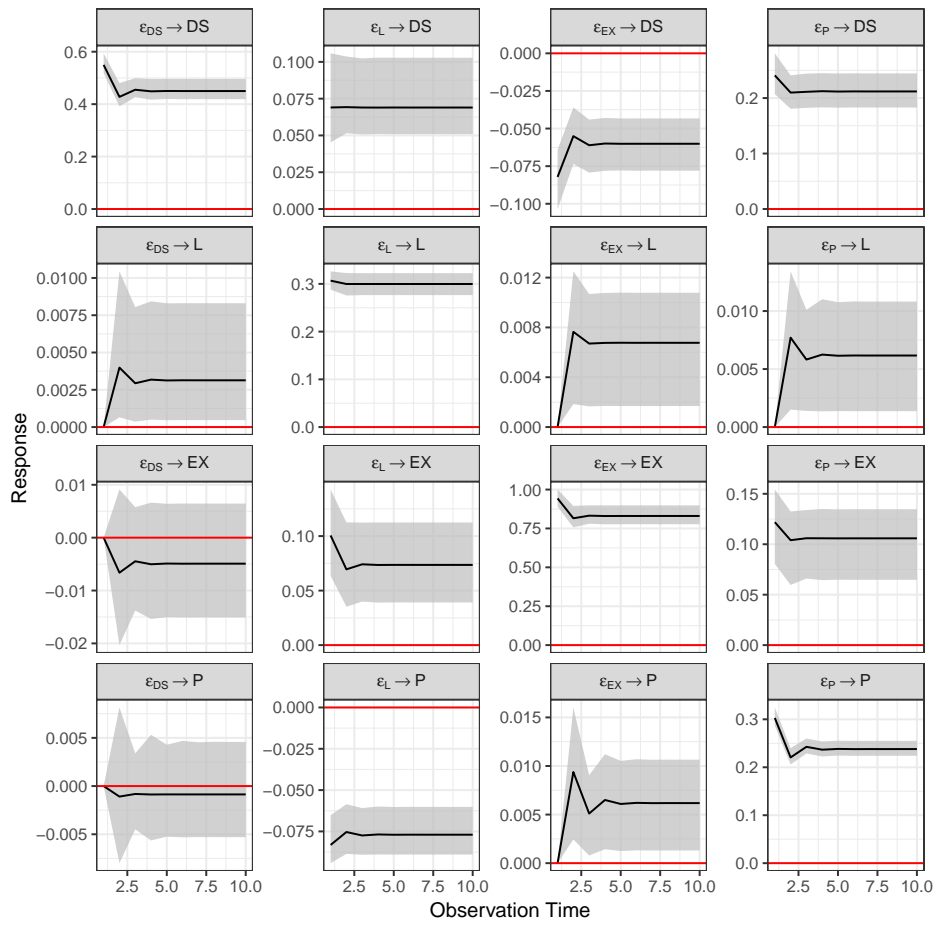


Figure 6: Impulse response functions with 90% confidence intervals obtained through bootstrap.




Figure 7: Cumulated impulse response functions with 90% confidence intervals obtained through bootstrap.

Appendix B: Further robustness checks

We perform several robustness studies. We report here below only results concerning the contemporaneous causal structure, which we summarise in a very schematic way. Notice that across several and diverse specifications productivity causally precedes exporting and not the other way around. Complete robustness results are available upon request.

Alternative TFP indicator

Our baseline productivity indicator sums up the productivity estimation in each sector, using Gross Value of Production and one overall material input variable for a sector sample with different employment categories (blue and white collar employees).

In this subsection, we use an alternative productivity variable that sums up the productivity estimation in each sector, using Gross Value of Production and one overall material input variable for the total sector sample.

The causal output of the VAR-LiNGAM algorithm is identical to the main study (see Figure 2) as regards the contemporaneous causal structure ($L \rightarrow P \rightarrow EX \rightarrow DS$), both for the 1-lag and 2-lag model. As regards the output of the VAR-LiNG algorithm, the contemporaneous causal matrix turns out to be almost full (which does not seem to be very informative): both for the 1-lag and 2-lag model there is only one zero (i.e. lack of contemporaneous causal relationship), from P to L .

Sector disaggregation

Firms are allocated into macro-sectors according to their technological content. These macro-sectors follow the Ferraz taxonomy of sectors, which is itself an adaptation of the Pavitt taxonomy to Latin American industries (following Ferraz, Kupfer and Haguenaer). Ferraz macrosector 1 corresponds to ‘commodities’ and ‘food commodities.’ Ferraz macrosector 2 corresponds to ‘durables’ and ‘auto industry.’ Ferraz macrosector 3 corresponds to ‘traditional sectors.’ Ferraz macrosector 4 corresponds to ‘technology diffusers/suppliers.’ We begin with estimations on subsamples of these four macro-sectors. However, one possible problem is that each macro-sector, taken individually, contains too few observations for a meaningful SVAR analysis. We therefore continue this sector disaggregation exercise (part 2, below) using a trick in Balasubramanian and Sivadasan (2011), which involves dropping one macro-sector at a time and performing robustness analysis on the remaining macro-sectors grouped together, in order to ensure that each subsample contains enough observations.

Ferraz macro-sector 1

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow EX \rightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow EX \rightarrow DS$ (identical to Figure 2).

VARLiNG output (1-lag model), zeros in the contemporaneous causal structures (i.e. lack of causal influence): from DS to EX and P , from EX to P , and from P to L . VARLiNG output (2-lag model): zeros in the contemporaneous causal structures (i.e. lack of causal influence): from DS to EX and P , from EX to P , and from P to L (but the algorithm is not completely stable under change of initial conditions).

Ferraz macro-sector 2

Not enough observations to run the analysis (sample size only 69).

Ferraz macro-sector 3

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNG output (1-lag model), zeros in the contemporaneous causal structures (i.e. lack of causal influence): from DS to EX and P , from EX to P , and from P to L . VARLiNG output (2-lag model), zeros in the contemporaneous causal structures (i.e. lack of causal influence): from EX to P and L , and from P to L .

Ferraz macro-sector 4

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNGAM output (2-lag model), contemporaneous causal structure: $DS \rightarrow L \rightarrow P \rightarrow EX$.

VARLiNG output (1-lag model): not stable under variations of initial conditions.

VARLiNG output (2-lag model): not stable under variations of initial conditions (sample size 231).

Sector disaggregation, part 2

Excluding Ferraz macro-sector 1

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNG output (1-lag model) contemporaneous causal structure: same as Figure 4, except for the presence of an extra arrow from DS to EX . VARLiNG output (2-lag model), only one zero in the contemporaneous causal structure: L is

not affected by any other variable (included therefore P).

Excluding Ferraz macro-sector 2

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow EX \longrightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow DS \longrightarrow EX$.

VARLiNG output (1-lag model) only one zero in the contemporaneous causal structure: L is not affected by any other variable. VARLiNG output (2-lag model): output not stable across algorithm initial conditions.

Excluding Ferraz macro-sector 3

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow EX \longrightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow EX \longrightarrow DS$ (identical to Figure 2).

VARLiNG output (1-lag model), zeros in the contemporaneous causal structures (i.e. lack of causal influence): from DS to EX and P , from EX to P , and from P to L . VARLiNG output (2-lag model): same as in the 1-lag model.

Excluding Ferraz macro-sector 4

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow EX \longrightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow DS \longrightarrow EX$.

VARLiNG output (1-lag model), zeros in the contemporaneous causal structures (i.e. lack of causal influence): from EX to P , and from P to L . VARLiNG output (2-lag model): zeros in the contemporaneous causal structures (i.e. lack of causal influence): from EX to P , and from P to L (but the algorithm is not completely stable under change of initial conditions).

Size disaggregation

In this subsection, firms are sorted according to whether their initial size (i.e. in the year 2001) is small (fewer than 50 employees), medium (50 to 249 employees) or large (250 or more employees).

Small initial size

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow EX \longrightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \longrightarrow P \longrightarrow EX \longrightarrow DS$ (identical to Figure 2).

VARLiNG output (1-lag model): a part from a causal feedback from DS to L , it is confirmed the recursive causal structure of VAR-LiNGAM. VARLiNG output (2-lag model): unstable.

Medium initial size

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNG output (1-lag model): unstable.

VARLiNG output (2-lag model): unstable.

Large initial size

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow EX \rightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow EX \rightarrow DS$ (identical to Figure 2).

VARLiNG output (1-lag model): unstable.

VARLiNG output (2-lag model): unstable.

Other checks

Firms are also sorted into subsamples according to their relative productivity (either above-median or below-median) in the initial year (i.e. 2001).

Below-median productivity

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow DS \rightarrow EX$.

VARLiNG output (1-lag model): unstable.

VARLiNG output (2-lag model): unstable.

Above-median productivity

VARLiNGAM output (1-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow EX \rightarrow DS$ (identical to Figure 2).

VARLiNGAM output (2-lag model), contemporaneous causal structure: $L \rightarrow P \rightarrow EX \rightarrow DS$ (identical to Figure 2).

VARLiNG output (1-lag model): unstable.

VARLiNG output (2-lag model): unstable.

OLS estimation

We also estimated the model using standard OLS regression instead of LAD. As regards the model with one lag, VAR-LiNGAM produces the same result as the one obtained from LAD estimation ($L \rightarrow P \rightarrow EX \rightarrow DS$). As regards

the model with two lags, the contemporaneous causal structure output of VAR-LiNGAM is $L \longrightarrow P \longrightarrow DS \longrightarrow EX$.

VAR-LiNG output (1-lag): same as LAD, with the addition of an edge from EX to L .

VAR-LiNG output (2-lag): same result as 1-lag model.

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). Export behavior and firm productivity in German manufacturing: A firm-level analysis. *Review of World Economics (Weltwirtschaftliches Archiv)* 141(2), 219–243.
- Atkin, D., A. K. Khandelwal, and A. Osman (2017). Exporting and firm performance: Evidence from a randomized experiment. *Quarterly Journal of Economics* 132(2), 551–615.
- 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.
- Balasubramanian, N. and J. Sivadasan (2011). What happens when firms patent? New evidence from us economic census data. *The Review of Economics and Statistics* 93(1), 126–146.
- 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.
- Blalock, G. and P. J. Gertler (2004). Learning from exporting revisited in a less developed setting. *Journal of Development Economics* 75(2), 397 – 416. 15th Inter American Seminar on Economics.
- 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.
- 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. and N. Grassano (2019). Firm growth and R&D investment: SVAR evidence from the world's top R&D investors. *Industry and Innovation* 26(5), 508–533.
- 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 Policy 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.
- 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. (2007). Do exports generate higher productivity? evidence from Slovenia. *Journal of International Economics* 73(1), 69 – 98.
- 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.
- Fernandes, A. M. and A. E. Isgut (2015). Learning-by-exporting effects: Are they for real? *Emerging Markets Finance and Trade* 51(1), 65–89.
- 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.
- Harris, R. and Q. Cher Li (2011). Export-market dynamics and firm-level productivity: evidence for UK tradable sectors. *Industrial and Corporate Change* 21(3), 649–670.
- Horn, R. A. and C. R. Johnson (2012). *Matrix analysis*. Cambridge university press.
- Hosono, K., D. Miyakawa, and M. Takizawa (2015). Learning by export: Does the presence of foreign affiliate companies matter? RIETI Discussion Paper Series 15-E-053, Research Institute of Economy, Trade and Industry.
- 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 deflatores. 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.
- Kilian, L. and H. Lütkepohl (2017). *Structural vector autoregressive analysis*. Cambridge University Press.
- 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.
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports* 66(1), 195–244.
- 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.
- Mooij, J., D. Janzing, J. Peters, and B. Schölkopf (2009). Regression by dependence minimization and its application to causal inference in additive noise models. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, New York, NY, USA, pp. 745–752. ACM.
- 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.
- 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.
- Stone, J. (2004). *Independent Component Analysis: A Tutorial Introduction*. A Bradford book. Cambridge, MA: 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.
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.