Exporting and productivity as part of the growth process: Results from a structural VAR*

Tommaso Ciarli
SPRU, University of Sussex
t.ciarli@sussex.ac.uk

Alex Coad
Pontificia Universidad Catolica del Peru
acoad@pucp.edu.pe

Alessio Moneta
Institute of Economics, Scuola Superiore Sant'Anna
a.moneta@santannapisa.it

January 24, 2018

Abstract

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 studied country, Chile. First, we focus on exporting firms only. Second, we do not compare them with non exporting firms. Third, we explicitly look at the co-evolution of the the two variables, productivity and growth, including the causal relation within the period and

*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 conference/seminar participants at Carnegie Mellon University, University of Jena, Queen Mary University of London, Free University of Berlin, University of Kassel for useful comments. Funding from the ESRC, TSB, BIS and NESTA on grants ES/H008705/1 and ES/J008427/1 as part of the IRC distributed projects initiative, as well as from the AHRC as part of the FUSE project, is gratefully acknowledged. The usual caveat applies.

with two lags. Our causal analysis is performed in the framework of a structural vector autoregression (SVAR). We identify such a model adopting a recent method which exploits non-Gaussianity in the data. Our findings suggest that exporting does not have any causal influence on the other variables. Instead, it seems to be determined by other dimensions of firm growth. Of interest to the inconclusive literature on learning by exporting (LBE) is that 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 numbers: L21; D24; F14

1 Introduction

Within the literatures on economic development and growth, openness to trade and industrialisation policies, such as 'infant industry' policies, have played a crucial role, and are still among the liveliest debates. Consider for instance the recent debate on new structural economics (Krueger, 2011; Lin, 2011; Lin and Chang, 2009; Rodrik, 2011), the growing literature on product spaces, specialisation and growth (Hausmann and Hidalgo, 2011, 2010; Hidalgo et al., 2007) and a few examples of the recent revitalisation of old debates (Alacevich, 2007; Hoff and Stiglitz, 2001; Ocampo, 2001).

Since the 1990s, part of this debate has focussed on micro-level evidence, looking at how firms take advantage from trade, particularly 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 from a large number of different countries. The results of this literature are relatively unambiguous in indicating a positive relation between exporting activity and productivity: 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). There are many factors that can explain this difference. Bernard and Jensen (1999) soon highlighted that in order to enter the global market firms need to be more productive than average. In other words, higher productivity may well be due to self-selection. We can then distinguish two main sets of factors: those explaining why firms need to increase productivity before entering the market – such as trade costs, stronger competition and investments to increase the scale, and why firms may increase productivity while exporting – such as learning from foreign buyers, use of excess capacity or stronger competition. Identifying which of the two sets of factors better explains the higher productivity of exporting firms has relevant policy implication for the debate on openness, industrialisation, and economic development.

A large amount of research has attempted to identify the direction of causality between export and growth. 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 techniques are used (control are firms as similar as possible to the treated ones, but which do not export) 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) report that results on LBE are non conclusive. Wagner (2012) reviews the large number of recent publications, reporting 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, with only a suggested tendency to find that firms are more likely to increase productivity by exporting when they focus on high income countries. Similar inconclusive results on LBE are found when analysing the service sector.

The focus of this paper is on the productivity dynamics of the relatively small number of exporting firms (Bernard et al., 2007). As anticipated, the evidence on LBE is still very mixed. International Study Group on Exports and Productivity (2008) use panel data from 14 different countries, finding no evidence of LBE. Baldwin and Yan (2012) find that, following changes in the real exchange rate, firms which are already in the export market experience a relatively larger gain in productivity than new entrants. 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. 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. However, Girma et al. (2004), who introduces the matching techniques, find a significant positive effect of export on productivity for UK manufacturing firms. And 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.

Crespi et al. (2008) use learning measures to estimate directly the effect of export on learning, which may affect productivity only in a second stage, and they do find evidence for LBE. Using exogenous shocks in the demand for exporting firms, Park et al. (2010) find evidence of LBE for Chinese firms, especially when exporting to high income countries. Following the De Loecker (2010) method to estimate productivity, Manjón et al. (2013) find evidence of LBE for Spanish manufacturing firms. Martins and Yang (2009) analyse the vast literature on LBE conducting a meta-analysis. They find that relative to high income countries, firms in developing countries enjoy a stronger impact of exporting on productivity.

In sum, the evidence on whether firms that are in the export market increase their positive productivity gap with respect to non-exporting firms is quite mixed.

In this paper we use a new method employing Independent Component Analysis to identify the casual relations among changing variables, a data-driven SVAR used for example by Moneta et al. (2013) to study firm growth and monetary policy. Focusing on the co-evolution of export and productivity growth we aim to shed more light on the direction of causality between the two variables (once firms have entered the export market). The use of Independent Component Analysis allows to identify causal effects of changes occurring in the same time periods as well as with a number of lags. We focus on Chile, a middle-income developing country. The advantages are that: Chile has a very well tested manufacturing firm survey, also with respect to the analysis of export and productivity (e.g. Alvarez and Crespi, 2007; Alvarez and López, 2005, 2008; López, 2009; Pavcnik, 2002); has had an open economy for three decades (we use the new data collected after 2001); and as a middle income developing country should show larger effect of export on productivity (Martins and Yang, 2009).

Another distinctive feature of our paper is that, focussing only on exporting firms, we analyse growth rates instead of levels. Although some works analyse the effect of exporting on productivity growth differences rather than levels, most works we are aware of do not analyse the effect of changes in export. Increase in export may indicate a larger scale, a larger number of clients and/or markets from which the firm could learn, or a relative increase in market shares. All mechanisms that may induce a positive effect on firm performance.

A few studies look at changes in exports. Park et al. (2010) who use exogenous shocks on the demand for exporting firms, and find evidence of LBE. Berman and Rebeyrol (2010), using data on French firms, find positive effects of export growth on productivity growth, contrary to firms that are on the international market but do not see their export increasing. Fernandes and Isgut (2005) 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. Finally, Antolín et al. (2013) use a different measure of export growth, the proportion of exports with respect to sales, and also find evidence of LBE.

In sum, the studies which focus on the growing relevance of firms export activity, seem to find more conclusive results on the presence of LBE, in line with the theory and evidence showing that foreign markets represent a relevant source of growth and development for those firms from low- and middle-income countries that manage to enter.

Our paper contributes to this literature, assessing the causal relation between exporting and productivity growth of Chilean firms. However, our preliminary results seem to indicate that, while changes in growth follow changes in productivity (simultaneous and lagged), there is no evidence of export growth causing productivity growth. This means that, at least once a firm has taken the binary decision of whether to export or not, changes in exporting have no detectable effect on productivity.

Section 2 presents the methodology, while Section 3 presents the dataset. The analysis is in Section 4. Section 5 contains a discussion of our results.

2 Econometric method

2.1 VAR and SVAR models

Vector autoregressive models were introduced by Sims (1980) to describe macroeconomic dynamics by treating all variables as potentially endogenous (for a survey cfr. Lütkepohl, 2013). The point of departure in VAR analysis is the specification and estimation of a reduced form model:

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{A}_1\mathbf{y}_{t-1} + \ldots + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{u}_t, \tag{1}$$

where $\mathbf{y}_t = (y_{1t}, \ldots, y_{kt})'$ is a vector of k time series variables, the $\mathbf{A}_i$ $(i = 1, \ldots, p)$ are $(k \times k)$ coefficient matrices, and $\mathbf{u}_t = (u_{1t}, \ldots, u_{kt})'$ is a k-dimensional zero mean white noise process with covariance matrix $E(\mathbf{u}_t\mathbf{u}_t') = \boldsymbol{\Sigma}_\mathbf{u}$. The vector $\boldsymbol{\mu}_t$ is a deterministic part, which in most cases is simply equal to a vector of constants.

Equation (1) is an approximate description of the unobserved data generating process (DGP), whose adequacy can be checked with the typical criteria of model selection and model checking, such as, for example, sequential testing procedure, Akaike or Schwarz's information criteria for selecting the VAR order (i.e. p), and tests for residual autocorrelation (cfr. Lütkepohl, 2013). An important feature of the reduced form model (1) is that it omits the fact that there might be mutual influences among the contemporaneous variables (within the period of observation) among $y_{1t}, \ldots, y_{kt}$. This omission is done in order to keep the variables on right hand side of the equation as pre-determined and hence getting consistent estimation of the coefficients through simple linear regression. But, precisely because of this omission, the coefficients being estimated cannot being interpreted as genuine causal influences.

Structural VAR analysis attempts to identify structural, i.e. causally meaningful, relations among the variables. The structural VAR model lets the (unobserved) coefficients that describe the contemporaneous causal influence to appear in the equation:

$$\mathbf{\Gamma}_0\mathbf{y}_t = \boldsymbol{\nu}_t + \mathbf{\Gamma}_1\mathbf{y}_{t-1} + \ldots + \mathbf{\Gamma}_p\mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \tag{2}$$

where $\mathbf{\Gamma}_0$ is a $(k \times k)$ matrix reflecting the instantaneous relations, and the $\mathbf{\Gamma}_i$ $(i = 1, \ldots, p)$ are the coefficient matrices of the lagged structural relations, reflecting causal influences present in the DGP. Again, the vector $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \ldots, \varepsilon_{kt})'$ is a k-dimensional zero mean white noise process with covariance matrix $E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = \mathbf{\Sigma}_{\boldsymbol{\varepsilon}}$ and $\boldsymbol{\nu}_t$ is the deterministic or constant part. In standard structural VAR analysis $\mathbf{\Sigma}_{\boldsymbol{\varepsilon}}$ is assumed to be diagonal, i.e. correlations among ε_{it} (over $i = 1, \ldots, k$) are zero. This is also equivalent to stating that the ε_{it} are orthogonal (conditional on the ε_{it} having a mean of zero). Besides assuming orthogonality and uncorrelatedness, we assume that ε_{it} is independent of ε_{jt} for each $i, j = 1, \ldots, k$ $(i \neq j)$. The independence assumption is consistent with the interpretation of the elements of $\boldsymbol{\varepsilon}_t$ as structural shocks, i.e. exogenous processes that affect each variable of the system at each time with the important feature that each term influences each variable in its own independent way. While in a setting with normally distributed error terms the distinction between independence and uncorrelatedness does not matter, in a non-Gaussian setting the further specification is crucial.

It is also assumed that the diagonal elements of $\mathbf{\Gamma}_0$ are equal to one (or that the system can always be rescaled in order to have ones in the main diagonal of $\mathbf{\Gamma}_0$). Let $\mathbf{B} = \mathbf{I} - \mathbf{\Gamma}_0$. Thus equation (2) can be rewritten as

$$\mathbf{y}_t = \boldsymbol{\nu}_t + \mathbf{B}\mathbf{y}_t + \mathbf{\Gamma}_1\mathbf{y}_{t-1} + \ldots + \mathbf{\Gamma}_p\mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t \tag{3}$$

Equation (3) cannot be directly estimated by linear regression because not all the variables on right hand side are predetermined or exogenous: some elements of $\mathbf{y}_t$ may instantaneously (i.e. within the period of observation) cause other elements of $\mathbf{y}_t$, without knowing which one. The relationship between the reduced form (cfr. equation 1) and the structural model (cfr. equation 2 or 3) is evident by pre-multiplying equation (2) or (3) by $\mathbf{\Gamma}_0^{-1}$ or $(\mathbf{I}-\mathbf{B})^{-1}$. We have that $\mathbf{\Gamma}_0^{-1}\mathbf{\Gamma}_i = \mathbf{A}_i$ $(i = 1, \ldots, p)$, $\mathbf{\Gamma}_0^{-1}\boldsymbol{\nu}_t = \boldsymbol{\mu}_t$ and $\mathbf{\Gamma}_0^{-1}\boldsymbol{\varepsilon}_t = \mathbf{u}_t$. The problem of identification consists in the fact that, having estimated the reduced form model we cannot directly recover the structural form model, because there are more parameters in equation (2) than in equation (1).

Structural VAR analysis is focused on imposing restrictions on $\mathbf{\Gamma}_0$ so that it can be retrieved from the data. The matrices $\mathbf{\Gamma}_i$ $(i = 1, \ldots, p)$, in turn, can be recovered from $\mathbf{\Gamma}_0$ joint with the $\mathbf{A}_i$ which are obtained from the estimation of equation (1). The restrictions on $\mathbf{\Gamma}_0$ are usually derived from theoretical or institutional knowledge, or placed on the basis of considerations about the long-run effects or signs of the shocks (Lütkepohl, 2013; Kilian, 2013). Other approaches aim at recovering $\mathbf{\Gamma}_0$ from an analysis of the conditional independence relations among the residuals $u_{1t}, \ldots, u_{kt}$ estimated in 1. Such conditional independence relations imply, under some assumptions about the restrictions that causality imposes on the probability structure, the presence of some causal relations and the absence of other. Search algorithms, such as those proposed by Spirtes et al. (2000) and Pearl (2009), exploit this information to find the class of admissible structures among $u_{1t}, \ldots, u_{kt}$ (cfr. Bessler and Lee, 2002; Demiralp and Hoover, 2003; Moneta, 2008; Bryant et al., 2009).

In this paper, we also apply an identification method focussed on inferring the instantaneous causal structure on the base of the study of the reduced-form residuals. But here we do not use conditional independence tests. The causal structure among the elements of $\mathbf{u}_t$ is captured by identifying their independent components. We apply here a technique first applied in econometrics by Moneta et al. (2013) which we extend here to a more general case, allowing the possibility of feedback loops among the contemporaneous variables. This method has, for the present case, a clear advantage with respect to the others because it exploits the non-Gaussian feature of the data[1] and in this manner permits us to further restrict the set of admissible causal structures.

2.2 Identification strategy

As in Moneta et al. (2013), our identification strategy is based on an approach for discovering linear non-Gaussian causal models which makes use of independent component analysis (ICA). The method searches for the appropriate matrix $\mathbf{\Gamma}_0$ that relates the vector of the structural shocks $\boldsymbol{\varepsilon}_t$ to the vector of reduced-form error terms $\mathbf{u}_t$ such that $\mathbf{u}_t = \mathbf{\Gamma}_0^{-1}\boldsymbol{\varepsilon}_t$. Given $\mathbf{u}_t$, ICA is able to find the latent sources which have been mixed to produce $\mathbf{u}_t$, under the assumption that the latent sources are independent and non-Gaussian. The underlying idea is to search for a mixture of the observed data (i.e. $\mathbf{u}_t$) 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, 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).

[1] As stated more explicitly below, non-Gaussianity and independence of the structural shocks are necessary conditions for the application of these methods.

No matter which algorithm is used, ICA leaves undetermined the scale, sign, and order of the latent sources or structural shocks. Further steps are needed to identify $\mathbf{\Gamma}_0$ and ε_t. We adopt here two different ICA-based search methods to identify the shocks and more in general the structural VAR model. The first one was proposed by Shimizu et al. (2006) and named LiNGAM (for linear, non-Gaussian, acyclic model). As applied to VAR models by Hyvärinen et al. (2008) and Moneta et al. (2013) is referred to as VAR-LiNGAM. The second one was proposed by Lacerda et al. (2012) and named LiNG (for linear, non-Gaussian model) (VAR-LiNG as applied to VAR). To our knowledge, this is the first time that the VAR-LiNG algorithm has been applied either in a VAR context 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 $\mathbf{P}$ ($\equiv \mathbf{\Gamma}_{ICA}^{-1}$) 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 underling 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 $\mathbf{\Gamma}_0$ must be equal to one, implies that if we find an ordering of the components $\hat{\varepsilon}_{1t}, \ldots, \hat{\varepsilon}_{kt}$ (output of the ICA algorithm) that produces a correspondence with the data $\hat{\varepsilon}_t = \tilde{\mathbf{\Gamma}}_0 \hat{\mathbf{u}}_t$ such that $\tilde{\mathbf{\Gamma}}_0$ has non-zero elements in its main diagonal, this ordering must be the correct one. Exploiting this fact, step 3 is devoted to find the permutation of the matrix $\mathbf{\Gamma}_{ICA}$ generating the independent components from $\hat{\mathbf{u}}_t$ which produces a correct matching between structural and reduced-form shocks. Step 4 solves the scale indeterminacy. This is simply done by normalizing the rows of $\tilde{\mathbf{\Gamma}}_0$, the correctly row-permuted version of $\mathbf{\Gamma}_{ICA}$), so that all diagonal elements equal unity. Let $\hat{\mathbf{\Gamma}}_0$ this row normalized matrix and $\hat{\mathbf{B}}_0 = I - \hat{\mathbf{\Gamma}}_0$ (step 5). Since it is assumed that there are no causal loops or feedbacks, there is a permutation (applied equally to columns and rows) of $\hat{\mathbf{\Gamma}}_\mathbf{0}$ which should be lower triangular. The same can be said for $\hat{\mathbf{\Gamma}}_\mathbf{0}^{-1}$ and $\hat{\mathbf{B}}_0$. 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 $\mathbf{\Gamma}_0$ has a zeroless diagonal, which is valid in the structural VAR model by construction. Step 3 tests which entries of the $\mathbf{\Gamma}_{ICA}$ are significantly different from zero. This can be done through a bootstrap procedure. Step 4 finds the permutation of the matrix $\mathbf{\Gamma}_{ICA}$ which produces a matrix $\tilde{\mathbf{\Gamma}}_{0,j}$ which has a zeroless diagonal. There might be several of such a matrix: we index each of them with $j = 1, \ldots, m$. 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. $\tilde{\mathbf{\Gamma}}_{0,j}$ such that $\tilde{\mathbf{\Gamma}}_{0,j}^{-1}$ has eigenvalues whose module 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 (1), obtaining estimates $\hat{\mathbf{A}}_i$ of the matrices $\mathbf{A}_i$ for $i = 1, \ldots, p$. Denote by $\hat{\mathbf{U}}$ the $k \times T$ matrix of the corresponding estimated VAR residuals, that is each column of $\hat{\mathbf{U}}$ is $\hat{\mathbf{u}}_t \equiv (\hat{u}_{1t}, \ldots, \hat{u}_{kt})'$, $(t = 1, \ldots, T)$. Check whether the u_{it} indeed are 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 $\hat{\mathbf{U}} = \mathbf{P}\hat{\mathbf{E}}$, where $\mathbf{P}$ is $k \times k$ and $\hat{\mathbf{E}}$ is $k \times T$, such that the rows of $\hat{\mathbf{E}}$ are the estimated independent components of $\hat{\mathbf{U}}$. Then validate non-Gaussianity and (at least approximate) statistical independence of the estimated components before proceeding.

3. Let $\mathbf{\Gamma}_{ICA} = \mathbf{P}^{-1}$. Find $\tilde{\mathbf{\Gamma}}_0$, the row-permuted version of $\mathbf{\Gamma}_{ICA}$ which minimizes $\sum_{i=1}^{k} 1/|\tilde{\mathbf{\Gamma}}_{0_{ii}}|$ 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 $\tilde{\mathbf{\Gamma}}_0$ by its diagonal element, to obtain a matrix $\hat{\mathbf{\Gamma}}_0$ with all ones on the diagonal.

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

6. Find the permutation matrix $\mathbf{Z}$ which yields a matrix $\hat{\mathbf{B}}_0 = \mathbf{Z}\tilde{\mathbf{B}}_0\mathbf{Z}'$ 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 $\hat{\mathbf{B}}_0$ to zero.

7. Calculate estimates of $\hat{\mathbf{\Gamma}}_i$ for lagged effects using $\hat{\mathbf{\Gamma}}_i = (\mathbf{I} - \hat{\mathbf{B}}_0)\hat{\mathbf{A}}_i$, 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 $\mathbf{\Gamma}_{ICA} = \mathbf{P}^{-1}$. Test which entries of $\mathbf{\Gamma}_{ICA}$ are zero. This can be done using a bootstrap procedure.
4. Find all admissible row-permuted matrices $\tilde{\mathbf{\Gamma}}_{0,1}, \dots, \tilde{\mathbf{\Gamma}}_{0,m}$ of $\mathbf{\Gamma}_{ICA}$ such that each $\tilde{\mathbf{\Gamma}}_{0,j}$ has zeroless diagonal for $j = 1, \dots, m$.
5. Divide each row of $\tilde{\mathbf{\Gamma}}_{0,j}$ by its diagonal element, to obtain a matrix $\hat{\mathbf{\Gamma}}_{0,j}$ with all ones on the diagonal, for each $j = 1, \dots, m$.
6. Let $\tilde{\mathbf{B}}_{0,j} = \mathbf{I} - \hat{\mathbf{\Gamma}}_{0,j}$, for each $j = 1, \dots, m$.
7. Calculate estimates of $\hat{\mathbf{\Gamma}}_{i,j}$ for lagged effects using $\hat{\mathbf{\Gamma}}_{i,j} = (\mathbf{I} - \hat{\mathbf{B}}_{0,j})\hat{\mathbf{A}}_{i,j}$, for $i = 1, \dots, p$, for $j = 1, \dots, m$.

---

3 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. We use the database that covers the period from 2001 to 2007.[2] 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 4 digits). For more information on the database see INE (2006, 2009a).

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

After some preliminary data cleaning[3], we create our SVAR variables. The variables used for the SVAR, as mentioned in Section 2 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. Year and sector dummies (respectively ANO and $Sector$) are also included. 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 $Privap$, $Privaf$, $Matvap$ and $Matvaf$). To compute value added at constant prices (Va) we use the generally preferred method of double deflation, and we remove initial inputs and add left overs at the end of the year: $Va = output - Material - (Privap + Matvap) + (Privaf + Matvaf)$.

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

Finally, we deflate the input variables used to compute productivity with the gross value of production ($output$): primary inputs, input materials purchased,

[3]We first check for inconsistencies in the data (Benavente and Ferrada (nd)), 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.

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 nothing that estimations of the TFP using value added and the whole sample of firms is highly correlated with labour productivity with a Spearman's correlation index of 0.96. However, for the sake of comparability with most other studies on the relation between export and productivity we use TFP estimations.

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

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

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

We create sub-samples for different size categories, based on number of employees: small (< 50), medium ($50 <= empl < 250$) 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 a couple of robustness checks, using a TFP estimated with no distinction between different types of employment ($tfp2$), leading to no significant differences in the relation between export 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 in the next section summarises the variables used for the analysis.

4 Analysis

Summary statistics are shown in Table 1, and correlations are shown in Table 2. Our four main variables are significantly correlated between them - although the magnitudes of the correlations are far below the value of 0.7 which is usually cited as the threshold above which problems of multicollinearity become severe. Of particular interest is the negative correlation between growth of exporting, and growth of TFP.

Table 3 presents the reduced-form VAR results, which describe the intertemporal associations between the variables but remain mute with regards to causal relations. A first observation is that the autocorrelation coefficients are generally negative, suggesting that increasing returns and sustained growth is not the norm. Of particular interest to our analysis is that lagged growth of exports is positively associated with subsequent growth of TFP in the 1-lag model (coefficient = 0.0192).

Before applying our identification method, we investigate whether the assumption of non-gaussian shocks is plausible. The evidence from the quantile-quantile plots in Figure 1 suggests that the shocks are far from Gaussian, providing support for our econometric strategy.

Table 4 shows the structural VAR estimated through algorithm 1 (VAR-LiNGAM), for both one-lag and two-lag models. Figure 2 shows both the contemporaneous and lagged causal relationships for the one-lag model.

Both the one-lag and two-lag models in Table 4 show that the *primus motor* is employment growth, which has large positive effects on growth of domestic sales and growth of exports. These can both be interpreted as 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

Table 1: Summary statistics

Variable	**Description**	**Obs**	**Mean**	**Std. Dev.**	**Min**	**Max**
Year	Year dummy	2303	2004.56	1.12	2003	2006
Sector	Sector dummy[a]	2303	3.47	2.39	0	9
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

[a] 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) Others.

[b] Estimated for different sectors, and differentiating between blue and 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. 2303 observations. All correlations significant at the 1% level, except for the Spearman rank correlation between gr_domsales and gr_exp (ρ=-0.0254, p-value=0.2223)

	gr_domsales	gr_empl	gr_exp	gr_tfp
gr_domsales	1	*0.1554*	*-0.0254*	*0.4679*
gr_empl	0.1431	1	*0.1670*	*-0.1870*
gr_exp	-0.0587	0.1356	1	*0.1508*
gr_tfp	0.4097	-0.2881	0.0885	1

Table 3: Reduced-form VAR estimated using median regression. Standard errors in parentheses. 1-lag VAR: 1389 observations in each regression. 2-lag VAR: 755 observations in each regression.

	first lag				second lag			
	l_gr_domsales	l_gr_empl	l_gr_exp	l_gr_tfp	l2_gr_domsales	l2_gr_empl	l2_gr_exp	l2_gr_tfp
gr_domsales	-0.149***	0.0345	0.0147	-0.0128				
	(0.0146)	(0.0277)	(0.0120)	(0.0299)				
gr_empl	0.00464	-0.00894	0.00993	0.0137				
	(0.00863)	(0.0165)	(0.00714)	(0.0182)				
gr_exp	0.0214	-0.0882**	-0.0602***	-0.0161				
	(0.0229)	(0.0433)	(0.0187)	(0.0479)				
gr_tfp	-0.0202*	-0.0125	0.0192**	-0.166***				
	(0.0108)	(0.0211)	(0.00912)	(0.0233)				
gr_domsales	-0.246***	0.108**	0.00958	0.101*	-0.208***	0.0667	0.0252	0.0660
	(0.0266)	(0.0443)	(0.0213)	(0.0578)	(0.0286)	(0.0541)	(0.0209)	(0.0522)
gr_empl	0.00866	-0.0282	0.0206*	0.0387	0.0207	0.0346	0.0214*	0.0130
	(0.0145)	(0.0250)	(0.0116)	(0.0319)	(0.0156)	(0.0304)	(0.0114)	(0.0286)
gr_exp	0.0418	-0.0708	-0.0470	-0.0720	0.00829	0.0672	-0.0157	-0.000986
	(0.0370)	(0.0628)	(0.0293)	(0.0814)	(0.0395)	(0.0764)	(0.0286)	(0.0721)
gr_tfp	-0.00988	-0.0221	0.0140	-0.234***	-0.0144	-0.00922	0.0165	-0.0874***
	(0.0166)	(0.0307)	(0.0144)	(0.0385)	(0.0189)	(0.0371)	(0.0140)	(0.0332)

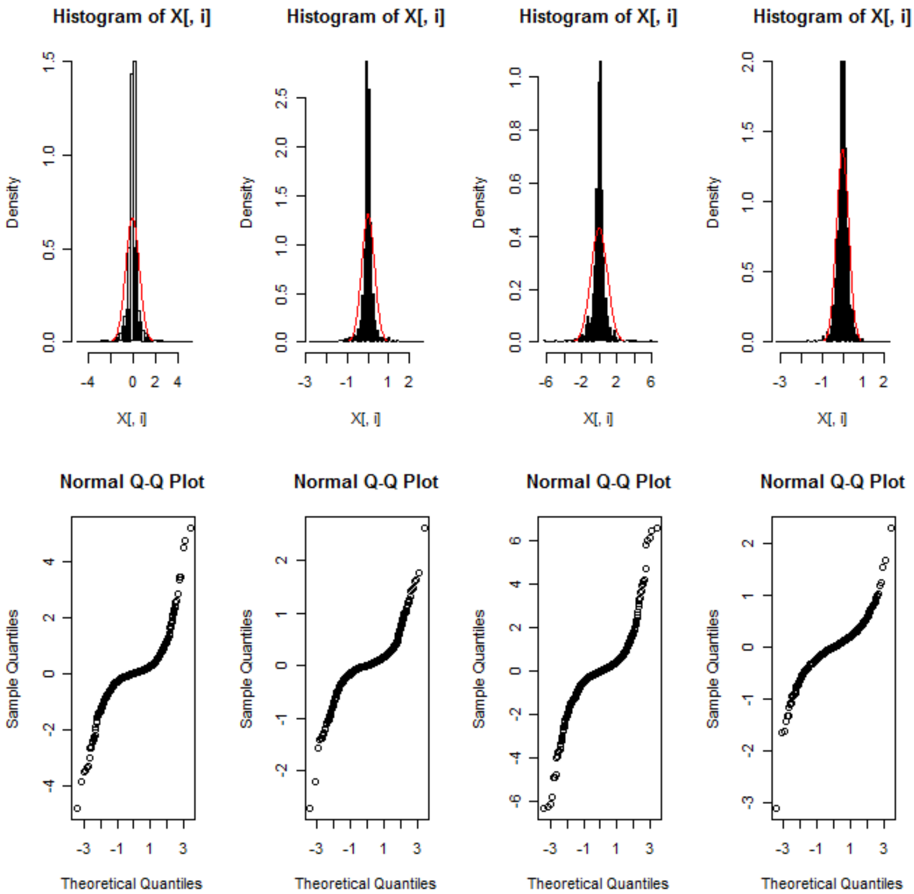

Figure 1: Quantile-quantile plots of the distributions of the four SVAR series, for the 2-lag model. The four variables are growth of domestic sales, growth of employment, growth of exports, and growth of TFP, respectively.

Table 4: VAR - LiNGAM estimates, 1- and 2-lag models.

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

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

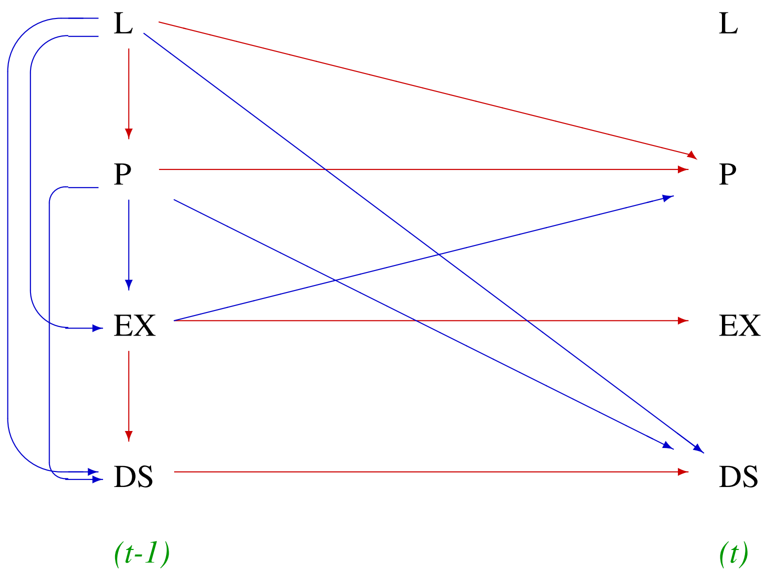

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

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 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

Figure 3: Bootstrap robustness analysis on the contemporaneous causal structure.

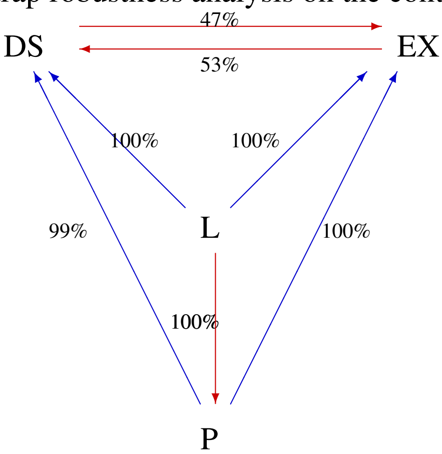

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

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 (significant at the 1% level in the 1-lag model, but slightly short of the 1% level in the 2-lag model).

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

Table 5: VAR - LiNG estimates, 1-lag model.

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

Table 5 reports the estimates of the application of algorithm 2 (VAR-LiNG, i.e. the algorithm which allows the possibility of feedback loops in the contemporaneous structure) 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

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

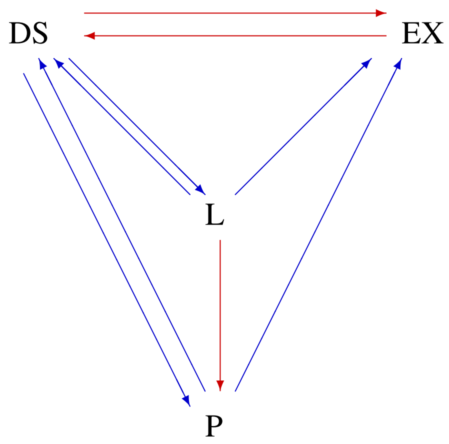

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

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 which we got from the first algorithm: 0.3782 vs. 0.4023 (standard errors are very similar as well: 0.0773 vs. 0.0707). But 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 in the results from algorithm one there was some (weak) influence from lagged exporting (growth) to current productivity (growth), in the results derived from algorithm 2 there is no such a causal link (see figure 5 for the lagged effects).

In further robustness analysis (not shown here), we use an alternative TFP indicator ("tfp2") which gives similar results.

Given the large number of heterogeneous effects found in the literature (Silva et al., 2012), we also plan to undertake further robustness analysis disaggregating

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

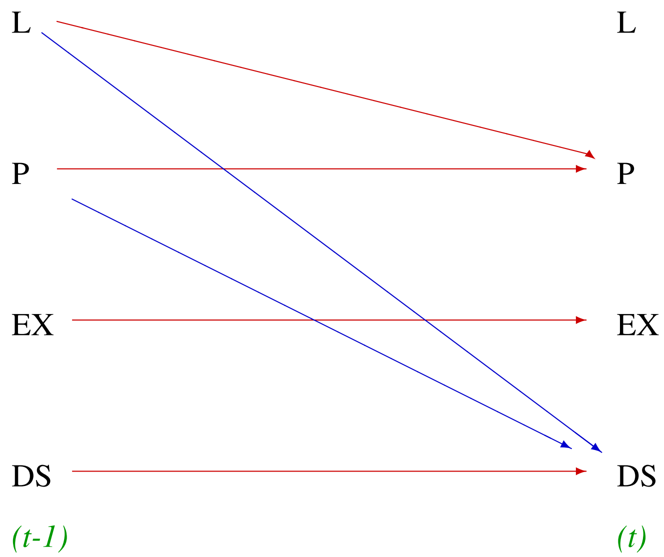

Note: L= *gr-empl*, P=*gr-tfp*, EX= *gr-exp*, DS= *gr-domsales*.
Blue edge = positive effect, red edge = negative effect.

by size groups and sectors, although we suspect that this may lead to problems of small sample sizes.

4.1 Robustness

Robustness can be investigated along a number of dimensions:

- Sector disaggregation (S1-S9)
- Ferraz sectors
- Size disaggregation (small/medium/large average size)
- New exporters vs old exporters
- Consider alternative estimations of TFP (beyond those already investigated), particularly the method suggested by De Loecker (2010) to take into account the effect that exporting has on productivity

5 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-studied country. First, we focus on exporting firms. Second, we do not compare them with non-exporting firms, but with other firms whose export grows 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 and with up to two lags.

Our most interesting finding is in relation with the extensively investigated LBE hypothesis. Adopting VAR-LiNGAM and VAR-LiNG, a class of SVAR models that estimates simultaneous directions of causality, it seems that exporting does not 'cause' any of the other variables. 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 so they 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, 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.

Apart from the robustness checks that we are still performing, there are a number of limitations to such a study. 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 extremely interested and curious to check how the method we used here would change earlier results on LBE 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. This is something we are willing to investigate in future work.

Third, our exercise on estimating TFP along different dimensions, shows that the measure is by no means robust. Apart from the critique of De Loecker (2010) to earlier LBE studies, there is a much general critique on the opportunity of using TFP (or productivity for that matters) as an indicator of performance, as somebody else has already pointed out.

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.

References

Alacevich, M. (2007, December). Early development economics debates revisited. Policy Research Working Paper Series 4441, The World Bank.

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, November). Exporting and performance: evidence from chilean plants. *Canadian Journal of Economics 38*(4), 1384–1400.

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

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

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

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 (nd). 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.

Bessler, D. A. and S. Lee (2002). Money and prices: Us data 1869–1914 (a study with directed graphs). *Empirical Economics 27*(3), 427–446.

Bryant, H. L., D. A. Bessler, and M. S. Haigh (2009). Disproving causal relationships using observational data. *Oxford Bulletin of Economics and Statistics 71*(3), 357–374.

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

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

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

De Loecker, J. (2010, November). A note on detecting learning by exporting. Working Paper 16548, National Bureau of Economic Research.

Demiralp, S. and K. D. Hoover (2003). Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and statistics 65*(1), 745–767.

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.

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.

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

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

Hausmann, R. and C. Hidalgo (2011). The network structure of economic output. *Journal of Economic Growth 16*, 309–342. 10.1007/s10887-011-9071-4.

Hausmann, R. and C. A. Hidalgo (2010, October). Country diversification, product ubiquity, and economic divergence. Working Paper Series rwp10-045, Harvard University, John F. Kennedy School of Government.

Hidalgo, C. A., B. Klinger, A.-L. Barabási, and R. Hausmann (2007). The product space conditions the development of nations. *Science 317*(5837), 482–487.

Hoff, K. and J. E. Stiglitz (2001). Modern economic theory and development. In G. M. Meier and J. E. Stiglitz (Eds.), *Frontiers of development economics: the future in perspective*, pp. 389–459. Oxford University Press.

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

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

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

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

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

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

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

Kilian, L. (2013). Structural vector autoregressions. In *Handbook of Research Methods and Applications in Empirical Macroeconomics*. Cheltenham, UK Edward Elgar Publishing, Inc.

Krueger, A. (2011). Comments on "new structural economics" by justin yifu lin. *The World Bank Research Observer 26*(2), 222–226.

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*.

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

Lin, J. and H.-J. Chang (2009, 09). Should industrial policy in developing countries conform to comparative advantage or defy it? a debate between justin lin and ha-joon chang. *Development Policy Review 27*(5), 483–502.

Lin, J. Y. (2011). New structural economics: A framework for rethinking development. *The World Bank Research Observer 26*(2), 1–29.

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

Lütkepohl, H. (2013). Vector autoregressive models. In *Handbook of Research Methods and Applications in Empirical Macroeconomics*. Cheltenham, UK Edward Elgar Publishing, Inc.

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, October). The impact of exporting on firm productivity: a meta-analysis of the learning-by-exporting hypothesis. *Review of World Economics 145*(3), 431–445.

Moneta, A. (2008). Graphical causal models and vars: an empirical assessment of the real business cycles hypothesis. *Empirical Economics 35*(2), 275–300.

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

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

Ocampo, J. A. (2001, December). Raúl prebisch and the development agenda at the dawn of the twenty–first century. *CEPAL Review 75*, 25–40.

Park, A., D. Yang, X. Shi, and Y. Jiang (2010, November). 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.

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.

Rodrik, D. (2011). Comments on “new structural economics” by justin yifu lin. *The World Bank Research Observer 26*(2), 227–229.

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.

Silva, A., O. Afonso, and A. P. Africano (2012). Learning-by-exporting: What we know and what we would like to know. *The International Trade Journal 26*(3), 255–288.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 1–48.

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

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

Wagner, J. (2007, 01). 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.