Venting Out: Exports during a Domestic Slump*

Miguel Almunia
CUNEF and CEPR

Pol Antràs
Harvard University, NBER and CEPR

David Lopez-Rodriguez
Banco de España

Eduardo Morales
Princeton University, NBER and CEPR

March 27, 2020

Abstract

We exploit plausibly exogenous geographical variation in the reduction in domestic demand caused by the Great Recession in Spain to document the existence of a robust, within-firm negative causal relationship between demand-driven changes in domestic sales and export flows. Spanish manufacturing firms whose domestic sales were reduced by more during the crisis observed a larger increase in their export flows, even after controlling for firms' supply determinants (such as labor costs). This negative relationship between demand-driven changes in domestic sales and changes in export flows illustrates the capacity of export markets to counteract the negative impact of local demand shocks. We rationalize our findings through a standard heterogeneous-firm model of exporting expanded to allow for non-constant marginal costs of production. Using a structurally estimated version of this model, we conclude that the firm-level responses to the slump in domestic demand in Spain could have accounted for around one-half of the spectacular increase in Spanish goods exports (the so-called 'Spanish export miracle') over the period 2009-13.

*Email: miguel.almunia@cunef.edu; pantras@fas.harvard.edu; david.lopezr@bde.es; ecmorale@princeton.edu. We thank María Jesús González Sanz and Xiang Zhang for research assistance, Rafael Frutos and Francesco Serti for help with the tax records of firm-level sales within Spain, Carlos Llano for help with the C-Intereg data, and Antoine Berthou, Rafael Dix-Carneiro, Fernando Leibovici, and Daniel Ramos for their detailed comments. We are also grateful to Olivier Blanchard, Kirill Borusyak, James Fenske, Clément Imbert, Wolfgang Keller, Enrique Moral, Michael Peters, Pedro Portugal, Tim Schmidt-Eisenlohr, Felix Tintelnot, Alberto Urtasun, and seminar audiences at the EEA meetings in Lisbon, the EIIT conference in Washington, D.C., Bank of Spain, Hitotsubashi, Warwick, CUNEF, Georgetown, Austin, IE Business School, University of Michigan, Banco de la República (Bogotá), CEMFI, Vienna, Vanderbilt, the NBER Summer Institute, UIBE (Beijing), UQAM (Montreal), Princeton, Harvard, the ECB, Erasmus University of Rotterdam, LSE, RIETI (Tokyo), Bilkent, Nebrija, Munich, Maryland, Boston College and the Philly Fed for useful comments. Finally, we particularly thank Óscar Arce for his continuous support throughout this project. All errors are our own. Appendices A and B appear at the end of the main text, while Appendices C–K are available online here. Any views expressed in this paper are only those of the authors and should not be attributed to the Banco de España or the Eurosystem.

1 Introduction

The Great Recession of the late 2000s and early 2010s shook the core of many advanced economies. Few countries experienced the consequences of the global downturn as intensively as the Southern economies of the European Monetary Union (EMU) did. Spain is a case in point. From its peak in 2008, Spain's real GDP fell by an accumulated 8.9% in the following five years, until bottoming out in 2013. During the same period, private final consumption contracted by 14.0%, and the unemployment rate shot up from 9.6% to 26.9%. Portugal and Greece also experienced marked contractions between 2008 and 2013, with their GDPs shrinking by 7.9% and 26.3%, respectively.

Despite these severe domestic slumps, merchandise exports in these economies demonstrated a remarkable resilience and partly contributed to mitigating the effects of the Great Recession. In the Spanish case, after tumbling by 11.5% in real terms during the global trade collapse of 2008-2009, Spanish merchandise exports quickly recovered and grew by 30.7% in real terms between 2009 and 2013.[1] Overall, real Spanish merchandise exports grew by an accumulated 15.6% during the 2008-2013 period, while real merchandise exports in the rest of the euro area increased by only 6.8% during the same years. As a result, and as shown in Figure 1, the share of euro area merchandise exports to non-euro area countries accounted for by Spain increased markedly during this period (especially in 2011-13), despite the contemporaneous decline in the relative weight of Spain's GDP in the euro area's GDP. Very similar patterns are observed for the cases of Portugal and Greece as well as for various other euro-area countries (see Appendix D.1).[2]

Figure 1: The Spanish Export Miracle

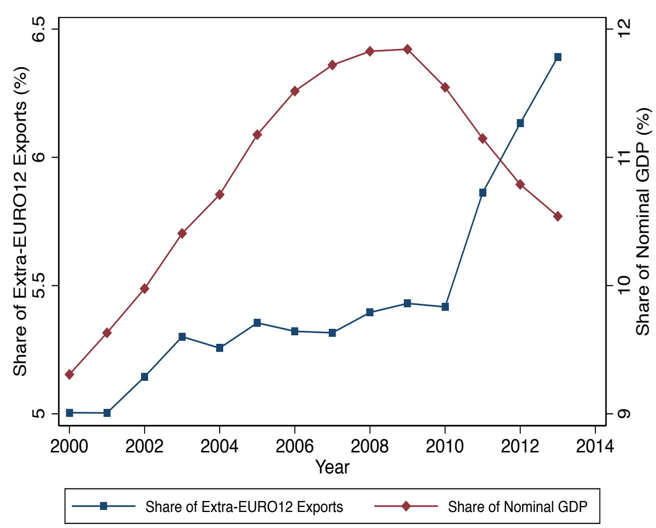

[1]The implied 6.9% annual growth in real exports from 2009 to 2013 almost doubled the 3.8% annual growth in real exports during the period 2000-2008.

[2]In Appendix D.1, we replicate Figure 1 for Portugal, Greece, Ireland, Italy, Germany, France, and the Netherlands. For Portugal and Greece, and less clearly for Germany and France, we observe a negative relationship between their GDP shares in the euro area and their shares in euro area exports of goods to other countries. See Appendix D.1 for a description of the data sources underlying these figures.

At first glance, this remarkable export performance appears to be consistent with the goals of the type of "internal devaluation" processes advocated by international organizations (such as the IMF, the ECB or the European Commission) since the onset of the crisis. According to this thesis, wage moderation coupled with a set of structural reforms (most notably labor market reforms) led to a fall in relative unit labor costs, allowing Southern European firms to reduce their relative export prices and increase their market shares abroad. Nevertheless, in the Spanish case, the adjustment in labor costs achieved via these policies was modest up to 2013 and this channel is believed to have had a limited contribution to export growth over the period 2010-13 (see, for instance, IMF, 2015, 2018; Salas, 2018).

What explains then the remarkable export growth in Spain, Portugal and Greece over the period 2010-2013? At least for the case of Spain, an often-invoked alternative explanation relates the growth in exports *directly* to the collapse in domestic demand. According to this hypothesis, the unexpected demand-driven reduction in Spanish firms' domestic sales, in combination with the irreversibility of certain investments in inputs, freed up capacity that these firms used to serve customers abroad.[3] More precisely, this explanation posits that, as domestic demand dropped, Spanish firms were able to cut their short-run marginal costs by reducing their usage of flexible inputs (e.g., temporary workers and materials) relative to their usage of fixed inputs (e.g., physical capital and permanent workers). This fall in short-run marginal costs translated into a gain in competitiveness in foreign markets and, consequently, to an increase in firms' exports.[4]

This alternative explanation resonates with the "vent-for-surplus" theory of the benefits of international trade, which has a long tradition in economics dating back to Adam Smith.[5] Despite its intuitive nature and distinguished lineage, the link between a domestic slump and export growth is hard to reconcile with modern workhorse models of international trade. The reason for this is that these canonical models – including those emphasizing product differentiation and economies of scale as in Krugman (1984) and Melitz (2003) – assume that firms face constant marginal costs of production, an assumption that implies that demand shocks in one market do not affect a firm's sales in another market.

In this paper, we leverage Spanish firm-level data from 2002 to 2013, and geographic variation

[3]See "La exportación como escape" in El País, 1/16/2016, for a journalistic account in Spanish with some specific case studies (https://elpais.com/economia/2016/01/14/actualidad/1452794395_894216.html). Further firm-level examples are provided in the more recent "El milagro exportador español" in El País, 5/27/2018 (http://elpais.com/economia/2018/05/25/actualidad/1527242520_600876.html), a newspaper article which was inspired by an early version of our paper.

[4]Generally, one can interpret this explanation as encompassing any mechanism that makes firms' short-run marginal cost curves increasing and that, thus, links the drop in firms' domestic demand to a downward movement *along* their supply curves. This effect is distinct from that of an "internal devaluation", which is associated with a downward *shift* in firms' marginal cost or supply curves (e.g., reductions in the price of factors or materials, or increases in productivity).

[5]In *The Wealth of Nations* (1776) Book II, Chapter V, Adam Smith writes "When the produce of any particular branch of industry exceeds what the demand of the country requires, the surplus must be sent abroad, and exchanged for something for which there is a demand at home. Without such exportation, a part of the productive labour of the country must cease, and the value of its annual produce diminish." The term "vent-for-surplus" was introduced by John Stuart Mill in his *Principles of Political Economy* (1848) and popularized by Williams (1929) and Myint (1958).

across Spanish regions in the reduction in domestic demand that took place during the Great Recession, to study the empirical relevance of the "vent-for-surplus" mechanism. To do so, we first divide our sample into a "boom" period (2002-08) and a "bust" period (2009-13), and measure the extent to which, at the firm level, a decline in the domestic sales in the bust period relative to the boom period is associated with an increase in export sales over the two periods. When measuring this association, we control for "boom-to-bust" changes in observed marginal cost shifters (i.e., measures of factor prices and productivity) to account for potential internal devaluation effects.

To further isolate demand-driven changes in domestic sales, we develop an instrumental variable approach that exploits the fact that the Great Recession affected different geographical areas in Spain differentially. In particular, we rely on municipality-level registration data on a major household durable consumption item, vehicles, as well as on municipality-to-municipality manufacturing sales data (from tax records of firm-level sales within Spain) to construct an instrument proxying the extent to which the Great Recession affected the domestic demand faced by Spanish manufacturing firms located in different municipalities. More precisely, we employ the municipality-to-municipality trade flows to estimate the extent to which firms in a given municipality are exposed to demand shocks in other municipalities, and we use changes in the municipality-level stock of vehicles per capita between 2002-08 and 2009-13 as a proxy for the changes in demand in each municipality caused by the Great Recession.

To understand the properties of our estimates of the causal impact of demand-driven changes in domestic sales on exports, we first rely on a commonly used model of firms' export behavior: a model à la Melitz (2003). For our purposes, this framework serves the role of identifying several empirical challenges that one encounters when measuring the relevance of the "vent-for-surplus" mechanism; i.e., when measuring the causal impact of changes in a firm's domestic sales due to changes in its domestic demand on exports.[6] We draw two main conclusions from our theoretical analysis. First, as long as firms' marginal cost shifters (i.e., firms' productivity and production factor costs) are not perfectly observable – and their unobserved component is not fully captured by various fixed effects – there will tend to be a *positive* spurious correlation between domestic sales and exports that is not informative about the impact of demand-driven changes in the former on the latter. Second, an instrumental variable approach identifies the causal impact of demand-driven changes in domestic sales on exports as long as the instrument satisfies two conditions: (i) it is a good predictor of the domestic sales of Spanish firms, and (ii) it is not correlated with firms' unobserved marginal cost or export-demand shifters.[7]

With these considerations in mind, we construct a theory-based instrument that, for each firm, weighs proxies for the local demand shocks experienced by municipalities other than the municipality of location of the firm, using as weights gravity-based estimates of the forces affecting

[6] The Melitz (2003) model assumes that firms face constant marginal costs of production, implying the null hypothesis of a zero effect of demand-driven changes in domestic sales on exports. However, as we show in section 7, the lessons we learn from this model in terms of the econometric challenges one faces when evaluating the "vent-for-surplus" mechanism are also applicable to more general models that feature increasing marginal costs of production.

[7] In our theoretical analysis, we also discuss the implications of measurement error whenever domestic sales are computed as the difference between firm-level total sales and exports (see also Berman et al., 2015).

trade flows between any two Spanish municipalities. We compute these gravity-based estimates using municipality-to-municipality sales data; these estimates reveal a significant amount of "home bias" *within* Spain, with shipments declining in distance with an elasticity of around -0.4 even after controlling for the discontinuity in sales observed when shipping outside a firm's municipality. These results are consistent with the findings of Hillberry and Hummels (2008) for the U.S. and of Díaz-Lanchas et al. (2013) for Spain.[8] To find a credible proxy for 'local demand', we invoke an extensive literature in empirical macroeconomics documenting that consumption of durable goods such as vehicles is strongly procyclical (see, for instance, the survey by Stock and Watson, 1999). More recently, Mian et al. (2013), Hausman et al. (2019), and Waugh (2019) have also documented a strong link between wealth (or income) shocks and vehicle consumption. Against this backdrop, we use as an instrument for a firm's boom-to-bust change in domestic sales the boom-to-bust change in a weighted-sum of the stock of vehicles per capita in all municipalities other than the municipality in which the firm is located, with weights obtained from our gravity-equation estimation. Our first-stage results indicate that our instrument is indeed relevant.

Armed with these first-stage results, we show that a larger *demand-driven* drop in domestic sales in the bust period relative to the boom period is associated with a significantly larger growth in export sales from boom to bust (conditional on exporting in both periods). Furthermore, these IV estimates are significantly larger than the OLS ones. This is consistent with the biases predicted by our baseline Melitz (2003)-type model in the plausible scenario in which our covariates only imperfectly control for a firm's supply determinants. Specifically, our IV estimates point at an intensive-margin elasticity of exports to domestic sales in the neighborhood of -1.6, while the OLS one is around -0.3. To assuage concerns about our long-differences approach to identification, we also present estimates that break the period 2002-2013 into four subperiods (2002-05, 2006-08, 2009-11, 2012-13), and which allow for the inclusion of municipality-specific time trends in the estimation. This robustness exercise shows that our estimates are not biased by the presence of underlying heterogeneous trends across regions in Spain. When estimating the effect of a demand-driven changes in domestic sales on the probability of exporting, we instead find quantitatively small effects that are somewhat sensitive to the way in which the extensive margin of exporting is measured.

As indicated above, a potential challenge to our identification approach is that the boom-to-bust changes in our instrument may be correlated with the extent to which unobserved shifters of the firm's marginal cost curve changed in the bust period relative to the boom period. Although we do not use information on demand changes in the firm's own municipality when constructing the instrument, one may still be concerned about spatial correlation in demand shocks posing a threat to identification. With that in mind, in section 6 we provide additional pieces of evidence that are consistent with the empirical relevance of the "vent-for-surplus" hypothesis and that address some specific sources of endogeneity that could affect the validity of our baseline instrument.

[8] The estimates in Hillberry and Hummels (2008) rely on survey data; those in Díaz-Lanchas et al. (2013) rely on a sample of road freight shipments.

First, an identification threat arises if differences in the severity of the contraction in vehicle purchases across Spanish municipalities are not exclusively a reflection of differences in demand shocks, but also partly a reflection of unobserved production costs affecting car manufacturers. According to this hypothesis, if a significant share of vehicles is sold in the near vicinity of where they are produced, municipalities that concentrate a significant share of firms operating in the auto industry could observe a correlation in the boom-to-bust changes in production costs and nearby purchases of new vehicles. Our results are robust to this identification threat. Both the relevance of our instrument as well as the finding of a sizable negative elasticity between domestic sales and exports are robust to excluding from the estimating sample: (a) all firms in the auto industry, no matter where they are located; (b) all firms located in any zip code that hosts at least one automaker employing more than 20 workers; (c) all firms located in any zip code that is geographically close to a zip code in which a significant share of manufacturing employment is in the auto industry; and (d) all firms producing in sectors that are either leading input providers or leading buyers of the vehicles manufacturing industry.

Second, although we control for firm-specific average wages in all of our specifications, compositional changes in the firm's workforce may have caused changes in effective labor costs that our wage measure does not capture correctly. An important feature of the Spanish labor market is the division of the workforce into permanent and temporary workers, the latter group being typically less productive than the former (see Dolado et al., 2002). We do indeed observe that firms whose share of temporary workers dropped by more in the bust relative to the boom experienced a smaller drop in their exports, consistently with the hypothesis that an increase in the ratio of permanent to temporary workers had an effect equivalent to a positive supply shock. The elasticity of exports with respect to domestic sales remains however largely unaffected when we control for the firm's change in the share of temporary workers. Similarly, controlling for the change in financial costs experienced by exporters or for proxies of trade credit available to firms does not change the second-stage estimate of the elasticity of exports with respect to domestic sales. In addition, in section 6 we explore the robustness of our results to alternative constructions of our instrumental variable and to an alternative approach to measuring firm-level total factor productivity.[9]

Having established a causal link between changes in domestic demand and exports that operates through firms' changes in domestic sales, we generalize our baseline model à la Melitz (2003) to allow for non-constant marginal costs of production. We rationalize this cost structure by including a pre-determined and fixed factor into the firm's production function, and show that the curvature of the firm's marginal cost function is related to the elasticity of output with respect to all flexible factors. Furthermore, we demonstrate how to estimate the curvature of the marginal cost function using a simple variant of our IV estimator. Consistently with our micro-foundation, we find that our estimate of this curvature is larger for sectors that make a more intensive use of fixed factors,

[9] Appendix G contains several additional estimates that illustrate the robustness of our baseline results. In Appendix H.2, we exploit alternative instrumentation strategies that focus on the deep roots of the differential fall in demand across Spanish regions (e.g., the bursting of a housing bubble, the collapse of the construction sector, and the decline in income from tourism).

though the statistical significance of this result is sensitive to which factors one classifies as "fixed".

Finally, we employ our model with increasing marginal costs and the corresponding IV estimates to quantitatively evaluate the importance of the "vent-for-surplus" mechanism in explaining the 2009-13 observed export miracle in Spain. More specifically, we implement a variance-decomposition exercise to determine the extent to which the domestic slump in Spain was driven by demand versus supply shocks. We then use our model to predict the boom-to-bust growth in Spanish exports that we would have observed if there had been no change in demand between the boom and bust periods. We find that, in this case, the growth in Spanish exports would have been 51.71% smaller than what we observe in the data and, thus, we conclude that slightly more than half of the Spanish export miracle of the period 2009-2013 can be attributed to the "vent-for-surplus" mechanism.

Our paper connects with several branches of the literature. As mentioned above, we relate the Spanish export miracle to Adam Smith's "vent-for-surplus" theory. The international trade literature has largely ignored this hypothesis as exemplified by the fact that we have only found one mention (in Fisher and Kakkar, 2004) of the term "vent-for-surplus" in all issues of the *Journal of International Economics*.[10] Nevertheless, there has been an active recent international trade literature focused on relaxing the assumption of constant marginal costs in the canonical (Melitz) model of firm-level trade, and has shown that, in the presence of increasing marginal costs, there is a natural substitutability between domestic sales and exports for which there is supporting empirical evidence. This literature includes the work of Vannoorenberghe (2012), Blum, Claro and Horstman (2013), Soderbery (2014), and Ahn and McQuoid (2017). Relative to this prior literature, our paper exploits plausibly exogenous variation in demand during a particularly salient episode to identify the causal effect of a drop in domestic sales on exports. Additionally, it provides an approach to identify and structurally estimate the slope of firms' short-run marginal cost curves. Relatedly, in contemporaneous work, Fan et al. (2018) exploit variation in the extent to which Chinese authorities enforce the collection of value-added taxes to establish a negative causal link between the profitability of domestic sales and firm-level exports. Conversely, using French data over the period 1995-2001, Berman et al. (2015) document a *positive* causal effect of changes in firm-level exports on firm-level domestic sales. Their identification strategy (based on exogenous variation in foreign demand conditions) is quite distinct from ours and so is their setting, since 1995-2001 was a tranquil period of sustained economic growth in France. In Appendix H.1, we use data on Spanish firms for the period 2002-07 to perform an analysis analogous to that in Berman et al. (2015), and we find no evidence supporting the positive causal relationship between exports and domestic sales that these authors previously found; on the contrary, for most specifications, we find a *negative* causal effect of (plausibly) exogenous changes in exports on domestic sales, in line with our core finding of substitution between exports and domestic sales.[11]

[10] A broader search to include top general-interest journals identified Neary and Schweinberger (1986), who provide a neoclassical rationale for the "vent-for-surplus" idea.

[11] Our paper also relates to a prior literature describing the behavior of firm-level exports in Spain around the Great Recession, including Antràs (2011), Myro (2015), Eppinger et al. (2017), and De Lucio et al. (2017a, 2017b).

The rest of the paper is structured as follows. In section 2, we lay out a baseline model of firm behavior in the spirit of Melitz (2003) and discuss its implications for the estimation of the causal impact of demand-driven changes in domestic sales on exports. In section 3, we introduce our firm-level data and, in section 4, we develop our core instrumental variable estimation approach. The results of this instrumental variable approach are presented in section 5. We present additional evidence in favor of the "vent-for-surplus" mechanism in section 6. In section 7, we generalize the baseline model à la Melitz (2003) to allow for non-constant marginal costs, and use this framework to quantify the importance of the "vent-for-surplus" channel in linking the slump in domestic demand to the growth in Spanish exports. We offer some concluding remarks in section 8.

2 Benchmark Model: Estimation Guidelines

As indicated in the Introduction, we aim to estimate the causal impact of within-firm demand-driven changes in domestic sales on firm-level exports. To guide our empirical analysis and our choice of an adequate estimator, we first consider the implications for this question of a model of exporting with heterogeneous firms along the lines of Melitz (2003), which is the canonical model of firm-level exports in the recent international trade literature. This model features the standard assumption of constant marginal costs. After presenting our evidence contradictory with this assumption, in section 7 we will develop an extension of this benchmark model that allows for non-constant marginal costs. Crucially, the lessons we learn in this section about the properties of different estimators will also apply in the more general model.

2.1 Benchmark Model: Estimating Equation

We index manufacturing firms producing in Spain by i, the sectors to which firms belong by s, and the two potential markets in which they may sell by $j = \{d, x\}$, with d denoting the domestic market and x denoting the export market. In principle, both the domestic and export market are an aggregate of several destinations, but due to data limitations, we focus in the main text on this dichotomous case (we develop a multi-destination extension of our model in Appendix E.2).

In any given period, firm i faces the following isoelastic demand in market j,

$$Q_{ij} = \frac{P_{ij}^{-\sigma}}{P_{sj}^{1-\sigma}} E_{sj} \xi_{ij}^{\sigma-1}, \qquad \sigma > 1, \tag{1}$$

where Q_{ij} denotes the number of units of output of firm i demanded in market j if it sets a price P_{ij}, P_{sj} is the sector s price index in j, E_{sj} is the total sectoral expenditure in market j expressed in units of the numeraire; and ξ_{ij} is a firm-market specific demand shifter.

Firm i's total variable cost of producing Q_{ij} units of output for market j is given by

$$c_{ij} Q_{ij} \qquad \text{with} \qquad c_{ij} \equiv \tau_{sj} \frac{1}{\varphi_i} \omega_i, \tag{2}$$

where c_{ij} denotes the marginal cost to firm i of selling one unit of output in market j, τ_{sj} denotes an iceberg trade cost, φ_i denotes firm i's productivity, and ω_i is the firm-specific cost of a bundle of inputs. Additionally, we assume that firm i needs to pay an exogenous fixed cost F_{ij} to sell a positive amount in market j.

Firm i chooses optimally the quantity offered in each market j, Q_{ij}, taking the price index, P_{sj}, and the size of the market, E_{sj}, as given. As the marginal production cost is independent of the firm's total output and the per-market fixed costs are independent of the firm's participation in other markets, the optimization problem of the firm is separable across markets. Specifically, conditional on selling to a market j, firm i solves the following optimization problem

$$\max_{Q_{ij}} \left\{ Q_{ij}^{\frac{\sigma-1}{\sigma}} P_{sj}^{\frac{1-\sigma}{\sigma}} E_{sj}^{\frac{1}{\sigma}} \xi_{ij}^{\frac{\sigma-1}{\sigma}} - \tau_{sj} \frac{1}{\varphi_i} \omega_i Q_{ij} \right\},$$

and sales by firm i to market j are thus $R_{ij} = P_{ij} Q_{ij} = \kappa \left((\xi_{ij}\varphi_i)/(\tau_{sj}\omega_i) \right)^{\sigma-1} E_{sj} P_{sj}^{\sigma-1}$, where κ is a function of σ. For the case of exports ($j = x$), and taking logs, we can rewrite this expression as:

$$\ln R_{ix} = \ln \kappa + (\sigma - 1)(\ln \xi_{ix} + \ln \varphi_i - \ln \omega_i) - (\sigma - 1)(\ln \tau_{sx} - \ln P_{sx}) + \ln E_{sx}. \tag{3}$$

The bulk of our empirical analysis will compare firm-level export behavior in a *bust* period, relative to a *boom* period. With that in mind, and letting $\Delta \ln X$ denote the log change in the cross-year average value of X from boom to bust, we can express the log change in exports from boom to bust as

$$\Delta \ln R_{ix} = (\sigma - 1)\left[\Delta \ln \xi_{ix} + \Delta \ln \varphi_i - \Delta \ln \omega_i\right] - (\sigma - 1)(\Delta \ln \tau_{sx} - \Delta \ln P_{sx}) + \ln \Delta E_{sx}. \tag{4}$$

In order to transition to an estimating equation, we model the change in firm-specific foreign demand, productivity and cost levels as follows:

$$\begin{aligned} \Delta \ln(\xi_{ix}) &= \xi_{sx} + u_{ix}^{\xi}, \\ \Delta \ln(\varphi_i) &= \varphi_s + \delta_\varphi \Delta \ln(\varphi_i^*) + u_i^{\varphi}, \\ \Delta \ln(\omega_i) &= \omega_s + \delta_\omega \Delta \ln(\omega_i^*) + u_i^{\omega}. \end{aligned} \tag{5}$$

Note that we are decomposing these terms into (i) a sector fixed effect, (ii) an *observable* part of these terms for the case of productivity (φ_i^*) and for input bundle costs (ω_i^*), and (iii) a residual term. We can thus re-write equation (4) as:

$$\Delta \ln R_{ix} = \gamma_{sx} + (\sigma - 1)\,\delta_\varphi \Delta \ln(\varphi_i^*) - (\sigma - 1)\,\delta_\omega \Delta \ln(\omega_i^*) + \varepsilon_{ix}, \tag{6}$$

where $\gamma_{sx} \equiv (\sigma - 1)\left[\xi_{sx} + \varphi_s - \omega_s - \ln \tau_{sx} + \ln P_{sx}\right] + \ln E_{sx}$, and where

$$\varepsilon_{ix} = (\sigma - 1)\,[u_{ix}^{\xi} + u_i^{\varphi} - u_i^{\omega}]. \tag{7}$$

Following analogous steps as above, we derive an expression for the change in domestic sales:

$$\Delta \ln R_{id} = \gamma_{sd} + (\sigma - 1)\, \delta_{\varphi} \Delta \ln(\varphi_i^*) - (\sigma - 1)\, \delta_{\omega} \Delta \ln(\omega_i^*) + \varepsilon_{id}, \tag{8}$$

where $\gamma_{sd} \equiv (\sigma - 1)\, [\xi_{sd} + \varphi_s - \omega_s - \ln \tau_{sd} + \ln P_{sd}] + \ln E_{sd}$, and where

$$\varepsilon_{id} = (\sigma - 1)\, [u_{id}^{\xi} + u_i^{\varphi} - u_i^{\omega}]. \tag{9}$$

We use equations (6) through (9) to generate predictions for the asymptotic properties of several estimators of the response of log exports to demand-driven changes in log domestic sales. The assumption of constant marginal costs implies that, according to this baseline model, the parameter of interest is zero: demand-driven changes in $\ln R_{id}$ have no causal effect on $\ln R_{ix}$. However, many estimators of the impact of log domestic sales on log exports based on observational data will yield estimates that differ from zero, even in large samples. We discuss here the asymptotic properties of different OLS and IV estimators.

Consider first using OLS to estimate the parameters of the following regression, which includes the change in log domestic sales as an additional covariate in equation (6):

$$\Delta \ln R_{ix} = \gamma_{sx} + (\sigma - 1)\, \delta_{\varphi} \Delta \ln(\varphi_i^*) - (\sigma - 1)\, \delta_{\omega} \Delta \ln(\omega_i^*) + \beta \Delta \ln R_{id} + \varepsilon_{ix}. \tag{10}$$

From equations (7), (9), and (10), the probability limit of the OLS estimator of the coefficient on domestic sales can be written as

$$plim(\hat{\beta}_{OLS}) = \frac{cov(\Delta \ln \mathcal{R}_{ix}, \Delta \ln \mathcal{R}_{id})}{var(\Delta \ln \mathcal{R}_{id})} = \frac{cov(u_{ix}^{\xi} + u_i^{\varphi} - u_i^{\omega}, u_{id}^{\xi} + u_i^{\varphi} - u_i^{\omega})}{var(u_{id}^{\xi} + u_i^{\varphi} - u_i^{\omega})}, \tag{11}$$

where we denote by $\Delta \ln \mathcal{X}$ the residual of a regression of a variable $\Delta \ln X$ on a set of sector fixed effects and the observable covariates $\Delta \ln \varphi_i^*$, and $\Delta \ln \omega_i^*$.

We draw two main conclusions from equation (11). First, as long as changes in productivity and production factor costs are not perfectly observable – and their unobserved component is not fully captured by the sector fixed effects – there will be a positive correlation between changes in exports and changes in domestic sales. Intuitively, unobserved productivity or factor cost changes will affect sales in the same direction in all markets in which a firm sells. In large samples, this will lead $\hat{\beta}_{OLS}$ to be positive and, thus, to be biased upwards as an estimate of the causal impact of demand-driven changes in domestic sales on exports. Second, even when one proxies for changes in productivity and factor costs perfectly (i.e., $u_i^{\varphi} = u_i^{\omega} = 0$), in the presence of a non-zero correlation in the change in residual demand faced by firms in domestic and foreign markets (i.e., $\text{cov}(u_{ix}^{\xi}, u_{id}^{\xi}) \neq 0$), the estimator $\hat{\beta}_{OLS}$ will also converge to a non-zero value. As this residual demand does not capture sector- and market-specific aggregate shocks (which are controlled by the sector fixed effects), it seems plausible that u_{ix}^{ξ} and u_{id}^{ξ} will be positively correlated in the data, leading $\hat{\beta}_{OLS}$ again to be biased upwards.

Consider next using an IV estimator of the parameters in equation (11). Specifically, consider instrumenting $\Delta \ln R_{id}$ with an observed covariate Z_{id} such that Z_{id} is either a proxy for $\Delta \ln \xi_{id}$ or has a causal impact on this firm-specific domestic demand shifter. In this case, the probability limit of the IV estimator of β is

$$plim(\hat{\beta}_{IV}) = \frac{cov(\Delta \ln \mathcal{R}_{ix}, \mathcal{Z}_{id})}{cov(\Delta \ln \mathcal{R}_{id}, \mathcal{Z}_{id})} = \frac{cov(u^{\xi}_{ix} + u^{\varphi}_{i} - u^{\omega}_{i}, \mathcal{Z}_{id})}{cov(u^{\xi}_{id} + u^{\varphi}_{i} - u^{\omega}_{i}, \mathcal{Z}_{id})}, \tag{12}$$

where, as above, we use $\mathcal{Z}_{id}$ to denote the residual from projecting Z_{id} on a vector of sector fixed effects and on the observable covariates $\Delta \ln \varphi^*_i$, and $\Delta \ln \omega^*_i$. The constant-marginal-cost model predicts that $\hat{\beta}_{IV}$ converges in probability to the true zero causal effect of demand-driven changes in domestic sales on exports as long as the variable $\mathcal{Z}_{id}$ satisfies two conditions: (a) it is correlated with the change in domestic sales of firm i after partialling out sector fixed effects as well as observable determinants of the firm's marginal cost; and (b) it is mean independent of the change in firm-specific unobserved productivity, u^{φ}_{i}, factor costs, u^{ω}_{i}, and export demand u^{ξ}_{ix}. As illustrated by the second equality in equation (12), an instrument can only (generically) verify conditions (a) and (b) if its effect on domestic sales works *exclusively* through the component of the change in domestic demand that is not accounted for by the sector fixed effects and the observable covariates included in the estimating equation, i.e., if it works exclusively through u^{ξ}_{id}.

Although our discussion above has centered around the role of unobserved supply and export demand factors in biasing estimates of β, Berman et al. (2015) emphasize that measurement error in both domestic sales and exports constitutes an additional source of possible bias when estimating the effect of exports on domestic sales (or vice versa). Because in many empirical settings – ours included – domestic sales are computed by subtracting exports from the total sales of firms, measurement error in firm total sales and exports will lead to a bias in the OLS estimate $\hat{\beta}_{OLS}$ that is likely to be of the opposite sign to that generated by the unobserved supply and export demand shocks accounted for by the residuals defined in equations (7) and (9). Consequently, as we detail in Appendix E.1 (see also Berman et al., 2015), negative values of $\hat{\beta}_{OLS}$ in large samples may be compatible with firms having constant marginal costs as long as the researcher's measures of either total sales or exports are affected by measurement error. Nevertheless, as we also show in Appendix E.1, if an instrument satisfies the same conditions (a) and (b) outlined above, and is also mean independent of the measurement error in exports, the IV estimator in equation (12) will still converge to zero in the presence of measurement error in total sales and exports.[12]

We have focused our discussion on the intensive margin of exports, namely the impact of domestic demand shocks on the level of exports conditional on exporting. In Appendix E.3, we show that an analysis of the extensive margin of exports modeled as a linear probability model delivers

[12] As mentioned above, in Appendix E.2, we generalize our model to incorporate multiple domestic and foreign markets. Theoretically, firms' choice over multiple export destinations may render instruments $\mathcal{Z}_{id}$ invalid, even if they satisfy the conditions outlined in the main text. However, as the model simulations presented in Appendix E.2.3 illustrate, the resulting potential bias in the corresponding TSLS estimates is for most possible parameter values generally small.

Figure 2: The Great Recession in Spain

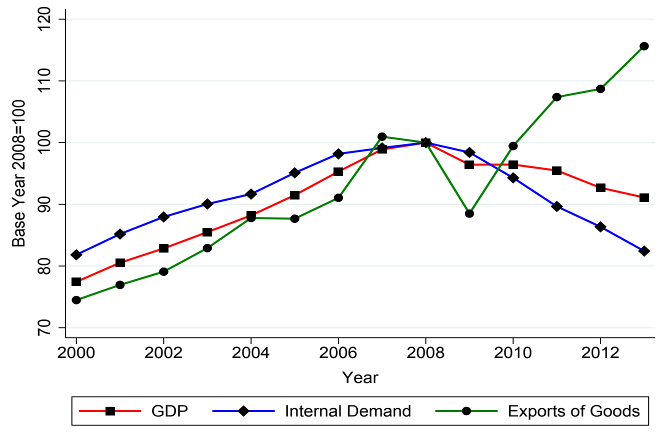

very similar insights. More specifically, when estimating the effect of demand-driven changes in domestic sales on the probability of exporting, even if the true effect were to be zero, one is likely to obtain a positive OLS estimate whenever productivity and production factor costs are not perfectly captured by sector fixed effects and observable controls, or whenever unobserved *firm-specific* demand shocks are positively correlated across the domestic and export markets. An instrument satisfying conditions (a) and (b) above will continue to effectively remove these biases as long as it satisfies the additional condition of being mean independent of the part of the change in the firm's fixed cost of exporting not captured by sector fixed effects and marginal cost proxies (see Appendix E.3 for more details).

3 Setting and Data

To construct a plausibly exogenous measure of the changes in domestic demand faced by firms, we exploit geographical variation in the severity of the impact of the Great Recession of the late 2000s and early 2010s in Spain. In this section, we describe the setting and data, and we defer a more detailed account of our identification strategy to section 4.

3.1 The Great Recession in Spain: Description

The macroeconomic history of Spain during the period 2000-2013 is a tale of a boom followed by a bust. As shown in Figure 2, between the year 2000 and the peak of the cycle in 2008, Spain's GDP and internal demand grew by approximately 20% in real terms.[13] In the five subsequent years until 2013, domestic demand decreased to the level of the year 2000, while real GDP fell by an accumulated 8.9%. In that same period, the unemployment rate shot up from 9% to 26%.

The particularly severe impact of the Great Recession in Spain is largely explained by the fact

[13]Internal demand is defined as final consumption expenditure by households and non-profit institutions serving households (NPISHs) plus investment plus acquisitions of public administrations minus imports.

that the economic boom of the early 2000s was primarily fueled by a real estate bubble. The construction sector accumulated an increasing share of GDP and employment.[14] For instance, in 2006, 658,000 new houses were built in Spain, a number corresponding to 80% of those built in Germany, Italy and the UK *combined*.[15] This real estate boom was in turn fostered by the increased availability of cheap credit to households, firms and real estate developers, which resulted from capital inflows related to the adoption of the euro in 2002 and the global savings glut (Santos, 2014). As a result, the ratio of mortgage credit to GDP went up from 40% in 2000 to 100% in 2008 (Basco and Lopez-Rodriguez, 2018). Importantly, the very high loan-to-value (LTV) ratios associated with residential mortgage credit were partly used by households to finance private consumption, particularly vehicle purchases (Masier and Villanueva, 2011).

The unraveling of the subprime mortgage market in the U.S. in the summer of 2007 had an immediate effect on the supply of credit in Spain. However, the effects were fully transmitted to the real economy only about one year later, coinciding with the fall of Lehman Brothers in September 2008, and the sudden stop in capital inflows (Basco and Lopez-Rodriguez, 2018). The recession officially started in the fourth quarter of 2008, and intensified during 2009 with a 3.6% annual drop in real GDP. The growth in the stock of vehicles in Spain, which had been stable at an average rate of 3.6% a year during the boom, suddenly came to a halt in 2008 (see Figure C.1 in Appendix C). In fact, in 2013, the national stock of vehicles per capita in Spain was *lower* than in 2008 by around 52,000 units.

Importantly for the identification strategy we describe in the next section, the real estate boom and subsequent bust featured significant geographic variation, affecting mainly some parts of the Mediterranean coast and medium-sized and large cities. As we shall document in section 4, this in turn translated into substantial geographic variation in the extent to which the Great Recession affected domestic demand and thus the domestic sales of Spanish firms.

3.2 The Spanish Export Miracle

As Figure 2 illustrates, the evolution of Spain's aggregate merchandise exports during the period 2008-2013 was significantly different from that of aggregate domestic demand. After a significant 11.5% drop in real terms during the global trade collapse of 2008-09, aggregate exports grew during the period 2009-2013 at an even faster rate than during the boom years. Specifically, while exports had grown by an accumulated 34% in the eight-year period 2000-2008, they grew by a very similar 31% in just the four years between 2009 and 2013. This acceleration in export growth occurred at a time during which all indicators of domestic economic activity were showing a significant decline. As a consequence, the fall in real GDP was significantly smaller than the fall in domestic demand, and the ratio of exports of goods to GDP grew from 15.1% in 2009 to 23.33% in 2013. In Appendix D.2, we use the firms in our sample to describe the dynamics of the exports-to-sales ratio by sector.

[14] The share of total employment in the construction sector peaked at 13.5% in the summer of 2007 and then collapsed, reaching 5.4% by early 2014, with a similar pattern for the contribution of this sector to Spain's GDP (12.4% in 2007 and 6.8% in 2014).

[15] See EU Buildings Database at https://ec.europa.eu/energy/en/eu-buildings-database.

One might wonder whether a depreciation in the euro could explain the growth in Spanish exports during the period 2009-2013. Figure 1 in the Introduction shows however that this could not have been the main explanation, as Spanish exports to non-euro area countries clearly outperformed those of other countries in the euro area (even though Spain's GDP dropped faster than the euro area average).[16] It has also been argued that Spain underwent an internal devaluation during this period (through wage moderation starting in 2009, and via a labor market reform in 2012), but there is little evidence that these policies had a significant effect on relative production costs before 2012. For instance, unit labor costs in Spain were only 2.2% lower in 2012 than in their peak in 2009 (OECD Statistics). Conversely, as we document in Appendix D.3, export prices (unit values from product-level export data) in Spanish manufacturing fell relative to export prices in other euro area countries from the onset of the crisis, before Spanish unit labor costs had started to fall.

Motivated by these facts, we will hereafter focus on an exploration of the "vent-for-surplus" mechanism, according to which the domestic slump, by freeing up production capacity, might have *directly* incentivized Spanish producers to sell their goods in foreign markets. More precisely, we hypothesize that the domestic slump led firms to move down along their short-run marginal-cost schedule, thereby allowing them to lower their export prices and gain market share in export markets.[17]

In principle, the associated growth in exports could have materialized along the intensive margin (with continuing exporters increasing their exports) or along the extensive margin (via net entry into the export market). Later in the paper, we will explore both margins, but descriptive evidence suggests that the bulk of the growth was driven by the intensive margin. Using detailed Spanish Customs data, De Lucio et al. (2017a) find that net firm entry (i.e., new exporters net of firms quitting exporting) contributed a mere 14% to the export growth between 2008 and 2013, while the remaining 86% was driven by continuing exporters. Similarly, in our sample of manufacturing firms, we find that continuers contributed 91% of the growth in exports between the boom and the bust periods, and the extensive margin only accounted for 9% of export growth.[18]

3.3 Data Sources

Our data cover the period 2000-2013 and come from various confidential administrative data sources. The first is the Commercial Registry (*Registro Mercantil Central*). It contains the annual financial statements of around 85% of registered firms in the non-financial market economy in

[16] An interesting aspect of Figure 1 is that most of the *relative* take-off of Spain occurred after 2010. The same is not true when looking at Spain's share in overall goods exports (including exports to euro area countries); in that case, Spain's share increased markedly already in 2009. This suggests that the increase in Spanish exports (relative to euro area countries) immediately following the Great Recession was largely driven by increased exports within the EU zone.

[17] This is in contrast with the type of downward shift in marginal costs associated with internal devaluations.

[18] De Lucio et al. (2017a) also show that a third of the contribution of continuing exporters is due to entry into new destination countries and products, while the other two thirds are due to growth in existing product-country combinations. Unfortunately, the nature of the export data available to us does not allow us to explore the firm-level extensive margin at the product or destination country level. See section 3.3 for a description of our data limitations.

Spain.[19] Among other variables, it includes information on the following: sector of activity (4-digit NACE Rev. 2 code), 5-digit zip code of location, net operating revenue, material expenditures (cost of all raw materials and services purchased by the firm in the production process), labor expenditures (total wage bill, including social security contributions), number of employees (full-time equivalent), and total fixed assets.[20]

The second dataset is the foreign transactions registry collected by the Bank of Spain (*Banco de España*). For both exports and imports, it contains transaction-level information on the fiscal identifier of the Spanish firm involved in the transaction, the amount transacted, the product code (SITC Rev. 4), the country of the foreign client, and the exact date of the operation (no matter when the payment was performed). Starting in 2008, however, the dataset's information on the product code and on the destination country became unreliable. The reason for this is that, to save on administrative costs, the entities reporting to the Bank of Spain were given the option of bundling a set of transactions together. In those cases, each entry reflects *only* the country of destination and product code of the largest transaction in that bundle (see Appendix B for more details). This feature of the dataset precludes us from studying exports at the firm-product-destination-year level during the crisis, but we can still reliably aggregate this transaction-level data to obtain information on total export volume by firm and year.

This international trade database has an administrative nature because *Banco de España* legally required financial institutions and external (large) operators to report this information for foreign transactions above a fixed monetary threshold. Until 2007, the minimum reporting threshold was fixed at 12,500 euros per transaction. Since 2008 until the end of the mandatory registry in 2013, information had to be reported for all transactions performed by a firm during a natural year as long as at least one of these transactions exceeded 50,000 euros. In order to homogenize the sample, for the period 2000 to 2007, we only record a positive export flow in a given year for firms that had at least one transaction exceeding 50,000 euros in that year (for more details see Appendix B). The foreign transactions registry collected by the Bank of Spain was discontinued in early 2014, which precludes us from extending our analysis past the year 2013.

In both datasets, a firm is defined as a business constituted in the form of a Corporation (*Sociedad Anónima*), a Limited Liability Company (*Sociedad Limitada*), or a Cooperative (*Cooperativa*). We merge both datasets using the fiscal identifier of each firm. Using the merged database, we define each firm's domestic sales as the difference between its total annual sales and its total export volume (see section 2 and Appendix E.1 for a discussion of measurement error in these variables).

[19] We obtain information on the Commercial Registry from two different sources: (i) the *Central de Balances* dataset, compiled by the Bank of Spain, and (ii) the *Sabi* dataset, compiled by Informa (a private company). For details on how we combine these two datasets, see Almunia et al. (2018).

[20] NACE (*Nomenclature générale des activités économiques dans les Communautés Européennes*) is the European statistical classification of economic activities. It classifies manufacturing firms into 24 different sectors. Some firms move to a different zip code or change their sectoral classification during the period of analysis. We assign to these firms a fixed zip code and sector using their most frequent value in each case. A firm's zip code corresponds to the location of its headquarters.

Figure 3: Output, Employment, Wage Bill and Export Dynamics

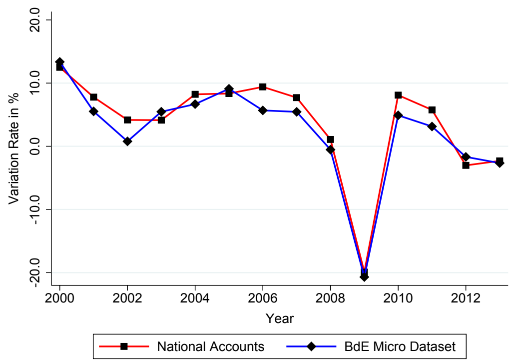

Panel (a): Output

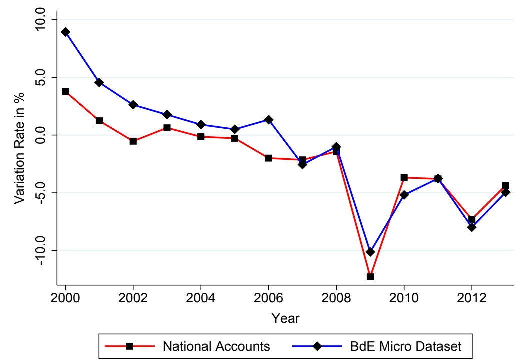

Panel (b): Employment

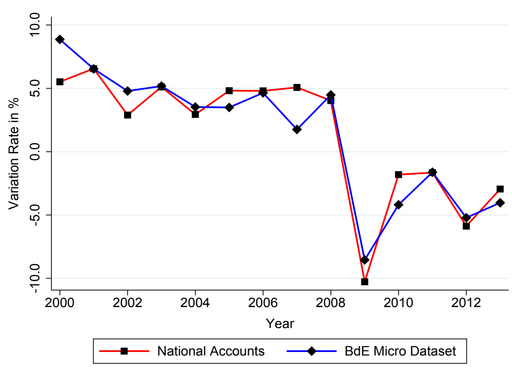

Panel (c): Wage Bill

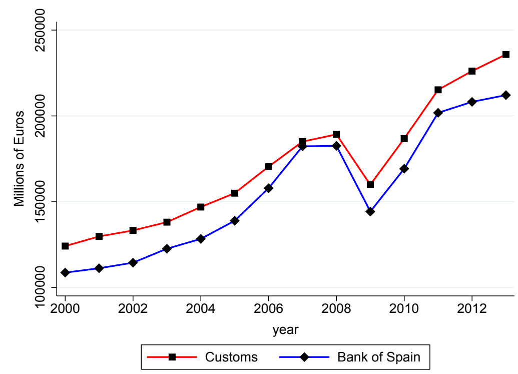

Panel (d): Exports

To confirm the validity of the information contained in the resulting dataset, we compare its coverage with the official publicly available aggregate data on output, employment and total wage bill (from National Accounts) and on goods exports (from Customs). Figure 3 shows that our dataset tracks nearly perfectly the aggregate evolution over time of output, employment, total payments to labor, and exports. Due to the reporting thresholds described above, aggregate exports in our sample naturally fall a bit short of aggregate exports in the Customs data, but note that the gap is very similar in the boom and bust periods (the average coverage is 91.8% in 2000-08 and 91.3% in 2009-13).[21]

We complement the firm-level data described above with yearly municipality-level data on the stock of vehicles per capita. The information on the stock of vehicles by municipality is provided by the Spanish Registry of Motor Vehicles, compiled by the General Directorate of Traffic (*Dirección General de Tráfico*), while the information on the population by municipality is provided by

[21] Most of the gap in coverage is explained by the fact that a nontrivial share of Spanish exports recorded by Customs is carried out by legal entities or individuals that are not registered as firms undertaking economic activity in Spain, and are thus exempted from submitting their financial statements to the Commercial Registry. The share of goods exports by non-registered entities was on average around 8% in 2010-2013 (own calculations based on public Customs data).

the Spanish National Statistical Office (*Instituto Nacional de Estadística*). When matching this municipality-level data with our firm-level data, we need to deal with the fact that the information on the location of firms is provided at the zip code level, and that the mapping between municipalities and zip codes is not one-to-one. More precisely, larger municipalities are often assigned multiple zip codes and, in a very small number of cases, a single zip code is assigned to more than one municipality. In the former case, we associate the same value for the stock of vehicles per capita to all firms located in the same municipality, independently of the zip code of location; for firms in zip codes containing multiple municipalities, we associate with them a stock of vehicles per capita constructed as an average the stock of vehicles per capita across these municipalities.

We also employ data on firm-level sales within Spain.[22] We obtained this data from the Spanish Tax Agency (*Agencia Estatal de Administración Tributaria*, AEAT), which collects this data from all firms (legal entities) and professionals (natural persons) that undertake economic activities in Spain (see more details in Appendix B). The AEAT was willing to share one year of data with us, so we work with data for the year 2006 because it is the first year for which a comprehensive digitization of the data is available. We obtained two datasets from the AEAT: first, aggregate data on municipality-to-municipality flows for *all* firms in the manufacturing sector, excluding sales of entitites in the auto industry; and second, firm-to-municipality sales for those manufacturing firms that exported in the boom as well as in the bust.

When exploring the robustness of our results, we use information on additional variables. The underlying sources for these variables are discussed in Appendix B.

4 Identification Approach

In this section, we first describe our identification approach, and later highlight various potential threats affecting this strategy and how we seek to address them.

4.1 A Geography-based Proxy of Demand Changes

As explained in section 3.1, a key characteristic of the Great Recession in Spain is that it affected different regions differently. Panel (a) in Figure 4 illustrates this fact. The figure plots the standardized percentage change in domestic sales for the average manufacturing firm located in each of the 47 Spanish peninsular provinces and operating in at least one year of the boom period (2002-2008) and at least one year of the bust period (2009-2013).[23] The provinces where the average firm experienced a reduction in domestic sales smaller than the national average are in darker color, while those where the average firm experienced a larger reduction in domestic sales are in lighter color. Specifically, Figure 4 illustrates that firms located in the Northern and Western regions saw

[22] We thank Francesco Serti for having brought to our attention the existence of these data.

[23] Figure C.2 in Appendix C.2 shows the yearly average number of firms and exporters by province for the period 2002-2008. Economic activity in Spain is concentrated mostly in the coast (Galicia, Basque Country, Catalonia, Valencian Community, Murcia and Andalusia) and in the center (Madrid). Exporting firms are concentrated in the center (Madrid) and in the Mediterranean coast (Catalonia and Valencian Community).

Figure 4: The Great Recession in Spain: Variation Across Provinces

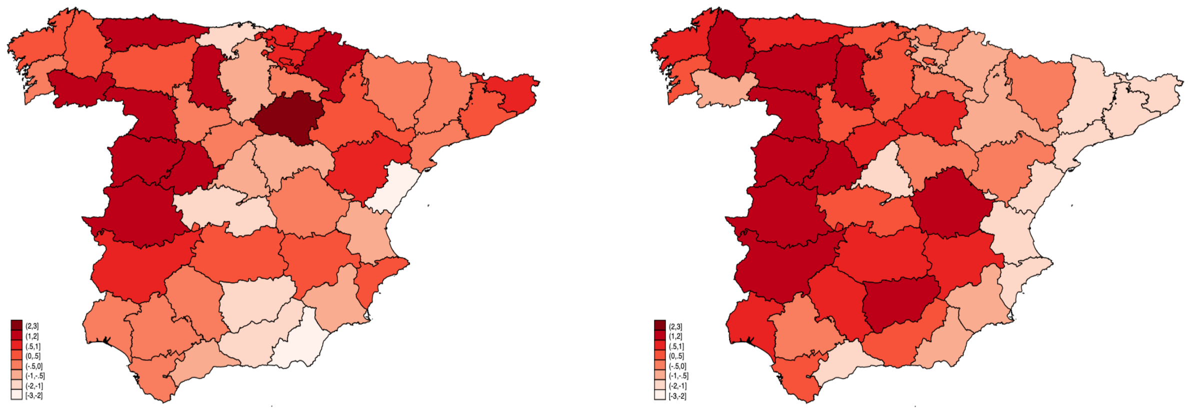

(a) Relative Change in Domestic Sales (b) Relative Change in Cars per Capita

Notes: For each province, panel (a) illustrates the standardized percentage change in average firm-level domestic sales between the period 2002-2008 and the period 2009-2013, where the average is computed across manufacturing firms active in at least one year in both periods. Therefore, if this variable takes value p for a given province, it means that the average firm located in this province experienced a relative change in average yearly domestic sales between 2002-2008 and 2009-2013 that was p standard deviations above the change experienced by the average province. Panel (b) illustrates the standardized percentage change in vehicles per capita between the period 2002-2008 and the period 2009-2013. Therefore, if this variable takes value p for a given province, it means that this province experienced a relative change in vehicles per capita between 2002-2008 and 2009-2013 that was p standard deviations above the change experienced by the average province.

changes in domestic sales larger (less negative) than the average, while firms located in the center of the country and in Southern and Eastern regions experienced relatively large domestic sales reductions. Furthermore, deviations from the national average are sizable in many cases.

The heterogeneity in the changes in domestic sales that we document in panel (a) of Figure 4 could have been caused by heterogeneity in supply factors or by heterogeneity in factors affecting local demand for manufacturing goods. We next propose an approach to measuring variation in local demand for manufacturing goods.

Our approach consists in proxying changes in local demand for manufacturing goods using observed changes in demand per capita for one particular type of manufacturing products: vehicles. Panel (b) in Figure 4 shows that there is substantial variation in the degree to which the number of vehicles per capita changed across provinces between the boom and the bust years.[24] Specifically, the provinces in the Northwest and in the Southwest experienced a relative increase in the number of vehicles per capita, while the region around Madrid and the provinces in the Northeast and along the Mediterranean coast experienced a relative reduction. As in panel (a) of Figure 4, the regional deviations from the national averages in panel (b) are large for many provinces.

By illustrating provincial averages, the maps in Figure 4 hide substantial spatial variation at the sub-province level (across 5-digit zip codes) in both the boom-to-bust changes in average firm-level

[24]Changes in the stock of vehicles per capita between the boom and the bust years could have been due either to purchases of new vehicles or to scrapping of old ones. We measure the change in the stock, rather than just new purchases, to avoid our measure of domestic demand for manufacturing firms from being contaminated by the effect of the "cash for clunkers" program (*Plan PIVE*) that the Spanish government put in place during the bust period.

domestic sales and in the boom-to-bust changes in the number of vehicles per capita. We illustrate this variation in Figure C.3 (see Appendix C.3) for the case of the two most populated provinces in Spain: Madrid and Barcelona.

Our core empirical strategy exploits the variation illustrated in Figures 4 and C.3 to identify the impact of domestic demand shocks on firms' exports operating through its effect on the firms' domestic (Spain-wide) sales. Specifically, we divide our sample into a "boom" period (2002-08) and a "bust" period (2009-13), and assess the extent to which a demand-driven decline in a firm's domestic sales in the bust period relative to the boom period is associated with a relative increase in its export sales between these two periods. We choose this 'long-differences' approach to avoid having to take a stance on the precise lag structure of the effect of domestic demand on exports, and also because the macroeconomic evidence in Figure 3 and the time series of the national stock of vehicles per capita in Figure C.1 cleanly identify the year 2009 as the break between two distinct periods. Having said this, we will demonstrate in section 6 that our results are quantitatively robust to other classifications of the boom and the bust periods, and also remain qualitatively similar when breaking the sample into four subperiods (2002-05, 2006-08, 2009-11, 2012-13).

To build a measure of the boom-to-bust change in domestic demand for each Spanish firm, we follow a two-step procedure. First, we use observed boom-to-bust changes in the stock of vehicles per capita at the municipality level as a proxy for the boom-to-bust changes in the demand for manufacturing goods in those municipalities. For this, we rely on a body of work documenting that durable goods consumption, and vehicle purchases in particular, are strongly procyclical and thus are a useful proxy for changes in 'local demand', i.e., the overall propensity of an area's inhabitants to consume (see Stock and Watson, 1999). Consistent with this notion, Mian et al. (2013) document how variation in the extent to which the U.S. subprime mortgage default crisis of 2007-10 affected household housing wealth in different areas in the United States translated into geographical variation in vehicle purchases.[25] There is also a broader literature documenting the large impact that housing wealth effects had on consumption more generally in the years around the Great Recession; e.g., Guren et al. (2020), Kaplan et al. (2020).

Second, with this measure of local demand at hand, we next take into account the different exposure of firms located in different municipalities to these local demand shocks in order to construct our baseline measure of the boom-to-bust change in domestic demand for each Spanish firm. To measure the different exposure of each firm to local demand shocks, we rely on information on the location of these firms together with municipality-to-municipality trade flows data (aggregated across *all* manufacturing firms) from our tax record data for the year 2006. We use this data to estimate municipality-to-municipality gravity regressions, and use the estimated coefficients for

[25]It would be interesting to tie the geographical variation in the change in the stock of vehicles per capita in Spain to the housing slump, but idiosyncratic features of the Spanish housing market complicate such an analysis (see Akin et al., 2014). In Appendix H.2, we revisit this issue and explore the robustness of our results to an alternative shifter of firms' domestic sales that uses a determinant of the housing supply elasticity in a given municipality as a proxy for the magnitude of the negative impact of the Great Recession on household housing wealth and consumption. This instrument is somewhat weak, but our second-stage elasticities are remarkably close to those obtained under our baseline instrument.

log population and log distance to back out the relevant weights needed to construct our baseline measure of firm-level exposure to local shocks. As a robustness check, we also exploit below our information on *firm-level* sales across destinations within Spain to estimate firm-to-municipality gravity equations and compute an alternative set of weights.

Estimates of this municipality-to-municipality gravity equation for Spanish manufacturing flows in the year 2006 are presented in columns 1 through 3 of Table 1. The first column presents a parsimonious specification with municipality of origin fixed effects, log population of the municipality of origin, and log distance between origin and destination. The results illustrate the relevance of gravity forces, with shipments declining with distance with an elasticity of -0.429. The inclusion in column 2 of dummies for own-municipality and own-province flows slightly reduces this distance elasticity, while these two dummies appear to have themselves a positive and significant effect on shipments. This suggests that part of the negative effect of distance on within-Spain municipality-to-municipality shipments is related to a discontinuous fall in shipments at the municipality border and at the province border. The extent of "home bias" at the municipality level is remarkably large: it implies that, *ceteris paribus*, shipments are $\exp(1.607) \approx 5$ times larger within a municipality than outside. The existence of such strong local home bias is in line with the findings of Hillbery and Hummels (2008) for the U.S., although the magnitude of this home bias is larger in our setting.[26] Column 3 presents estimates for a third specification analogous to that in column 1 but employing a set of distance dummies to capture the effect of distance on sales. The results corroborate the fact that, controlling for municipality of origin fixed effects and for the population of the municipality of destination, sales decay monotonically with distance.

We combine the estimates in Table 1 with data both on the distance between any two municipalities i and j and on the population of each possible municipality of destination j to predict the share of domestic (within-Spain) sales that any firm located in municipality i will sell in municipality j. We use these shares to compute, for each municipality of origin i, a weighted sum of the local demand level (i.e., vehicles per capita) in all destinations j *other than* i. We do so for the boom and for the bust periods (taking the average number of vehicles per capita within each period as our period-specific measure of demand in each municipality) and then we take the log difference of these period-specific average demand levels as our proxy for the demand shock that manufacturing firms located in any municipality i experienced between the boom and the bust period. These demand proxies play a key role in our identification approach, as they will serve in section 5 as instruments for firm-specific domestic sales in the estimation of regression equations analogous to that in equation (10); i.e., the proxy for municipality-specific boom-to-bust demand

[26] The analysis in Hillbery and Hummels (2008) is at the zip-code level, which admittedly makes the comparison of its results with those in our municipality-level analysis imperfect. Díaz-Lanchas et al. (2013) also estimate a very sizeable 'zipcode effect' using a micro-database based on a random sample of shipments by road within Spain during the period 2003-2007 (C-Intereg). Although we do not have access to this micro-database, we have obtained province-to-province shipments from that database. In Appendix D.4, we compare some aggregate statistics on firms' within-Spain sales from our sales data based on tax records and from the C-Intereg dataset. The extent of provincial home bias is very similar in both datasets.

Table 1: Estimates from Gravity Equations at Municipal Level

Dependent Variable:	Ln(Bilateral Trade Flows between Municipalities)				
	(1)	(2)	(3)	(4)	(5)
Ln(Population)	0.493^a	0.490^a	0.485^a	0.300^a	0.322^a
	(0.031)	(0.031)	(0.030)	(0.012)	(0.015)
Ln(Distance)	-0.429^a	-0.378^a		-0.150^a	-0.145^a
	(0.011)	(0.019)		(0.021)	(0.019)
Dummy for own-municipality flows		1.607^a			
		(0.111)			
Dummy for own-province flows		0.131^b			
		(0.065)			
Dummy for distance 5-10Km			-1.159^a		
			(0.079)		
Dummy for distance 10-50Km			-1.924^a		
			(0.090)		
Dummy for distance 50-100Km			-2.465^a		
			(0.114)		
Dummy for distance 100-200Km			-2.718^a		
			(0.113)		
Dummy for distance 200-500Km			-2.962^a		
			(0.120)		
Dummy for distance 500-1000Km			-3.237^a		
			(0.120)		
Dummy for distance >1000Km			-3.590^a		
			(0.141)		
Observations	417,936	417,936	417,936	675,715	675,589
R-squared	0.30	0.31	0.30	0.15	0.31
Municipality of origin FE	Yes	Yes	Yes	Yes	No
Sector FE	No	No	No	Yes	No
Firm FE	No	No	No	No	Yes

Note: [a] denotes 1% significance, [b] denotes 5% significance, [c] denotes 10% significance. Standard errors clustered at the province (of origin) level are reported in parenthesis. The data on municipality-level trade flows for manufacturing firms is for the 2006 fiscal year. Ln(Population) denotes the log of the population of the destination municipality in 2006. Ln(Distance) denotes the log of the distance, in kilometers, between the two municipalities in each pair. The estimates in columns 1 to 3 use municipality-to-municipality sales data; the estimates in columns 4 and 5 use firm-to-municipality data.

shocks described here will play the role of the instrument Z_{id} introduced in section 2.1. [27]

Although our baseline instrument builds on aggregate municipality-to-municipality trade flows, as mentioned in section 3.3 the Spanish tax authority also shared with us 2006 data on firm-

[27] Our instrument is inspired by the "market access" term in Redding and Venables (2004). There are however two important differences. First, we measure the change in demand in a location not as the change in a destination fixed effect estimated from a gravity equation, but as the change in the stock of vehicles per capita in that location. Second, as the value of our instrument for a municipality excludes the change in vehicles per capita in such municipality, our instrument excludes what Redding and Venables (2004) denote the domestic market access term. The first difference is a consequence of the fact that we observe within-Spain flows only for the year 2006; the second one aims to make our instrument more plausibly exogenous. Similar market access terms have been defined in Harris (1954), Hanson (2005) and Donaldson and Hornbeck (2016).

destination level shipments within Spain for a sample of around $8,000$ continuing exporters (i.e., firms that exported both in the boom as well as in the bust). We use this data in columns 4 and 5 of Table 1 to run gravity equations at the firm-destination level. For simplicity, we focus in column 4 on the parsimonious specification in column 1, extended in column 5 to account for firm fixed effects. The results are in line with those in column 1, but with somewhat smaller log population and log distance coefficients, as one would expect given that these specifications do not account for extensive margin variation in the set of municipalities firms sell to. As a robustness check, we present in section 6.3 results that rely on instruments analogous to that described above, but built using predicted sales shares based on the estimates of these firm-to-municipality gravity specifications.[28]

4.2 Threats to Validity of the Instrument

The main concern with our identification approach is that our municipality-level measure of demand changes between boom and bust might be correlated with changes in supply shocks affecting the firms located in the corresponding municipalities. This exclusion restriction is central to the validity of our strategy, so we next outline how it might be violated and how we deal with these potential threats to identification.

First, we control in our specifications for sector fixed effects. Thus, we base our identification on observing how domestic sales and exports changed between the boom and the bust for different firms operating in the same sector but located in regions experiencing different exposure to local demand changes. More specifically, by controlling for sector fixed effects we control for sector-specific foreign demand shocks, sector-specific trade cost shocks, and domestic supply shocks affecting Spanish firms (see definition of γ_{sx} in equation (6)). For example, these sector fixed effects control for shocks such as the expiration of the Multi Fiber Arrangement on January 1, 2005, which eliminated all European Union quotas for textiles imported from China and which increased the competition that Spanish textile manufacturers faced both in the domestic and foreign markets.[29]

Admittedly, sector fixed effects may not effectively control for all heterogeneity across firms in their export demand shocks; specifically, firms located in different Spanish regions may be differentially affected by export demand shocks even if they operate in the same sector. A possible source of this heterogeneity in demand shocks is the different exposure of firms located in different Spanish regions to changes in demand in different foreign countries (e.g., firms located in Southern regions are more exposed to demand changes in Northern African countries than firms located in the north of Spain); in Appendix E.4, we provide suggestive evidence that this type of heterogeneity in export demand shocks is not, in practice, damaging for our instrument.

[28]We have also estimated gravity equations that use the firm-to-destination data and that include the border and distance dummies introduced in columns 2 and 3. Notably, we again find a remarkably large coefficient of 1.308 for the own-municipality dummy variable.

[29]If sector fixed effects had not been included in our specifications and textile firms were to be on average located in Spanish regions that suffered larger negative local demand shocks, our estimates would confound the impact of the Multi Fiber Arrangement expiration and the negative local shocks.

Second, as different firms operating in the same sector may experience different supply shocks, we also control for firm-specific measures of productivity and labor costs. By controlling for changes in wages and productivity at the firm level, we aim to identify the effect that changes in local demand had on firms' exports through channels other than the internal devaluation channel. More specifically, these controls help address the concern that the reduction in unit labor costs observed in Spain during the period 2009-13 might have been heterogeneous across different Spanish regions in a manner that is correlated with our demand measure.

Our third approach to assuage endogeneity concerns is motivated by the fact that our various fixed effects and proxies for firm-level productivity and wage costs might not perfectly capture supply-side factors, and that unobserved *residual* supply shocks might be correlated with our proxy for changes in local demand. For instance, this will be the case if changes in the stock of vehicles per capita in a municipality are not entirely due to changes in the supply of credit to households but also partly due to supply shocks affecting either the firms located in such municipality or in neighboring municipalities. This would be the case if these supply shocks (e.g., changes in labor payments not properly controlled for by our measure of firm-level wages) impact the purchasing power of consumers living in the corresponding municipalities. To deal with this threat to identification, in all regressions presented in this paper, we do not use information on the change in the number of vehicles per capita in the municipality of location of a firm when constructing our baseline instrument, which is thus a weighted sum of demand levels in municipalities *other than the one* where the firm is located. In section 6, we also present estimates from regression specifications in which we control for additional measures of municipality-specific boom-to-bust changes in economic conditions (e.g., changes in financial costs, changes in the share of temporary workers, etc.).

A fourth concern relates to the presence of a nontrivial number of car manufacturers in our sample of Spanish manufacturing firms. These firms' supply shocks are especially likely to have impacted the boom-to-bust changes in the stock of vehicles per capita in their own municipality and in geographically close ones. More specifically, if a disproportionate share of cars in Spain was sold in municipalities that are geographically close to where the car was manufactured, supply shocks in these firms may affect the number of vehicles per capita not just by affecting the purchasing power of consumers in certain municipalities, but by affecting directly the supply of cars in those municipalities. Roughly three quarters of all cars purchased in Spain are imported (as indicated by data from the Spanish National Institute of Statistics); thus, supply shocks affecting car manufacturers are likely to have a limited impact on the total amount of cars in Spain. To deal with this threat to identification, we exclude all firms operating in the auto industry (NACE Rev. 2 code 29) in all the regressions we present. Additionally, in section 6, we also explore how our results are impacted when excluding from our sample: (i) all firms located in a zip code that hosts at least one firm in the auto industry employing more than 20 workers; (b) all firms located in a zip code or in the neighborhood of a zip code with a significant share of manufacturing employment accounted for by the auto industry; and (c) all firms producing in sectors that are either leading input providers or leading buying industries of the vehicles manufacturing industry.

Finally, it is important to remark that, as illustrated in equation (12), any unobserved factor costs that are negatively correlated with our instrument will cause our IV estimator to be positively biased. For example, if the tightening of the credit supply in a region caused firms' marginal production costs to increase and consumers' demand to fall, the resulting endogeneity of our instrument would bias our IV estimator upwards. Thus, negative IV estimates of the elasticity of firm-level exports with respect to a firm's domestic sales would still reflect patterns in the data that would be inconsistent with the constant marginal cost model described in section 2.

5 Baseline Results

In this section, we present our baseline results on the impact of demand-driven changes in domestic sales on firms' behavior in the export market. Specifically, in section 5.1, we present evidence on the impact of the Great Recession on Spanish firms' intensive margin of exports. In section 5.2, we present analogous evidence of its impact on the extensive margin.

5.1 Intensive Margin

Table 2 presents OLS estimates of the elasticity of boom-to-bust changes in firms' export flows with respect to boom-to-bust changes in domestic sales for continuing exporters – i.e., firms that exported both in the boom as well as in the bust. There are $8,009$ such firms in our dataset.

As discussed in section 2, when no instrument for the change in domestic sales is used in the estimation, even in a model with constant marginal costs, unobserved (residual) supply factors tend to make the OLS estimate of a firm's change in foreign sales on its change in domestic sales positive. Conversely, measurement error in both total sales and exports tends to make this OLS estimate negative. As illustrated in column 1 of Table 2, when no controls are included, we estimate an OLS elasticity of export flows with respect to domestic sales that is very close to zero. Consistently with the expected biases in the OLS estimator, as we control for various sources of marginal cost heterogeneity across firms in the remaining columns of Table 2, the OLS estimates become negative. Specifically, we control in column 2 for the change in firms' productivity (estimated following the procedure in Gandhi et al., 2016, as detailed in Appendix F), and in column 3 for the change in the firm's average wages (reported by the firm in its financial statement). Consistent with the discussion in section 2, controlling for these supply shocks reduces the OLS estimate of the coefficient on domestic sales. In fact, the coefficient turns negative (-0.298), indicating that, once we control for the observable part of firms' supply shocks, domestic sales and exports are negatively correlated. Columns 4, 5 and 6 aim to control for additional unobserved determinants of firms' marginal costs that are time varying. To do so, and motivated by the specification in equation (10), we sequentially add sector fixed effects (in column 4) and location fixed effects (in columns 5 and 6). In the latter case, we first include province fixed effects and, in column 6, we instead include municipality fixed effects. The resulting estimates continue to be negative and indicate that a 1% decrease in a firm's domestic sales, keeping its productivity and average wages constant, is

Table 2: Intensive Margin: Ordinary Least Squares Estimates

Dependent Variable:	ΔLn(Exports)					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔLn(Domestic Sales)	0.063	-0.209^a	-0.298^a	-0.292^a	-0.284^a	-0.271^a
	(0.044)	(0.049)	(0.043)	(0.032)	(0.032)	(0.036)
ΔLn(TFP)		1.142^a	1.448^a	1.535^a	1.522^a	1.514^a
		(0.043)	(0.046)	(0.057)	(0.055)	(0.057)
ΔLn(Average Wages)			-0.744^a	-0.723^a	-0.712^a	-0.706^a
			(0.062)	(0.072)	(0.070)	(0.067)
Observations	8,009	8,009	8,009	8,009	8,009	7,502
R-squared	0.001	0.100	0.126	0.162	0.171	0.278
Sector FE	No	No	No	Yes	Yes	Yes
Province FE	No	No	No	No	Yes	No
Municipality FE	No	No	No	No	No	Yes

Notes: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered at the province level are reported in parenthesis. For any X, $\Delta\text{Ln}(X)$ is the difference in $\text{Ln}(X)$ between its average in the 2009-2013 period and its average in the 2002-2008 period. The estimation sample includes all firms exporting in at least one year in the period 2002-2008 and in the period 2009-2013.

associated with close to a 0.3% increase in its overall export flows.

In Table 3, we turn to our baseline two-stage least squares (TSLS) estimates of the elasticity of the firm's boom-to-bust change in exports with respect to its boom-to-bust demand-driven change in domestic sales. As discussed in section 2, if the instrument is orthogonal to both unobserved supply factors and to the measurement error in both total sales and exports, the corresponding TSLS estimator should converge to zero in large sample, as long as firms' marginal costs are constant. The first-stage estimates (reported in columns 1 to 4 and plotted in panel (a) of Appendix Figure C.4) reveal that firms located in municipalities that experienced a larger drop in the distance- and population-weighted municipality-level stock of vehicles per capita also suffered a larger decline in their domestic (Spain-wide) sales. This relationship is robust to controlling for our measures of firms' changes in productivity and labor costs and for sector and province fixed effects: the statistic of an F-test for the null hypothesis that the change in the municipality-specific average of the stock of vehicles per capita across all other municipalities that we use as our instrument has no impact on the domestic sales of the firms located in that municipality is in all specifications above threshold values generally applied to detect weak instrument problems, the only exception being the value of 7.85 in column 1.

The second-stage estimates (reported in columns 5 to 8) reveal elasticities of exports with respect to domestic sales that are significantly larger (in absolute value) than the OLS elasticities reported in Table 2.[30] This is true regardless of whether one controls for sector and province fixed effects as well as for changes in our measures of the firm's productivity and labor costs. Our

[30]We illustrate the reduced-form relationship between the log change in firms' exports and the log change in the demand proxy we use as our baseline instrument in panel (b) of Appendix Figure C.4.

Table 3: Intensive Margin: Two-Stage Least Squares Estimates

Dependent Variable:	ΔLn(Domestic Sales)				ΔLn(Exports)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔLn(Domestic Sales)					-10.068^a	-2.081^a	-1.751^a	-1.607^a
					(3.454)	(0.319)	(0.238)	(0.248)
ΔLn(Dist-Pop-Weighted	0.339^a	1.194^a	1.346^a	1.312^a				
Vehicles p.c.)	(0.121)	(0.145)	(0.135)	(0.119)				
ΔLn(TFP)		0.829^a	1.031^a	1.023^a		2.623^a	2.876^a	2.810^a
		(0.028)	(0.029)	(0.028)		(0.241)	(0.222)	(0.213)
ΔLn(Average Wages)			-0.621^a	-0.526^a			-1.620^a	-1.387^a
			(0.037)	(0.047)			(0.174)	(0.151)
F-statistic	7.85	67.47	99.84	122.44				
Observations	8,009	8,009	8,009	8,009	8,009	8,009	8,009	8,009
Sector FE	No	No	No	Yes	No	No	No	Yes

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered by province appear in parenthesis. For any X, $\Delta\text{Ln}(X)$ is the log difference between the average of X in 2009-2013 and its average in 2002-2008. ΔLn(Dist-Pop-Weighted vehicles p.c.) is the instrument constructed using data on vehicles per capita at the municipal level and applying the weights from the gravity equation reported in column 1 of Table 1. Columns 1-4 contain first-stage estimates; columns 5-8 contain second-stage estimates. F-statistic denotes the corresponding test statistic for the null hypothesis that the coefficient on Ln(Dist-Pop-Weighted Vehicles p.c.) equals zero.

preferred estimate in column 8 indicates an elasticity of exports with respect to domestic sales of around -1.6.[31] It is clear that this elasticity is significantly more negative than the OLS one, which we take to be a validation of the hypothesis, formalized in equation (11), that, even after controlling for sector fixed effects and for firm proxies of productivity and average labor costs, there still remains substantial unobserved determinants of firms' marginal costs that induce a spurious positive correlation between their sales in the domestic and foreign markets. Also, the threats to the internal validity of the instrument notwithstanding (see section 4.2 for a detailed account of these threats), the fact that the TSLS estimate is negative is suggestive of the firm's marginal cost function not being flat.

One might be concerned that, because firms' total sales are a key input in the computation of our TFP measure, our empirical results are just unveiling a mechanical negative correlation between exports and domestic sales once one holds total sales revenue constant (by controlling for it). Although log TFP and log total sales are obviously positively correlated (as one would expect in light of our model), the correlation is far from perfect, particularly when considering log changes in these variables. More specifically, the correlation between log changes in TFP and log changes in total sales in data is 0.31 at the yearly level, while it is 0.56 when looking at boom-to-bust 'long differences' in these variables. To further assuage this concern, in section 6.4 we explore the robustness of our results to an alternative measure of log firm TFP computed using value-added that features a much lower correlation with log firm sales.

[31] The unrealistically high point estimate and the corresponding large standard error reported in column 5 should be discounted on the basis that, as shown in column 1, our instrument is weak in this regression specification.

In terms of the quantitative relevance of our results, it is worth emphasizing that an elasticity of -1.6 does not necessarily imply a more-than-complete substitution of exports for domestic sales. For a firm with an initial export share of $100 \times \chi\%$, a demand-driven drop of €100 in their domestic sales would lead to a €$160 \times (\chi/\,(1-\chi))$ increase in exports. For example, for every €100 of lost domestic sales, a firm with an export share of 25% would able to recoup €53.3 via exports, while a firm with an export share of one-third would be able to recoup €80.[32]

In terms of the statistical significance of our results, it is important to remark that, unless otherwise noted, all standard errors presented in this paper are clustered by province. It is worth highlighting that although our instrument has a shift-share flavor, it does not strictly fall in the class of shift-share instruments studied recently by Adão et al. (2018), Borusyak et al. (2019) and Goldsmith-Pinkham et al. (2020), among others. The reason for this is that our instrument does not simply use changes in a weighted-average of vehicles per capita across municipalities but log changes in such weighted average.[33] Although the functional form of our instrumental variable prevents us from computing the standard errors according to the formulas introduced in Adão et al. (2018), it is conceivable that our standard error estimates may suffer from the downward bias that, as these authors illustrate, typically affect clustered standard errors when the instrument is of the shift-share type. We revisit this question in section 6.3, where we provide suggestive evidence showing that the bias affecting our standard errors, if present, is likely to be very small.

5.2 Extensive Margin

We next turn to studying the causal impact of demand shocks on the extensive margin of exporting. As in our intensive margin regressions, we divide the sample period into a boom (2002-08) and a bust period (2009-13), and explore how demand-driven changes in domestic sales affect firms' probability of exporting in each of these two periods. More specifically, we compute a TSLS estimates of a linear probability model in which a firm's dummy capturing positive exports in a given period (boom or bust) is regressed on firm, sector-period and province-period fixed effects, the log of firm-level average TFP in that period, the log of average wages in that period, and the log of average domestic sales in that period, with log domestic sales in a given period instrumented with the log of the same weighted sum of the stock of vehicles per capita around a firm's municipality as in Table 3.[34] Besides this linear probability model, we also estimate analogous specifications in which we substitute the dependent variable by a variable capturing the proportion of years in a given period (boom or bust) for which a firm exports.

The results for the $62,527$ firms in our sample that are active in the domestic market in both the boom and the bust are presented in Table 4. Column 1 reports the first-stage estimates. As in

[32] The median export share among the 8,009 firms exporting in both boom and bust periods is 16.2%.

[33] The reason for employing log changes to build our instrument is that, according to the multi-destination model in Appendix E.2, log changes in domestic sales are linearly related to log changes in a weighted average of local demand shifters across the different domestic markets.

[34] Our results are similar under an alternative specification in which the left-hand-side variable is a dummy variable that treats a firm as an 'exporter' only if it exports for two or more years in a given period.

Table 4: Extensive Margin: Two-Least Squares Estimates

Dependent Variable:		Export Dummy		Proportion of Years	
	1st Stage	OLS	2nd Stage	OLS	2nd Stage
	(1)	(2)	(3)	(4)	(5)
Ln(Domestic Sales)		0.021[a]	-0.099[a]	0.008[b]	0.040[b]
		(0.005)	(0.034)	(0.004)	(0.019)
Ln(Dist-Pop-Weighted Vehicles p.c.)	1.024[a]				
	(0.110)				
Ln(TFP)	1.169[a]	0.068[a]	0.204[a]	0.062[a]	0.024
	(0.018)	(0.007)	(0.039)	(0.005)	(0.020)
Ln(Average Wages)	-0.589[a]	-0.046[a]	-0.114[a]	-0.041[a]	-0.022[b]
	(0.015)	(0.007)	(0.022)	(0.004)	(0.010)
Observations	125,054	125,054	125,054	125,054	125,054
R-squared	0.983	0.842	-0.017	0.920	0.012
Firm FE	Yes	Yes	Yes	Yes	Yes
Sector-Period FE	Yes	Yes	Yes	Yes	Yes
F-statistic	86.32				

Note: [a] denotes 1% significance, [b] denotes 5% significance, [c] denotes 10% significance. Standard errors clustered by province reported in parenthesis. Ln(Dist-Pop-Weighted vehicles p.c.) is the instrument constructed using data on vehicles per capita at the municipal level and applying the weights from the gravity equation reported in column 1 of Table 1. F-statistic denotes the corresponding test statistic for the null hypothesis that the coefficient on Ln(Dist-Pop-Weighted Vehicles p.c.) equals zero. All specifications include firm fixed effects, sector-period fixed effects, and province-period fixed effects. The estimation sample includes all firms selling in the domestic market in at least one year in the period 2002-2008 and in the period 2009-2013.

Table 3, the results indicate that domestic sales fell more for firms located in municipalities that, according to our measure, experienced a larger drop in domestic demand. The F-stat (86.32) is, as in our intensive margin specifications, well above standard threshold values. Columns 2 and 3 then present OLS and IV estimates of the link between domestic sales and export status, while columns 4 and 5 report OLS and IV estimates of the link between domestic sales and the proportion of years exported. The results of these two specifications deliver statistically significant estimates of opposite signs but, in both cases, quantitatively very small. First, the OLS estimates in columns 2 and 4 suggest a positive relationship between domestic sales and the propensity to export. When isolating demand-driven variation in domestic sales, the coefficient in column 3 turns negative and suggests that a 1% drop in domestic sales leads to a 0.099% increase in the probability of exporting. Nevertheless, the results in column 5 indicate a *positive* effect of domestic demand shocks on the proportion of years exported, though the estimated effect is modest (a 1% drop in domestic sales leads to a 0.04% drop in the proportion of export year) and only significant at the 5% level.

Taken together, the muted results in Table 4 lead us to conclude that the vent-for-surplus mechanism does not appear to operate particularly via the extensive margin (i.e., via entry and exit from the export market). This result is perhaps not entirely surprising in light of the fact, discussed in section 3.2, that more than 90% of the growth in Spanish exports during the bust period was explained by continuing exporters. More substantively, there are at least two potential explanations that make our muted extensive margin results not surprising, even in the face of the sizeable intensive margin effects discussed in section 5.1. The first explanation relates to the fact that we only have data on aggregate exports, and thus changes in the extensive margin in our context refer to entry and exit from export markets *altogether*, which is a decision involving much larger investments than entry and exit from specific export markets.[35] Second, despite the static nature of our model, in practice, the relevant changes in profitability that shape the intensive and extensive margin of trade may be quite different from each other. For instance, for the extensive margin, expectations of future earnings may be much more relevant than for the intensive margin (see Dickstein and Morales, 2018). Similarly, it is well-know that sunk costs of exporting tend to generate substantial hysteresis in exporting status.

Regardless of the economic factors explaining it, the empirical evidence in Table 4, and the fact that most of the growth in Spanish exports in the years following the Great Recession was due to firms that were already exporting during the boom, strongly suggest that the relationship between demand-driven changes in domestic sales and the extensive margin of exports in Spain in this period was not quantitatively important. Consequently, we focus in the remaining of this paper on exploring the robustness of the intensive margin results in Table 3.

6 Robustness

In this section, we complement our baseline results with additional evidence that further supports the empirical relevance of the "vent-for-surplus" mechanism. Specifically, we present estimates of regression specifications that address some specific sources of endogeneity that could affect the validity of our baseline identification strategy.

6.1 Panel Specification with Four Periods

One could be concerned about the exogeneity of our instrument in the regression specifications discussed in section 5 being polluted by the presence of underlying trends that are heterogeneous across regions in Spain and that affect both of the demand and supply conditions of the firms located in those regions. To assuage this concern, in Table 5, we report TSLS estimates for a panel specification that breaks the period 2002-2013 into four superiods: two boom subperiods (2002-05 and 2006-08), and two bust subperiods (2009-11 and 2012-13). The key advantage of this panel specification is that it permits the inclusion of municipality-specific time trends, apart from firm fixed effects and sector-period fixed effects. The various columns of Table 5 present results under alternative sets of fixed effects. The F-statistic for the null that the coefficient on our instrument in the first-stage regressions equals zero is always above 20, and while the second-stage elasticities of exports to demand-driven changes in domestic sales are slightly larger than our baseline one in

[35] It is important to remark that our findings do not rule out the possibility that, in reaction to the drop in domestic demand, continuing exporters increased the set of exported products or the set of countries to which they exported.

Table 5: Two-Stage Least Squares Estimates with Four Periods

Dependent Variable:	ΔLn(Exports)			
	(1)	(2)	(3)	(4)
Ln(Domestic Sales)	-1.773^a	-2.085^a	-1.853^a	-2.070^a
	(0.315)	(0.483)	(0.447)	(0.363)
Ln(TFP)	2.952^a	3.131^a	3.064^a	3.238^a
	(0.306)	(0.463)	(0.418)	(0.350)
Ln(Average Wages)	-1.951^a	-1.928^a	-1.371^a	-1.435^a
	(0.244)	(0.277)	(0.219)	(0.187)
1st-Stage Coefficient	0.731^a	1.206^a	1.168^a	1.457^a
F-statistic	104.89	23.13	24.13	47.80
Observations	24,036	24,036	24,036	23,995
Firm FE	Yes	Yes	Yes	Yes
Period FE	No	Yes	No	No
Sector-Period FE	No	No	Yes	Yes
Municipality-specific trend	No	No	No	Yes

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered at the province level are reported in parenthesis. The four periods considered are 2002-05, 2006-08, 2009-11 and 2012-13. The estimation sample includes all continuing exporters with complete data in at least three of the four subperiods.

column 8 of Table 3, they all lie within one standard deviation of it.

In Appendix J, we also present results from panel specifications analogous to those in Table 5 but in which each period corresponds to a year. Thus, these specifications estimate the causal link between demand-driven changes domestic sales and exports at a *yearly frequency*. The OLS results are very much in line with those obtained in Table 2, but the TSLS estimates are more sensitive to the type of fixed effects included in the estimation. More specifically, when municipality-specific time trends are included in the regression specification, either in combination with firm fixed effects or with both firm and sector-period fixed effects, the instrument remains relevant and the elasticity of exports with respect to demand-driven changes in domestic sales is still negative and statistically significant, although smaller than in our baseline boom-to-bust analysis.[36] The smaller elasticity is not surprising, as it is likely that the full impact of domestic demand shocks on exports will not manifest itself at a yearly frequency, since substituting from domestic to foreign markets often requires non-trivial investments on the part of firms, and these investments may take one, two or more years to materialize. When accounting for municipality-specific time trends, firm fixed effects and sector-*year* (instead of sector-period) fixed effects in the same regression specification, our instrument becomes much weaker (F-stat of 2.2), making the second-stage estimates unreliable in this case. Given that this same specification yielded a strongly statistically significant first-stage coefficient when each period was longer than a year, we again attribute the lack of statistical

[36]For the purpose of this specification, 'sector-period' fixed effects correspond to two sets of sector fixed effects, one set specific to the boom period and the other specific to the bust period.

Table 6: Intensive Margin: Robustness to Excluding Zip Codes Linked to Auto Industry

	Panel A: Exclude zipcodes w/ high auto employment share			*Panel B: Exclude zipcodes with at least one sizeable auto maker*		
Dependent Variable:	ΔLn(Exp)	ΔLn(DSales)	ΔLn(Exp)	ΔLn(Exp)	ΔLn(DSales)	ΔLn(Exp)
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	1st Stage	2nd Stage	OLS	1st Stage	2nd Stage
ΔLn(Domestic Sales)	-0.305^a		-1.693^a	-0.294^a		-1.663^a
	(0.036)		(0.277)	(0.040)		(0.371)
ΔLn(Dist-Pop-Wght. Vehicles p.c.)		1.290^a			1.372^a	
		(0.119)			(0.166)	
ΔLn(Average Wages)	-0.691^a	-0.511^a	-1.372^a	-0.714^a	-0.505^a	-1.364^a
	(0.075)	(0.052)	(0.166)	(0.114)	(0.070)	(0.222)
ΔLn(TFP)	1.519^a	1.022^a	2.867^a	1.505^a	1.006^a	2.801^a
	(0.063)	(0.034)	(0.243)	(0.072)	(0.050)	(0.300)
Observations	7,180	7,180	7,180	4,595	4,595	4,595
F-statistic		118.30			68.58	
	Panel C: Exclude zipcodes 'neighboring' zipcodes in Panel A			*Panel D: Exclude sectors w/ I-O links to automakers*		
Dependent Variable:	ΔLn(Exp)	ΔLn(DSales)	ΔLn(Exp)	ΔLn(Exp)	ΔLn(DSales)	ΔLn(Exp)
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	1st Stage	2nd Stage	OLS	1st Stage	2nd Stage
ΔLn(Domestic Sales)	-0.277^a		-1.693^a	-0.255^a		-1.864^a
	(0.036)		(0.317)	(0.032)		(0.343)
ΔLn(Dist-Pop-Wght. Vehicles p.c.)		1.311^a			1.238^a	
		(0.132)			(0.130)	
ΔLn(Avgerage Wages)	-0.664^a	-0.490^a	-1.325^a	-0.718^a	-0.491^a	-1.459^a
	(0.083)	(0.053)	(0.183)	(0.079)	(0.049)	(0.165)
ΔLn(TFP)	1.465^a	1.008^a	2.817^a	1.464^a	1.010^a	3.011^a
	(0.063)	(0.037)	(0.277)	(0.064)	(0.036)	(0.285)
Observations	6,131	6,131	6,131	6,072	6,072	6,072
F-statistic		98.78			91.09	

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered by province in parenthesis. For any X, $\Delta\text{Ln}(X)$ is the log difference between the average of X in 2009-2013 and its average in 2002-2008. 'Exp' denotes exports, and 'DSales' denotes domestic sales. 'ΔLn(Dist-Pop-Wght. vehicles p.c.)' denotes the baseline instrument constructed using data on vehicles per capita at the municipal level and applying the weights from the gravity equation reported in column 1 of Table 1. 'F-statistic' denotes the corresponding statistic for the null hypothesis that the coefficient on the ΔLn(Dist-Pop-Wght. vehicles p.c.) covariate is equal to zero. All regressions include sector fixed effects.

significance in this case to the lack of contemporaneous response of manufacturing demand to the type of purchasing power shocks captured by our instrument.

6.2 Further Purges of the Auto Industry

While the sample used to compute the estimates in Table 3 excludes firms classified in the manufacturing of motor vehicles sector, one might still be concerned that the salient presence of firms in that industry in a given municipality might lead to a negative association between the boom-to-bust changes in the stock of vehicles per capita and in the unobserved *residual* marginal costs shifters of the firms located in that municipality (even if they operate in other industries). This would be the case if the boom-to-bust drop in the number of vehicles per capita in a municipality was caused by an exogenous increase in marginal costs affecting the firms in the motor vehicles industry, and this negative supply shock was transmitted to other firms within the same municipality, reducing aggregate labor demand and, thus, aggregate consumption in this municipality.[37] Notice however that, as discussed in more detail in section 4.2, this source of endogeneity in our instrument would cause the TSLS estimates presented in Table 3 to be upward biased, as unobserved shocks that increase firms' marginal costs would have a negative impact on their exports. In order to evaluate the robustness of our estimates to this concern, we report in Table 6 TSLS estimates for regressions specifications analogous to those in columns 4 and 8 of Table 3, but for four alternative samples. In Panel A, we exclude from our sample all firms located in a zip code that ranks in the top 25% of zip codes by share of manufacturing employment accounted for by motor-vehicles producers (as computed from our micro-level data). In Panel B, we further restrict the sample relative to Panel A by excluding all firms located in a zip code in which at least one motor-vehicles producer with more than 20 workers operates. In Panel C, we exclude all firms from zip codes neighboring a zip code that ranks in the top 25% of zip codes by share of manufacturing employment in motor-vehicles producers.[38] Finally, in Panel D, we exclude all firms producing in sectors that are either one of the two top leading input providers or two top leading buying industries of the vehicles manufacturing industry. The results in all panels point at slightly larger estimated elasticities (in absolute value), consistently with the notion that these sample restrictions attenuate concerns about our estimates being up upward biased.

6.3 Alternative Instruments

As described in detail in section 4.1, the value of our baseline instrument for each municipality is computed as the log change in a weighted average of vehicles per capita in every other municipality, where the weight attached to a municipality depends on its population and the distance to the municipality of origin, and where the elasticities of this weight with respect to these covariates correspond to the estimates reported in column 1 of Table 1. In columns 1 to 3 of Table 7, we test the robustness of our results to instruments constructed similarly to our baseline instrument, but with weights that depend on distance and population in different ways. In particular, after reproducing our baselines estimates in column 1, columns 2 to 4 present results corresponding to instruments based on the gravity specifications in columns 2 to 4 of Table 1. As a reminder, the first of these includes own-municipality and own-province dummies, the second one uses a more flexible specification to estimate the impact of distance on municipality-to-municipality shipments,

[37]For example, the post-2009 trade collapse may have increased the input costs for firms in the motor vehicles industry, which may have passed these higher costs through to their buyers, which may be other firms located in the same municipality but operating in different industries.

[38]We identify two zip codes as neighboring each other if they share the first four digits of their 5-digit code.

Table 7: Alternative Instruments and Overidentification Tests

Dependent Variable:	ΔLn(Domestic Sales)					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔLn(Dist-Pop-Wght. Vehicles p.c.) Gravity: mun-mun flows (Baseline)	1.312^a (0.119)					
ΔLn(Dist-Pop-Wght. Vehicles p.c.) Baseline incl. own mun. dummy		1.321^a (0.118)				
ΔLn(Dist-Pop-Wght. Vehicles p.c.) Gravity: distance dummies			0.954^a (0.128)			
ΔLn(Dist-Pop-Wght. Vehicles p.c.) Gravity: firm-mun flows				1.303^a (0.106)		
ΔLn(Weighted Vehicles p.c.) Weights: firm-level mun. shares					0.521^a (0.111)	
Δ(Dist-Pop-Wght. Vehicles p.c.) Baseline in levels with AKM s.e.'s						0.967^a (0.115) (0.208)
F-statistic	122.44	125.46	55.90	151.03	21.82	70.94 21.57
	ΔLn(Exports)					
	(1)	(2)	(3)	(4)	(5)	(6)
ΔLn(Domestic Sales)	-1.607^a (0.248)	-1.613^a (0.214)	-1.336^a (0.228)	-1.642^a (0.245)	-2.162^a (0.412)	-1.701^a (0.262) (0.271)
Observations	8,009	8,009	8,009	7,906	7,850	8,009

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered by province reported in parentheses. In column 6, the we first report standard errors clustered by province, and below we report standard errors that apply the formula in Adão et al. (2019). All specifications include firm-level log differences in TFP and log differences in average wages as additional controls (estimates not included to save space). Additionally, all specifications also include sector fixed effects. For a detailed description of each of the instruments, see text.

and the third one estimates a gravity equation that is very similar to the one we use to compute our baseline instrument, but exploits firm-to-municipality shipment flows instead of municipality-to-municipality ones. Because our instrument excludes predicted own-municipality flows, it is not surprising that the results in column 2 are almost identical to those in column 1. The results in column 3 are more dissimilar, with a second-stage elasticity slightly lower (-1.336) than our baseline one. Finally, despite the fact that the instruments used in column 4 are based on the smaller distance and population elasticities estimated with our firm-destination trade flow data in column 4 of Table 1, the resulting second-stage elasticities are very similar to the baseline ones (-1.607 vs -1.642).[39]

In column 5 of Table 7, we present results based on instruments computed in an analogous

[39]The results are also virtually identical when using the estimated distance and population elasticities reported in column 5 (rather than column 4) of Table 1.

manner as in our baseline specification, but with the difference being that, instead of using weights that only vary bilaterally between municipalities and are predicted by a gravity equation, we use weights that vary across firms and that correspond to the *actual* domestic sales share of each firm in each Spanish municipality (other than the firm's own municipality) in 2006. This specification results in a larger (in absolute terms) elasticity of exports to domestic sales, but note that this coefficient is approximately only one standard deviation away from our baseline of -1.6.

Finally, in column 6 we present results associated with a variant of our baseline instrument that, instead of computing a log change in our weighted sum measure of demand, computes simply a weighted sum of the change in vehicles per capita at the municipality level. This instrument belongs to the category of shift-share instruments considered by Adão et al. (2019), thus allowing us to compute standard errors in the manner recommended in that paper. Given that both the baseline instrument and the shift-share instrument used in column 6 rely on the same identification assumptions, it is reassuring that the second-stage point estimates they yield are very similar (-1.607 vs. -1.701). To evaluate the possible bias of standard errors clustered by province in our empirical application, we compute in column 6 both standard errors clustered by province (first number in parenthesis) and the standard errors suggested in Adão et al. (2019) (second number in parenthesis). As illustrated in column 6 of Table 7, the downward bias affecting the standard errors based on clustering by province is very small; while the second-stage standard error that clusters by province is 0.262, the standard error computed according to the formula introduced in Adão et al. (2019) is 0.271.[40]

In section 6.6, we discuss additional alternative instrumentation strategies related to the deep roots of the differential fall in demand across Spanish regions.

6.4 Controlling for Additional Confounding Factors

In spite of the controls included in our baseline specification, one may still be concerned that this specification might not be accounting for the effect of some marginal cost shifters that could be correlated with our instrument, thus biasing our estimates. The fact that our instrument is constructed based on municipality-to-municipality trade flows aggregated across all manufacturing firms and excludes predicted sales in a firm's own municipality should diminish concerns related to unobserved supply factors, but it is nonetheless important to assess the robustness of our results to the inclusion of additional proxies for firm- and municipality-level cost shifters. More specifically, one may be concerned that our firm-level measures of average wages and TFP are too crude to fully capture changes in firm-level supply conditions, even when additionally controlling for sector fixed effects.

For instance, the dual nature of the Spanish labor market, with large differences in dismissal costs between temporary- and permanent-contract workers, might have led firms to shed a dispro-

[40]The similarity in our application between the standard errors that cluster by province and those suggested in Adão et al. (2019) is hardly surprising: municipalities located in the same province assign similar weights to every other municipality and, thus, firms located in the same province are, according to our shift-share instrument, similarly exposed to changes in the stock of vehicles per capita across all Spanish municipalities.

portionate number of temporary workers during the bust. If so, given the average differences in skill and experience between both types of workers, firms shedding temporary workers may have undergone skill- and experience-upgrading that changed firms' marginal production costs in a way that is not properly accounted for by our firm-level measures of TFP and average wages. Similarly, our baseline specifications do not include any proxies for factor costs other than labor costs, yet it is likely that financial costs faced by firms (explicit via interest rates, or implicit via rationing) were also significantly impacted by the Great Recession.

After reproducing our baseline estimates in column 1, columns 2 to 4 of Table 8 present variants of our baseline specification that add several covariates with the aim of controlling for various alternative confounding factors. More specifically, in column 2, we additionally control for the firm-level change in the share of temporary workers. The negative and statistically significant point estimate indicates that firms that shed a disproportionate number of temporary workers during the bust period experienced a larger increase in exports, which is in line with our hypothesis above. The IV estimate of the causal effect of demand-driven changes in domestic sales on exports is however only slightly increased (elasticity of -1.639).[41] In columns 3 and 4, we introduce municipality-level controls for local labor market conditions. Column 3 includes the same change in the ratio of temporary workers over total employment as in column 2, but computed at the municipality level. In column 4, we further control for a municipality-level measure of the change in the manufacturing employment per capita. The inclusion of these two controls has a negligible impact on the main coefficient of interest, and only the second one has a significant effect on exporting.[42]

In columns 5 to 7 of Table 8, we study potential confounding effects related to financial costs. We construct a measure of the financial costs that each firm faces in each period as the within-period average ratio of financial expenditures over total outstanding debt with financial institutions (both measures are annually reported by firms in their financial statements). In column 5, we add the log change in this firm-level measure of financial costs as an additional control; the impact of this measure on firms' changes in exports is statistically different from zero only at the 10% level, and including this variable only has a marginal effect on the estimate of the elasticity of exports to domestic sales (which becomes -1.678). In columns 6 and 7, we explore the possibility that the relevant increase in the financial costs faced by firms in the bust relative to the boom happened through credit rationing, instead of via explicit interest rates. Although we do not have measures of firms' credit applications and whether these were denied, one may conjecture that firms whose financial costs were larger in the boom were more likely to suffer credit rationing in the bust. Regardless of whether we measure financial costs in the boom using each firm's financial information (column 6) or as the average financial costs of all other firms located in the same

[41] We obtain similar results when instead controlling for the (initial) firm-level share of temporary workers during the boom period. Specifically, firms that entered the bust period with a larger share of temporary workers (and thus had a larger potential to affect their skill composition when transitioning to the bust period) experienced higher export growth in the bust relative to the boom, but the estimate of our parameter of interest is largely unaffected.

[42] Firms located in municipalities with larger declines in manufacturing employment per capita experienced higher export growth, potentially due to workers' extra effort in reaction to the reduction in employment opportunities in their municipality.

Table 8: Confounding Factors

Dependent Variable:	ΔLn(Exports)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ΔLn(Domestic Sales)	-1.607^a	-1.639^a	-1.618^a	-1.632^a	-1.678^a	-1.680^a	-1.611^a
	(0.248)	(0.251)	(0.263)	(0.259)	(0.251)	(0.257)	(0.264)
ΔShare of Temp. Workers (firm level)		-0.250^a					
		(0.089)					
ΔShare of Temp. Workers (munic. level)			-0.019				
			(0.170)				
ΔManufacturing Empl. p.c. (munic. level)				-0.272^a			
				(0.050)			
ΔLn(Financial Costs) (firm level)					-0.027^c		
					(0.014)		
Financial Costs in Boom (firm level)						-0.008	
						(0.015)	
Financial Costs in Boom (munic. level)							-0.039
							(0.041)
Observations	8,009	7,640	7,743	7,745	6,879	6,945	7,741
F-Statistic	122.44	131.97	138.65	136.93	88.43	89.27	139.26

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered by province reported in parentheses. In all specifications, ΔLn(Domestic Sales) is instrumented by ΔLn(Distance-Population-Weighted Vehicles per capita), defined as in previous tables. All specifications include firm-level log changes in TFP and in log wages as additional controls (coefficients not included to save space), and sector fixed effects.

municipality (column 7), our results indicate that either credit rationing had little impact on firms' exports or our conjecture that it may be measured through the firms' financial costs in the boom has little empirical support.[43]

6.5 Alternative Productivity Estimates

We next test the robustness of our results to alternative approaches to measuring firms' productivity. Columns 1 and 2 in Table 9 replicate our baseline OLS and IV estimates presented in column 6 of Table 2 and column 8 of Table 3, while columns 3 and 4 of Table 9 present estimates of specifications that differ exclusively on the productivity measure.

Consistently with the model described in section 2, both productivity measures exploit the assumption that firms: (a) face a CES demand function and are monopolistically competitive in both the domestic and the foreign market; (b) take all factor prices as given. The two approaches

[43] In Table G.8, we additionally control for the change in the number of bank offices per capita in the municipality of location of a firm and for the change in firm-level short-term liabilities over total liabilities. We interpret the first of these two variables as an alternative proxy for firms' financial constraints and the second one as a way of partly capturing the potential role of international trade credit in facilitating the growth of exports in municipalities that were hard-hit by the financial crisis. In Table G.8, we also experiment with the inclusion of the change in the value of land (measured at the municipality level). The estimate of our key parameter is robust to the inclusion of these additional controls.

we implement differ however on the assumptions we impose on the shape of the production function. In our baseline approach, we assume that the firm's production function is a Leontief aggregator of materials and a translog function of labor and capital (as in Ackerberg et al., 2015). Given these assumptions, we describe our estimation procedure in detail in Appendix F. A possible concern with this estimation approach is that, if it were to be the case that materials are not perfect complements with the output of labor and capital, our measure of the firm's productivity would automatically incorporate a measure of the firm's materials' usage. This would be problematic for our identification approach, as firms may adjust their materials' usage directly in reaction to a demand-driven change in domestic sales, which would result in our instrument being endogenous. To address this possible concern, the second approach assumes instead that the production function is a Cobb-Douglas aggregator of materials and the same translog function of labor and capital employed in our baseline approach (see Bilir and Morales, 2020, for details on the estimation procedure). Thus, while our baseline approach imposes that material inputs have a zero elasticity of substitution with the output of labor and capital, the second approach imposes instead a unit elasticity of substitution.

In both estimation approaches, we invoke optimality conditions for the static inputs (labor and materials) in order to estimate the relevant parameters of the production function and, in this sense, both approaches are specific cases of the general estimation framework in Gandhi et al. (2016). Both estimation approaches do however use different outcome measures; while the approach that assumes a Leontief production function exploits data on the firm's sales revenue, the approach that assumes a Cobb-Douglas production function uses information on the firm's value added. We thus refer in Table 9 to the two measures of productivity that we obtain as "TFP Sales" and "TFP Value Added", respectively.

A general concern with our productivity estimates is that, if they do not correctly account for the impact of different factors of production on the firm's total sales, they may just become an imperfect proxy of these total sales, which would cause our estimate of the elasticity of exports with respect to demand-driven changes in domestic sales to be biased downwards. We should however point out that our measures of productivity are far from being perfectly correlated with the firm's total sales; specifically, this correlation is 0.56 for our baseline approach and 0.22 for our alternative approach. The higher correlation of our baseline approach is consistent with it partly accounting for the firm's usage of material inputs.

Perhaps reflecting the lower correlation between our alternative productivity proxy and the firm's total sales, the OLS estimator in column 3 reveals a *positive* partial correlation between exports and domestic sales. However, the IV elasticity in column 4 is again negative and, though it is significantly lower in absolute value than in our baseline specification (see column 2), it still implies a sizeable substitution between domestic sales and exports at the firm level.

A second concern with our productivity estimates is that, because we do not observe separately prices and quantities for each firm, they may capture not only the firm's actual productivity but also the firm's demand shifter. Specifically, this would be a concern if our productivity estimates

Table 9: Alternative TFP Measures

Dependent Variable:	ΔLn(Exports)			
	(1)	(2)	(3)	(4)
	OLS	IV	OLS	IV
ΔLn(Domestic Sales)	-0.292^a	-1.607^a	0.020	-1.057^a
	(0.032)	(0.248)	(0.038)	(0.197)
ΔLn(Average Wages)	-0.723^a	-1.387^a	-0.753^a	-1.024^a
	(0.072)	(0.151)	(0.070)	(0.100)
ΔLn(TFP Sales): Baseline	1.535^a	2.810^a		
	(0.057)	(0.213)		
ΔLn(TFP Value-Added)			1.014^a	1.338^a
			(0.076)	(0.096)
Observations	8,009	8,009	8,009	8,009
F-Statistic		122.44		66.29

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered at the province level are reported in parenthesis. For any X, $\Delta\text{Ln}(X)$ is the difference in $\text{Ln}(X)$ between its average in the 2009-2013 period and its average in the 2002-2008 period. All specifications include sector fixed effects.

were implicitly already controlling for the impact of our instrument. There is however no empirical evidence of this being relevant in our data: the correlation between our instrument for location-specific demand shocks and log changes in our productivity measures is actually negative (it is -0.18 for our baseline approach and -0.33 for our alternative approach).

6.6 Additional Robustness Tests

We finally succinctly discuss a number of additional robustness tests. To save space, these results are reported in Appendices G and H. First, we study in columns 1 to 3 of Table G.1 how our results are affected when restricting our estimation sample according to different criteria. First, we exclude Spanish subsidiaries of foreign multinationals, since both their exports and domestic sales may react to local demand shocks differently than other firms in Spain. Second, we restrict our attention either to firms with a single manufacturing establishment or even to firms with a single establishment, as multi-establishment firms might have production plants in locations other than the one where they are incorporated and, consequently, may react to local demand shocks in the headquarters' location very differently from single-establishment firms. No matter which of these sample restrictions we implement, the results are not significantly affected.[44] In columns 4 and 5 of Table G.1, we also verify that our results are not materially affected when defining the bust period as 2010-2013 or 2011-13, instead of 2009-13.[45]

[44]Relatedly, we have used information on the location of all car assembly plants in Spain (many of which part of multi-plant firms), and have confirmed that results are robust to excluding all provinces where these plants are located.

[45]The motivation for exploring these alternative definitions of the bust period is that, relatively to prior years, the acceleration in the growth rate of Spanish exports starts in 2010 (see Figure 2) and, relative to other countries in the euro area, this acceleration starts in 2011 (see Figure 1).

In Tables G.2, G.4, and G.6, we modify our baseline regression specification so as to: (i) include province or province-sector fixed effects; (ii) cluster standard errors at various levels other than province, and (iii) weight the observations according to different criteria (number of years exporting, log average sales during the boom period, log average employment during the boom period, and log average assets during the boom period). The inclusion of province and province-sector fixed effects and the weighting of observations have a very small effect on our estimates. Some forms of clustering (particularly two-way clustering by province and sector) tend to yield larger our standard errors, but our key estimates remain significant at the 1% level.

Next, in Tables G.8 and G.10, we report specifications that control for a variety of additional firm- or municipality-level supply factors, and specifications that control for firm-specific log average sales, log average employment, log average assets or average export-to-sales ratio during the boom period. When introducing these controls, the estimates of the elasticity of interest are never lower than -1.4 or higher than -1.9. In Tables G.12 and G.13, we reproduce our findings when aggregating the firm-level data at the municipality-sector level. When estimating our baseline specification at the municipality-sector level, we obtain lower (in absolute value) elasticities of exports with respect to domestic sales, but the qualitative nature of our findings remains unaffected.

In Appendix H, we explore alternative strategies to identify the potential substitutability of domestic and export markets. First, in Appendix H.1, we follow the identification approach implemented Berman et al. (2015) to estimate the causal impact of demand-driven changes in exports on domestic sales, swapping then the role that these two variables play in our main specification. Due to data restrictions (see section 3.3), we can only carry out such analysis for the period 2002-07. Consistently with our main findings, and contrary to the results in Berman et al. (2015) using French data, we find a *negative* effect of demand-driven changes in exports on domestic sales. Second, in Appendix H.2, we re-estimate the parameters of our main regression specification but explore alternative instrumentation strategies that focus on the deep roots of the differential fall in demand across Spanish regions. More specifically, we posit that, relative to the boom years, municipality-level demand shocks were larger (i) in municipalities with lower housing supply elasticities (in which house prices grew disproportionately during the boom years), (ii) in municipalities with a larger pre-crisis contribution of the construction sector to total labor income, and (iii) in provinces that experienced larger declines in tourism during the bust years.[46] We then weigh the municipality-specific demand shocks in (i) and (ii) following the same procedure as in our baseline IV approach. The second-stage estimates of our parameter of interest computed using these alternative instruments are all negative and generally a bit lower in absolute value than our baseline one; however, one should interpret these estimates with caution, as these instruments are generally not as strong as our baseline one.[47]

[46] The construction and tourist sectors are among the ones that experienced the largest reduction in total sales and employment in the bust relative to the boom. Regions more exposed to these sectors are likely to have experienced a larger drop in demand for manufactured goods.

[47] When adding each of these instruments one by one to our baseline instrument, standard tests of overidentifying restrictions fail to reject at typically used significance levels the null hypothesis that our instruments are jointly valid.

7 Structural Interpretation and Quantification

Our empirical results suggesting a negative impact of demand-driven changes in domestic sales on changes in exports are in contradiction with the theoretical framework described in section 2. In section 7.1, we show how a simple extension of that framework incorporating non-constant marginal costs can rationalize our empirical results. In section 7.2, we use this extended framework to provide a quantitative assessment of the importance of the domestic slump for the observed export miracle in Spain during the period 2009-13.

7.1 Structural Interpretation

The theoretical environment we consider here is identical to that in section 2, except that the cost structure in equation (2) is now replaced with a total variable cost of producing Q_{id} units of output for the domestic market and Q_{ix} units of output for the foreign market given by

$$\frac{1}{\widetilde{\varphi}_i}\omega_i\frac{1}{\lambda+1}\left(\tau_d Q_{id}+\tau_x Q_{ix}\right)^{\lambda+1}, \quad \lambda \geq 0, \tag{13}$$

where $\tau_d Q_{id}+\tau_x Q_{ix}$ denotes firm i's total output in the presence of iceberg trade costs in the domestic (τ_d) and foreign (τ_x) markets. Notice that the parameter λ governs how steeply marginal costs increase with output. When $\lambda = 0$, marginal costs are constant and equation (13) reduces to our previous expression in equation (2). We show in Appendix A that the cost function in equation (13) can derived in a model in which the firm's production function is a Cobb-Douglas aggregator of a fixed or pre-determined input and a flexible and static input; without loss of generality, we can refer to these two inputs as capital and labor, respectively. Under this micro-foundation, the parameter λ is decreasing in the elasticity of output with respect to the flexible factor, and λ is equal to 0 when this elasticity is equal to one. Note also that we denote firm productivity with the new notation $\widetilde{\varphi}_i$ (rather than φ_i), as the microfoundation in Appendix A shows that this productivity level $\widetilde{\varphi}_i$ depends not only on the TFP level φ_i but also on the stock of fixed factors.

Solving for the optimal level of exports by firm i under the cost function in equation (13), and taking log differences, we obtain

$$\begin{aligned}\Delta\ln R_{ix} \;&=\; (\sigma-1)\left[\Delta\ln\xi_{ix}+\Delta\ln\widetilde{\varphi}_i-\Delta\ln\omega_i\right]-(\sigma-1)\left(\Delta\ln\tau_{sx}-\Delta\ln P_{sx}\right)+\Delta\ln E_{sx}\\ &\quad-(\sigma-1)\,\lambda\Delta\ln\left(\tau_d Q_{id}+\tau_x Q_{ix}\right),\end{aligned} \tag{14}$$

which is analogous to equation (4) except for the last term, which reflects the effect of total output on marginal production costs.[48] Next, note that, due to constant mark-up pricing, we can write

$$\ln\left(\tau_d Q_{id}+\tau_x Q_{ix}\right)=\ln\left(\frac{\tau_d R_{id}}{P_{id}}+\frac{\tau_x R_{ix}}{P_{ix}}\right)=\ln\left(R_{id}+R_{ix}\right)-\ln\left(\frac{\sigma\omega_i\left(\tau_d Q_{id}+\tau_x Q_{ix}\right)^{\lambda}}{(\sigma-1)\widetilde{\varphi}_i}\right). \tag{15}$$

[48] The expression in equation (14) as well as all other expressions in this section implicitly assume that there are only two markets, one domestic market and one foreign market. For an extension of this model to a setting with multiple domestic and foreign markets, see Appendix E.2.

Table 10: Intensive Margin with Total Sales

Dependent Variable:	ΔLn(Exp)	ΔLn(TotSales)	ΔLn(Exp)	ΔLn(TotSales)	ΔLn(Exp)
	(1)	(2)	(3)	(4)	(5)
	OLS	1st Stage	2nd Stage	1st Stage	2nd Stage
ΔLn(Total Sales)	0.724^a		-2.374^a		-2.590^a
	(0.050)		(0.526)		(0.606)
ΔLn(TFP)	0.509^a	1.063^a	3.690^a	1.015^a	3.739^a
	(0.055)	(0.026)	(0.482)	(0.027)	(0.539)
ΔLn(Average Wages)	-0.217^a	-0.509^a	-1.750^a	-0.493^a	-1.801^a
	(0.063)	(0.043)	(0.250)	(0.041)	(0.282)
ΔLn(Dist-Pop-Weighted Vehicles p.c.)		0.888^a		0.838^a	
		(0.103)		(0.107)	
ΔLn(Stock of Capital)				0.101^a	0.382^a
				(0.009)	(0.067)
Observations	8,009	8,009	8,009	8,009	8,009
Sector FE	Yes	Yes	Yes	Yes	Yes
F-statistic		75.00		61.30	

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered at the province level are reported in parenthesis. For any X, $\Delta\text{Ln}(X)$ is the difference in $\text{Ln}(X)$ between its average in the 2009-2013 period and its average in the 2002-2008 period. All specifications include sector fixed effects.

Solving for $\ln(\tau_d Q_{id} + \tau_x Q_{ix})$, plugging this expression into equation (14), and imposing the same decomposition as in equation (5), we then find that:

$$\Delta \ln R_{ix} = \gamma_{sx} + \frac{(\sigma-1)}{1+\lambda}\delta_\varphi \Delta \ln(\widetilde{\varphi}_i^*) - \frac{(\sigma-1)}{1+\lambda}\delta_\omega \Delta \ln(\omega_i^*) - \frac{(\sigma-1)\,\lambda}{1+\lambda}\Delta \ln(R_{id}+R_{ix}) + \varepsilon_{ix}, \quad (16)$$

where $\varepsilon_{ix} \equiv u_{ix}^{\xi} + ((\sigma-1)/(1+\lambda))(u_i^{\varphi} - u_i^{\omega})$. This equation is analogous to equation (10), except that it features the log difference of *total* sales (and not just domestic sales) on the right-hand side, and that it calls for the inclusion of the firm's stock of fixed factors as an additional control. The intuition for the need to include the change in total sales rather than in domestic sales as an explanatory variable is straightforward: marginal costs of production are increasing in total output, not just output destined for the domestic market.

Estimating equation (16) via OLS is problematic not just for the reasons identified in section 2, but also because the fact that the log change in total sales naturally depends on the log change in exports implies that any unobserved determinant of exports accounted for by the regression residual will be correlated with our covariate of interest and, thus, will bias the OLS estimate of $(\sigma-1)\lambda/(1+\lambda)$. Importantly, because the regression residual in equation (16) depends on the same terms as that in equation (10) (i.e., unobserved productivity, factor costs and export demand shifters), a TSLS estimator based on our instrument will deliver consistent estimates of this regression coefficient as long as the identification assumptions outlined in section 2.1 hold. Consequently, the threats to the validity of our instrument discussed in section 4.2 also apply here.

In Table 10, we present OLS and TSLS estimates of the regression coefficients in equation (16). In columns 1 to 3, we include in the regression specification the same controls as in previous tables. In columns 4 and 5, we additionally control for the change in the stock of capital, as indicated by the micro-foundation in Appendix A. As expected, the OLS estimates in column 1 indicate a strong positive correlation between exports and total sales, even when controlling for sector fixed effects and for our measures of firms' average wages and TFP. The first-stage results in column 2 indicate that our baseline instrument is a strong predictor of a firm's total sales (not just its domestic sales), with an F-stat of 75. The second-stage elasticity of exports to total sales in column 3 is negative and significant and stands at a value of -2.374. Adding the boom-to-bust log change in the firm's stock of physical capital does not affect significantly the first-stage nor the second-stage results. Thus, henceforth, we treat the estimates in column 3 as our baseline estimates.

To understand the magnitude of our estimates, take a firm with an initial export share of 16.2% (which corresponds to the median export share during the boom in our sample of 8,009 continuing exporters). Suppose that, due to a drop in demand, this firm experiences a 1% drop in its domestic sales. Our estimated elasticity of exports to domestic sales in Table 3 indicates that, other things equal, the firm should see its exports increase by 1.6%. This also implies that the firm's total sales will decrease by $83.8\% \times 1\% + 16.2\% \times (-1.6\%) = 0.58\%$. For this change in total sales, our estimated elasticity of exports to total sales in Table 10 suggests an implied increase in exports of $0.58\% \times 2.374 = 1.4\%$, which is close to the 1.6% increase predicted by the estimates in Table 3. This demonstrates that our IV results in Tables 3 and 10 deliver congruent estimates for the response of exports to local demand shocks.[49]

With an estimate of the demand elasticity σ in hand, it is easy to infer an estimated value of λ from the estimates in Table 10. Specifically, given the estimates in column 3, we can compute an estimate of λ as $\hat{\lambda} = 2.374/\left(\sigma - 1 - 2.374\right)$. For $\sigma = 6$, we obtain $\hat{\lambda} = 0.90$, which indicates a significant departure from constant marginal costs.

We conclude this section with additional tests of several implications of the vent-for-surplus mechanism. These tests are based on the idea that one should expect the increase in exports in reaction to a common demand-driven drop in domestic sales to be larger for those firms whose short-run marginal cost function is steeper or, equivalently, for those firms whose elasticity of output with respect to flexible inputs is lower.[50] The results in Table 11 confirm that the elasticity of the change in exports to changes in total sales is higher for firms having lower output elasticities with respect flexible inputs. More specifically, the elasticity of exports with respect to total sales is lower in sectors with higher elasticities with respect to materials (column 1), in sectors with a higher elasticity with respect to labor (column 2), and especially in sectors with a higher elasticity of

[49]In Appendix I, we present results for specifications analogous to those in Tables 2 to 9, with the only difference being that the boom-to-bust log change in total sales is included as right-hand-side variable instead of the corresponding log change in domestic sales. The conclusions discussed in section 6 are generally corroborated by the results reported in Appendix I.

[50]See Appendix A for a formalization of the link between the slope of the short-run marginal cost function and the elasticity of output with respect to flexible inputs. We rely on our production function estimates in Appendix F.2 to measure these output elasticities.

Table 11: Heterogeneous Effects with Total Sales: Second Stage

Dependent Variable:	ΔLn(Exports)		
	(1)	(2)	(3)
Δ Ln(Total Sales)	-2.508[a]	-3.635[a]	-2.688[a]
	(0.755)	(1.134)	(0.579)
Δ Ln(Total Sales) $\times$ High Output elasticity wrt Materials	0.462		
	(0.782)		
Δ Ln(Total Sales) $\times$ High Output elasticity wrt Labor		2.205[c]	
		(1.143)	
Δ Ln(Total Sales) $\times$ High Output elast. wrt Labor $\times$ Temp. Ratio			3.963[a]
			(1.350)
Observations	8,009	8,009	7,889
Sector FE	Yes	Yes	Yes
P-value for $H_0 : \beta_1 + \beta_2 = 0$	0.00	0.00	0.23

Note: ^{a} denotes 1% significance, ^{b} denotes 5% significance, ^{c} denotes 10% significance. Standard errors clustered at the province level reported in parenthesis. The output elasticities with respect to inputs are estimated with the same production function we use to estimate TFP.

output with respect to the use of temporary workers (column 3). Notice, however, that only the two latter results are statistically significant at standard levels. Interestingly, the estimates in column 3 imply that we cannot rule out that the elasticity of exports with respect to a domestic demand-driven change in total sales equals 0 for firms that satisfy two conditions: (a) their elasticity of output with respect to labor equals 1; (b) their share of temporary workers in their total workforce also equals 1. This prediction is consistent with the model described in section 7.1 and the micro-foundation in Appendix A, as these firms would have short-run constant marginal costs according to this micro-foundation.

7.2 Quantification

In this final section, we evaluate the quantitative importance of the "vent-for-surplus" channel for explaining the remarkable growth in Spanish exports during the period 2009-13. Specifically, we use the model described in section 7.1 to compute the counterfactual boom-to-bust change in the aggregate exports that would have taken place in the absence of the marginal-cost reductions driven by the fall in domestic demand.

To implement our counterfactual analysis, we need to make admittedly imperfect inferences from our data, and impose several assumptions beyond those embedded in the framework described in section 7.1. First, we need to assess the extent to which Spanish domestic demand fell between the boom and the bust periods. To do so, we perform a variance decomposition of firms' changes in total sales, use this decomposition to infer the contribution of domestic demand factors to the observed boom-to-bust drop in Spanish aggregate sales, and select a specific demand shock in the model that can be targeted to match this inferred contribution. Second, given this targeted demand shock, we need to make assumptions about how different equilibrium variables featuring in the model adjust to it. We provide more details on these choices below but, to build intuition, we first present a simple back-of-the-envelope calculation.

Back-of-the-Envelope Quantification. Let us assume the existence of a "representative" firm in Spain, with boom-to-bust log changes in total sales and exports corresponding to the aggregate ones we see in our sample of manufacturing firms. Building on equation (16), we can back out any counterfactual change in aggregate exports that is due to changes in total sales driven exclusively by domestic demand shocks as

$$\ln\left(\frac{R'_{x1}}{R_{x0}}\right) = \ln\left(\frac{R_{x1}}{R_{x0}}\right) - \frac{(\sigma-1)\,\lambda}{1+\lambda} \times \left(\ln\left(\frac{R'_1}{R_0}\right) - \ln\left(\frac{R_1}{R_0}\right)\right), \tag{17}$$

where, for any variable A (either exports R_x or total sales R), we denote its actual value in the boom by A_0, its actual value in the bust by A_1, and its counterfactual value in the bust by A'_1.

Imagine that R'_1/R_0 were to indicate the fall in total sales we would have observed if domestic demand in Spain had remained constant between the boom and the bust periods. Then, the expression in equation (17) relates the counterfactual log change in exports to the actual one observed in the data minus a term capturing the effect of the lower fall in marginal costs due to total sales falling by less once the demand-driven decline is removed. For our sample of manufacturing firms, total sales fell in the data from boom to bust by 10.23%, while exports increased by 11.99%. To compute R'_1/R_0, we assume for now that 40% of the observed fall in total sales was driven by demand factors (as we show below, this value is close to what the we infer from the data). With this assumption and the estimate of $((\sigma-1)\lambda)/(1+\lambda)$ in column 3 of Table 10, we obtain a counterfactual growth in exports of only 0.74%:

$$\frac{R'_x}{R_{x0}} = \exp\big(\ln(1+0.1199) - 2.374 \times (\ln(1-0.1023\times(1-0.4)) - \ln(1-0.1023))\big) = 1.0074.$$

Thus, this counterfactual analysis implies that 94% ((11.99%−0.74%)/11.99%) of the boom-to-bust growth in exports was accounted for by the vent-for-surplus mechanism.

There are several problems with this back-of-the-envelope calculation. The most important of these are that, by treating all firms as identical, it imposes the same contribution of demand factors to each firm's change in total sales, and that it does not distinguish between purely domestic firms and exporters (while, naturally, only the latter's demand-driven changes in domestic sales got translated into exports). As we show below, these shortcomings result in a large overestimate of the importance of the vent-for-surplus mechanism.

Baseline Quantification. To address the limitations of the back-of-the-envelope calculation, we perform a quantification exercise that makes fuller use of the heterogeneous-firm model described in section 7.1. Specifically, through the lens of this model, we conceptualize the domestic demand changes that affected Spanish firms between the boom and the bust as changes in the sectoral demand shifters $\{Q_{sd}\}_{s=1}^{S}$. According to our model, changes in Q_{sd} for a sector s determine changes in the residual demand function that each firm in s faces and, thus, from the perspective of each individual firm, are purely demand shifters.[51]

By focusing on the counterfactual impact of changes in the demand shifters $\{Q_{sd}\}_{s=1}^{S}$, we overcome one of the limitations of the back-of-the-envelope calculation described above, as we allow for the fact that firms faced a variety of idiosyncratic supply and demand shocks between the boom and the bust periods and, consequently, for the fact that the contribution of the changes in these aggregate demand shifters to each firm's change in total sales was heterogeneous across firms. Unfortunately, taking the changes in the sectoral demand shifters $\{Q_{sd}\}_{s=1}^{S}$ as the primitive in our counterfactual analysis implies that equation (17) is not enough to determine the outcome of this analysis: we need to resort to additional model-implied expressions that map changes in these shifters to changes in firms' domestic and export sales. Through these expressions, we take a stand on which equilibrium market variables we allow to respond to our counterfactual shocks of interest.

More specifically, we compute our counterfactual analysis by implementing the following three steps.

Step 1: Computing counterfactual changes in aggregate domestic sales and exports for given changes in the sectoral demand shifters $\{Q_{sd}\}_{s=1}^{S}$. We first derive the system of equations that we use to map relative counterfactual boom-to-bust changes in sectoral demand shifters to changes in aggregate domestic sales and exports (and, thus, total sales). Importantly, our aim when deriving this system of equations is to exclusively capture the impact that demand changes have on firms' sales through the 'vent-for-surplus' mechanism; i.e., holding firms' marginal cost shifters constant, but letting firms move along their marginal cost curve.

When deriving the system of equations that we use to carry out our counterfactual analysis, we impose the following restrictions. First, we maintain the boom-to-bust changes in Spanish firms' supply parameters $(\varphi_i, \omega_i, \tau_{sx}, \tau_{sd})$ and idiosyncratic demand shifters (ξ_{id}, ξ_{ix}) at their realized values.[52] Second, we assume that Spain is a small open economy and, thus, counterfactual changes in the Spanish aggregate demand shifters $\{Q_{sd}\}_{s=1}^{S}$ do not affect: (a) the boom-to-bust change in the foreign price indices and aggregate demand shifters $\{P_{sx}, Q_{sx}\}_{s=1}^{S}$; and (b) foreign firms'

[51]In a general equilibrium model, changes in Q_{sd} could be due to a variety of factors. For example, in a model with consumers that have Cobb-Douglas preferences over the S manufacturing sectors we consider in our analysis and other 'outside' sectors, decreases in the share of spending on manufacturing goods (e.g., a decrease in the Cobb-Douglas parameters associated with manufacturing sectors) would reduce the demand shifters $\{Q_{sd}\}_{s=1}^{S}$. In a more general model in which consumers own houses, these would also be affected by wealth effects associated to changes in housing prices, which may themselves reflect changes in expectations of future housing demand and supply shocks.

[52]We do not need to estimate these parameters for every firm in our sample. Following an approach analogous to that popularized by Dekle et al. (2008), we use data on the observed boom-to-bust changes in the domestic sales and exports of every Spanish firm to proxy for the changes in the functions of these supply and idiosyncratic demand parameters that are relevant for our counterfactual exercise. It is important to remark that, in general, certain supply shifters (e.g., the wage level ω_i that every firm i faces) will endogenously react to changes in demand shifters. However, as indicated above, the aim of our counterfactual is to capture the impact that demand changes have on firms' sales through channels *other than shifts in firms' marginal cost curves,* and, hence, we maintain the boom-to-bust changes in supply parameters at their realized values.

marginal production costs.[53] Third, consistently with the estimates presented in Table 4, we assume that firms do not change their export status in reaction to counterfactual changes in the domestic demand shifters. Although, according to these restrictions, we impose that certain variables do not react to counterfactual changes in the aggregate demand shifters of interest, firms' prices will obviously change in reaction to demand-driven movements along their marginal cost curves and, thus, the system of equations we use to implement our counterfactual analysis will indeed take into account that counterfactual changes in $\{Q_{sd}\}_{s=1}^{S}$ may imply counterfactual changes in the price indices $\{P_{sd}\}_{s=1}^{S}$.

More precisely, our counterfactual analysis relies on three sets of equations. We describe in Appendix K.1 the step-by-step derivation of these equations starting from the assumptions of the model described in section 7.1 and those imposed in the previous paragraph. The first set computes the counterfactual change in exports of every firm i that exports in both the boom and the bust periods. Each equation in this first set is analogous to equation (17), with the only difference that we add a subscript i to illustrate that it holds for every firm in our sample (and not just for a hypothetical representative firm), and that we decompose the counterfactual change in total sales as a function of the counterfactual change in domestic sales and exports:

$$\ln\left[\frac{R'_{ix1}}{R_{ix0}}\right] = \ln\left[\frac{R_{ix1}}{R_{ix0}}\right] - \frac{(\sigma-1)\,\lambda}{1+\lambda}\left[\ln\left(\frac{R'_{ix1}}{R_{ix0}}\chi_{i0} + \frac{R'_{id1}}{R_{id0}}(1-\chi_{i0})\right) - \ln\left[\frac{R_{i1}}{R_{i0}}\right]\right], \tag{18}$$

where $\chi_0 \equiv R_{ix0}/(R_{id0} + R_{ix0})$ denotes the initial export share of firm i, and both

$$\frac{R'_{ix1}}{R_{ix0}}\chi_{i0} + \frac{R'_{id1}}{R_{id0}}(1-\chi_{i0}) \qquad \text{and} \qquad \frac{R_{i1}}{R_{i0}}$$

denote, respectively, the counterfactual and observed change in firm i's total sales.

The second set of equations computes the counterfactual change in domestic sales of every firm i that is active in the domestic market in both boom and bust periods. Following analogous steps to those followed to derive equation (18), we obtain

$$\begin{aligned}\ln\left[\frac{R'_{id1}}{R_{id0}}\right] &= \ln\left[\frac{Q'_{sd1}}{Q_{sd0}}\left(\frac{Q_{sd1}}{Q_{sd0}}\right)^{-1}\right] + \sigma\ln\left[\frac{P'_{sd1}}{P_{sd0}}\left(\frac{P_{sd1}}{P_{sd0}}\right)^{-1}\right] + \ln\left[\frac{R_{id1}}{R_{id0}}\right] \\ &\quad - \frac{(\sigma-1)\,\lambda}{1+\lambda}\left[\ln\left(\frac{R'_{ix1}}{R_{ix0}}\chi_{i0} + \frac{R'_{id1}}{R_{id0}}(1-\chi_{i0})\right) - \ln\left[\frac{R_{i1}}{R_{i0}}\right]\right],\end{aligned} \tag{19}$$

where

$$\frac{Q'_{sd1}}{Q_{sd0}}\left(\frac{Q_{sd1}}{Q_{sd0}}\right)^{-1} \qquad \text{and} \qquad \frac{P'_{sd1}}{P_{sd0}}\left(\frac{P_{sd1}}{P_{sd0}}\right)^{-1}, \tag{20}$$

denote the counterfactual change (relative to the actual change) in the sectoral demand shifter and price index, respectively.

The third set of equations computes the counterfactual change (relative to the actual change)

[53] While part (a) implies that Spain is a small exporter to the rest of the world; part (b) implies that Spain is a small importer from the rest of the world.

in the price index of every sector s. Noting that $Q_{sd} = E_{sd}/P_{sd}$, and that total Spanish spending in a sector must equal total domestic sales by Spanish firms plus imports, one can derive:

$$\ln\left[\frac{P'_{sd1}}{P_{sd0}}\left(\frac{P_{sd1}}{P_{sd0}}\right)^{-1}\right] = \ln\left((1-s^{\mathcal{D}}_{sd0})\frac{Q'_{sd1}}{Q_{sd0}}\left(\frac{Q_{sd1}}{Q_{sd0}}\right)^{-1}\left(\frac{P'_{sd1}}{P_{sd0}}\left(\frac{P_{sd1}}{P_{sd0}}\right)^{-1}\right)^{\sigma}\frac{R^{\mathcal{X}}_{sd1}}{R^{\mathcal{X}}_{sd0}} + s^{\mathcal{D}}_{sd0}\sum_{i\in\mathcal{D}_s} s^{\mathcal{D}}_{id0}\frac{R'_{id1}}{R_{id0}}\right) - \ln\left(s^{\mathcal{D}}_{sd0}\sum_{i\in\mathcal{D}_s} s^{\mathcal{D}}_{id0}\frac{R_{id1}}{R_{id0}} + (1-s^{\mathcal{D}}_{sd0})\frac{R^{\mathcal{X}}_{sd1}}{R^{\mathcal{X}}_{sd0}}\right) - \ln\left[\frac{Q'_{sd1}}{Q_{sd0}}\left(\frac{Q_{sd1}}{Q_{sd0}}\right)^{-1}\right], \tag{21}$$

where $s^{\mathcal{D}}_{sd0} \equiv R_{sd0}/(R_{sd0}+R^{\mathcal{X}}_{sd0})$ denotes the share of *total* expenditure in sector s spent in varieties produced by Spanish firms, $s^{\mathcal{D}}_{id0} \equiv R_{id0}/R_{sd0}$ denotes firm i's share of the total domestic sales of Spanish firms in sector s, and $R^{\mathcal{X}}_{sd1}/R^{\mathcal{X}}_{sd0}$ denotes the observed boom-to-bust change in Spanish imports in sector s.

Equations (18), (19) and (21) depend only on two parameters, σ and $((\sigma-1)\lambda)/(1+\lambda)$. We set the former to 5, which is a central value in the range of estimates used in the international trade literature (see Head and Mayer, 2014), and the latter to the estimated value of 2.374 (see Table 10). Given these parameter values and, for every sector s, a value of the relative counterfactual change in the demand shifter described in equation (20), we use equations (18), (19) and (21) to compute the counterfactual change in domestic sales R'_{id1}/R_{id0} for every firm in the economy, and the counterfactual change in exports R'_{ix1}/R_{ix0} for every firm that exports a positive amount in the boom and in the bust periods.[54] We then aggregate these firm-specific counterfactual changes to construct counterfactual changes in aggregate domestic sales and exports, R'_{d1}/R_{d0} and R'_{x1}/R_{x0}.

Among all the counterfactual changes in $\{Q_{sd}\}_{s=1}^{S}$ we could consider, we focus on relative counterfactual changes in aggregate demand shifters that are constant across sectors; i.e., $(Q'_{sd1}/Q_{sd0})(Q_{sd1}/Q_{sd0})^{-1} = \Gamma_Q$. In Figure 5, we plot the changes in total exports, domestic sales, and total sales predicted by our model when Γ_Q takes different values between 0.5 and 1.5. For the set of firms that we use in our counterfactual analysis, aggregate domestic sales dropped 15.91% between the boom and the bust periods. As mentioned before, exports grew by 11.99% and aggregate total sales dropped by 10.23%. These are the values that our counterfactual analysis naturally generates when we set the relative counterfactual change in the aggregate demand shifter of every sector s to equal 1; i.e., $\Gamma_Q = 1$. If the value of the aggregate demand shifters in the bust had been 50% smaller than they actually were (i.e., $\Gamma_Q = 0.5$), our model predicts that aggregate domestic sales would have dropped by 56.64% and aggregate exports would have increased by 60.1%. In this case, aggregate total sales would have dropped by 32.87%. Conversely, if it had been 50% larger (i.e., $\Gamma_Q = 1.5$), aggregate domestic sales would have grown by 17.82% and aggregate exports would have dropped by 14.76%, and the result would have been an 11.18% growth in total sales.

Step 2: Decomposing the variance of total sales. Our ultimate goal is not to indicate how aggregate exports, domestic and total sales react to arbitrary counterfactual domestic demand

[54] Our system also yields a value for the relative counterfactual changes in price indices described in equation (20). Note that $(Q'_{sd1}/Q_{sd0})(Q_{sd1}/Q_{sd0})^{-1} = Q'_{sd1}/Q_{sd1}$ and, thus, we can interpret the choice of values for the relative counterfactual changes in the sector-specific demand shifters as a choice over the counterfactual values of these shifters in the bust period.

Figure 5: Impact of Aggregate Demand Shocks

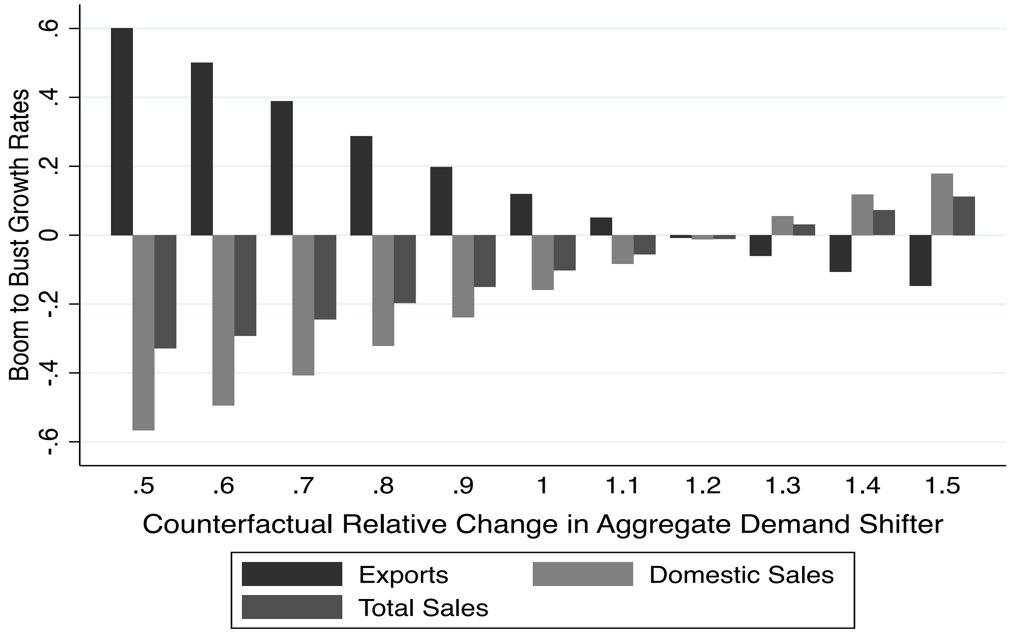

Notes: The horizontal axis indicates the value of Γ_Q. The export and domestic sales growth rates indicated in the vertical axis correspond to those predicted by equations (20) and (21). Given these counterfactual growth rates in export and domestic sales, we compute the counterfactual growth rate in total sales as $(R'_{ix1}/R_{ix0})\chi_{i0} + (R'_{id1}/R_{id0})(1 - \chi_{i0})$.

changes, but to predict the change in these variables that we would have observed if demand shifters in Spain had remained constant between the boom and the bust periods. Doing so requires measuring the extent to which Spanish domestic demand actually fell between these two periods. With this aim, we follow an approach analogous to that of Autor et al. (2013) and base our measure on a variance decomposition of observed boom-to-bust changes in firms' total sales. Specifically, we use equation (16) to decompose the variance of $\Delta \ln(R_i)$, with $R_i \equiv R_{id} + R_{ix}$, into a component due to firms' marginal cost and export demand shifters and a component attributed to factors orthogonal to these shifters (see Appendix K.2 for details). When performing this decomposition, we find the contribution of the combination of marginal cost and export demand shifters to be 59%, and that of factors orthogonal to it to be 41%. On the basis of this number, we predict that 41% of the 10.23% drop in total sales between the boom and the bust periods was due to domestic demand factors.

Our procedure to measure the contribution of demand changes to the domestic slump is admittedly imperfect. First, we use the results of a variance decomposition to infer the contribution of different factors to the change over time in an aggregate variable. Second, our decomposition reveals that 41% of the variance of the changes in firms' total sales is due to *any* residual factor that is orthogonal to firms' marginal cost shifters and export demand shocks. Thus, the conclusion that changes in demand shifters explain 41% of the variance in firms' changes in total sales relies on assuming that these shifters are the only determinants of firms' total sales whose changes between the boom and bust are orthogonal to contemporaneous changes in firms' marginal cost shifters and export demand shocks.

Step 3 and Quantification results. Given the result in *Step 2* that 41% of the 10.23% drop in total sales between the boom and the bust periods was due to changes in demand, we use the counterfactual results computed in *Step 1* to find the value of Γ_Q for which our model predicts a drop in total sales that is equal to $6.04\% = (1-41\%) \times 10.23\%$. Intuitively, this is the drop in total sales that we would have observed if aggregate demand shifters had not changed between the boom and the bust periods. Our model predicts a drop in total sales of 6.04% if Γ_Q=1.09.[55] For this value of the parameter Γ_Q, our model predicts that the growth in exports would have been 5.79%, and the drop in domestic sales would have been 9.10%. Given that the observed growth in exports was 11.99% (for $\Gamma_Q = 1$), our analysis indicates that the vent-for-surplus mechanism explains $(11.99\%-5.79\%)/11.99\% = 51.71\%$ of the total growth in exports. This is a significant contribution, but it is noticeably smaller than the one delivered by our back-of-the-envelope quantification.

Additional results. We discuss here results for two additional counterfactual exercises. First, we quantify how much more would the total sales of Spanish firms have dropped if firms had faced an increase in export costs in the bust period, when trying to substitute export markets for domestic markets (see Appendix K.3 for details on the system of equations we use to compute this counterfactual). Our results show that even moderate increases in trade costs would have had a very large impact, accentuating severely the drop in total sales that Spanish firms would have experienced; for example, if trade costs in the bust had been 10% larger than they actually were, holding the changes in every other variable constant, then total sales of Spanish firms would have dropped by 13.97% (while they had dropped by 10.23% in the data).

Second, we explore how robust our results are to different values of the parameter $((\sigma-1)\lambda)/(1+\lambda)$. Specifically, we compute the predicted contribution of the vent-for-surplus mechanism for values of $((\sigma-1)\lambda)/(1+\lambda)$ that are at the boundaries of the 95% confidence interval for this parameter implied by the estimates reported in column 3 of Table 10. For $((\sigma-1)\lambda)/(1+\lambda) = 1.343$ (i.e., the lower boundary of the confidence interval), the predicted contribution of vent-for-surplus mechanism to the observed boom-to-bust change in total exports is 35.03%; for $((\sigma-1)\lambda)/(1+\lambda) = 3.405$ (i.e., the higher boundary of the interval), this predicted contribution is 64.39%.[56]

8 Conclusion

In this paper, we provide evidence suggesting that export and domestic sales decisions are interdependent at the firm level. Faced with a severe domestic slump during the Great Recession, Spanish producers appear to have experienced a decline in their short-run marginal production costs, with

[55] If it were true that, as we impose in our analysis, the boom-to-bust change in the aggregate demand shifters $\{Q_{sd}\}_{s=1}^{S}$ was the same across all manufacturing sectors $s = 1, \ldots, S$, then we can infer from our analysis that the aggregate demand shifters fall between the boom and the bust in $1 - 1/\Gamma_Q^* = 1 - 1/1.09 = 1 - 0.9174 = 8.26\%$.

[56] Table K.1 in Appendix K.3 presents results on the implied contribution of the vent-for-surplus mechanism to the boom-to-bust growth in exports for several additional values of the parameter $((\sigma-1)\lambda)/(1+\lambda)$.

this gain in competitiveness translating into an increase in their sales in foreign markets. We circumvent the inherent difficulties associated with establishing a causal link between demand-driven changes in domestic sales and exports by exploiting geographic variation in the incidence of the Great Recession in Spain.

Our empirical findings are inconsistent with international trade models featuring constant markups and technologies with constant marginal costs of production. We rationalize and interpret our results through the lens of a model with increasing marginal costs, and show that the "vent-for-surplus" mechanism is powerful enough to explain approximately half of the growth in Spanish exports in the period 2009-13. Although there are a few singular aspects of the Spanish experience in the years around the Great Recession that may make the "vent-for-surplus" mechanism particularly important in this context (e.g., the large boom in investment preceding the bust, or the sclerotic nature of the Spanish labor market), we think that some of the insights and methodology in our paper can be transported to other countries that experienced severe domestic slumps.

Our paper also offers a new perspective on the literature studying interdependencies in the extensive and intensive margins of trade (e.g., Antràs et al. (2017) or Morales et al. (2019). Due to data limitations, we have restricted our analysis to the study of interdependencies between the domestic market and a single (aggregate) export destination, and we have modeled these interdependencies as arising exclusively from an increasing marginal cost function. In such a case, the firm's profit function is submodular in the extensive margin of trade, a feature that would be preserved in a model with multiple export markets, as shown in Appendix E.2. With access to data on firms' exports by destination market, and borrowing tools from Arkolakis and Eckert (2017), one could estimate the key parameters of a multi-country version of our model, and thus explore interdependencies also in the intensive and extensive margin of trade *across export markets*. Even more ambitiously, with better data on prices, one could potentially expand our analysis to explore the extent to which endogenous markup adjustments (see De Loecker et al., 2016) or price stickiness in both the domestic and export prices (Gopinath and Rigobon, 2008; Nakamura and Steinsson, 2008, 2013) affect the way in which a firm's exports react to domestic demand shocks. We leave the study of these questions for future research.

References

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales (2019), "Shift-Share Designs: Theory and Inference," The Quarterly Journal of Economics, Vol. 134, No. 4, pp. 1949-2010.

Antràs, Pol, Teresa C. Fort, and Felix Tintelnot (2017), "The Margins of Global Sourcing: Theory and Evidence from US Firms." *American Economic Review*, Vol. 107, No. 9, pp. 2514-2564.

Ahn, JaeBin, and Alexander F. McQuoid (2017), "Capacity Constrained Exporters: Identifying Increasing Marginal Cost," Economic Inquiry, Vol. 55, No. 3, pp. 1175-1191.

Akin, Ozlem, José García Montalvo, Jaume García Villar, José-Luis Peydró, and Josep María Raya (2014), "The Real Estate and Credit Bubble: Evidence from Spain," SERIEs, Vol. 5, No. 2–3, pp. 223–243.

Almunia, Miguel, David Lopez-Rodriguez, and Enrique Moral-Benito (2018), "Evaluating the Macro-representativeness of a Firm-level Dataset: An Application for the Spanish Economy," mimeo Bank of Spain.

Antràs, Pol (2011), "El Comportamiento de las Exportaciones Españolas," Apuntes Fedea, Vol 2, pp. 1-11.

Arkolakis, Costas, and Fabian Eckert (2017), "Combinatorial Discrete Choice," mimeo Yale University.

Autor, David H., David Dorn, and Gordon H. Hanson (2013), "The China Syndrome: Local Labor Market Effects of Import Competition in the United States,"American Economic Review, Vol. 103, No. 6, pp. 2121-2168.

Basco, Sergi and David Lopez-Rodriguez (2018), "Credit Supply, Education and Mortgage Debt: The BNP Securitization Shock in Spain," mimeo Bank of Spain.

Berman, Nicolas, Antoine Berthou, and Jèrôme Héricourt (2015), "Export Dynamics and Sales at Home," Journal of International Economics, Vol. 96, No. 2, pp. 298-310.

Bilir, L. Kamran, and Eduardo Morales (2020), "Innovation in the Global Firm," Journal of Political Economy, Vol. 128, No. 4, pp. 1566-1625.

Blum, Bernardo S., Sebastian Claro, and Ignatius J. Horstmann (2013), "Occasional and Perennial Exporters," Journal of International Economics, Vol. 90, No. 1, pp. 65-74.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2019), " Quasi-Experimental Shift-Share Research Designs," mimeo University College London.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008), "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," IMF Staff Papers, Vol. 55, No. 3, pp. 511-540.

De Loecker, Jan, Pinepoli K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik (2016), "Price, Markups, and Trade Reform," Econometrica, Vol. 84, No. 1, pp. 445-510.

De Lucio, Juan, Raúl Mínguez, Asier Minondo, and Francisco Requena (2017a), "Los Márgenes del Crecimiento de las Exportaciones Españolas Antes y Después de la Gran Recesión," Estudios de Economía Aplicada, Vol. 35, No. 1, pp. 43-62.

De Lucio, Juan, Raúl Mínguez, Asier Minondo, and Francisco Requena (2017b), "Is Spain Experiencing an Export Miracle?", Funcas SEFO Vol. 7, No. 4, pp. 53-66.

Díaz-Lanchas, Jorge, Carlos Llano, and José Luis Zofío (2013), "Trade Margins, Transport Cost Thresholds and Market Areas: Municipal Freight Flows and Urban Hierarchy," Working Papers in Economic Theory 2013/10, Universidad Autónoma de Madrid (Spain), Department of Economic Analysis (Economic Theory and Economic History).

Dickstein, Michael, and Eduardo Morales (2018), "What do Exporters Know?" The Quarterly Journal of Economics, Vol. 133, No. 4, 1753-1801.

Dolado, Juan J., Carlos García-Serrano, and Juan F. Jimeno (2002), "Drawing Lessons from the Boom of Temporary Jobs in Spain," Economic Journal, Vol. 112, No. 480, pp. 270-295.

Donaldson, Dave, and Richard Hornbeck (2016), "Railroads and American Economic Growth: A "Market Access" Approach", The Quarterly Journal of Economics, Vol. 131, No. 2, pp. 799-858.

Eppinger, Peter S., Nicole Meythaler, Marc-Manuel Sindlinger, and Marcel Smolka (2017), "The Great Trade Collapse and the Spanish Export Miracle: Firm-level Evidence from the Crisis," The World Economy, Vol. 41, No. 2, pp. 457-493.

Fan, Haichao, Yu Liu, Larry Qiu, and Xiaoxue Zhao (2018),"Export to Elude," mimeo University of Hong Kong.

Fisher, Eric O'N, and Vikas Kakkar (2004), "On the Evolution of Comparative Advantage in Matching Models," Journal of International Economics, Vol. 64, No. 1, pp. 169-193.

Gandhi, Amit, Salvador Navarro, and David Rivers (2016), "On the Identification of Production Functions: How Heterogeneous is Productivity?" mimeo University of Pennsylvania.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020), "Bartik Instruments: What, When, Why and How," mimeo Stanford University.

Gopinath, Gita, and Roberto Rigobon (2008), "Sticky Borders',' Quarterly Journal of Economics, Vol. 123, No. 2, 531-575.

Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson (2020), "Housing Wealth Effects: The Long View", Review of Economic Studies, forthcoming.

Hanson, Gordon H. (2005), "Market potential, Increasing Returns and Geographic Concentration," Journal of International Economics, Vol. 67, pp. 1-24.

Harris, Chauncy D. (1954), "The Market as a Factor in the Localization of Industry in the United States," Annals of the Association of American Geographers, Vol. 54, No. 4, pp. 315-348.

Hausman, Joshua K., Paul W. Rhode, and Johannes F. Wieland (2019), "Recovery from the Great Depression: The Farm Channel in Spring 1933," American Economic Review, Vol. 109, No. 2, pp. 427-472.

Head, Keith and Thierry Mayer (2014), "Gravity Equations: Workhorse, Toolkit, and Cookbook,"Chapter 3 in Gopinath, G, E. Helpman and K. Rogoff (eds.), vol. 4 of the Handbook of International Economics, Elsevier, pp. 131-195.

Hillberry, Russell and David Hummels (2008), "Trade Responses to Geographic Frictions: A Decomposition Using Micro-Data." European Economic Review, Vol. 52, No. 3, pp. 527-550.

International Monetary Fund (2015), Crisis Program Review.

International Monetary Fund (2018), Spain, IMF Country Report No. 18/330.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante (2020), "Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data", mimeo University of Chicago.

Krugman, Paul R. (1984) "Import Protection as Export Promotion: International Competition in the Presence of Oligopoly and Economies of Scale," in H. Kierzkowski (ed.), Monopolistic Competition and International Trade, Clarendon Press, Oxford.

Masier, Giacomo and Ernesto Villanueva (2011), "Consumption and Initial Mortgage Conditions" ECB Working Paper #1297, February.

Melitz, Marc J. (2003), "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," Econometrica, Vol. 71, No. 6, pp. 1695-1725.

Mian, Atif, Kamalesh Rao, and Amir Sufi (2013), "Household Balance Sheets, Consumption, and the Economic Slump." The Quarterly Journal of Economics, Vol. 128, No. 4, pp. 1687-1726.

Morales, Eduardo, Gloria Sheu, and Andrés Zahler (2019), "Extended Gravity," The Review of Economic Studies, Vol. 86, No. 6, pp. 2668-2712.

Myint, Hla (1958), "The 'Classical Theory' of International Trade and the Underdeveloped Countries," Economic Journal, Vol. 68, No. 270, pp. 317-337.

Myro, Rafael (2015), "España en la Economía Global: Claves del Éxito de las Exportaciones Españolas", RBA Libros.

Nakamura, Emi, and Jón Steinsson (2008) "Five Facts About Prices: A Reevaluation of Menu Cost Models," Quarterly Journal of Economics, Vol. 123, No. 4, 1415-1464.

Nakamura, Emi, and Jón Steinsson (2013) "Price Rigidity: Microeconomic Evidence and Macroeconomic Implications," Annual Review of Economics, Vol. 5, 133-163.

Neary, Peter J. and Albert G. Schweinberger (1986), "Factor Content Functions and the Theory of International Trade." Review of Economic Studies, Vol. 53, No. 3, pp. 421-432.

Redding, Stephen, and Anthony J. Venables (2004), "Economic Geography and International Inequality", Journal of International Economics, Vol. 62, pp. 53-82.

Salas, Jorge (2018), "Drivers of Spain's Export Performance and the Role of Labor Market Reforms." IMF Working Paper.

Santos, Tano (2014), "Antes del Diluvio: The Spanish Banking System in the First Decade of the Euro," mimeo Columbia Business School.

Soderbery, Anson (2014), "Market Size, Structure, and Access: Trade with Capacity Constraints." European Economic Review, Vol. 70, pp. 276-298.

Stock, James H. and Mark W. Watson (1999), "Chapter 1: Business Cycle Fluctuations in US Macroeconomic Time Series." Handbook of Macroeconomics, Elsevier, Volume 1, Part A, pp. 3-64.

Vannoorenberghe, Gonzague (2012), "Firm-level Volatility and Exports." Journal of International Economics, Vol. 86, No. 1, pp. 57-67.

Waugh, Michael (2019), "The Consumption Response to Trade Shocks: Evidence from the US-China Trade War," mimeo New York University.

Williams, J. H. (1929), "The Theory of International Trade Reconsidered," Economic Journal, Vol. 39, No. 154, pp. 195-209.

A Convexity of the Short-run Marginal Cost Function

Suppose a firm's production function depends on fixed or pre-determined input K_i and a flexible and static input L_i. Let us refer to the former as capital and the latter as labor. Assuming a Cobb-Douglas technology in capital and labor, the cost minimization problem of a firm with productivity φ_i seeking to produce a total amount of output Q_i can be expressed as:

$$\begin{aligned} \min \quad & \omega_i L_i \\ s.t. \quad & \varphi_i K_i^{\alpha_K} L_i^{\alpha_L} \geq Q_i, \end{aligned}$$

where ω_i denotes the nominal wage that firm i faces, and α_K and α_L denote then the output elasticities with respect to capital and labor, respectively. The first-order condition of the cost-minimization problem of the firm delivers

$$\begin{aligned} \omega_i &= \mu \alpha_L \frac{Q_i}{L_i} \\ \varphi_i K_i^{\alpha_K} L_i^{\alpha_L} &= Q_i, \end{aligned}$$

where μ denotes the Lagrange multiplier on the constraint $\varphi_i K_i^{\alpha_K} L_i^{\alpha_L} = Q_i$. After solving for L_i in the second of these equalities, we can rewrite the short-run total costs as a function of output, Q_i, as follows

$$\omega_i L_i = \omega_i \left(\varphi_i K_i^{\alpha_K}\right)^{-\frac{1}{\alpha_L}} \left(Q_i\right)^{\frac{1}{\alpha_L}} .$$

Note that, unless $\alpha_L = 1$, this short-run total cost function is not linear in the total output Q_i. More specifically, as long as $0 < \alpha_L < 1$, this short-run cost function will be convex in Q_i. Using $\widetilde{\varphi}_i$ to denote a shifter of the short-run costs, and using λ to denote the deviation of the output elasticity of the short-run cost function relative to the case in which this function is linear in Q_i, i.e.,

$$\begin{aligned} \widetilde{\varphi}_i &= \alpha_L \left(\varphi_i K_i^{\alpha_K}\right)^{\frac{1}{\alpha_L}} \\ \lambda &= \frac{1-\alpha_L}{\alpha_L}, \end{aligned}$$

we can rewrite the short-run total costs as

$$\omega_i L_i = \frac{1}{\widetilde{\varphi}_i} \omega_i \frac{1}{1+\lambda} (Q_{it})^{1+\lambda}.$$

The elasticity of the short-run total costs with respect to output is thus

$$\frac{\partial \ln(\omega_i L_i)}{\partial \ln(Q_{it})} = 1 + \lambda.$$

Note that, the lower the value of α_L (i.e., the lower the elasticity of output with respect to the flexible input), the larger the elasticity with respect to output of the short-run total cost function. The curvature of the total cost schedule is thus crucially shaped by the parameter determining the elasticity of output with respect to the flexible input.

B Data Appendix

B.1 Macroeconomic Data

Data on Spanish unemployment, real GDP, internal demand, private final consumption expenditure and exports of goods come from the Spanish National Statistical Office (*Instituto Nacional de Estadística*). Data on merchandise exports and real GDP shares for the countries that belong to the European Monetary Union come from the AMECO Dataset (i.e., annual macro-economic database of the European Commission's Directorate General for Economic and Financial Affairs). Data on unit labor costs in the manufacturing sector for Spain and the European Monetary Union were obtained from the Bank of Spain (*Banco de España*), the Eurosystem, and the OECD dataset on Productivity and ULC by main economic activity. We use the input-output tables produced by the Spanish National Statistical Office for the year 2005 to identify the interlinkages across industries (e.g., the two top leading input providers or two top leading buying industries of the vehicles manufacturing industry discarded in the robustness analysis described in Table 6).

B.2 Construction of the Commercial Registry Dataset

As described in section 3.3, our main source of firm-level data is the Commercial Registry (*Registro Mercantil Central*), which contains annual financial statements of around 85% of registered firms in the non-financial economy. We collect data from two separate sources to construct our own firm-level dataset: (i) the *Central de Balances* dataset from the Bank of Spain, and (ii) *Sabi*, from Informa, a private company. The Bank of Spain made an effort to expand and treat the information for small firms gathered in the Commercial Registry but the Central de Balances dataset does not cover the universe of private-sector firms. In particular, this dataset excludes, mainly, medium and large firms that submit information after the regular submission deadline or that do not use a digital support. Conversely, the Informa dataset puts special emphasis on compiling information on large and medium-sized firms that submit their statements either late or on paper. We combine the information in these two datasets to take advantage of their complementarities in order to maximize the coverage of the resulting database. A detailed description of how we combine the two sources to construct our firm-level dataset can be found in Almunia et al. (2018).

B.3 Foreign Transactions Dataset

As described in section 3.3, until 2014 the Bank of Spain required all financial institutions and a set of large companies to report all foreign transactions, including imports, exports and other financial transactions. Until 2007, there is information for each transaction on the country of destination (or origin). However, from 2008 to 2013, the Bank of Spain relaxed this requirement and allowed reporting institutions to group multiple transactions into a single reported transaction. In those cases, the country of destination (or origin) reflected in the data corresponds to the country of the largest transaction in that group. Similarly, the product code reported corresponds to the largest transaction as well. This implies that one cannot analyze changes in exports or imports by country of destination (or origin) nor by product in a consistent way for periods spanning around 2008. The foreign transactions registry collected by the Bank of Spain was discontinued in early 2014. Since then, the Bank of Spain's monitoring of foreign transactions mainly relies on aggregate data built from transaction-level information that is provided by the Spanish tax administration.

B.3.1 Minimum Reporting Threshold

Between 2001 and 2007, all foreign transactions of more than €12,500 had to be reported to the Bank of Spain. In order to reduce the compliance costs for reporting institutions, the minimum reporting threshold was updated in 2008 to €50,000. From that year onwards, a firm appears in the dataset if it has at least one transaction larger than €50,000 in that year. In order to create a homogeneous sample for the period 2002-2013, we apply the post-2008 minimum reporting threshold to the data from 2002 to 2007, meaning that we only record a positive export flow in a given year for firms that have at least one transaction exceeding €50,000 in that year. This adjustment reduces substantially the number of exporting firms that appear in the data, but the impact on the aggregate amount exported is small.

B.4 Instruments: Vehicles per Capita and Tax Records of Firm-Level Sales within Spain

The information on the stock of vehicles by both municipality and province is provided by the Spanish Registry of Motor Vehicles. According to Spanish Law, vehicles have to be registered in the municipality where the owner has her permanent residence. This residence status should match the one reported in the municipal census (*Padrón*). In the case of legal entities (business or institutions), vehicles must be registered in the municipality where they undertake their main activity, which should match the location reported to the tax authorities. The census of vehicles at the municipal level is maintained by the General Directorate for Traffic (*Dirección General de Tráfico*). Each city council has the capacity to levy a small fee on the registered vehicles in its municipality (*Impuesto sobre Vehículos de Tracción Mecánica*). This fee usually depends on several criteria such as vehicle power, type of vehicle, pollution level, etc. In aggregate terms, these fees collectively raised tax revenue equivalent to around 0.2% of GDP in 2016.

The information on population both at the municipality and province levels is provided by the Spanish National Statistical Office.

Regarding our data on firm-level sales within Spain, Spanish Tax Law obligates all firms (legal entities) and professionals (natural persons) that undertake economic activities to report detailed information on the transactions with their trading partners. In 2006, this information is collected in Form 347, officially called "Annual information return on transactions with third parties". We work with data for the year 2006 because it is the first year for which a precise and consistent comprehensive digitization of the data is available. In particular, each business must report the monetary value of its individual sales to each trading partner. The reported transactions include all domestic sales to businesses, households and the public sector. The law uniquely exempts mandatory reporting of individual transactions when the annual aggregate sales to a trading partner do not exceed €3,005.06. This tax record of sales to third-parties is a fundamental tool of tax enforcement for both the VAT and the corporate income tax given that transactions included in Form 347 must be reported consistently in both tax returns.

The Spanish Tax Agency (*Agencia Estatal de Administración Tributaria*, AEAT) shared with the Bank of Spain aggregate data on municipality-to-municipality flows for firms in the manufacturing sector, excluding sales of businesses in the auto industry, for the year 2006. In particular, for each municipality where manufacturing sellers are located, the Spanish Tax Agency computed the total amount of sales to any municipality where purchases are made. We thus have access to a matrix of bilateral flows of manufactured goods between Spanish municipalities. This matrix contains data on 485,565 municipality-to-municipality flows, with 2,305 municipalities of origin and 6,623 municipalities of destination. After restricting the sample of municipalities to those observed in our subsample of continuing exporters, the dataset contains a matrix of 412,500 bilateral municipality flows, with 1,224 municipalities of origin and 6,587 municipalities of destination.

Apart from the matrix of aggregate municipality-to-municipality flows, the Spanish Tax Agency also provided us, for each firm in our sample of 8,009 continuing exporters, the 2006 share of its total domestic sales going to each Spanish municipality.