

Predicting Exporters with Machine Learning

Francesca Micocci*

Armando Rungi†

This version: June 2021

Abstract

In this contribution, we exploit machine learning techniques to predict out-of-sample firms' ability to export based on the financial accounts of both exporters and non-exporters. Therefore, we show how forecasts can be used as exporting scores, i.e., to measure the distance of non-exporters from export status. For our purpose, we train and test various algorithms on the financial reports of 57,021 manufacturing firms in France in 2010-2018. We find that a Bayesian Additive Regression Tree with Missingness In Attributes (BART-MIA) performs better than other techniques with a prediction accuracy of up to 0.90. Predictions are robust to changes in definitions of exporters and in the presence of discontinuous exporters. Eventually, we argue that exporting scores can be helpful for trade promotion, trade credit, and to assess firms' competitiveness. For example, back-of-the-envelope estimates show that a representative firm with just below-average exporting scores needs up to 44% more cash resources and up to 2.5 times more capital expenses to reach full export status.

Keywords: exporting; machine learning; trade promotion; trade finance; competitiveness.

JEL Codes: F17; C53; C55; L21; L25

We want to thank Tommaso Aquilante, Kristina Bluwstein, Falco Bargagli Stoffi, Mahdi Ghodsi, Andreas Joseph, Massimo Riccaboni, Michele Ruta, Beata Smarzynska Javorcik, Maurizio Zanardi, and participants to the seminar series jointly organized by FIW and WüW in Vienna for valuable comments. Armando Rungi claims financial support from Artes 4.0 - Industry 4.0 Competence Center on Enabling Digital Technologies and Systems.

*Mail to: francesca.micocci@imtlucca.it. Laboratory for the Analysis of Complex Economic Systems, IMT School for Advanced Studies, piazza San Francesco 19 - 55100 Lucca, Italy.

†Mail to: armando.rungi@imtlucca.it. Laboratory for the Analysis of Complex Economic Systems, IMT School for Advanced Studies, piazza San Francesco 19 - 55100 Lucca, Italy.

1 Introduction

Building trade capacity is a purpose of many international and national agencies. The World Trade Organization provides special support programs for developing countries to better integrate into the multilateral trading system. On the other hand, many developing and developed economies prefer to establish their facilitative agencies to provide firms with information, technical advice, marketing services, and policy advocacy about access to foreign markets. The general idea is that there are opportunities for gains from trade, yet not all firms have the same ability to sell their goods and services abroad. Exporting entails beach-head costs when handling different regulatory environments, meeting different consumer tastes, and establishing marketing and logistics channels. Only some more productive firms may have the ability to self-select into exporting status, while others may not have the necessary skills or resources to start with¹. Hence, the necessity to resort to trade promotion programs to fill the gap and help firms build trade capacity to take advantage of open markets. Eventually, openness to trade is a determinant of economic growth insofar as it allows exploiting differential comparative advantages and economies of scale while tapping into foreign technology and raising aggregate productivity in the home countries².

Against the previous background, our simple intuition is to adopt machine learning techniques to predict exporters and non-exporters based on the assumption that firms' accounts convey non-trivial information on firm-level trade capacity. In other words, we propose to train an algorithm on in-sample financial statements to predict out-of-sample firms' ability to start exporting. We perform a predictive exercise on a sample of French manufacturing firms that may have exported or not in 2010-2018. Thus, we randomly partition the dataset in an 80-20 proportion in training vs. testing sets. First, we train different models on training sets armed with a battery of 52 predictors derived from financial statements, and then we predict export status on testing sets. Eventually, we find that we can correctly separate firms with different export status with an accuracy of about 90%. The latter is a figure that we obtain from a horse race among different algorithms, after which we find that the winner is a Bayesian Additive Regression Tree with Missingness not at Random (BART-MIA) (Kapelner and Bleich, 2015). The BART is a regression tree with a Bayesian component for regularization through a prior specification that allows flexibility in fitting a

¹For a review of the arguments according to which only the most efficient firms are able to self-select into an export status and the consequences on the sources of gains from trade, see among others [Bernard and Jensen \(1999\)](#); [Bernard et al. \(2012\)](#); [Melitz and Redding \(2014\)](#); [Hottman et al. \(2016\)](#)

²Seminal works identify macroeconomic linkages between trade openness, technological progress, and economic growth. See [Grossman and Helpman \(1990\)](#), [Rivera-Batiz and Romer \(1991\)](#), [Romer \(1994\)](#), [Barro and Sala-i Martin \(1997\)](#).

variety of regression models while avoiding strong parametric assumptions (Hill et al., 2020). In particular, the BART-MIA variant exploits additional predictive power from non-random missing values on predictors. The latter is a feature that is especially useful in catching business dynamics when coverage of financial accounts is likely to be correlated with other dimensions, e.g., firms’ size or productivity, which in turn can correlate with firms’ export status. Crucially, considering missing values as predictors helps in increasing prediction accuracy about 14.4%. Eventually, we make sure that prediction accuracies are robust to different definitions of exporters and we test the the model performance when we consider cases of heterogeneous exporting patterns and discontinuous exporters (Geishecker et al., 2019; Békés and Muraközy, 2012). Our framework is also robust to different cross-validation strategies since we obtain similar performance by randomly picking training and testing subsets in different ways, albeit from a unique sample. Finally, we test that different subsets of predictors would not bring the same high levels of prediction accuracies after we perform a Least Absolute Shrinkage and Selection Operator (LASSO) for dimensionality reduction in predictors (Belloni et al., 2013, 2014, 2016; Ahrens et al., 2020).

In the second part of the paper, we discuss how our prediction exercise can be useful to assess the distance of a non-exporter from export status, i.e., how far a firm is from becoming an exporter. We suggest looking at baseline predictions to attribute a probabilistic exporting score to a firm, i.e., a score summarising how similar a non-exporter is to benchmark exporters on a scale from 0 to 1. We believe that such exporting scores could be helpful for trade promotion or trade finance programs. Therefore, to illustrate the utility of exporting scores, we classify firms into risk categories and provide a simple back-of-the-envelope calculation for how much cash resources and capital expenses they would need to reach export status. We find that increasing cash and capital is needed to reduce the distance from export status. For example, in the case of medium-risk firms, when they have just below 50% probability of exporting, we show a need of up to 44% more cash resources and up to 246% more capital expenses to reach full export status.

Finally, we show how exporting scores can be used as an additional tool to describe trade competitiveness. Once we consider the French case study, we observe that there is high heterogeneity in the potential for exporting across industries and regions. For example, the North-West hosts a relatively higher number of potential exporters than the rest of the country. A variety of industries presents a high share of potential exporters, including refineries, producers of rubber, plastic, paper products, and manufacturers of basic metals.

The remainder of the paper is organized as follows. In Section 2 we relate to previous literature. We introduce data and sample coverage in Section 3, whereas Section 4 discusses the empirical strategy. Results are commented in Section 5, and a proposal for using exporting

scores is offered in Section 6. Section 7 concludes.

2 Related literature

Most countries around the world implement trade promotion programs. Thus, it is hardly surprising that there have been concerns about the efficacy and effectiveness of those support programs. Interestingly, [Volpe Martincus and Carballo \(2008\)](#) show how export promotion actions are associated with increased exports by already trading firms and traded products, i.e., the intensive margin. In terms of extensive margins, i.e., the increase of firms and products crossing national borders, [Volpe Martincus et al. \(2010\)](#) show that an influential role is often played by the establishment of diplomatic representations, especially in the case of producers of homogeneous goods. In general, the activation of new trading relationships may require a variety of services bundled together into export promotion programs ([Volpe Martincus and Carballo, 2010b](#)). Eventually, a majority of studies investigate how effective a policy is on the *ex-post* exporting performance while controlling for cherry-picking, as in [Volpe Martincus and Carballo \(2010a\)](#). In general, [Van Biesebroeck et al. \(2016\)](#) demonstrate that trade promotion programs have been a vital tool to overcome crises, as in the case of recovery after the global recession in 2008-2009.

In this context, our contribution focuses explicitly on the trade extensive margin since we aim to predict firms' ability to start exporting. From this perspective, we propose a pure prediction exercise based on the intuition that exporters are statistically different from non-exporters. In this sense, we rely on a two-decades-long strand of research that has established a connection between firms' heterogeneity and trading status ([Bernard and Jensen, 1999](#); [Melitz, 2003](#); [Melitz and Ottaviano, 2008](#); [Bernard et al., 2012](#); [Melitz and Redding, 2014](#); [Hottman et al., 2016](#)). Our intuition is that a prediction on export status is possible because we have prior knowledge that exporters do have different cost structures than non-exporters. After all, they have to sustain the fixed costs to gain access to foreign markets, where regulations and consumer tastes can be much different from home, and where shipping is costly. Thus, we demonstrate that starting from a comprehensive battery of economic and financial predictors allows indeed separating exporters from non-exporters with a relatively high prediction accuracy, up to 90%.

Please, however, note that ours is not a classic policy evaluation exercise because we do not assess whether any specific policy design works to support would-be exporters. Ours is a simple scoring exercise in the fashion of what one can find in previous literature about credit scoring, where there is a long tradition to try and spot firms in financial distress based on the disclosure of financial accounts. See seminal attempts with Z-scores by [Altman](#)

(1968); Altman et al. (2000) and Distance-to-Default by Merton (1974), where some specific thresholds is set as a rule of thumbs to say whether a firm is financially sound and worthy of credit. Nowadays, most financial institutions adopt predictive models to evaluate credit risk, including machine learning (Uddin, 2021). See also the exercises made on firm-level correlations to spot investment-to-cash-flow sensitivities and assess time-varying financial constraints (Fazzari et al., 1988; Almeida et al., 2004; Chen and Chen, 2012). The additional difficulty in our exercise is that we want to score success, i.e., the ability of a firm to outreach across national borders, whereas credit risk analyses take as reference previous firms' failures, i.e., the distance to default. Yet, from our perspective, the problem can take a similar approach: to get as benchmark firms (and their financial accounts) that realized an outcome, in our case export status, and thus measure how far we are from that benchmark.

Eventually, routine access to trade finance is needed, and well-functioning financial markets are crucial to export performance (Manova, 2012). External finance helps to gain and keep access to foreign markets despite the high beach-head costs they entail, especially in the case of smaller producers who have a reduced ability to provide collateral for trade credit (Chor and Manova, 2012). In this context, exporting scores can be as useful to financial institutions as to trade promotion agencies. As in credit scoring literature, we believe our perspective can be potentially valuable to better target credit policies by financial institutions in a familiar way, e.g., by considering credit risk classes. Hence, to better grasp our intuition, we propose a back-of-the-envelope exercise that estimates *ceteris-paribus* how much cash resources and capital expenses firms need to switch across low, medium and high-risk classes.

Moreover, from a macroeconomic viewpoint, one can use firms' scoring as yet another indicator of the competitiveness of an economy (or lack thereof). Inspired by so-called growth diagnostics, international and national statistics offices have developed frameworks for assessing the potential of countries, regions and industries to compete on international markets. See, for example, the work by the World Bank on measuring trade competitiveness (Reis et al., 2010; Gaulier et al., 2013). In the case of French manufacturing, we show how potential exporters are unevenly distributed across industries and regions. We believe there is no reason why an indicator like ours about the potential of extensive margins should not find room in a standard trade diagnostic kit.

Finally, we want to remark how ours is one of the first attempts to exploit statistical learning techniques in international economics. As far as we know, there are only a few notable efforts in progress, including Gopinath et al. (2020) and Breinlich et al. (2021). Yet, we firmly believe that statistical learning exercises have great potential and should find their way in a field where one often needs to extract information from big and complex data sets,

which can be dealt with by a combination of predictive tasks and standard causal inference exercises (Athey, 2018; Mullainathan and Spiess, 2017).

3 Data

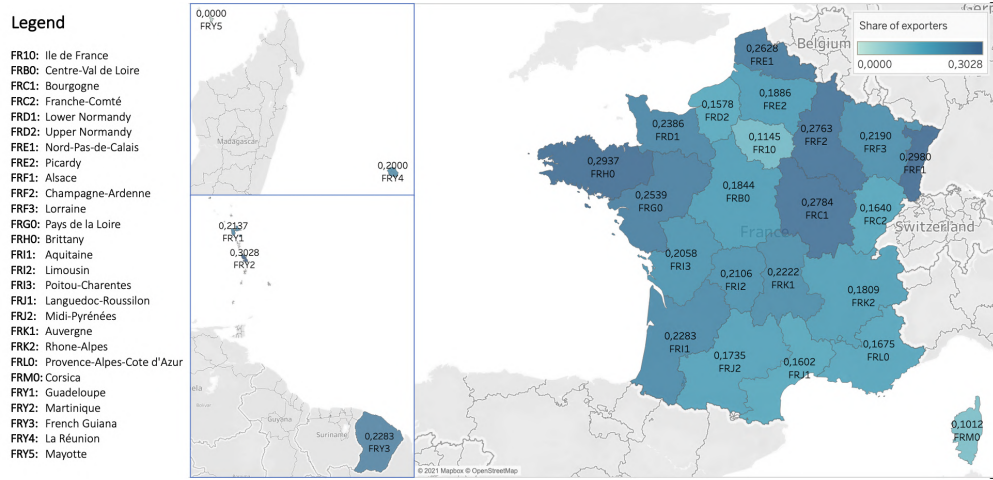
We source firm-level information from ORBIS³, compiled by the Bureau Van Dijk. Notably, France is a much-explored case study of firm-level trade data that allows us to confront previous literature. See among others Crozet et al. (2011) and Fontagné et al. (2018). Our main outcome of interest is the export status of a firm, which we derive from information on export revenues⁴. *Prima facie*, we will consider a firm as an exporter if it reports positive export revenues. Then, in Sections 5.2 and 5.4, we will challenge our baseline definition to comply with the phenomenon of temporary trade and discontinuous exporters (Békés and Muraközy, 2012), when it is optimal for firms to export every once in a while. As for firm-level predictors of exporting status, we employ a battery of 52 indicators elaborated on original financial accounts that we use to train our models. Further details on our choice are discussed in Section 4.2, while we include the list of predictors with a complete description in the Data Appendix.

To grasp the coverage of our sample, we draw Figure 1 and Table 1. Figure 1 shows how relevant exporters are in every NUTS-2 region in France, as from our sample. Table 1 compares sample industry coverage with the one provided by EUROSTAT census in 2018. We do find that we have a fair coverage by 2-digit industries since the correlation by industry shares is about 0.90. Yet, our sample covers 32.6% of firms' population, which, however, represents about 75% of total operating revenues in France according to Eurostat business demographics. As largely expected, we cannot retrieve financial accounts of smaller firms, because they are not required to comply with accounting regulations in the same way as medium and larger ones. See also a comparison by class categories with EUROSTAT in Appendix Table B1. In the following paragraphs, we will show how our baseline analysis can handle non-random missing values in financial information.

³The ORBIS database has become a standard source for global firm-level financial accounts. For a previous usage of this database, among others, see Gopinath et al. (2017), Cravino and Levchenko (2016), Del Prete and Rungi (2017), and Rungi and Del Prete (2018). It complements financial accounts with other information from different sources on ownership, corporate governance, and intellectual property rights, which we also use for predictions in the following analyses.

⁴Interestingly enough, French firms must report the amount of revenues from exports separately, as from the subsequently amended *Règlement n. 99-03 du Comité de la réglementation comptable*.

Figure 1: Sample coverage: exporters by region



Note: Unitary shares indicate exporters on total firms in NUTS 2-digit regions.

Table 1: Sample coverage by industry

NACE rev. 2	code	Population		Sample			
		Eurostat	(%)	non-exporters	exporters	total	(%)
Food products	10	51,288	0.29	13,057	1,429	14,486	0.25
Beverages	11	3,853	0.02	1,176	395	1,571	0.03
Textiles	13	5,076	0.03	919	389	1,308	0.02
Wearing apparel	14	9,694	0.06	1,060	336	1,396	0.02
Leather and related products	15	3,243	0.02	374	142	516	0.01
Wood and of products of wood and cork	16	9,956	0.06	2,203	509	2,712	0.05
Paper and paper products	17	1,292	0.01	455	362	817	0.01
Printing and reproduction of recorded media	18	15,316	0.09	2,995	584	3,579	0.06
Coke and refined petroleum products	19	35	0.01	17	14	31	0.01
Chemicals and chemical products	20	2,515	0.01	958	705	1,663	0.03
Basic pharmaceutical products and pharmaceutical preparations	21	252	0.01	151	148	299	0.01
Rubber and plastic products	22	3,205	0.02	1,436	931	2,367	0.04
Other non-metallic mineral products	23	7,803	0.04	1,929	393	2,322	0.04
Basic metals	24	599	0.01	354	267	621	0.01
Fabricated metal products, except machinery and equipment	25	18,460	0.11	8,135	2,540	10,675	0.19
Computer, electronic and optical products	26	2,295	0.01	965	605	1,570	0.03
Electrical equipment	27	2,048	0.01	789	495	1,284	0.02
Machinery and equipment	28	4,534	0.03	1,938	1,194	3,132	0.05
Motor vehicles, trailers and semi-trailers	29	1,635	0.01	748	424	1,172	0.02
Other transport equipment	30	1,107	0.01	330	186	516	0.01
Furniture	31	9,356	0.05	1,416	249	1,665	0.03
Other manufacturing	32	21,338	0.12	2,796	518	3,314	0.06
Total		174,890	1.00	44,201	12,815	57,016	1.00

Note: French manufacturing firms are sourced from Orbis, by Bureau Van Dijk. On the third column a comparison with Eurostat census. On columns 5 and 6, we separate exporters and non-exporters in our sample. When we look at shares on columns 4 and 8, we find our sample is well balanced by industry if compared with the population.

4 The empirical strategy

Our main intuition is that we can predict out-of-sample exporters based on the in-sample experience of both exporters and non-exporters. Thus, we can make use of the generic predictive model for firms' export status in the form:

$$f(\mathbf{X}_i) = Pr(Y_i = 1 \mid \mathbf{X}_i = x) \quad (1)$$

where Y_i is the binary outcome that assumes value 1 if the i th firm is exporting, and zero otherwise. \mathbf{X}_i is a P -dimensional matrix that includes a full battery of firm-level predictors, which we discuss in detail in the following Section 4.2. Please note that, at this stage, we do not consider the time dimension, i.e., we train the predictive model considering the export status of a firm in relation with present predictors. In this baseline model, it is entirely possible that a firm is considered as an exporter in one year and a non-exporter in another year. See Section 5.4 where we introduce the time dimension, thus looking at heterogeneous exporting patterns.

The functional form that links predictors to outcomes is *ex-ante* unknown and looked for by the generic supervised machine learning technique. We provide an overview of the different methods we use in Section 4.1. The advantage of any of them is to catch non-linearities that may be present in the association between export status and its predictors. Briefly, the generic predictive model has to pick the best in-sample loss-minimizing function in the form:

$$\arg \min \sum_{i=1}^N L(f(x_i), y_i) \quad \text{over } f(\cdot) \in F \quad \text{s. t.} \quad R(f(\cdot)) \leq c \quad (2)$$

where F is a function class from where to pick the specific function $f(\cdot)$. Importantly, $R(f(\cdot))$ is the generic regularizer that summarizes the complexity of $f(\cdot)$. The latter is a tool that allows us to solve the common trade-off between an as high as possible in-sample fit and an as high as possible flexibility of the prediction model, able to take on board new out-of-sample information. It is the solution to the so-called bias-variance trade-off. The set of regularizers, R 's, will change following standards proposed by each method that we will compare in the following paragraphs. Eventually, any method shall minimize the constrained loss function represented in eq. 2, while searching for the function that can be better used to process new out-of-sample information.

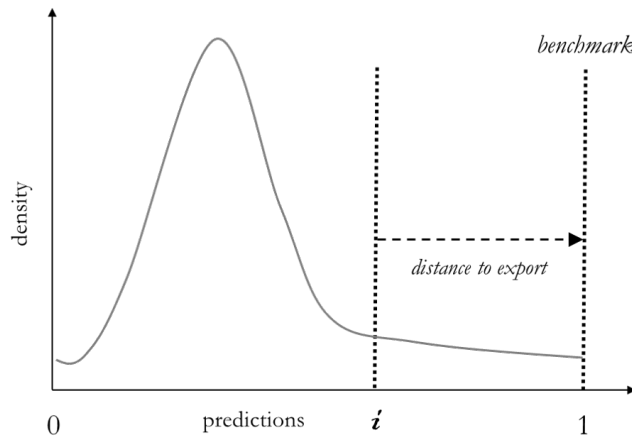
As a common strategy across our different models, we will pick at random 80% of our French firms to be considered as in-sample information and use it to train the generic statistical learning algorithm while keeping the remaining 20% as out-of-sample information

to predict export status. Hence, we will be able to assess the accuracy of our predictions within the limit of our data sources. As it is standard in similar exercises, we perform a cross-validation described in Section 5.2, to check that a specific segment of the sample does not affect predictions and related accuracies despite the initial random 80 – 20 partition.

Thus, once we assess the method that assures the best predictive accuracy with the minimum numbers of false positives and false negatives (see Section 5.1), we propose to use predictions to attribute each out-of-sample firm an exporting score bounded by construction in an interval from 0 to 1. Our main intuition is that we can use prediction scores to catch the *distance to export* of non-exporters, i.e., how suitable each out-of-sample firm is to access foreign markets. We further discuss switching from binary to continuous predictions on exporting status in Section 6. In Figure 2, we report a visual fictional representation of our intuition.

Assuming that we did a good job in training and that prediction accuracy is acceptable, we can reasonably locate actual exporters at the end of the right tail of the distribution of exporting predictions. Thus, any i th non-exporting firm located on the left of predicted exporters will come with a positive distance from exporters, which will convey non-trivial information on how viable that firm is to start exporting. In other words, we take as a reference point the maximum exporting scores a firm can obtain, and thus check how far we are from that reference point, where we checked that a firm is certainly fit for export.

Figure 2: Visual intuition of an exporting score.



Note: We represent a fictional distribution of predictions of export status that is by definition bounded in an interval $[0, 1]$. Along the distribution, we could spot an i -th non-exporting firm. We reasonably assume that actual exporters locate at the end of the right tail. By definition, non-exporters are less and less likely to start exporting at an increasing distance from predicted exporters.

4.1 Methods

To get our best predictions, we train and compare different statistical learning techniques. In the following paragraphs, we show how a specific variant of the Bayesian Additive Regression Tree (BART) performs better than others, because it is able to consider the presence of non-random missing values as a further predictor for the outcome. The variant we use is the BART with Missingness In Attributes (BART-MIA). For more details, see also [Kapelner and Bleich \(2015\)](#). For a previous application to firms' dynamics, see [Bargagli-Stoffi et al. \(2020\)](#).

In general, any regression tree \mathcal{T} is built on *if-then* statements that split the training data according to the observed values of predictors, allowing for non-linear relationships between the predictors and the outcomes. Thus, the generic algorithm for the construction of a regression tree, \mathcal{T} , is based on a top-down approach that recursively splits the main sample into non-overlapping sub-samples (i.e. the nodes and the leaves). Therefore, the tree is pruned iteratively with the generic regularizer R to improve its predictive ability while avoiding overfitting in case trees develop along too many layers ⁵.

As in the baseline version ([Chipman et al., 2010](#)), BART-MIA elaborates a sum-of-trees model by imposing a prior that regularizes the fit by keeping the individual trees' effects small in an adaptive way. The Bayesian component of the technique is a prior that helps in iterations constructing and fitting successive residuals. The result is a sum of trees, each of which explains a small and different portion of the predictive function. The BART-MIA variant we adopt can be expressed as:

$$\mathbb{P}(Y = 1|\mathbf{X}) = \Phi \left(\mathcal{T}_1^{\mathcal{M}}(\mathbf{X}) + \dots + \mathcal{T}_q^{\mathcal{M}}(\mathbf{X}) \right), \quad (3)$$

where Φ denotes the cumulative density function of the standard normal distribution and the q distinct binary trees are denoted by \mathcal{T} , each being a single tree coming with an entire structure made of nodes and leaves. The sum-of-trees model serves as an estimate of the conditional probit at \mathbf{x} , which can be easily transformed into a conditional probability estimate of $Y = 1$. The Bayesian component of the BART includes three priors that have demonstrated to use efficiently the data at disposal:

1. the prior on the probability that a node will split at depth k is $\beta(1 + k)^{-\eta}$, where $\beta \in (0, 1)$, $\eta \in [0, \infty)$, and the hyper-parameters are chosen to be $\eta = 2$ and $\beta = 0.95$;
2. the prior on the probability distribution in the leaves is a normal distribution with zero

⁵It is beyond the scope of this paper to get into further details of single techniques. For a deeper introduction to statistical learning, we refer to [Hastie et al. \(2017\)](#).

mean: $\mathcal{N}(0, \sigma_q^2)$, where $\sigma_q = 3/d\sqrt{q}$ and $d = 2$;

3. the prior on the error variance is $\sigma^2 = 1$.

In addition to the Bayesian component, the BART-MIA variant augments the original algorithm by exploiting information on missing values and splitting on *missingness* features that are used as additional predictors in each binary-tree component.

Eventually, the BART-MIA is chosen in the following paragraphs as the baseline method after a comparison with four other alternatives. At first, we compare with a simple logistic regression (LOGIT) as the latter is a classical econometric technique for binary outcomes with a specific *ex-ante* assumption on the functional form linking predictors with the outcome. Then, we perform three other methods based on regression trees, namely a Classification and Regression Tree (CART) (Breiman et al., 1984), a Random Forest (RF) (Breiman, 2001), and the original unaugmented BART. CART is the most basic regression tree, while RF is an ensemble method that aggregates different regression trees to get a stronger predictive power, as the BART does, but without a Bayesian framework. Finally, we compare previous regression trees' models with the Least Absolute Shrinkage and Selection Operator (LASSO), in the form:

$$\arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2N} \sum_{i=1}^N \left(y_i(x_i^T \beta) - \log(1 + e^{(x_i^T \beta)}) \right)^2 \quad \text{subject to } \|\beta\|_1 \leq k. \quad (4)$$

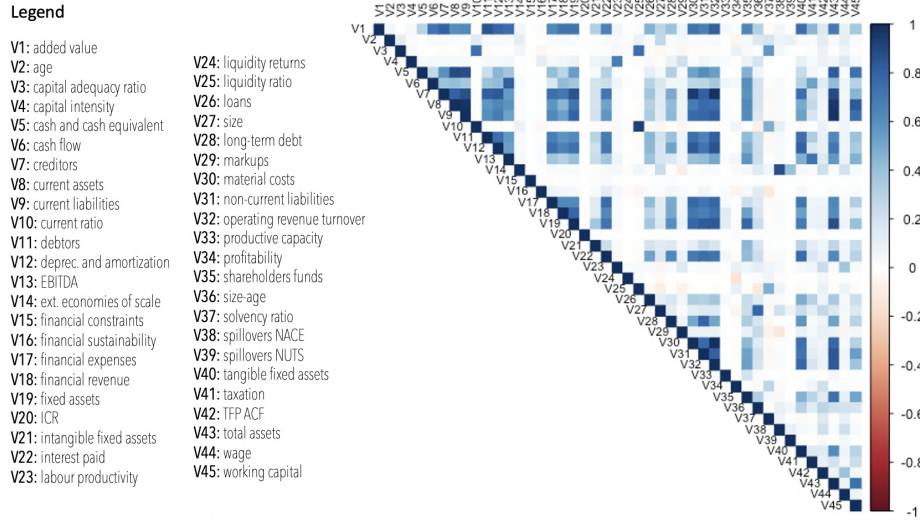
where y_i is a binary variable equal to one if a firm i is an exporter and zero otherwise. Any x_i is a predictor chosen in \mathbb{R}^p , whereas $\|\beta\|_1 = \sum_{j=1}^p |\beta_j|$ and $k > 0$. The constraint $\|\beta\|_1 \leq k$ limits the complexity of the model to avoid overfitting, and k is chosen, following Ahrens et al. (2020), as the value that maximises the Extended Bayesian Information Criteria (Chen and Chen, 2008). To account for the potential presence of heteroskedastic, non-Gaussian and cluster-dependent errors, we adopt the rigorous penalization introduced by Belloni et al. (2016).

4.2 Predictors

To increase predictability, we include a full battery of 52 predictors that are derived from firms' balance sheets and profit and loss accounts. A detailed description is reported in the Data Appendix. Broadly speaking, we choose to include:

1. original financial accounts without any elaboration;

Figure 3: Correlation matrix of predictors



Note: We report a correlation matrix of the predictors we use. Non-numeric predictors are excluded yet included in following analyses: NUTS-2 locations, NACE Rev.2 industries, a categorical variable for consolidated accounts, patents' dummy, inward FDI, outward FDI, and corporate control. Positive correlations are reported as shades of blue, while negative correlations are reported as shades of red.

2. financial ratios and other proxy indicators (e.g., productivity, economies of scale, spillovers) that are based on financial accounts and that we expected to correlate with the ability of exporting;
3. firms' locations, ownership status, and industry affiliations, which can help in spotting categories of firms at a competitive advantage or disadvantage.

Usefully, in Figure 3, we show a correlation matrix including all numeric predictors. Please note how many of them are indeed much cross-correlated with values well above 0.6. In a context of a pure predictive exercise, we do not know *ex-ante* which financial information can convey the highest predictive power. In principle, we can be informed from previous theory and empirical analyses that some variables are more associated than others with export status, i.e., they can be drivers of exporters, as for example in the case of productivity, firm size, financial constraints or ownership status. Yet, we prefer to keep all of them as they altogether allow us to reach high levels of prediction accuracy. See also a specific robustness check in Section 5.2. Of course, we are well aware that our list of predictors entails a great deal of endogeneity among variables that are otherwise studied in different structural relationships, e.g., financial constraints and productivity.

Yet, from a pure predictive perspective, we do not want to leave any available information unexploited, even if it contributed only marginally to increase our prediction accuracy. In

Section 5.3, we further discuss the limits and benefits of a pure predictive exercise when it comes to interpretability of predictors. Here, we just want to highlight once again that we are neither interested in studying causality. Eventually, what is relevant for the scope of our research question is just to obtain the minimum number of false negatives and false positives. In a trade-off between higher interpretability of parameters and better prediction accuracy, we decide to lean exclusively on the latter. We will devote a specific robustness check in Section 5.2 to show that if we selected only a subset of (best) predictors, we would obtain a worse predictive performance. See also Section 5.3 for an assessment of contribution of predictors to predictions.

5 Results

5.1 Models' horse race

In Table 2, we compare measures of standard prediction accuracy across the methods we test. For details on how they are constructed, please see Appendix C. In our case, Sensitivity focuses on the ability to predict exporters, i.e., the amount of *true positives*, while Specificity focuses on the ability to predict non-exporters, i.e., the amount of *true negatives*. Balanced Accuracy is just an average of Sensitivity and Specificity values. AUC (Area Under the Curve) is derived from evaluation of the performance at different classification thresholds, as reported in Figure 4, and it is our baseline measure of performance across different models. Finally, Precision-Recall is of help in assessing the trade-off between returning accurate results (high precision) vis á vis returning a majority of positive results (high recall).

Table 2: Prediction accuracies

	Specificity	Sensitivity	Balanced Accuracy	AUC	PR	N. obs.
LOGIT	0.6642	0.7776	0.7210	0.7940	0.8053	86,754
LOGIT-LASSO	0.6606	0.7722	0.7164	0.7847	0.7891	86,754
CART	0.5700	0.7896	0.6796	-	-	86,754
Random Forest	0.6078	0.8276	0.7178	0.7947	0.8010	86,754
BART	0.6272	0.8048	0.7158	0.7911	0.7998	86,754
BART-MIA	0.9064	0.6496	0.7782	0.9054	0.7375	382,606

Note: We report standard measures of prediction accuracies (by column) for different methods we train (by row). For details on how prediction accuracies are constructed, see Appendix C. Any observation is a firm-year present in the sample. All methods but BART-MIA do not train or test on observations when at least one predictor is missing. Hence, a larger number of observations in testing BART-MIA.

Immediately, we notice that BART-MIA performs better as it shows an AUC equal to 0.9054, which is considerably higher than in the case of other methods. BART-MIA is in general more able than others to predict both exporters and non-exporters. Its overall ability is confirmed by a high value of Balanced Accuracy (0.77).

Yet, when we look at Specificity *vis á vis* Sensitivity values, we realize it predicts relatively better non-exporters rather than exporters. The reason is that the boost in overall prediction accuracy by BART-MIA is largely due to an efficient use of the non-random missing information on smaller firms reporting incomplete financial accounts. Yet, as largely expected, smaller firms with partial information are also the ones that are more likely to be non-exporters. Therefore, BART-MIA is able to include them in predictions, while other methods simply drop them for lack of data. Thus, we observe an increase in Specificity that corresponds to a decrease in Sensitivity⁶.

Finally, a simple comparison between the prediction accuracies of BART and BART-MIA allows us to quantify what is the gain in considering also missing values. Overall, we observe a 14.4% increase in AUC, our baseline measure of prediction accuracy.

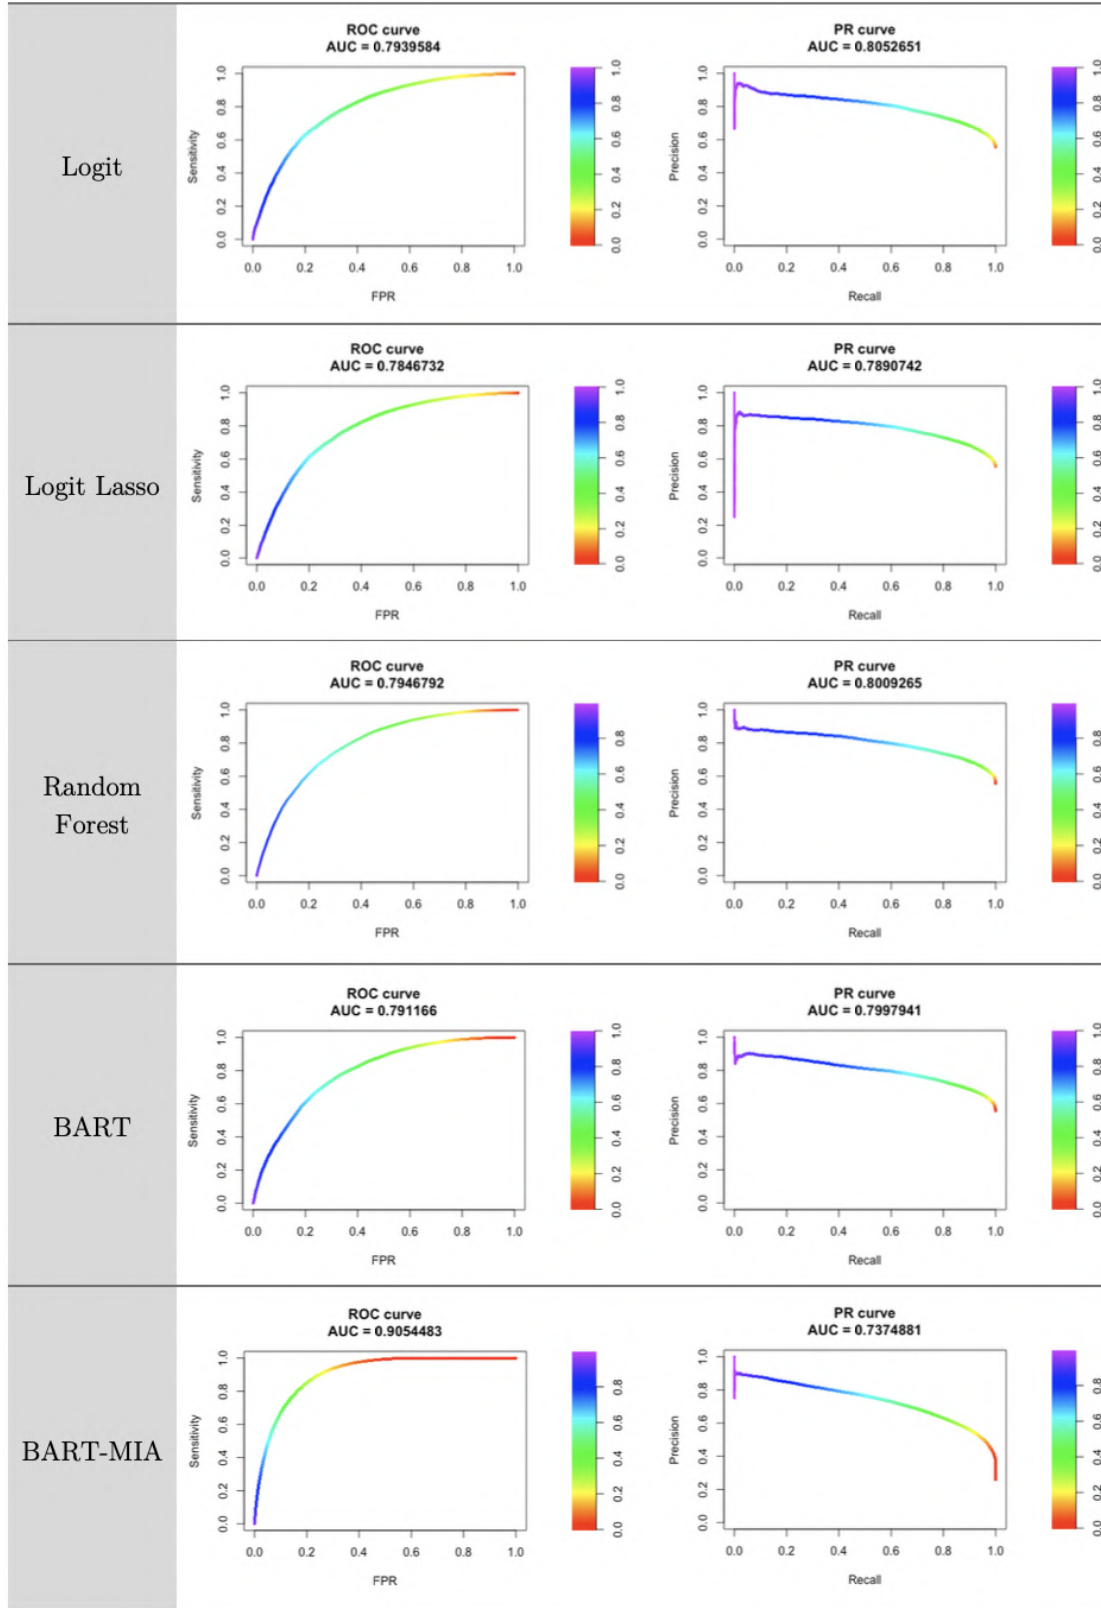
Eventually, the reason why BART-MIA performs better in Specificity and, in turn, on AUC and Balanced Accuracy is that smaller firms are more likely non-exporters, thus our relative number of *true negatives* (i.e., non-exporters) is higher than *true positives* with BART-MIA. We will further discuss the trade-off between Specificity and Sensitivity once we challenge our results in Section 5.4. Suffice it to say here that, in general, predicting true exporters is made difficult by the presence of heterogeneous exporting patterns, when firms export in some years and not in others, hence some uncertainty as summarized by Sensitivity values in Table 2.

5.2 Robustness and sensitivity

So far, we adopted a relatively standard 80 – 20 random partition of the firms in the sample at our disposal when training our model (Athey et al., 2021). Therefore, our first concern here is to cross-validate our choice by repeating the prediction exercise other four times with a similar random partition. Any time, we train on a random 80% of the dataset that we consider as in-sample information, then we test the accuracy of our predictions on the rest 20%, which we take as out-of-sample information. We show in Table B2 how we obtain similar performance scores across all exercises, and we pick BART-MIA once again as the most predictive algorithm. We conclude that previous results had not been driven by a specific selection of training *vis á vis* testing data.

⁶Indeed, we notice that the share of exporters is 56% when we exclude firms with missing predictors, while we find 26% of exporters in the entire sample tested by BART-MIA.

Figure 4: Out-of-sample Goodness-of-Fit



Our second concern is that prediction accuracies are robust to different definitions of exporters. So far, we defined an exporter as any firm with positive exporting revenues. Here, we will define an exporter as a firm whose export share over total revenues is higher than a specific minimum threshold, to make our results robust to the presence of so-called *passive exporters* (Geishecker et al., 2019), i.e., domestic firms that engage in one-off exporting events.

In the first case, models’ performance scores are similar across all periods but much worse than in our baseline, as evident in Appendix Table B4, therefore pointing to the necessity of a less volatile definition of exporters. Appendix Table B6 shows prediction accuracies after we run simulations by excluding from the category of exporters those firms that report export shares lower than the first, second, and fifth percentile, respectively. Prediction accuracies are similar in magnitude to those of our benchmark definition. Latter evidence suggests that baseline predictions are not affected by the presence of a few less proactive firms.

A third concern we have is to verify the robustness to changes in predictors. Our problem here is whether we could obtain similar prediction accuracy with a minor effort, once neglecting variables that contribute with a relatively little predictive power. For this purpose, we perform a Logit-LASSO exercise before running again the models described in 4.1. As in standard applications (Belloni et al., 2017), the Logit-LASSO selects a subset of best predictors (in our case, 23 out of 52) to contribute relatively more to predict export status. Once again, BART-MIA outperforms other statistical learning techniques. However, when we perform BART-MIA including only such a subset of predictors, we obtain lower accuracy than baseline results, as reported in Appendix Table B3. Yet, we gather there is no reason to exclude available predictors despite the high cross-correlations we observed in Figure 3.

A fourth concern we have is to check whether the time of training and testing matters for predictions. So far, we considered firms and their export status throughout the entire period at our disposal, between 2010 and 2018. In Appendix Table B4, we train and test our predictive model separating each year. It is evident how predictions do not change dramatically over the timeline.

Finally, we report Spearman’s rank correlations in Table 3 to test whether rankings in predictions are sensitive to the choice of predictive models. Please note how, by construction, the Spearman’s rank correlations can be performed only on the subset of the data where every technique obtains predictions, excluding firms with missing values tested only by BART-MIA. We get relatively high rank-correlations with a minimum of 0.87 and a maximum of 0.96. In general, models do not dramatically alter the relative positions of firms on the distribution of predictions.

However, please note that rank-correlation is about 0.92 between the simpler BART and

its variant with missingness-not-at-random, the BART-MIA. The inclusion of firms with partial information does alter the ranking in predictions even if we compare across the same observations. The latter is a significant result that allows us to further qualify the difference between the simpler BART and its variant. The bottom line is that information from firms with missing values in predictors allows BART-MIA to identify different thresholds on predictors’ distributions, which in turn change the relative positions of firms on the distribution of predictions.

Table 3: Spearman’s rank correlations of predicted probabilities from different models

	LOGIT	LOGIT-LASSO	Random Forest	BART	BART-MIA
LOGIT	1	0.9657	0.8773	0.8841	0.9012
LOGIT-LASSO		1	0.8925	0.9030	0.9118
Random Forest			1	0.9112	0.9167
BART				1	0.9179
BART-MIA					1

Note: We report a Spearman’s rank correlation among out-of-sample predictions to show how rankings in export status are sensitive to changes in predictive models. All models, including BART-MIA, are thus trained and tested on the same observations.

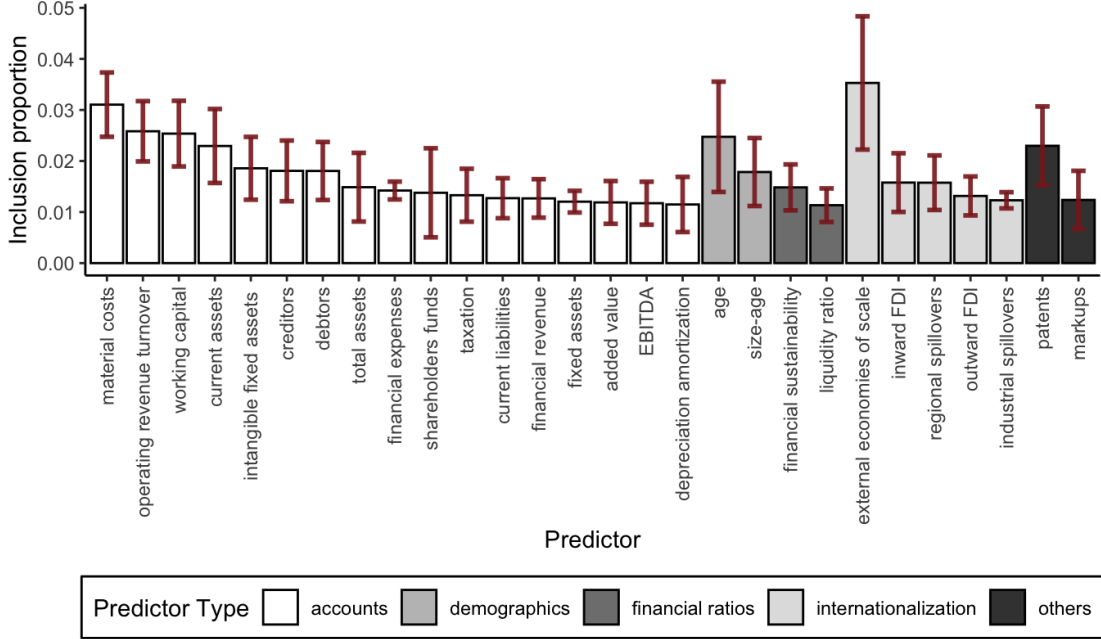
5.3 Predictors’ power

In line with our empirical strategy, we focused so far on prediction accuracy while neglecting the role of single predictors and their contributions. We discussed in Section 4 how our choice is driven by the necessity to maximize prediction accuracy; therefore we use information from an as complete as possible list of predictors. Yet, we are aware that our selection brings to a list of predictors that includes a compound of endogenous variables that are also highly cross-correlated, as shown in Figure 3.

This section wants to show how predictors do have different predictory power, which we can discuss without implicating any direction of causality. We measure so-called Variable Inclusion Proportions (VIP) after testing our baseline BART-MIA. We report visualization of their relative influence on Figure 5 while providing a standard deviation measured after running five random tests. Please note how averaging across multiple trials allows us to improve the stability of estimates, as in [Kapelner and Bleich \(2013\)](#). For a different choice of method to catch relative importance of predictors, see also [Joseph \(2020\)](#). For the sake of visualization, we report only the predictors that register an inclusion proportion that is at least 1%.

When we look at Figure 5, we find that the best predictor of exporting status is a proxy of external economies of scale based on the presence of firms in the same industry and the

Figure 5: Variable inclusion proportions after BART-MIA



Note: We report the proportion of times each predictor is chosen for a splitting rule in BART-MIA, collecting by main type of predictor. Of all the predictors in baseline, we visualize those with an inclusion proportion higher than 1%. Red bars represent standard deviations of inclusion proportions obtained by replicating BART-MIA on the same random training set but five different times.

same region, following suggestions by [Bernard et al. \(1995\)](#). As we are in a pure prediction framework, we cannot say whether external economies of scale, measured in this way, are indeed a determinant of export status. We cannot exclude reversal causality. On the one hand, it is indeed possible that local spillovers help neighbouring firms to start exporting after, for example, sharing infrastructures or intangible knowledge about foreign markets. On the other hand, we cannot exclude that firms in industries at a comparative advantage located in proximity before becoming exporters. It is beyond the scope of our analysis to unravel the endogeneity of this specific relationship or any other we find among predictors and the outcome. Suffice it to say that an industrial concentration of exporting firms in a geographical area is a good albeit not unique predictor of export status for the representative firm located in that area.

In general, we observe in Figure 5 how original accounts contribute best to identify export status. Yet, no predictor contributes more than 5% in any of the tests we performed. To name just a few of the first predictors, we have material costs, turnover, working capital, and current assets. Yet, considering the overall distribution of predictive power and their standard deviations, we conclude that there is no unique indicator that alone can predict the firm’s status. All convey non-trivial information on the ability to export. Besides

financial accounts, business demography is also essential: firm age and size have an inclusion proportion higher than 2%. It also makes perfect sense that the activities of multinational enterprises play a role in export status. Being either a foreign subsidiary (inward FDI) or owning a subsidiary abroad (outward FDI) is associated with a higher probability of exporting. As expected, the ability to innovate and register patents is also related to the likelihood of becoming an exporter.

Please note, however, that a much-studied determinant of export status, Total Factor Productivity (TFP), is completely missing from Figure 5. Our educated guess is that its role is captured by the sample variation in raw financial accounts, including turnover, costs of materials, and other variables needed to estimate the production function from which one would extract Total Factor Productivity.

5.4 Heterogeneous exporting patterns

The biggest challenge in predicting exporters is that exporting is an event that can be repeated with some heterogeneity over the timeline. Firms can export for some time and then lay idle for a while before re-proposing on foreign markets. Yet, the statistical learning techniques we have been using in previous analyses rely on classifications of an outcome that is simple and binary: based on their observed characteristics at time t , firms are either exporters or not.

The training strategy so far relies on the assumption that a successful exporter is one that at some point could afford the searching costs entailed to access a foreign market. Hence, based on the general intuition that exporters are statistically different from non-exporters in financial accounts, we trained on in-sample information to test on out-of-sample firms and predict whether they were exporters or not. The standard methodology that we adopted did not allow for intermediate alternatives.

Here we want to test the sensitivity of predictions to heterogeneous exporting patterns, including the case of discontinuous exporters. For our purpose, we perform a separate sensitivity check by classifying firms into five categories:

1. firms that always export, which we call *constant exporters*;
2. firms that never export, which we call *non-exporters*;
3. firms that start exporting at some period t and always export afterwards, which we call *switching exporters*;

4. firms that export all periods until t and never export afterwards, which we call *switching non-exporters*⁷;
5. *discontinuous exporters*, which export with an irregular pattern over our timeline, with more than one gap along the timeline.

In Table 4, we report separate prediction accuracies for the previous categories. On the one hand, we observe that our predictive model performs quite well in separating constant exporters and non-exporters, where Sensitivity and Specificity are about 0.86 and 0.95, respectively.⁸ On the other hand, our predictions are less reliable when we start looking at out-of-sample information on firms that show gaps along the timeline. In general, we have ROCs of about 0.86 and 0.81, respectively, in the case of *switching exporters* and *switching non-exporters*. Interestingly enough, the quality of predictions is proportional to the number of years that the firms actually exported. We are more able to predict the export status of firms that started (stopped) exporting sooner (later) in our data.

With a similar approach, we focus on discontinuous exporters at the bottom of Table 4. Here, we find a relatively lower prediction accuracy (ROC: 0.80) if compared with constant exporters and non-exporters. Evidently, in this case, we are less and less able to predict the export status of firms observed exporting fewer years over the timeline.

Eventually, we compare previous exercises with the more liberal definitions proposed by Békés and Muraközy (2012), according to whom firms with at least four years of consecutive exporting can be considered as *permanent exporters* vis á vis other *temporary exporters*. As largely expected, we find in Table B5 that prediction accuracies for *permanent exporters* are relatively higher (AUC: 0.849; PR: 0.934) than in the case of temporary exporters. In particular, the model fails at predicting the export status of temporary exporters, i.e., it reports a relatively lower true positives' rate, as shown by the low scores on sensitivity, PR and AUC. From our viewpoint, it makes sense that exporters with irregular exporting patterns represent intermediate cases somewhere between firms that always export and firms that never export. Therefore, classification algorithms struggle to separate intermediate cases on a binary outcome. Based on financial accounts, such firms can be seen neither as fit for exporting as constant exporters nor as unfit as non-exporters. Yet, it is more likely that such intermediate cases are of less interest in policy applications because trade promoters or

⁷Please note how we may have had more switching non-exporters if we were able to zoom out on a longer timeline. We cannot exclude that firms that do not export in our sample did in previous unobserved periods. The latter is an element of imperfection that we cannot expunge from our prediction accuracy.

⁸Please note that we cannot estimate other measures of prediction accuracy when we focus exclusively on either positive or negative outcomes. See Appendix C for a definition of different measures of prediction accuracies.

financial institutions need instead to understand whether a firm that never exported needs some support or not.

Table 4: Prediction accuracies and exporting patterns

Firm category	Sensitivity	Specificity	Balanced Accuracy	ROC	PR	Num. Obs.
Constant Exporters	0.856	-	-	-	-	21,834
Non-exporters	-	0.951	-	-	-	158,625
Switching Exporters	0.629	0.849	0.739	0.864	0.764	15,084
<i>Start in 2011</i>	<i>0.749</i>	<i>0.682</i>	<i>0.716</i>	<i>0.794</i>	<i>0.954</i>	<i>1,980</i>
<i>Start in 2012</i>	<i>0.729</i>	<i>0.694</i>	<i>0.712</i>	<i>0.808</i>	<i>0.914</i>	<i>1,296</i>
<i>Start in 2013</i>	<i>0.711</i>	<i>0.751</i>	<i>0.731</i>	<i>0.838</i>	<i>0.888</i>	<i>1,179</i>
<i>Start in 2014</i>	<i>0.618</i>	<i>0.806</i>	<i>0.712</i>	<i>0.832</i>	<i>0.821</i>	<i>1,215</i>
<i>Start in 2015</i>	<i>0.582</i>	<i>0.796</i>	<i>0.689</i>	<i>0.812</i>	<i>0.73</i>	<i>1,323</i>
<i>Start in 2016</i>	<i>0.585</i>	<i>0.819</i>	<i>0.702</i>	<i>0.823</i>	<i>0.638</i>	<i>1,683</i>
<i>Start in 2017</i>	<i>0.463</i>	<i>0.835</i>	<i>0.649</i>	<i>0.804</i>	<i>0.45</i>	<i>2,187</i>
<i>Start in 2018</i>	<i>0.262</i>	<i>0.903</i>	<i>0.583</i>	<i>0.792</i>	<i>0.251</i>	<i>4,221</i>
Switching non-exporters	0.599	0.802	0.7	0.819	0.786	27,891
<i>Stop in 2011</i>	<i>0.269</i>	<i>0.81</i>	<i>0.539</i>	<i>0.643</i>	<i>0.152</i>	<i>3,915</i>
<i>Stop in 2012</i>	<i>0.376</i>	<i>0.745</i>	<i>0.561</i>	<i>0.65</i>	<i>0.291</i>	<i>2,511</i>
<i>Stop in 2013</i>	<i>0.419</i>	<i>0.725</i>	<i>0.572</i>	<i>0.689</i>	<i>0.443</i>	<i>2,124</i>
<i>Stop in 2014</i>	<i>0.479</i>	<i>0.737</i>	<i>0.608</i>	<i>0.733</i>	<i>0.599</i>	<i>2,412</i>
<i>Stop in 2015</i>	<i>0.508</i>	<i>0.815</i>	<i>0.662</i>	<i>0.816</i>	<i>0.757</i>	<i>2,844</i>
<i>Stop in 2016</i>	<i>0.563</i>	<i>0.925</i>	<i>0.744</i>	<i>0.929</i>	<i>0.924</i>	<i>5,409</i>
<i>Stop in 2017</i>	<i>0.664</i>	<i>0.843</i>	<i>0.754</i>	<i>0.877</i>	<i>0.931</i>	<i>3,996</i>
<i>Stop in 2018</i>	<i>0.742</i>	<i>0.813</i>	<i>0.778</i>	<i>0.874</i>	<i>0.97</i>	<i>4,680</i>
Discontinuous	0.547	0.807	0.677	0.796	0.686	85,023
<i>exporting years: 1</i>	<i>0.216</i>	<i>0.873</i>	<i>0.544</i>	<i>0.686</i>	<i>0.171</i>	<i>19,152</i>
<i>exporting years: 2</i>	<i>0.313</i>	<i>0.823</i>	<i>0.568</i>	<i>0.702</i>	<i>0.334</i>	<i>12,816</i>
<i>exporting years: 3</i>	<i>0.387</i>	<i>0.796</i>	<i>0.592</i>	<i>0.718</i>	<i>0.483</i>	<i>10,962</i>
<i>exporting years: 4</i>	<i>0.478</i>	<i>0.736</i>	<i>0.607</i>	<i>0.719</i>	<i>0.595</i>	<i>8,910</i>
<i>exporting years: 5</i>	<i>0.519</i>	<i>0.74</i>	<i>0.63</i>	<i>0.753</i>	<i>0.72</i>	<i>9,297</i>
<i>exporting years: 6</i>	<i>0.593</i>	<i>0.721</i>	<i>0.657</i>	<i>0.755</i>	<i>0.808</i>	<i>8,460</i>
<i>exporting years: 7</i>	<i>0.662</i>	<i>0.7</i>	<i>0.681</i>	<i>0.774</i>	<i>0.886</i>	<i>7,758</i>
<i>exporting years: 8</i>	<i>0.757</i>	<i>0.658</i>	<i>0.708</i>	<i>0.781</i>	<i>0.951</i>	<i>7,668</i>
Total	0.6491	0.9080	0.7785	0.9048	0.7383	308,457

Note: We report prediction accuracies after BART-MIA for firms with different exporting patterns. For switching-exporters and switching-non-exporters we identify the year when they are observed changing status, i.e., the year when the firm passes from never exporting to always exporting, and vice versa. For discontinuous exporters we distinguish by number of exporting years over the sample timeline.

6 From predictions to firms' scoring

From our perspective, a pure prediction exercise for firms' exporting ability is helpful to assess the distance of non-exporters to export status. Based on the prior knowledge that exporters and non-exporters are statistically different across many attributes, we can use baseline predictions and build a continuous indicator that gives a score to indicate the potential to successfully propose on foreign markets.

Briefly, we can get a basic and simple export (probabilistic) score for any non-exporting i th firm that we can indicate as a distance from the export status as:

$$distance_i = 1 - Pr(Y_i = 1 \mid \mathbf{X}_i = x) \quad (5)$$

which is by definition bounded in a range $(0, 1)$, and made conditional on the entire set of predictors \mathbf{X}_i . In a nutshell, after we successfully train on previous in-sample information, we can just plug new out-of-sample information in and get a continuous (probabilistic) score as a distance from one, which is the value at which we can find exporters on the prediction distribution.

We believe that such a score can be a valuable tool to design target-specific policies. For example, one can design better programs to promote firms' access to foreign markets. One can assess credit worthiness of potential exporters when they ask for financial resources to outreach foreign consumers. From a broader perspective, one can adopt exporting scores as indicators of competitiveness, aggregating them on a subset of firms, let's say an industry or a region, to monitor which segments of an economy have the potential to export successfully and which segments have not. In the following paragraphs, we discuss the benefits and limits of possible applications with the help of some descriptive statistics and back-of-the envelope calculations.

6.1 Financial constraints and trade promotion

Exporting requires routine access to financial resources; thus, well-functioning financial markets are crucial to support exporters (Manova, 2012). Since they incur high fixed costs to access distant foreign markets, exporters depend relatively more on external resources than domestic producers. Therefore, the presence of financial market imperfections constrains opportunities for trade, all the more when firms are heterogeneous in the ability to provide collateral (Chor and Manova, 2012).

Against this background, national and international agencies establish trade promotion programs to fill the gap in financial markets' imperfections and develop skills that help catch

business opportunities on global markets ⁹. Export promotion programs are effective tools in helping firms reach new destination countries and introduce new differentiated products (Volpe Martincus and Carballo, 2010a). They facilitated the recovery after the global recession of 2009 (Van Biesebroeck et al., 2016).

If we focus on firms’ financial constraints, financial institutions and trade promotion agencies all face a common credit scoring problem when firms ask for their support. Both scholars and practitioners have developed several tools to reduce the informative gap between borrowers and lenders from disclosed financial accounts. Usually, the main idea is to check how far a company is from a situation of financial distress using some combination of financial ratios ¹⁰.

As far as we know, there has been no previous attempt to score the exporting ability of a firm starting from financial accounts. In this context, we believe that our prediction exercise could be useful to catch the sustainability of firms’ internationalization strategies. For example, after looking at the entire distribution we obtain from French non-exporters in Figure 6, one could design an intervention based on how distant a company is from an ideal benchmark of exporters that we could easily locate on the right tail.

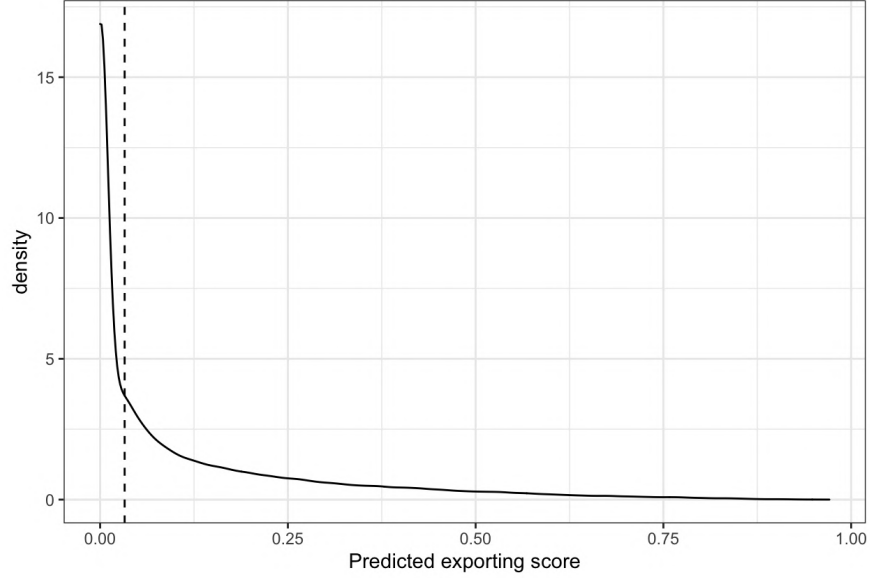
Interestingly, Figure 6 shows that a majority of French non-exporters is located on a heavily thick left tail, thus showing to be much different from what an exporter would look like. In general, some non-exporters more than others may be proximate to reaching the right tail’s goal. Thus, one could calibrate the financial support to focus on the aspects that need it most.

To illustrate our idea, we perform back-of-the-envelope estimates of how many capital expenses and cash resources a representative firm needs to climb risk categories. We can classify firms in different risk categories based on a simple partition of exporting scores as if we were a financial institution. By construction, probabilistic exporting scores obtained from baseline BART-MIA are in a range $(0, 1)$. Let us consider all firms included in a segment of predictions as belonging to the same risk category. Obviously, the higher the distance from

⁹A variety of services are provided to firms that apply for trade support programs, ranging from training to financial resources. International organizations specifically support firms in less advanced countries to fill the gap in global markets. See, for example, the experience of the Inter-American Development Bank and the International Trade Center.

¹⁰For example, Z-scores (Altman, 1968; Altman et al., 2000) and Distance-to-Default (Merton, 1974) have been first tools used to assess the viability of a firm based on a combination of financial accounts, which could indicate financial distress. Recent advances in predictive models for bankruptcies also include machine learning methods. See, for example, Bargagli-Stoffi et al. (2020).

Figure 6: Distributions of exporting scores of non-exporters after BART-MIA



Note: We report the distribution of the score after implementing BART-MIA on the entire sample and selecting all non-exporting firms. The vertical line identifies the median non-exporting firm.

export status, $1 - Pr(Y_i)$, the higher the risk for trade credit. For simplicity, let us assume that we can identify up to ten main categories of firms. The analyst could find a rationale for a different partition of risk classes. For the moment, let us just rely on symmetric segments of length equal to 0.1, i.e., about ten percentage points of lower risk in each following category when approaching export status. Therefore, we can run the following simple specification:

$$\log Y_{it} = \beta_0 + \sum_{risk=1}^{10} \theta_{risk} + \beta_1 x_{it} + \phi_t + \delta_s + \eta_r + \epsilon \quad (6)$$

where Y_{it} is either cash resources or fixed assets for firm i at time t , and x_{it} is its firm-level size. We will always control for time (ϕ_t), four-digit NACE sector (δ_t), and two-digit NUTS region (η_r) fixed effects. We cluster standard errors at the firm level.

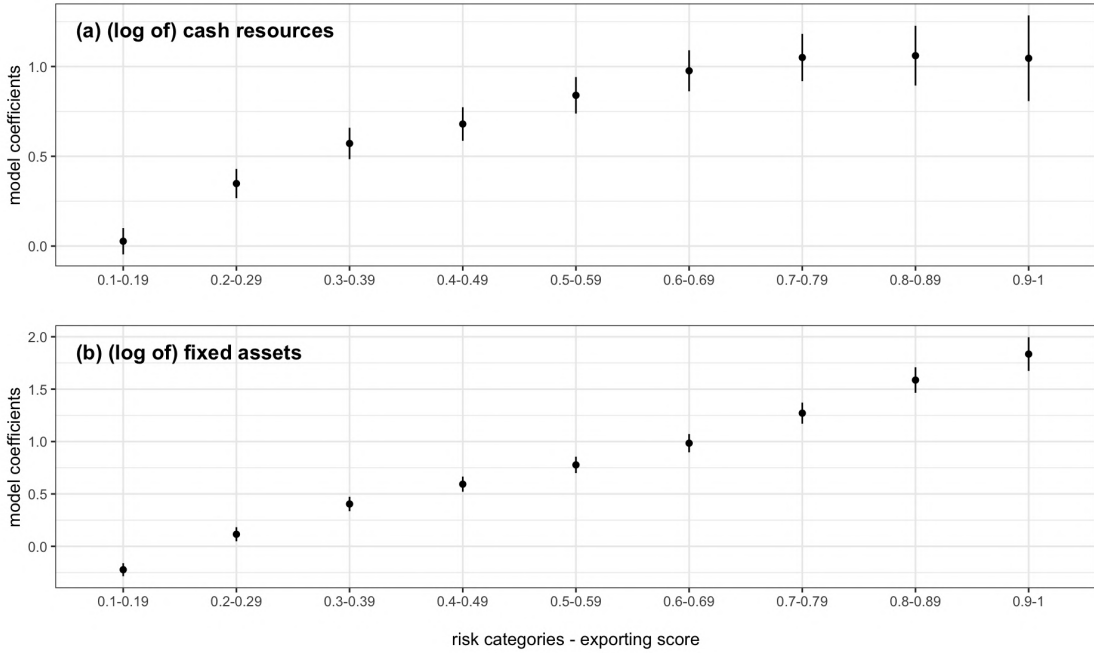
Crucially, our coefficients of interest are the ones on θ_{risk} , as these are risk classes built on exporting scores. We report them in decreasing order of risk in Figure 7 together with 99% confidence intervals. Once we omit the first segment $[0, 0.09]$, the estimated intercepts of eq. 6 will indicate (logs of) cash resources and fixed assets needed by a representative firm that is more distant from export status. Therefore, to obtain what is needed by following categories, we can just consider (log) premia with respect to the first segment.

For example, the representative firm with exporting scores lower than 0.1 operates with $\exp(\hat{\beta}_0) = \exp(11.6338) \approx 112,850$ euro of cash resources and $\exp(\hat{\beta}_0) = \exp(13.4027) \approx$

661,790 euro of fixed assets. Firms in the fifth category, when exporting scores are in a range $[0.4, 0.5)$, will need $\exp(\hat{\beta}_0 + \hat{\theta}_5) = (11.6338 + 0.6797) \approx 222,690$ euro of cash resources and $\exp(\hat{\beta}_0 + \hat{\theta}_5) = \exp(13.4027 + 0.5933) \approx 1,197,800$ euro of fixed assets. To put it differently, we can say that a firm that is in a medium-risk category needs about 97% more cash resources and about 81% more fixed assets if compared with a firm with the lowest exporting scores.

On the other hand, if we look at firms in a comfort zone with exporting scores in a range $[0.9, 1]$, we see that they operate with $\exp(\hat{\beta}_0 + \hat{\theta}_{10}) = \exp(11.6338 + 1.0459) \approx 321,160$ euro of cash and $\exp(\hat{\beta}_0 + \hat{\theta}_{10}) = \exp(13.4027 + 1.8348) \approx 4,145,360$ euro of fixed assets. Please note that the higher the probability that a firm starts exporting, the higher the cash resources and the capital expenses it needs. In the latter case, if we compare with average exporting scores in the fifth risk class, we find that medium-risk firms need 44% more cash resources and up to 246% more capital expenses to look like firms that have been classified under the lowest risk category.

Figure 7: Premia on relevant firm dimensions across exporting scores



Note: Fixed effects on segments of exporting scores after linear regressions where the outcomes are (log of) cash resources and (log of) fixed assets, respectively. We always control for firm size, NUTS 2-digit regions, NACE 2-digit industries, and time fixed effects. Errors are clustered at the firm level.

In terms of trade credit, we observe that there is an increasing need for financial resources to climb risk categories and reduce the distance from export status. Based on predictions made on the experience of both exporters and non-exporters, a financial institution could

evaluate whether it's worth the effort of investing in internationalization and, in case, how much resources a firm needs to reach its target.

6.2 Export competitiveness

Openness to international trade is a determinant of economic growth. Thanks to differential comparative advantages and economies of scale, consumers can gain from trade. Both developed and developing economies have benefited from integration into the global economy through export growth and diversification. Thus, export performance has been long used as yet another proxy for measuring countries' competitiveness by a consolidated tradition in economic literature and by international organizations ([Leamer and Stern, 1970](#); [Richardson, 1971a,b](#); [Gaulier et al., 2013](#)).

In this context, we believe that predictive models like ours could help further understanding the export competitiveness of a country, a region or an industry, specifically focusing on the potential for extensive margins, i.e., by looking at the number of firms that could become exporters given the right conditions. Take the case of Figure 8. Once we focus on French NUTS 2-digit regions, we spot where are the non-exporting firms with exporting scores above the median of the overall national distribution we observe in Figure 6. This is the segment of firms where we can assume that there is a high potential for exporting. Interestingly, we find that a relative majority share of 15.41% is in Île-de-France followed by a 15.22% share that operates in Rhône-Alpes. The third most trade competitive region is Provence-Alpes-Côte d'Azur with however just 8.52% of firms with a score above the national median. Comprehensibly, we mainly find French overseas territories at the bottom of the ranking.

Clearly, absolute numbers in Figure 8 are also higher in some regions, like Rhône-Alpes and Île-de-France, because this is where we find a higher density of manufacturing activities. To control for concentrations of business activity, we follow a dartboard approach as in [Ellison and Glaeser \(1997\)](#) and propose location quotients in Figure 9. See Appendix D for further details on computations. Regions with location quotients greater than one are the ones where potential exporters are more concentrated than what one would expect given the underlying distribution of manufacturing activities. Eventually, we do find a geographic pattern in Figure 9, since non-exporters with the highest potential are mainly present in North-Eastern regions, while Southern regions and overseas territories lag behind in trade potential.

In Figure 10, we observe that there is a high variation of exporting scores for non-exporters at the NACE 2-digit industry-level, which can also be much informative for the

Legend

FRY5
0,03

FRY1
0,41
FRY2
0,20

FRY3
0,17

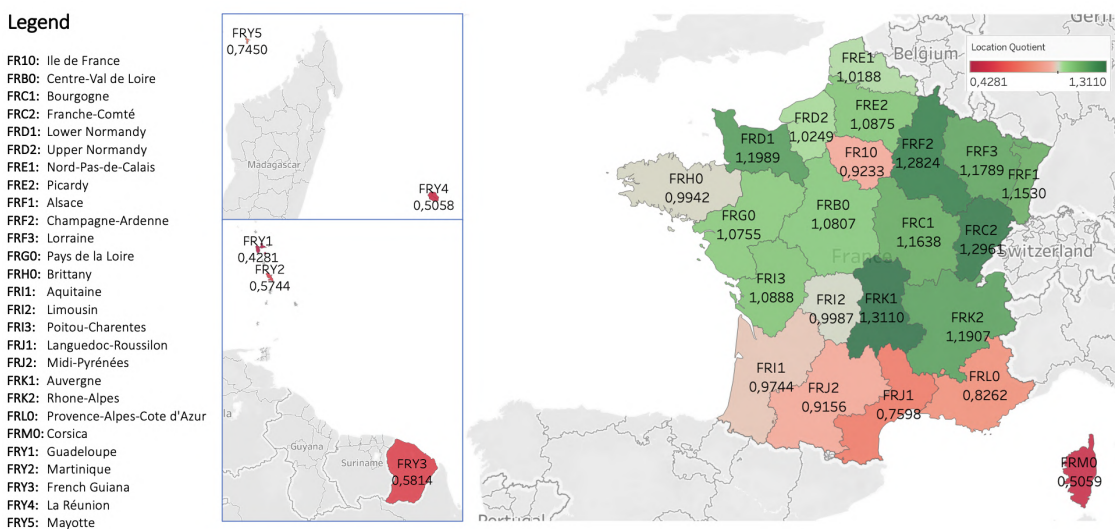
FRY4
0,58

Share (in % pts)

0,03 15,41

Region	Share (in % pts)
FRY1	15,41
FRY0	8,52
FRY2	5,70
FRY3	0,17
FRY4	0,58
FRY5	0,03

Figure 9: Location quotient of non-exporters with exporting scores above the national median

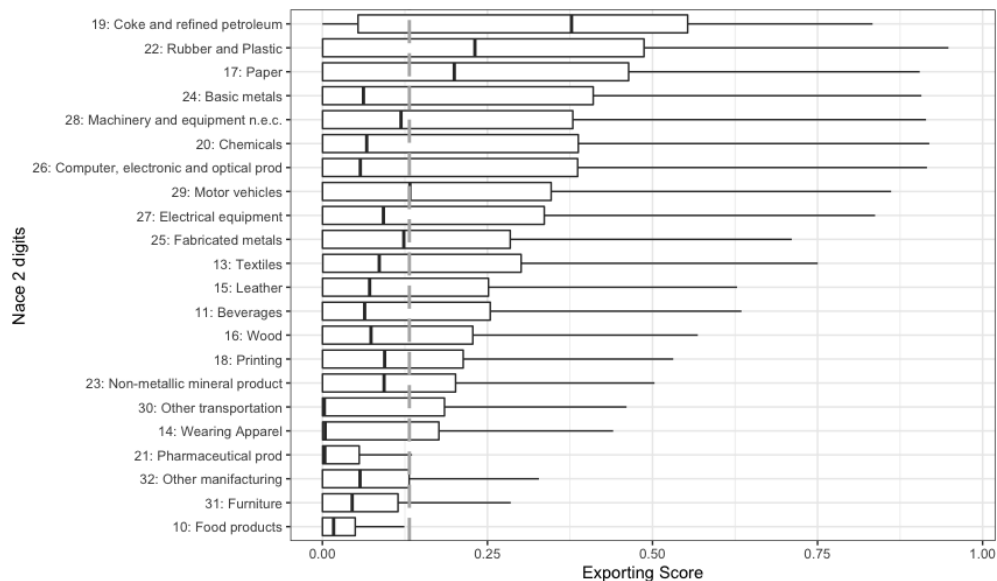


26

policymaker. Industries do report different dispersion values across the industry medians. Thick bars on boxplots indicate industry-level medians. Interestingly, the Coke and Refined Petroleum (NACE 19) report the highest industry median, followed by Rubber and Plastic (NACE 22), Paper (NACE 17), and Basic Metals (NACE 24) industries. Notably, Food Products (NACE 10) is the industry with the minimum dispersion and median.

Eventually, more sophisticated analyses on the distribution of exporting scores in industries and regions can be performed to evaluate trade potential. For example, one could exploit the variation in time to understand how much competitive in trade a region or an industry is evolving. One could compare across countries to check whether there is potential for trade beyond actual export performance. We believe any of them could be a useful tool in the kit of the analyst that aims at assessing the trade competitiveness of an economy.

Figure 10: Exporting scores by industry



Note: On boxplots, we report the distributions of exporting scores for non-exporters after BART-MIA by NACE 2-digit industries. The grey line vertically crossing industry bars corresponds to the median of the overall distribution of non-exporters. Thicker black vertical bars represent industry medians.

7 Conclusions

This paper exploits statistical learning techniques to predict the ability of firms to export. After showing how financial accounts convey non-trivial information to separate exporters from non-exporters, we propose predictions as a tool that can be useful for targeting trade promotion programs, trade credit, and assessing firms' competitiveness.

The central intuition is that exporters and non-exporters are statistically different in their

financial structures since they have to sustain the sunk costs of gaining access to foreign markets, where regulations and consumer tastes differ. On this, we rely on the long-established literature that connects firm heterogeneity with self-selection into exporting. Thus, we train and test various algorithms on a dataset of French firm-level data from 2010-2018. Eventually, we find that the Bayesian Additive Regression Tree with Missingness In Attributes (BART-MIA) outperforms other models due to an efficient use of the non-random missing information on smaller firms reporting incomplete financial accounts. Moreover, prediction accuracy is rather high, up to 90%, and robust to changes in the definition of exporters and different training strategies. Interestingly enough, our framework allows handling cases of discontinuous exporters, as they show up as intermediate cases between permanent exporters and non-exporters. Eventually, the more firms export over the timeline, the more likely we correctly classify them as actual exporters.

In the second part of our contribution, we discuss how export predictions can be used as scores to catch the sustainability of firms' internationalization strategies and their creditability. For example, imitating what a financial institution would professionally do, we order firms along exporting scores in different risk classes. Thus, we show back-of-the-envelope estimates of how much cash resources and capital a firm would need to climb those risk classes. In our case study, we show that a French non-exporter that has just half the exporting score needs up to 44% more cash and 246% more capital assets to reach full export status.

To conclude, we argue that exporting scores obtained as predictions from firm-level financial accounts can be yet another useful tool in the analyst kit to evaluate trade potential at different levels of aggregations. As we show in the case of France, for which we provide summary statistics where a high heterogeneity of trade potential is detected across regions and industries.

References

- Akerberg, D.A., Caves, K., Frazer, G., 2015. Identification properties of recent production function estimators. *Econometrica* 83, 2411–2451.
- Ahrens, A., Hansen, C.B., Schaffer, M.E., 2020. lassopack: Model selection and prediction with regularized regression in stata. *arXiv preprint arXiv:1901.05397* .
- Almeida, H., Campello, M., Weisbach, M.S., 2004. The cash flow sensitivity of cash. *The Journal of Finance* 59, 1777–1804.
- Altman, E.I., 1968. Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *The Journal of Finance* 23, 589–609.

- Altman, E.I., et al., 2000. Predicting financial distress of companies: revisiting the z-score and zeta models. Stern School of Business, New York University , 9–12.
- Athey, S., 2018. The impact of machine learning on economics, in: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Athey, S., Imbens, G., Metzger, J., Munro, E., 2021. Using wasserstein generative adversarial networks for the design of monte carlo simulations. *Journal of Econometrics* .
- Bargagli-Stoffi, F., Riccaboni, M., Rungi, A., 2020. Machine Learning for Zombie Hunting. Firms’ Failures and Financial Constraints. EIC Working Paper Series 1/2020. IMT School for Advanced Studies.
- Barro, R.J., Sala-i Martin, X., 1997. Technological diffusion, convergence, and growth. *Journal of Economic Growth* 2, 1–26.
- Békés, G., Muraközy, B., 2012. Temporary trade and heterogeneous firms. *Journal of International Economics* 87, 232–246.
- Belloni, A., Chernozhukov, V., Fernández-Val, I., Hansen, C., 2017. Program evaluation and causal inference with high-dimensional data. *Econometrica* 85, 233–298.
- Belloni, A., Chernozhukov, V., Hansen, C., 2014. High-dimensional methods and inference on structural and treatment effects. *Journal of Economic Perspectives* 28, 29–50.
- Belloni, A., Chernozhukov, V., Hansen, C., Kozbur, D., 2016. Inference in high-dimensional panel models with an application to gun control. *Journal of Business & Economic Statistics* 34, 590–605.
- Belloni, A., Chernozhukov, V., et al., 2013. Least squares after model selection in high-dimensional sparse models. *Bernoulli* 19, 521–547.
- Bernard, A.B., Jensen, J.B., 1999. Exceptional exporter performance: cause, effect, or both? *Journal of international economics* 47, 1–25.
- Bernard, A.B., Jensen, J.B., 2004. Why some firms export. *Review of economics and Statistics* 86, 561–569.
- Bernard, A.B., Jensen, J.B., Lawrence, R.Z., 1995. Exporters, jobs, and wages in us manufacturing: 1976-1987. *Brookings papers on economic activity. Microeconomics* 1995, 67–119.
- Bernard, A.B., Jensen, J.B., Redding, S.J., Schott, P.K., 2012. The empirics of firm heterogeneity and international trade. *Annual Review of Economics* 4, 283–313.
- Breiman, L., 2001. Random forests. *Machine learning* 45, 5–32.
- Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J., 1984. *Classification and regression trees*. Belmont, CA: Wadsworth & Brooks .
- Breinlich, H., Corradi, V., Rocha, N., Ruta, M., Santos Silva, J., Zylkin, T., 2021. Machine Learning in Gravity Models: An Application to Agricultural Trade. *Research Working*

- Paper 9629. World Bank.
- Caballero, R.J., Hoshi, T., Kashyap, A.K., 2008. Zombie lending and depressed restructuring in japan. *American Economic Review* 98, 1943–77.
- Chen, H.J., Chen, S.J., 2012. Investment-cash flow sensitivity cannot be a good measure of financial constraints: Evidence from the time series. *Journal of Financial Economics* 103, 393–410.
- Chen, J., Chen, Z., 2008. Extended bayesian information criteria for model selection with large model spaces. *Biometrika* 95, 759–771.
- Chipman, H.A., George, E.I., McCulloch, R.E., et al., 2010. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics* 4, 266–298.
- Chor, D., Manova, K., 2012. Off the cliff and back? credit conditions and international trade during the global financial crisis. *Journal of International Economics* 87, 117–133. Symposium on the Global Dimensions of the Financial Crisis.
- Cravino, J., Levchenko, A.A., 2016. Multinational Firms and International Business Cycle Transmission*. *The Quarterly Journal of Economics* 132, 921–962.
- Crozet, M., Head, K., Mayer, T., 2011. Quality Sorting and Trade: Firm-level Evidence for French Wine. *The Review of Economic Studies* 79, 609–644.
- De Loecker, J., Warzynski, F., 2012. Markups and firm-level export status. *American Economic Review* 102, 2437–71.
- Del Prete, D., Rungi, A., 2017. Organizing the global value chain: A firm-level test. *Journal of International Economics* 109, 16–30.
- Ellison, G., Glaeser, E., 1997. Geographic concentration in u.s. manufacturing industries: A dartboard approach. *Journal of Political Economy* 105, 889–927.
- Fazzari, S.M., Hubbard, R.G., Petersen, B.C., 1988. Financing Constraints and Corporate Investment. *Brookings Papers on Economic Activity* 19, 141–206.
- Fontagné, L., Secchi, A., Tomasi, C., 2018. Exporters’ product vectors across markets. *European Economic Review* 110, 150–180.
- Gaulier, G., Santoni, G., Taglioni, D., Zignago, S., 2013. In the Wake of the Global Crisis. Evidence from a New Quarterly Database of Export Competitiveness. Policy Research Working Paper 9629. World Bank.
- Geishecker, I., Schröder, P.J., Sörensen, A., 2019. One-off export events. *Canadian Journal of Economics/Revue canadienne d’économie* 52, 93–131.
- Gopinath, G., Kalemli-Özcan, Ş., Karabarbounis, L., Villegas-Sanchez, C., 2017. Capital allocation and productivity in south europe. *The Quarterly Journal of Economics* 132, 1915–1967.
- Gopinath, M., Batarseh, F.A., Beckman, J., 2020. Machine Learning in Gravity Models: An

- Application to Agricultural Trade. Working Paper 27151. National Bureau of Economic Research.
- Grossman, G.M., Helpman, E., 1990. Comparative advantage and long-run growth. *The American Economic Review* 80, 796–815.
- Hastie, T., Tibshirani, R., Friedman, J., 2017. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- Hill, J., Linero, A., Murray, J., 2020. Bayesian additive regression trees: A review and look forward. *Annual Review of Statistics and Its Application* 7, 251–278.
- Hottman, C.J., Redding, S.J., Weinstein, D.E., 2016. Quantifying the Sources of Firm Heterogeneity *. *The Quarterly Journal of Economics* 131, 1291–1364.
- Joseph, A., 2020. Parametric inference with universal function approximators. Bank of England Staff Working Paper no. 784.
- Kapelner, A., Bleich, J., 2013. bartmachine: Machine learning with bayesian additive regression trees. arXiv preprint arXiv:1312.2171 .
- Kapelner, A., Bleich, J., 2015. Prediction with missing data via bayesian additive regression trees. *Canadian Journal of Statistics* 43, 224–239.
- Leamer, E., Stern, R., 1970. *Quantitative International Economics*. Allin and Bacon, Boston. chapter 7.
- Manova, K., 2012. Credit Constraints, Heterogeneous Firms, and International Trade. *The Review of Economic Studies* 80, 711–744.
- Melitz, M.J., 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *econometrica* 71, 1695–1725.
- Melitz, M.J., Ottaviano, G.I.P., 2008. Market Size, Trade, and Productivity. *The Review of Economic Studies* 75, 295–316.
- Melitz, M.J., Redding, S.J., 2014. Chapter 1 - heterogeneous firms and trade, in: Gopinath, G., Helpman, E., Rogoff, K. (Eds.), *Handbook of International Economics*. Elsevier. volume 4 of *Handbook of International Economics*, pp. 1–54.
- Merton, R.C., 1974. On the pricing of corporate debt: The risk structure of interest rates. *The Journal of finance* 29, 449–470.
- Mullainathan, S., Spiess, J., 2017. Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31, 87–106.
- Nickell, S., Nicolitsas, D., 1999. How does financial pressure affect firms? *European Economic Review* 43, 1435–1456.
- Reis, J.G., Wagle, S., Farole, T., 2010. *Analyzing Trade Competitiveness : A Diagnostics Approach*. The World Bank.
- Richardson, J., 1971a. Constant-market-shares analysis of export growth. *Journal of Inter-*

- national Economics 1, 227–239.
- Richardson, J., 1971b. Some sensitivity tests for a "constant-market-shares" analysis of export growth. *The Review of Economics and Statistics* 53, 300–304.
- Rivera-Batiz, L.A., Romer, P.M., 1991. International trade with endogenous technological change. *European Economic Review* 35, 971–1001.
- Romer, P., 1994. New goods, old theory, and the welfare costs of trade restrictions. *Journal of Development Economics* 43, 5–38.
- Rungi, A., Del Prete, D., 2018. The smile curve at the firm level: Where value is added along supply chains. *Economics Letters* 164, 38–42.
- Uddin, M.S., 2021. Machine learning in credit risk modeling: Empirical application of neural network approaches. *The Fourth Industrial Revolution: Implementation of Artificial Intelligence for Growing Business Success* , 417–435.
- Van Biesebroeck, J., Konings, J., Volpe Martincus, C., 2016. Did export promotion help firms weather the crisis? *Economic Policy* 31, 653–702.
- Volpe Martincus, C., Carballo, J., 2008. Is export promotion effective in developing countries? firm-level evidence on the intensive and the extensive margins of exports. *Journal of International Economics* 76, 89–106.
- Volpe Martincus, C., Carballo, J., 2010a. Entering new country and product markets: does export promotion help? *Review of World Economics* 146, 437–467.
- Volpe Martincus, C., Carballo, J., 2010b. Export promotion: Bundled services work better. *The World Economy* 33, 1718–1756.
- Volpe Martincus, C., Estevadeordal, A., Gallo, A., Luna, J., 2010. Information barriers, export promotion institutions, and the extensive margin of trade. *Review of World Economics* 146, 91–111.

Appendix A: Data

Table A1: Panel (B): List of predictors

Variable	Description
Value Added, Depreciation, Creditors, Current Assets, Current liabilities, Non-current liabilities, Current ratio, Debtors, Operating Revenue Turnover, Material Costs, Costs of Employees, Taxation, Financial Revenues, Financial Expenses, Interest Paid, Number of Employees, Cash Flow, EBITDA, Total Assets, Fixed Assets, Intangible Fixed Assets, Tangible Fixed Assets, Shareholders' Funds, Long-Term Debt, Loans, Sales, Solvency Ratio, Working Capital	Original financial accounts expressed in euro.
Corporate Control	A binary variable equal to one if a firm belongs to a corporate group.
Dummy Patents	equal to 1 if the firm issued any patent, and 0 otherwise.
Consolidated Accounts	A binary variable equal to one if the firm consolidates accounts of subsidiaries
NACE rev. 2	A 2-digit industry affiliation following the European Classification
NUTS 2-digit	The region in which the company is located following the European classification.
Productive Capacity	It is an indicator of investment in productive capacity computed as $\frac{Fixed\ Assets_t}{Fixed\ Assets_{t-1} + Depreciation_{t-1}}$
Capital Intensity	It is a ratio between fixed assets and number of employees for the choice of factors of production.
Labour Productivity	It is a ratio between value added and number of employees for the average productivity of labor services.
Interest Coverage Ratio (ICR)	It is a ratio between EBIT and Interest Expenses, as yet another proxy of financial constraints as in Caballero et al. (2008) .
TFP	It is the Total Factor Productivity of a firm computed as in Akerberg et al. (2015) .
Financial Constraints	It is a proxy of financial constraints as in Nickell and Nicolitsas (1999) , calculated as a ratio between interest payments and cash flow

Table A1: Panel (B): List of predictors

Variable	Description
Markup	It an estimate of a firm's markup following De Loecker and Warzynski (2012) .
ROA	It is a ratio of EBITDA on Total Assets for returns on assets.
Financial Sustainability	It is a ratio between Financial Expenses and Operating Revenues.
Size-Age	It is a synthetic indicator proposed by Hadlock and Pierce (2010), computed as $(-0.737 \cdot \log(\text{totalassets})) + (0.043 \cdot \log(\text{totalassets}))^2 - (0.040 \cdot \text{age})$ to catch the non-linear relationship between financial constraints, size and age.
Capital Adequacy Ratio	It is a ratio of Shareholders' Funds over Short and Long Term Debts.
Liquidity Ratio	A ratio between Current Assets minus Stocks and Current Liabilities.
Liquidity Returns	It is a ratio between Cash Flow and Total Assets
Regional Spillovers	It is a proxy proposed by Bernard and Jensen (2004) computed as a share of exporting plants out of total plants in a region.
Industrial spillovers	It is a proxy proposed by Bernard and Jensen (2004) computed as a share of exporting plants on total plants in a 2-digit industry.
External Economies of Scale	It is a proxy proposed by Bernard and Jensen (2004) computed as a share of exporting plants out of the total in an industry-region cell.
Size	Measure of firm size computed as (log of) number of employees.
Average Wage Bill	It is computed as (log of) costs of employees divided by number of employees.
Inward FDI	It is a binary variable with value 1 if the firm has foreign headquarters and 0 otherwise.
Outward FDI	It is a binary variable with value 1 if the firm has subsidiaries abroad and 0 otherwise.

Appendix B: Figures and Tables

Table B1: Sample coverage - size classes

NACE rev.2	Sample - N. employees						Population - N. employees					
	0-9	10-19	20-49	50-249	250+	Total	0-9	10-19	20-49	50-249	250+	Total
10	1,649	711	611	488	172	3,631	45,798	3,225	1,382	679	204	51,288
11	233	105	93	59	21	511	3,397	205	147	76	28	3,853
13	93	76	107	80	7	363	4,586	209	151	113	17	5,076
14	117	51	49	47	22	286	9,391	140	89	57	16	9,694
15	43	24	36	47	16	166	3,038	70	69	45	21	3,243
16	274	182	178	93	8	735	8,869	560	337	168	21	9,956
17	48	64	105	129	39	385	865	123	121	120	62	1,292
18	381	144	167	86	6	784	14,455	445	277	123	17	15,316
19	1	3	4	6	5	19	NA	NA	3	3	7	25
20	134	109	177	223	87	730	NA	NA	190	219	99	2,515
21	16	18	36	58	61	189	NA	NA	31	50	55	252
22	192	173	274	279	53	971	1,963	405	431	319	86	3,205
23	348	135	161	136	59	839	7,094	266	234	136	72	7,803
24	39	33	53	122	51	298	377	60	56	70	35	599
25	988	792	869	571	75	3,295	13,917	2,174	1,498	734	136	18,460
26	134	113	136	154	70	607	1,700	219	157	171	49	2,295
27	106	83	120	123	64	496	1512	169	168	136	63	2,048
28	281	171	320	319	101	1,192	2,983	455	536	399	160	4,534
29	84	62	103	157	98	504	1,092	156	160	152	75	1,635
30	36	22	30	70	41	199	838	57	63	95	55	1,107
31	148	55	78	66	9	356	8,976	164	134	68	13	9,356
32	311	121	108	102	26	668	20,551	394	217	133	44	21,338
Total	5,656	3,248	3,816	1,091	3,415	17,226	151,402	9,496	6,451	4,066	1,335	174,898

Note: French manufacturing firms are sourced from Orbis, by Bureau Van Dijk. Sample coverage by number of employees in 2018 (left panel) is compared with information on population sourced from EUROSTAT Structural Business Statistics. Please note that number of employees may report missing values from sample data, thus number of observations do not sum up to sample totals.

Table B2: Prediction accuracies after cross-validating training and testing sets

Measure	Sample 1	Sample 2	Sample 3	Sample 4	Sample 5
Sensitivity	0.649	0.647	0.654	0.65	0.648
Specificity	0.911	0.904	0.905	0.905	0.907
Balanced Accuracy	0.780	0.775	0.780	0.778	0.778
ROC	0.909	0.903	0.907	0.903	0.908
PR	0.739	0.738	0.742	0.732	0.739
N.Obs	103,540	102,748	102,169	102,028	101,712

Note: We report prediction accuracies of BART-MIA after cross-validating the algorithm on five different random training and testing sets. Our aim is to check whether predictions are robust against data sampling.

Table B3: Prediction accuracies with a subset of predictors

Model	Sensitivity	Specificity	Balanced Accuracy	ROC	PR
Logit-Lasso	0.668	0.768	0.718	0.786	0.785
CART	0.512	0.907	0.710	-	-
Random forest	0.810	0.627	0.719	0.791	0.793
BART	0.807	0.629	0.718	0.790	0.791
BART-MIA	0.623	0.914	0.768	0.902	0.725

Note: We report prediction accuracies after reducing the battery of predictors from 52 to 23 variables selected by a robust LASSO ([Ahrens et al., 2020](#)).

Table B4: Prediction accuracies after training and testing on separate years

Measure	2011	2012	2013	2014	2015	2016	2017	2018
Sensitivity	0.907	0.896	0.885	0.896	0.901	0.918	0.924	0.928
Specificity	0.637	0.632	0.641	0.627	0.639	0.651	0.652	0.654
Balanced Accuracy	0.772	0.764	0.763	0.761	0.770	0.784	0.788	0.791
ROC	0.903	0.889	0.886	0.888	0.894	0.910	0.919	0.930
PR	0.759	0.718	0.725	0.723	0.722	0.729	0.734	0.727
N.Obs	11,375	11,377	11,378	11,383	11,386	11,392	11,388	11,387

Note: We report prediction accuracies of BART-MIA after training and testing on separate years. Our aim is to check whether predictions are robust along the timeline.

Table B5: Prediction accuracies of exporters defined *à la* [Békés and Muraközy \(2012\)](#)

Exporter Class	Sensitivity	Specificity	Balanced Accuracy	ROC	PR	Num. Obs.
Permanent Exporters	0.723	0.779	0.751	0.849	0.934	76,185
Temporary Exporters	0.421	0.820	0.621	0.755	0.447	73,647
Non-Exporters		0.949				158,625
Total	0.650	0.9066	0.7783	0.9048	0.7383	232,272

Note: We report prediction accuracies after BART-MIA for firms classified according to [Békés and Muraközy \(2012\)](#): i) *permanent exporters* are firms that export at least four consecutive years; ii) *temporary exporters* are remaining firms that export at least once; iii) *non-exporters* are firms that never export.

Table B6: Prediction accuracies after an exporters' definition based on thresholds of the share of export revenues over total revenues

Measure	1 st Percentile	2 nd Percentile	5 th Percentile	Benchmark
Sensitivity	0.652	0.641	0.625	0.658
Specificity	0.835	0.837	0.852	0.833
Balanced Accuracy	0.744	0.739	0.738	0.745
ROC	0.836	0.835	0.836	0.836
PR	0.737	0.731	0.724	0.738
N.Obs	41,911	41,911	41,911	41,911

Note: We report prediction accuracies of BART-MIA after defining as exporters the firms with share of export revenues over total revenues above some specific thresholds, at the 1st, 2nd, and 5th percentiles of the distribution of the share of export revenues over total revenues.

Appendix C: Metrics of prediction accuracy

Different metrics are used to evaluate prediction accuracy of machine learning algorithms. Briefly, prediction accuracy metrics compare the classes predicted by the algorithm with the actual ones. In the case of a binary outcome, the comparison generates four classes of results:

- **True Positives:** cases when the actual class of the data point is 1 (Positive) and the predicted is also 1 (Positive);
- **False Positives:** cases when the actual class of the data point is 0 (Negative) and the predicted is 1 (Positive);
- **False Negatives:** cases when the actual class of the data point is 1 (Positive) and the predicted is 0 (Negative);
- **True Negatives:** cases when the actual class of the data point is 0 (Negative) and the predicted is also 0 (Negative);

In an ideal scenario we want to minimize the number of False Positives and False Negatives.

Table B1: Confusion Matrix

		Actual	
		<i>Positives</i> (1)	<i>Negatives</i> (0)
Predicted	<i>Positives</i> (1)	True Positives (TP)	False Positives (FP)
	<i>Negatives</i> (0)	False Negatives (FN)	True Negatives (TN)

The metrics we use to evaluate prediction accuracy in our exercises are based on the relationship between the sizes of the above classes.

Sensitivity (or Recall) Sensitivity (or Recall) is a measure of the proportion of correctly Predicted Positives, out of the total Actual Positives.

$$Sensitivity = \frac{True\ Positives}{True\ Positives + False\ Positives}$$

Specificity Specificity is a measure that catches the proportion of correctly Predicted Negatives, out of total Actual Negatives.

$$Specificity = \frac{True\ Negatives}{True\ Negatives + False\ Negatives}$$

Balanced Accuracy (BACC) The Balanced Accuracy (BACC) is a combination of Sensitivity and Specificity. It is particularly useful when classes are imbalanced, i.e., when a class appears much more often than the other. It is computed as the average between the rate of True Positives and the rate of True Negatives.

$$BACC = \frac{Sensitivity + Specificity}{2}$$

Receiving Operating Characteristics (ROC) The ROC curve is a graph showing the performance in classification at different thresholds, expressed in terms of the relationship between True Positive Rate (TPR) and False Positive Rate (FPR), defined as follows:

$$True\ Positive\ Rate = \frac{True\ Positives}{True\ Positives + False\ Negatives}$$

$$FalsePositiveRate = \frac{False\ Positives}{False\ Positives + True\ Negatives}$$

The Area Under the Curve (AUC) of ROC is then useful to evaluate performance in a bounded range between 0 and 1, where 0 indicates complete misclassification, 0.5 corresponds to an uninformative classifier, and 1 indicates perfect prediction.

Precision-Recall (PR) The PR curve is a graph showing the trade-off between Precision and Recall at different thresholds. Note that Precision and Recall are defined as follows:

$$Precision = \frac{True\ Positives}{True\ Positives + False\ Positives}$$

$$Recall = \frac{True\ Positives}{True\ Positives + False\ Negatives}$$

As for the ROC curve, the PR AUC is used to evaluate the classifier performance. A High AUC represents both high recall and high precision, thus meaning the classifier is returning accurate results (high precision), as well as returning a majority of all the positive results (high recall).

Appendix D: Location Quotients

Let us define $\mathcal{I} = \{1, \dots, n\}$ the set of non-exporting firms and $\mathcal{R} = \{1, \dots, r\}$ the set of regions (NUTS 2-digit). The r partitions of \mathcal{I} by region $j \in \mathcal{R}$ are defined as:

$$I_j \subset \mathcal{I}, j = 1, \dots, r \quad s.t. \quad \bigcup_{j=1}^r I_j = \mathcal{I}$$

Let \mathcal{P} be the set of non-exporting firms whose exporting score e is above the one of the median firm in the total distribution of non-exporters, i.e.:

$$\mathcal{P} \subset \mathcal{I} = \{i \in \mathcal{I} : e_i > median(e)\}$$

Again we can define the r partitions of \mathcal{P} by region $j \in \mathcal{R}$ as

$$P_j \subset \mathcal{P}, j = 1, \dots, r \quad s.t. \quad \bigcup_{j=1}^r P_j = \mathcal{P}$$

The location quotient, for each region $j = 1, \dots, r$ is computed as

$$LQ_j = \frac{\#P_j / \#I_j}{\#\mathcal{P} / \#\mathcal{I}} \quad (7)$$

In our case, location quotients (LQ) detect concentration of potential exporters in excess of what one would expect from the national distribution. If, for example, region j has $LQ_j = 1.5$, it implies that firms with a high trade potential are 1.5 times more concentrated in such region than the average.