Matrix completion of world trade: an analysis of interpretability through Shapley values

Matrix completion of world trade: an analysis of interpretability through Shapley values

Giorgio Gnecco[1], Federico Nutarelli[2], Massimo Riccaboni[1]
[1] IMT - School for Advanced Studies
Piazza San Francesco, 19 - 55100 Lucca, Italy
{giorgio.gnecco,massimo.riccaboni}@imtlucca.it
[2] Bocconi University, Milan, 20136, Italy
Via Sarfatti, 25 - 20100 Milan, Italy
federico.nutarelli@unibocconi.it

May 11, 2026

Abstract

In recent years, economic complexity and machine learning have become popular approaches for analyzing international trade. However, for effective use of machine learning in relation to economic complexity and policy-making, it is important to understand what are the key features for predictions. In this framework, this article addresses the issue of the interpretability of results obtained with a machine learning technique – namely, matrix completion – when applied to economic complexity, specifically in predicting Revealed Comparative Advantages (RCAs) of countries in different product categories. Shapley values are used to measure the role each country plays in predicting the RCAs of other countries. Countries relevant for prediction may differ from countries whose RCA values are similar to those of the country of interest when a standard similarity measure such as cosine similarity is used. We demonstrate the usefulness of our approach to identifying comparable countries by focusing our analysis on export diversification into complex goods of selected European countries.

Keywords: World trade, revealed comparative advantage, machine learning, interpretability, Shapley value.

1 Introduction

Economic complexity has recently become a popular tool for analysis in various fields such as economic geography, international development, innovation, and trade (Hidalgo, 2021), as it has been found that countries that diversify their exports into more complex goods grow faster (Hidalgo and Hausmann, 2009). Economic complexity applies machine learning and network science techniques to predict and explain the product upgrading of countries, regions, or cities. Recent applications in the context of international trade can be found, e.g., by Albora et al. (2023), Estmann et al. (2022), Gnecco et al. (2022), Kannen (2020), Morrison et al. (2017), Sciarra et al. (2020), and Tacchella et al. (2021).

However, for effective use of machine learning for economic complexity analysis, it is essential to understand what are the most relevant features (e.g., key countries, products, or institutional factors) involved in making predictions through economic complexity methods. This requires the application of techniques from the field of interpretable machine learning, also known as explainable artificial intelligence, an important area of research in artificial intelligence (Molnar, 2022). The main idea behind interpretability is to understand how much each feature contributes to the prediction generated by a machine learning technique. While this is trivial in linear regression or classification models – where the magnitude of each learned coefficient is a measure of the importance of the associated feature, while the distinction between learned null and non-null coefficients allows for feature selection – it is usually quite difficult to investigate in more complex models such as deep neural networks. Interpretability is of paramount importance in various applications of machine learning, as higher interpretability of a given machine learning method increases user confidence in its predictions (Bhatt et al., 2019) and is, therefore, a prerequisite for its application in critical decision areas, such as trade policy.

In this paper, we use a general interpretability technique – specifically, the Shapley value, introduced by Shapley (1953) in the context of cooperative game theory – as a tool to improve the interpretability of supervised machine learning techniques in the context of their application to the analysis of economic complexity of world trade data. We further develop the matrix completion (MC)-based approach presented in our earlier work (Gnecco et al., 2022), which did not address the interpretability issue. The idea of this work is to use MC to extract information on Revealed Comparative Advantages (RCAs) of countries in exporting certain products. This information is presented for each year in a matrix, $\mathbf{RCA} \in \mathbb{R}^{C \times P}$ (where C denotes the number of countries included in the analysis and P denotes the number of products at a given level of aggregation). The element of the matrix $\mathbf{RCA}$ in position (c, p) is defined as the RCA of country c in product category p.[1] With the goal of analyzing the topological information associated with the matrix $\mathbf{RCA}$, the economic complexity literature typically also uses the incidence matrix $\mathbf{M} \in \mathbb{R}^{C \times P}$ (Hidalgo, 2021) whose generic element in position (c, p) is equal to 1 if $RCA_{c,p} \geq 1$ and 0 otherwise.

Gnecco et al. (2022) used MC to predict RCA values associated with pairs of countries c and products p that were not considered in the training phase of the learning machine. This approach is useful for predicting the future performance of countries in trade in complex goods. The importance of the application of MC to economics is demonstrated by the recent work of Athey et al. (2021) on MC -based causal inference. In general, the interest of economists in MC and related techniques has increased recently (Athey and Imbens, 2019; Liao et al., 2018).

However, the literature has recognized the need to advise governments, particularly in less developed countries, on their choice of trade development strategies. The economic complexity literature has emphasized the importance of product relatedness in predicting future trade (Hidalgo, 2021). More recently, machine learning techniques have been successfully used to predict new products traded by a given country (Albora et al., 2023; Gnecco et al., 2022).

This paper departs from Gnecco et al. (2022) by addressing the issue of interpretability of results obtained by MC in predicting countries' RCAs in specific product categories. Shapley values are used to identify countries whose export diversification is useful in predicting the future export performance of a country of interest. To illustrate our approach, it may be useful to consider an example of applying MC to predict future trade patterns. In this context, MC takes as input a portion of a matrix whose row/column indices refer to countries/products, respectively, while the matrix elements refer to the value of exports. As a result, MC will predict for all remaining country/product pairs (c, p) whether $RCA_{c,p} \geq 1$ or not. Then, two countries (e.g., A and B) may have quite similar export patterns (hence high cosine similarity). Another country (e.g., country C) may export different products than country A, but at the same time have very similar skills in certain product categories to country A. Thus, if there are several countries whose export profiles are similar to country B's, but none similar to country C's, it may be that country C's average marginal utility in predicting country A's future export is greater than country B's. In this paper, such average marginal utility is expressed as the average increase in the probability of correctly classifying an appropriate binary classifier derived from the output of MC.

It follows from the above discussion that there is currently a need for recommender systems (Hastie et al., 2015) that can provide additional guidance to firms and countries for upgrading and diversifying exports.[2] Hence, the interpretability of the predictions obtained through MC proves to be an important issue. In this context, the ranking of countries resulting from the Shapley values (with the aim of predicting the RCA values of a given country) may be significantly different from the ranking resulting from the cosine similarities. This conclusion is supported by the fact that the cosine similarities of any two countries A and B are symmetric by construction, whereas the Shapley values are not symmetric in our context, in the following sense: the Shapley value of country A in explaining the MC predictions with respect to country B is generally different from the Shapley value of country B in explaining the MC predictions with respect to country A. The Shapley value of country A in explaining the predictions with respect to country B is generally different from the Shapley value of country B in explaining the predictions with respect to country A. That is, for example, if a developing country is similar to an industrialized country in terms of cosine similarity, one could use the export of the developing country to predict the future export of the industrialized country and vice versa because of the symmetry property of cosine similarity. However, such a property is unlikely to be useful in some applications in economics, and in particular in predicting future trade, because while it might be useful for a developing country to mimic a developed country, the opposite might not be true. In summary, Shapley values in this context provide a measure that is more directly related to the specific prediction task than, for example, cosine similarity, and can help policymakers better interpret economic complexity indices based on machine learning. In this sense, our work fits into the framework of the recent literature on international trade (see, e.g., Zhu et al. (2018)), in which novel similarity measures are defined to better capture the relationships between countries. It is worth noting that the country rankings obtained in our framework (induced by Shapley values) are tailored to each individual country of interest. Therefore, different countries would automatically receive different indications of which other countries they should compare themselves to in terms of product improvement. This should increase the potential interest of our analysis for policy makers. For example, based on the results of an analysis similar to ours, each country of interest could be suggested to look at export diversification of comparable countries (automatically determined by their individual similarity measure using the specific ranking resulting from the Shapley values). This would be particularly important for trade policies targeting less developed countries (Gnangnon, 2019).

In conclusion, this paper contributes to economic complexity theory based on machine learning (see, e.g., Hidalgo (2021)) by focusing on interpretability issues. As stated above, the use of the Shapley value can allow researchers

[1] It should be noted that $RCA_{c,p} = NaN$ is obtained when the value (expressed in international dollars) of country c's exports with respect to product p is not known. This value is typically replaced by 0 (also in the present work) as a preprocessing step.

[2] For an example of such an application see `https://beta.trade.gov/marketdiversification`.

to further deepen the predictions that arise from complexity indices based on machine learning. Regardless of a country's complexity level, the Shapley value is better than traditional similarity measures at identifying which countries contribute most to explaining predictions of future export patterns for selected countries of interest.

The article is organized as follows. Section 2 describes our methodological approach. Section 3 summarizes our findings. Finally, section 4 concludes the paper. Additional results are provided in the Appendix.

2 Methods

Our approach combines machine learning with Shapley values to make the analysis of economic complexity more interpretable. First, Section 2.1 introduces the notion of Shapley values and explains how they are used in explainable machine-learning applications. Then, section 2.2 introduces the specific formulation of MC considered in this work. Section 2.3 describes the proposed combination of MC and Shapley values for the case where MC is applied to a matrix of (suitably preprocessed) RCA values. Finally, section 2.4 compares rankings of countries based on Shapley values and cosine similarity, respectively.

2.1 Shapley value

A cooperative game with transferable utility (TU) is a pair (N, v), where N is a finite set of $n = |N|$ players and $v : 2^N \to \mathbb{R}$ is a characteristic function that assigns a value (or utility) $v(S)$ to each subset S of players (including the empty set), also called the coalition. It also satisfies the condition $v(\varnothing) = 0$. The coalition $S = N$ is called a grand coalition. The use of the term 'transferable utility' is motivated by the fact that in this type of cooperative game, players have the ability to transfer utility from one player to the other.

The Shapley value is one of the most important pointwise solution concepts in cooperative game theory. It provides a means of sharing the utility $v(N)$ of the grand coalition in a fair way among all players (such fairness can be formulated rigorously according to several well-known axiomatizations of the Shapley value: see, e.g., Maschler et al. (2013)). In other words, the Shapley value represents a measure of the importance of each player $i \in N$ for a given TU game. Formally, it is defined as follows:

$$\varphi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S| - 1))!}{n!} \left[v(S \cup \{i\}) - v(S) \right] . \tag{1}$$

It can be shown that Eq. (1) represents the average marginal utility of player i when joining the random coalition S formed by all players preceding i in a random permutation of the set of players N (assuming that all possible $n!$ permutations are equally likely). It is worth noting that the number of coalitions involved S in Eq. (1) increases exponentially with the number of players, which usually makes the exact evaluation by Eq. (1) intractable. For this reason, it is common in the literature to approximate the Shapley value, e.g., by the Monte Carlo approach described below (initiated by Mann and Shapley (1960)). In this approach, one generates a subset $\mathcal{P}_m$ of m random permutations and approximates $\varphi_i(v)$ with the empirically averaged marginal utility $\hat{\varphi}_i(v)$ of player i when joining the coalition S formed by all players preceding i in each of these permutations. It is well known that the obtained estimate of the Shapley value is unbiased and its variance is inversely proportional to the number of permutations considered (Castro et al., 2009). Nevertheless, in real-world applications of the Shapley value, it is often necessary to restrict the number of players in order to obtain a reliable Monte Carlo approximation. This is particularly the case when applying TU games to network problems (Hadas et al., 2017; Gnecco et al., 2021), where the evaluation of the characteristic function v for a given coalition S can be computationally intensive, as it involves solving an optimization (sub) problem for each of these coalitions (Passacantando et al., 2021a).

The Shapley value has several classical applications in economics, such as in taxation and redistribution, in the production of nonexclusive public goods, and in fixed-price economies, as discussed by Aumann (1994). A more recent application in the context of trade is presented by Yeung et al. (2021). More recently, the Shapley value has also found applications in machine learning, the relevance of which stems from the growing importance of machine learning applications in economics (Athey and Imbens, 2019). In recent years, the Shapley value has appeared in the literature on the interpretability of machine learning (Molnar, 2022) as an important (and quite generally applicable) means of evaluating the importance of features in machine learning, e.g., binary classification problems (Štrumbelj and Kononenko, 2014). For this purpose, one interprets each feature as a player of a TU game in which N is the set of all considered features, and the value $v(S)$ taken by the characteristic function corresponding to a subset of features S represents the probability of correct classification (or accuracy) of a binary classifier (with a certain functional form) based on this subset of features. In practice, the Monte Carlo approach mentioned above can be used to approximate the Shapley value in this context as well. One difference, however, is that since the probability of correct classification is usually not known precisely, one replaces $v(S)$ with its estimate $\hat{v}(S)$, obtained, for example, from a test set. In this framework, in this article, we compute the Shapley value for several TU games (indexed by j), where N is a set of countries. The characteristic function v_j of the j-th game represents the probability of correctly classifying the binarized RCA values associated with a given country j using a binary classifier based on MC, which is described in the next section 2.3, after details of the specific MC formulation used in this paper are given in section 2.2.

2.2 Matrix completion

This subsection provides an overview of Matrix Completion (MC). For further details, the reader is referred to Hastie et al. (2015). Roughly speaking, MC refers to the automatic reconstruction of a matrix that has been only partially observed by an appropriate machine learning algorithm. Typically, when using MC, a low-rank structure of the reconstructed matrix is specified to avoid the inadequacy of the associated MC optimization problem. In the context of the present work, MC is applied to the reconstruction of a partially observed matrix derived from the **RCA** matrix. Its low-rank structure (justified in the Supplementary Material of Gnecco et al. (2022)) is motivated by the presence of similarities of different pairs of rows (countries) or pairs of columns (products) in the **RCA**-matrix. These are due, for example, to the adoption of similar technologies or the presence of a set of unobservable skills in different countries. The importance of MC and related techniques for economics and, in particular, for international trade has already been discussed in the introduction (see also Hidalgo (2021)).

In the present article, MC formulated in terms of the following optimization problem (studied by Mazumder et al. (2010)):

$$\underset{\mathbf{Z}\in\mathbb{R}^{C\times P}}{\text{minimize}} \left(\frac{1}{2} \sum_{(c,p)\in\Omega^{\text{tr}}} \left(A_{c,p} - Z_{c,p}\right)^2 \lambda \|\mathbf{Z}\|_* \right) . \tag{2}$$

In the above representation, Ω^{tr} represents a training subset of index pairs (c, p) corresponding to the positions of the elements of the partially observed matrix $\mathbf{A} \in \mathbb{R}^{C\times P}$ provided as inputs to the learning machine, $\mathbf{Z} \in \mathbb{R}^{C\times P}$ represents the reconstructed (or completed) matrix (suitably selected by solving the above optimization problem), $\lambda \geq 0$ represents a regularization parameter (selected by a suitable validation procedure), while $\|\mathbf{Z}\|_*$ stands for the nuclear norm of the matrix $\mathbf{Z}$, i.e. the summation of all singular values of this matrix.

In this article, the optimization problem (2) is solved by applying the soft impute algorithm[3] developed by Mazumder et al. (2010) by considering several possibilities for the training set Ω^{tr}. For each such MC application, two additional subsets of index pairs (c, p) are considered. These are called the validation subset Ω^{val} and the test subset Ω^{test}. By construction, there is no overlap between any two of the three sets Ω^{tr}, Ω^{val}, and Ω^{test}. The validation and test sets are used to find an optimal choice for the regularization parameter λ and to provide a final fair evaluation for this value of the regularization parameter of the MC predictions for elements of the matrix $\mathbf{A}$ that were not used in either the training phase of MC or the validation phase.

More precisely, with the optimization problem (2) is solved for each choice of the training set Ω^{tr} 2) is solved in correspondence with multiple choices λ_k for λ distributed as $\lambda_k 2^{(k-1)/2}$ for $k1, \ldots, 30$ to avoid so-called overfitting problem.[4] The elements of the final matrix are appropriately thresholded to fit the same value interval as the elements of the matrix $\mathbf{A}$ associated with the training set. The resulting matrix is then denoted as $\mathbf{Z}_{\lambda_k}$. First, for each λ_k and each validation set Ω^{val}, the root mean square error (RMSE) of the matrix reconstruction for this set is defined as

$$RMSE^{\text{val}}_{\lambda_k} := \sqrt{\frac{1}{|\Omega^{\text{val}}|} \sum_{(c,p)\in\Omega^{\text{val}}} \left(A_{c,p} - Z_{\lambda_k,c,p}\right)^2} . \tag{3}$$

Second, the specific choice λ_{k° that minimizes $RMSE^{\text{val}}_{\lambda_k}$ for $k = 1, \ldots, 30$ is obtained. Third, the RMSE of matrix reconstruction on the related test set is evaluated, in correspondence with the optimal value λ_{k°, as

$$RMSE^{\text{test}}_{\lambda_{k^\circ}} := \sqrt{\frac{1}{|\Omega^{\text{test}}|} \sum_{(c,p)\in\Omega^{\text{test}}} \left(A_{c,p} - Z_{\lambda_{k^\circ},c,p}\right)^2} . \tag{4}$$

2.3 Combining matrix completion and Shapley values for interpretable analysis of revealed comparative advantage

In the following, we compute the Shapley value (presented in Section 2.1) for several TU games (indexed by j), where N is a set of countries. The characteristic function v_j of the j-th game represents, for different coalitions of players (countries), the probabilities of correctly classifying the binarized RCA values associated with a given country j using appropriate binary classifiers based on MC. This function is a byproduct of the numerical search

[3]The reader is referred to Mazumder et al. (2010) or to the supplementary material of Gnecco et al. (2022) for details about the soft-impute algorithm and its convergence properties.

[4]To limit the computational effort, 30 such selections are made in the present application of MC.

for the optimal solutions to various MC optimization problems, which are also formulated and solved in Section 2.2. Such problems are distinguished by the choice of three sets Ω^{tr}, Ω^{val} and Ω^{test}. Our idea is to rank the countries that differ from j according to their importance (as measured by their – suitably approximated – Shapley values associated with the j-th game) in explaining the MC predictions of binarized RCA elements associated with country j. As explained in the introduction, this use of the (approximated) Shapley value is useful in finding subsets of countries that are thought to be relevant in predicting the products exported by country j. Moreover, it is crucial in recommending that country j consider the products exported by these countries for its future trade. Based on the results of the combined MC -Shapley value analysis, country j is expected to have similar export capabilities.

In the following, we describe the algorithm proposed in the present subsection for the approximate computation of the Shapley values for the characteristic functions v_j introduced in the above paragraph. The algorithm takes as input a matrix $\mathbf{A} \in \mathbb{R}^{C \times P}$, which is obtained by discretizing and shifting the original **RCA** matrix,[5] such that the elements of $\mathbf{A}$ are symmetrically distributed around 0 (this shift is justified by the fact that the soft-impute algorithm typically underestimates the matrix elements due to the presence of the regularization parameter λ in the MC optimization problem (2)). Due to the high computational cost required for the approximate evaluation of Shapley values, we consider a reduced number C of countries compared to Gnecco et al. (2022), in which the MC analysis was not combined with Shapley values. See section 3 for details on the possible choices of the number C and the set of C associated countries. As for the number of products, in this paper it is chosen equal to $P1,243$, which refers to an analysis at the 4-digit level according to the Harmonized System Codes 1992 (HS -1992).

The proposed algorithm, which combines MC and the Shapley values for an interpretable analysis of revealed comparative advantage, is described as follows (comments on the various steps are reproduced later after the description).

1. One generates Q' random permutations $\sigma'_1, \sigma'_2, \ldots, \sigma'_Q$ of the set of all C countries. Then, for each such permutation, other C permutations are generated, by applying a circular shift C times (since this is done C times, the circular shift could be made indifferently either to the left or to the right). In total, $Q = Q'C$ permutations are obtained, denoted as $\sigma_1, \sigma_2, \ldots, \sigma_Q$.

2. For each permutation σ_q (for $q = 1, 2, \ldots, Q$), one generates other $C-1$ permutations (named $\sigma_q^1, \sigma_q^2, \ldots, \sigma_q^{C-1}$) by circularly shifting (for a total number of $C - 1$ times, again moving either to the left or to the right) the elements $\sigma_q(1), \sigma_q(2), \ldots, \sigma_q(C-1)$, while keeping the last element $\sigma_q(C)$ always in the last position. In this way, each country different from $\sigma_q(C)$ appears in the first position exactly once (within the subset of permutations $\sigma_q^1, \sigma_q^2, \ldots, \sigma_q^{C-1}$).

3. For each permutation σ_q^r (with $r = 1, 2, \ldots, C-1$), one generates $C-1$ coalitions (subsets) of countries, named $S^1_{\sigma_q^r}, S^2_{\sigma_q^r}, \ldots, S^{C-1}_{\sigma_q^r}$, with respective cardinalities $|S^1_{\sigma_q^r}| = 1, |S^2_{\sigma_q^r}| = 2, \ldots, |S^{C-1}_{\sigma_q^r}| = C-1$. Such coalitions are defined as $S^u_{\sigma_q^r} = \{\sigma_q^r(s)\}_{s=1}^u$, for $u = 1, 2, \ldots, C-1$. In other words, they collect, respectively, the first u elements in the permutation σ_q^r.

4. The MC optimization problem (2) is solved in correspondence of each subset $S^u_{\sigma_q^r}$ generated according to Step 3 above. For that application of MC, the training set is constructed as follows. First, one inserts in that training set the positions of all the elements in the rows of the matrix $\mathbf{A}$ (i.e., of the discretized and shifted version of the **RCA** matrix) that belong to $S^u_{\sigma_q^r}$ and that are not associated with originally NaN RCA values. Moreover, the training set also contains a subset of the elements that belong to the other rows of the matrix $\mathbf{A}$ and that are not associated with originally NaN RCA values. Such elements are inserted in the training set, randomly and independently, with probability $p_{\text{insertion}} = 0.7$. All the other elements of the matrix $\mathbf{A}$ are hidden to MC. In particular, the positions of the hidden elements in the row $j = \sigma_q(C)$ and that are not associated with originally NaN RCA values form the test set. The same test set is considered for the various subsets $S^u_{\sigma_q^r}$ obtained starting from the same permutation σ_q^r. Similarly, the validation set (used to select the optimal value of λ) is formed by the positions of the hidden elements in the other rows (different from j, and not belonging to $S^u_{\sigma_q^r}$) that are not associated with originally NaN RCA values. Each of the matrices reconstructed by MC in this step is thresholded likewise in Section 2.2. Then, the resulting matrix, associated with the optimal value of the regularization parameter, is denoted as $\mathbf{Z}^{(S^u_{\sigma_q^r})}$ (for simplicity, its dependence on the regularization parameter is not reported in the notation).

5. For each matrix $\mathbf{Z}^{(S^u_{\sigma_q^r})}$, one constructs a binary classifier as follows. Every time an entry $Z_{c,p}^{(S^u_{\sigma_q^r})}$ of the matrix $\mathbf{Z}^{(S^u_{\sigma_q^r})}$ is in the test set, such an entry is attributed to the "negative" class (which is associated with $0 \leq RCA < 1$) in case its MC prediction is smaller than 0, otherwise, it is attributed to the "positive" class (which is associated with $RCA \geq 1$).

[5]The details of the discretization and shifting are the same as those given in the supplementary material of Gnecco et al. (2022).

6. The value assumed by the characteristic function $v_j(S^u_{\sigma^r_q})$ in correspondence of the coalition $S^u_{\sigma^r_q}$ is approximated by the average accuracy on the test set of the MC-based binary classifier constructed in Step 5. For each subset $S^u_{\sigma^r_q}$, this average is obtained by repeating N_{rep} times the procedure detailed in Step 4. The country $j = \sigma_q(C)$ associated with the test set is considered as a so-called null player for the game associated with the characteristic function v_j.

7. For each $j = 1, 2, \ldots, C$, one considers the set $\mathcal{P}^{(j)}_m$ of permutations composed of all permutations σ^r_q for which $\sigma^r_q(C) = j$ (i.e., for which the test set is associated with the specific country j). Then, the Shapley values $\varphi_i(v_j)$ (for $i \neq j$) are computed approximately according to the Monte Carlo procedure reported in Section 2.1, setting $\mathcal{P}_m = \mathcal{P}^{(j)}_m$. Finally, one takes $\varphi_i(v_j) = 0$ (being j a null player for the game associated with the characteristic function v_j).

It is worth explaining in less technical terms the motivations behind the main steps reported above. Due to Step 1, each country is associated with the test set the same number of times (as it turns out to be the last country in the permutation for exactly Q' of the permutations σ_q). In this way, the number of permutations used to estimate the Shapley values $\varphi_i(v_j)$ in Step 7 is the same for each country j (avoiding subsampling for some countries with respect to other countries). It is worth remarking that the number Q' is inversely proportional to the variance of the estimate of each Shapley value (Castro et al., 2009). So, increasing Q' reduces that variance, but at the same time, it increases the computational time needed to generate the estimate. Moreover, the motivation behind the common choice of the test set in Step 4 for the various subsets $S^u_{\sigma^r_q}$ obtained starting from the same permutation $S^u_{\sigma^r_q}$ is that, in this way, the estimated marginal utilities of players are less dependent on the variability of the test set. Finally, due to Step 2, for each j, when $j = \sigma_q(C)$, each country $i \neq j$ enters the same number of times as the first player, to form the coalitions $S^1_{\sigma^r_q}, S^2_{\sigma^r_q}, \ldots, S^{C-1}_{\sigma^r_q}$ used to estimate $\varphi_i(v_j)$, according to the Monte Carlo procedure of Step 7. This is motivated by the fact that the first player that enters a coalition may have an advantage over the other players. Indeed, its marginal utility is the estimated accuracy of a binary classifier (since the utility of the empty coalition is always 0, by the definition of characteristic function), whereas the marginal utilities of the subsequent players are differences of estimated accuracies of binary classifiers. By making all the countries enter the first position the same number of times, this possible first-player advantage is reduced (with respect to the case in which, due to random sampling, one country appears as the first player a much larger number of times with respect to another country). Moreover, by construction, due to Step 2, each country $i \neq j$ enters the same number of times also as second player, third player, and so on.

2.4 Ranking of countries according to Shapley values and cosine similarities

The methodology presented in section 2.3, based on the (numerical approximations of the) Shapley values $\varphi_i(v_j)$, makes it possible to identify the different characteristics" (i.e., the different countries i whose corresponding rows in the matrix $\mathbf{A}$ can only contain elements of the training set) according to how important they are in generating the MC -based classifications of the RCA values of each specific country j (assigned to the test set). Each ranking is created by ordering the Shapley values $\varphi_i(v_j)$ for a fixed j.

Another (this time unsupervised) way to rank countries i for fixed j is to replace the Shapley values $\varphi_i(v_j)$, for example, with the cosine similarities between the RCA lines $\mathbf{a}_i$ and $\mathbf{a}_j$ associated with the countries i and j (after removing from both rows all elements that appear in columns for which at least one of the two elements is originally a NaN value). Recall here that the cosine similarity between two (nonzero) row vectors $\mathbf{a}_i$ and $\mathbf{a}_j$ of the same dimension is defined as $S_C(\mathbf{a}_i, \mathbf{a}_j) \frac{\mathbf{a}_i \cdot \mathbf{a}_j}{\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2}$, where $\cdot$ denotes the inner product.

It is argued here that the approximate Shapley value-based ranking of countries is more appropriate than, for example, their cosine similarity-based ranking when examining the interpretability of the performance of MC. This is inspired by the following machine learning analogy: features selected by linear discriminant analysis (a supervised classification technique) tend to outperform features selected by principal component analysis (an unsupervised dimensionality reduction technique) when the machine learning task under study is supervised (Belhumeur et al., 1997). The issue of comparison between approximate Shapley value-based country rankings and cosine similarity-based rankings is discussed in more detail in section 3 using numerical results.

3 Results

This section contains some illustrative results of the methods described in section 2. The RCA values used in the following analyzes are calculated starting from the trade data contained in the CEPII - BACI dataset, which is freely distributed under the Etalab Open License 2.0 and can be accessed at the following hyperlink: `http://www.cepii.fr/cepii/en/bdd_modele/bdd.asp`. For the analyzes presented in this section, the year 2018 is considered.

Due to the high computational effort needed for the approximate evaluation of the Shapley values, only a subset of European countries is considered in the following, choosing $C = 14$ for the number of countries. In more detail, this choice for the number of countries is made with the aim of achieving a good compromise between the representativeness of the subset of European countries considered in the analysis, and the computational time needed for the approximation of their Shapley values, which in the present application increases quite rapidly with respect to the number of players (countries), as it can be deduced by the dependence of eq. (1) from that number. It is worth observing that, compared with our previous work Gnecco et al. (2022) – in which it was possible to include in the MC analysis a much larger number of countries (i.e., 119) – in the present work the MC application is much more expensive from a computational point of view because of the presence in the analysis of the additional cooperative game-theoretical component. Indeed, for the present analysis, one needs to vary not only the choice of the regularization parameter and of the positions of the hidden matrix elements but also of the specific coalition of countries (which can be in principle any subset of the set of 14 selected countries).

The specific 14 countries are chosen in such a way to span geographically across Europe and to explain at the same time a large portion of its Gross Domestic Product (GDP). Such countries are Albania (ALB), Austria (AUT), Belgium (BEL), Bulgaria (BGR), Germany (DEU), Spain (ESP), France (FRA), United Kingdom (GBR), Greece (GRC), Hungary (HUN), Italy (ITA), Poland (POL), Russia (RUS), and Sweden (SWE). Moreover, the parameters Q' and N_{rep} are set to 50 and 10, respectively. In this way, the total number of permutations σ_q^r is $Q = Q'C = 50 \cdot 14 = 700$. Hence, for each country j associated with the test set and for each choice of the regularization parameter λ, the number of applications of MC is $Q'C(C-1)N_{\text{rep}} = 50 \cdot 14 \cdot 13 \cdot 10 = 91,000$.

Table 1 reports the resulting approximations $\hat{\varphi}_i(v_j)$ of the Shapley values. The elements on the main diagonal of the table are not relevant, hence they are not reported. It is worth remarking that, by construction, the approximate Shapley values can be either positive, negative, or zero. In the case of the table, they turn out to be all positive, meaning that the inclusion of a player on average increases the (estimated) probability of correct classification.

	Country j													
Country i	ALB	AUT	BEL	BGR	DEU	ESP	FRA	GBR	GRC	HUN	ITA	POL	RUS	SWE
ALB	.	0.054	0.053	**0.057**	0.056	0.050	0.052	0.054	0.058	**0.060**	0.049	0.055	**0.068**	0.057
AUT	0.058	.	0.052	0.054	**0.060**	0.050	0.053	0.054	0.058	0.058	0.051	0.055	0.061	0.061
BEL	0.055	0.054	.	0.054	0.055	0.052	0.052	0.055	0.057	0.057	0.049	0.055	0.062	0.056
BGR	0.059	0.054	0.052	.	0.056	0.048	0.052	0.053	0.056	0.058	0.047	0.054	0.061	0.055
DEU	0.058	**0.060**	0.050	0.055	.	0.050	0.051	**0.058**	0.059	0.059	0.048	0.055	0.062	0.058
ESP	0.054	0.056	0.054	0.054	0.054	.	0.051	0.056	0.059	0.058	**0.054**	0.055	0.061	0.057
FRA	0.054	0.055	0.053	0.051	0.056	0.052	.	0.057	0.056	0.056	0.049	0.053	0.061	0.054
GBR	0.056	0.055	0.054	0.056	0.059	0.050	**0.055**	.	0.055	0.058	0.048	0.053	0.063	0.058
GRC	0.058	0.055	0.053	0.056	0.054	**0.056**	0.051	0.058	.	0.059	0.052	0.053	0.067	0.057
HUN	0.055	0.054	0.052	0.057	0.055	0.049	0.053	0.054	0.057	.	0.046	**0.060**	0.063	0.056
ITA	0.059	0.054	**0.055**	0.054	0.056	0.054	0.051	0.055	0.057	0.056	.	0.054	0.063	0.058
POL	0.056	0.057	0.053	0.055	0.057	0.050	0.050	0.056	0.058	0.060	0.047	.	0.060	0.058
RUS	**0.062**	0.054	0.050	0.056	0.054	0.051	0.051	0.055	**0.064**	0.059	0.053	0.054	.	**0.063**
SWE	0.057	0.059	0.051	0.052	0.057	0.050	0.052	0.056	0.057	0.057	0.050	0.051	0.065	.

Table 1: Approximations of the Shapley values $\hat{\varphi}_i(v_j)$ in 2018 for the selected 14 European countries (i=row index, j=column index) considered in Section 3. For each country j, the country $i \neq j$ presenting the largest approximate Shapley value $\hat{\varphi}_i(v_j)$ is highlighted.

Then, for each country j, Table 2 orders the other countries according to the so-obtained approximate Shapley values $\hat{\varphi}_i(v_j)$ (from the largest to the smallest). Country j (for which the Shapley value is 0) is excluded from such rankings. The rankings range from 1^{st} (the best) to 13^{th} (the worst).

	Country j													
Ranking	ALB	AUT	BEL	BGR	DEU	ESP	FRA	GBR	GRC	HUN	ITA	POL	RUS	SWE
1^{st}	RUS	DEU	ITA	ALB	AUT	GRC	GBR	DEU	RUS	ALB	ESP	HUN	ALB	RUS
2^{nd}	BGR	SWE	ESP	HUN	GBR	ITA	AUT	GRC	DEU	POL	RUS	ALB	GRC	AUT
3^{rd}	ITA	POL	GBR	GBR	POL	BEL	HUN	FRA	ESP	DEU	GRC	AUT	SWE	DEU
4^{th}	AUT	ESP	ALB	GRC	SWE	FRA	ALB	ESP	ALB	GRC	AUT	BEL	GBR	GBR
5^{th}	DEU	FRA	FRA	RUS	ALB	RUS	BEL	POL	AUT	RUS	SWE	DEU	HUN	ITA
6^{th}	GRC	GBR	GRC	DEU	BGR	ALB	BGR	SWE	POL	AUT	ALB	ESP	ITA	POL
7^{th}	SWE	GRC	POL	POL	FRA	AUT	SWE	BEL	BEL	BGR	BEL	BGR	BEL	ALB
8^{th}	GBR	ALB	AUT	AUT	ITA	DEU	DEU	ITA	HUN	ESP	FRA	ITA	DEU	ESP
9^{th}	POL	BEL	BGR	BEL	BEL	GBR	ESP	RUS	ITA	GBR	DEU	RUS	AUT	GRC
10^{th}	BEL	BGR	HUN	ESP	HUN	POL	GRC	ALB	SWE	BEL	GBR	FRA	BGR	BEL
11^{th}	HUN	HUN	SWE	ITA	ESP	SWE	ITA	AUT	BGR	SWE	BGR	GBR	ESP	HUN
12^{th}	ESP	ITA	DEU	SWE	GRC	HUN	RUS	HUN	FRA	FRA	POL	GRC	FRA	BGR
13^{th}	FRA	RUS	RUS	FRA	RUS	BGR	POL	BGR	GBR	ITA	HUN	SWE	POL	FRA

Table 2: For each column j: ranking of countries $i \neq j$ induced by the approximate Shapley values $\hat{\varphi}_i(v_j)$ in 2018 reported in Table 1.

In a similar way, Table 3 reports the cosine similarities $S_C(\mathbf{a}_i, \mathbf{a}_j)$. The elements on the main diagonal of the table are not relevant, hence they are not reported. It is worth remarking that, by construction, the cosine similarities can be either positive, negative, or zero. In the case of the table, they turn out to be all positive. By a comparison of Tables 1 and 3, it is worth observing that, while the cosine similarities are always symmetric, i.e., $S_C(\mathbf{a}_i, \mathbf{a}_j) = S_C(\mathbf{a}_j, \mathbf{a}_i)$, this is not the case for the approximate Shapley values, i.e., one gets typically $\hat{\varphi}_i(v_j) \neq \hat{\varphi}_j(v_i)$. This depends on the fact that, in the case of $\hat{\varphi}_i(v_j)$, the test set is constructed based on country j, whereas in the case of $\hat{\varphi}_j(v_i)$, it is built based on country i. In this context, it is easy to see how the symmetry property of the cosine similarity limits its potential use to identify countries to be used to predict future trade. For instance, the cosine similarity predicts that Poland is very similar to Germany in terms of the RCA values and that, symmetrically, Germany is very similar to Poland (still in terms of the RCA values). The literature, however, suggests that the cosine similarity may not be the most appropriate measure, given that Germany and Poland differ in several aspects in terms of export capabilities (Adigwe, 2022; Bojnec and Ferto, 2016; Kuźnar, 2016).

	Country j													
Country i	ALB	AUT	BEL	BGR	DEU	ESP	FRA	GBR	GRC	HUN	ITA	POL	RUS	SWE
ALB	.	0.023	0.055	**0.414**	0.038	0.056	0.048	0.022	0.113	**0.446**	0.087	0.333	0.089	0.015
AUT	0.023	.	0.262	0.101	0.447	0.227	0.324	0.274	0.0912	0.183	0.280	0.273	0.135	0.261
BEL	0.055	0.263	.	0.112	0.495	0.305	0.455	0.424	0.156	0.206	0.289	0.365	0.168	0.261
BGR	0.414	0.101	0.112	.	0.150	0.122	0.144	0.127	0.094	0.132	0.174	0.143	0.184	0.070
DEU	0.038	**0.447**	**0.495**	0.150	.	0.383	0.401	**0.547**	0.169	0.343	**0.489**	**0.505**	0.186	**0.387**
ESP	0.056	0.227	0.305	0.1212	0.383	.	0.401	0.287	**0.394**	0.189	0.437	0.343	0.115	0.195
FRA	0.048	0.324	0.455	0.144	0.401	0.401	.	0.486	0.150	0.293	0.415	0.360	0.177	0.305
GBR	0.022	0.274	0.424	0.127	**0.547**	0.287	**0.486**	.	0.147	0.214	0.374	0.338	0.168	0.287
GRC	0.113	0.091	0.156	0.094	0.169	0.394	0.150	0.147	.	0.091	0.261	0.184	0.075	0.105
HUN	**0.446**	0.183	0.206	0.132	0.344	0.189	0.293	0.214	0.091	.	0.210	0.316	0.105	0.149
ITA	0.087	0.280	0.289	0.174	0.489	**0.437**	0.415	0.374	0.261	0.210	.	0.318	0.101	0.235
POL	0.333	0.273	0.365	0.1427	0.505	0.343	0.360	0.338	0.184	0.316	0.318	.	0.229	0.270
RUS	0.089	0.135	0.168	0.184	0.186	0.115	0.177	0.168	0.075	0.105	0.101	0.229	.	0.238
SWE	0.015	0.261	0.261	0.070	0.387	0.195	0.305	0.287	0.105	0.149	0.235	0.270	**0.238**	.

Table 3: Cosine similarities $S_C(\mathbf{a}_i, \mathbf{a}_j)$ in 2018 for the selected 14 European countries (i=row index, j=column index) considered in Section 3. For each country j, the country $i \neq j$ presenting the largest cosine similarity $S_C(\mathbf{a}_i, \mathbf{a}_j)$ with respect to j is highlighted.

Then, for each country j, Table 4 orders the other countries according to the so-obtained cosine similarities $S_C(\mathbf{a}_i, \mathbf{a}_j)$ (from the largest to the smallest). Again, country j (for which the cosine similarity is 1 by default) is excluded from such rankings.

	Country j													
Ranking	ALB	AUT	BEL	BGR	DEU	ESP	FRA	GBR	GRC	HUN	ITA	POL	RUS	SWE
1^{st}	HUN	DEU	DEU	ALB	GBR	ITA	GBR	DEU	ESP	ALB	DEU	DEU	SWE	DEU
2^{nd}	BGR	FRA	FRA	RUS	POL	FRA	BEL	FRA	ITA	DEU	ESP	BEL	POL	FRA
3^{rd}	POL	ITA	GBR	ITA	BEL	GRC	ITA	BEL	POL	POL	FRA	FRA	DEU	GBR
4^{th}	GRC	GBR	POL	DEU	ITA	DEU	DEU	ITA	DEU	FRA	GBR	ESP	BGR	POL
5^{th}	RUS	POL	ESP	FRA	AUT	POL	ESP	POL	BEL	GBR	POL	GBR	FRA	BEL
6^{th}	ITA	BEL	ITA	POL	FRA	BEL	POL	ESP	FRA	ITA	BEL	ALB	BEL	BEL
7^{th}	ESP	SWE	AUT	HUN	SWE	GBR	AUT	SWE	GBR	BEL	AUT	ITA	GBR	RUS
8^{th}	BEL	ESP	SWE	GBR	ESP	AUT	SWE	AUT	ALB	ESP	GRC	HUN	AUT	ITA
9^{th}	FRA	HUN	HUN	ESP	HUN	SWE	HUN	HUN	SWE	AUT	SWE	AUT	ESP	ESP
10^{th}	DEU	RUS	RUS	BEL	RUS	HUN	RUS	RUS	BGR	SWE	HUN	SWE	HUN	HUN
11^{th}	AUT	BGR	GRC	AUT	GRC	BGR	GRC	GRC	AUT	BGR	BGR	RUS	ITA	GRC
12^{th}	GBR	GRC	BGR	GRC	BGR	RUS	BGR	BGR	HUN	RUS	RUS	GRC	ALB	BGR
13^{th}	SWE	ALB	ALB	SWE	ALB	ALB	ALB	ALB	RUS	GRC	ALB	BGR	GRC	ALB

Table 4: For each column j: ranking of countries $i \neq j$ induced by the cosine similarities $S_C(\mathbf{a}_i, \mathbf{a}_j)$ in 2018 reported in Table 3.

Furthermore, for each $j = 1, 2, \ldots, 14$, Table 5 shows the Kendall τ correlation between the two rankings of countries $i \neq j$ induced by $\hat{\varphi}_i(v_j)$ and $S_C(\mathbf{a}_i, \mathbf{a}_j)$. These rankings have been already reported in Tables 2 and 4, respectively. As expected, for each j, the two rankings differ, meaning that a country i with a high cosine similarity with respect to the country j is not necessarily the most important country for what concerns the prediction of the RCA values of that country j.

	Country j													
Kendall τ correlation	ALB	AUT	BEL	BGR	DEU	ESP	FRA	GBR	GRC	HUN	ITA	POL	RUS	SWE
Shapley value vs. cosine similarity	0.099	-0.011	-0.231	0.143	-0.033	0.275	-0.121	-0.165	0.033	-0.099	0.319	0.033	-0.033	-0.121

Table 5: For each column j: Kendall τ correlation between the respective rankings of countries $i \neq j$ induced by the approximate Shapley values and by the cosine similarities in 2018, already reported in Tables 2 and 4.

In the following, the economic relevance of the proposed approach – which combines MC with the Shapley value – is clarified, by showing that the approximate Shapley value-based rankings of countries reported in Table 2 are more appropriate than their cosine similarity-based rankings reported in Table 4, in the sense that they allow selecting subsets of countries associated with better MC performance (please see the end of Section 2.4 for an explanation of the reason why this is expected). In other words, such countries allow one to better classify RCA values of a country of interest.

In more detail, for each country j associated with the test set, we apply MC several times likewise in Step 4 of the algorithm described in Section 2.3, by replacing, for $r = 1, \ldots, \lfloor (C-1)/2 \rfloor = 6$, the subset $S^u_{\sigma^r_q}$ considered therein with a subset of countries made, respectively, by:

1. the first r countries in the ranking of countries induced by the approximate Shapley values, reported in column j of Table 2;
2. the first r countries in the ranking of countries induced by the cosine similarities, reported in column j of Table 4;
3. the last r countries in the ranking of countries induced by the approximate Shapley values, reported in column j of Table 2;
4. the last r countries in the ranking of countries induced by the cosine similarities, reported in column j of Table 4.

Then, we evaluate, for each of these cases, the average performance of MC on the test set, in terms of the estimated probability of correct classification. More precisely, for a more reliable evaluation, for each country j this average is performed now by considering $N_{\text{rep}} = 20$ MC applications with various independently generated test sets associated with the same country j, then the estimated probability of correct classification is averaged over all such test sets. The expectation is that for a generic value of r:

– first, the selection of the subset of countries obtained in Case 1 (countries with the r largest approximate Shapley values) is on average better than the one obtained in Case 2 (countries with the r largest cosine similarities);

– second, the selection of the subset of countries obtained in Case 1 (countries with the r largest approximate Shapley values) is on average better than the one obtained in Case 3 (countries with the r smallest approximate Shapley values);

– third, the selection of the subset of countries obtained in Case 4 (countries with the r smallest cosine similarities) is on average better than the one obtained in Case 3 (countries with the r smallest approximate Shapley values).

The results of this further comparison are reported in Table 6 and confirm that the approximate Shapley value-based rankings of countries look, indeed, more appropriate than their cosine similarity-based rankings for the present application of MC to international trade data. In particular, by comparing the corresponding entries in the table, it follows that:

– Case 1 produces better results than Case 2 about 74% of the time;

– Case 1 produces better results than Case 3 about 88% of the time;

– Case 4 produces better results than Case 3 about 68% of the time.

It is worth remarking that the problem of finding the best subset of r players (here, the one associated with the smallest estimated probability of correct classification) is a difficult combinatorial problem, whose optimal solution is approximated here by the subset of r players with the r largest approximate Shapley values. Although this is expected to be typically a good approximation to that optimal solution (Narayanam and Narahari, 2010), the fact of being just an approximation is the reason why the percentages above, despite being high, are not near 100%.

	Country j													
r	ALB	AUT	BEL	BGR	DEU	ESP	FRA	GBR	GRC	HUN	ITA	POL	RUS	SWE
1	0.736	0.706	0.674	0.712	0.721	0.647	0.672	0.711	0.744	0.743	0.622	0.697	0.798	0.742
2	0.746	0.719	0.686	0.711	0.718	0.648	0.668	0.712	0.745	0.746	0.631	0.698	0.803	0.741
3	0.745	0.718	0.674	0.711	0.729	0.658	0.677	0.720	0.745	0.750	0.643	0.710	0.811	0.742
4	0.744	0.717	0.669	0.702	0.727	0.652	0.669	0.714	0.760	0.744	0.648	0.705	0.820	0.747
5	0.745	0.714	0.672	0.708	0.724	0.661	0.680	0.716	0.751	0.759	0.635	0.707	0.820	0.746
6	0.748	0.722	0.682	0.712	0.728	0.664	0.670	0.727	0.747	0.755	0.639	0.701	0.817	0.754
Case 1. First r countries in the rankings induced by the approximate Shapley values $\hat{\varphi}_i(v_j)$														
1	0.727	0.720	0.674	0.696	0.718	0.652	0.667	0.700	0.750	0.745	0.624	0.694	0.798	0.733
2	0.731	0.721	0.680	0.714	0.700	0.658	0.682	0.708	0.736	0.739	0.633	0.694	0.797	0.729
3	0.715	0.713	0.674	0.704	0.716	0.658	0.669	0.711	0.743	0.752	0.622	0.686	0.796	0.733
4	0.727	0.715	0.680	0.708	0.719	0.651	0.670	0.712	0.735	0.744	0.623	0.696	0.806	0.735
5	0.742	0.715	0.677	0.712	0.726	0.658	0.677	0.713	0.748	0.754	0.628	0.701	0.799	0.745
6	0.734	0.717	0.678	0.706	0.716	0.657	0.673	0.708	0.743	0.747	0.630	0.695	0.805	0.741
Case 2. First r countries in the rankings induced by the cosine similarities $S_C(\mathbf{a}_i, \mathbf{a}_j)$														
1	0.715	0.704	0.676	0.693	0.705	0.649	0.671	0.704	0.731	0.737	0.626	0.695	0.788	0.738
2	0.731	0.704	0.675	0.706	0.711	0.645	0.661	0.707	0.741	0.747	0.626	0.690	0.784	0.730
3	0.720	0.709	0.679	0.701	0.706	0.637	0.659	0.705	0.739	0.743	0.630	0.697	0.789	0.730
4	0.724	0.708	0.675	0.705	0.716	0.639	0.672	0.702	0.743	0.742	0.623	0.692	0.794	0.737
5	0.726	0.706	0.671	0.700	0.707	0.641	0.671	0.718	0.745	0.744	0.627	0.694	0.792	0.737
6	0.731	0.704	0.672	0.703	0.710	0.633	0.675	0.705	0.743	0.741	0.6295	0.693	0.789	0.733
Case 3. Last r countries in the rankings induced by the approximate Shapley values $\hat{\varphi}_i(v_j)$														
1	0.719	0.709	0.677	0.702	0.706	0.637	0.674	0.709	0.728	0.743	0.629	0.697	0.795	0.728
2	0.717	0.711	0.673	0.699	0.707	0.654	0.673	0.708	0.748	0.747	0.629	0.689	0.793	0.737
3	0.729	0.708	0.669	0.704	0.713	0.650	0.667	0.694	0.739	0.747	0.633	0.695	0.811	0.728
4	0.725	0.705	0.668	0.707	0.705	0.645	0.676	0.706	0.743	0.760	0.629	0.708	0.807	0.734
5	0.734	0.704	0.671	0.707	0.711	0.645	0.663	0.712	0.744	0.740	0.636	0.697	0.810	0.737
6	0.741	0.718	0.680	0.711	0.708	0.633	0.670	0.702	0.747	0.746	0.633	0.697	0.812	0.733
Case 4. Last r countries in the rankings induced by the cosine similarities $S_C(\mathbf{a}_i, \mathbf{a}_j)$														

Table 6: Average estimated probabilities of correct classification on the test set for the Cases 1–4 described in Section 3. The results refer to the year 2018 and to the 14 selected European countries considered in Section 3.

Moreover, Figures 1 (a) and (b) show a colored visualization of the average ranking of each country involved in the analysis presented in this section, based on the results reported in Tables 2 and 4, which are related, respectively, to the approximate Shapley values and to the cosine similarities. In more detail, for each of the 14 countries, one averages its 13 rankings reported in the various columns of each of the two Tables 2 and 4, excluding the column associated with the country itself.

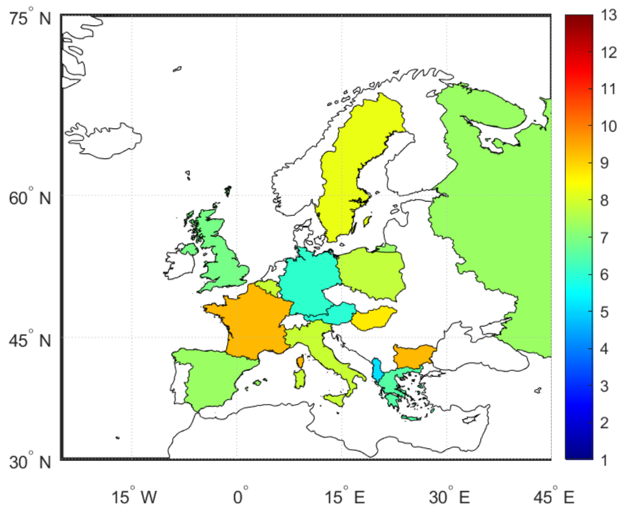

(a) Average rankings induced by the approximate Shapley values in 2018 for the 14 selected European countries.

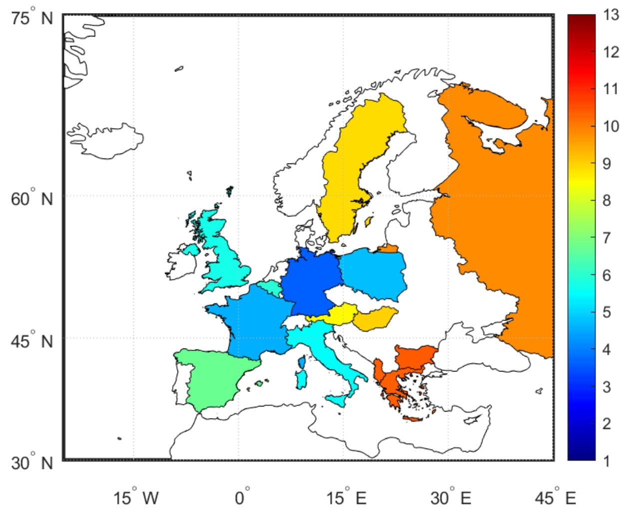

(b) Average rankings induced by the cosine similarities in 2018 for the 14 selected European countries.

Figure 1: Visual representations of the average rankings (from 1st to 13th) induced by the approximate Shapley values and by the cosine similarities in 2018 for the 14 selected European countries considered in Section 3.

It is worth remarking that the average ranking reported in Figure 1 (a) has to be interpreted not as a novel economic complexity index, but – taking into account the discussion of the results contained in Table 6 – as the average importance of each country in producing the MC-based classifications of the RCA values of the other 13 countries. Analogously, the average ranking reported in Figure 1 (b) represents how much each country is similar on average (according to the cosine similarity) to the other 13 countries considered in the analysis. As expected, the average rankings obtained in the two cases differ substantially. In particular, for the dataset analyzed, it turns out that the average ranking obtained by using the approximate Shapley values has a narrower distribution with respect to the one obtained by using the cosine similarities.

Finally, in order to further illustrate the potentialities of the proposed approach, an additional comparison between the average rankings induced, respectively, by the approximate Shapley values and by the cosine similarities is provided in the Appendix, with reference to the 2004 expansion of the European Union and with a slightly different choice of the set of countries considered in the analysis.

4 Conclusions

In this work, the interpretability of a machine-learning method whose application has been recently introduced in the context of economic complexity has been addressed. More specifically, a general-purpose interpretability technique - based on the Shapley value - has been applied to address the interpretability of the results of matrix completion, when it is applied to the reconstruction of the incidence matrix associated with the matrix of Revealed Comparative Advantages (RCAs) of various country/sector pairs, with the aim of addressing the degree of predictability of the elements of the incidence matrix related to each country. Indeed, economic complexity indices based on matrix completion (such as the one studied by Gnecco et al. (2022), in which matrix completion was applied with the aim of addressing the degree of predictability of the elements of the incidence matrix associated with each country) could be used to define suitable economic and financial recommender systems, by identifying countries and associated economic sectors that show relevant growth opportunities, on the basis of the predictions made by matrix completion itself. This calls for an interpretability analysis of such predictions, which in this work has been performed in terms of the Shapley value. This use of interpretability techniques is expected to be extendable to other applications of matrix completion to the analysis of international trade data, such as the one investigated by Metulini et al. (2022), in which matrix completion was combined with hierarchical clustering for the reconstruction of portions of Input/Output (I/O) tables.

Our methodological framework is expected to be beneficial for policymakers, in the sense that, based on a similar analysis as ours, each country of interest for a trade policy analysis (e.g., a developing country) could be suggested to look at the export capabilities' history of its most similar countries. Such similarity would be automatically measured on the basis of the proposed approximate Shapley values, which are expected to represent a more reliable similarity measure (in view of its application to international trade) than traditional approaches such as cosine similarity.

Still, since the Shapley value has a high computational cost in this context, it has been necessary to limit the analysis to a small number of countries. In order to extend it to a significantly larger number of countries, possibly at a continent or worldwide level, suitable parallelization techniques have to be adopted, to speed up the approximate computation of Shapley values. For instance, for each application of matrix completion, the search for the optimal value of the regularization parameter can be speeded up by assigning a specific subset of its values to each core (to avoid bottleneck effects, the size of each such subset should be approximately the same for every core). The maximum theoretical speed-up achieved by proceeding in this way is at most equal to the number of cores, although in practice a smaller speed-up is obtained, due to Amdahl's law of parallel computing (Amdahl, 1967). Based on the idea outlined above, together with an initial sequential implementation of matrix completion, we also considered its parallel-processing implementation with 8 cores, achieving a speed up approximately equal to 7. Another research direction would consist in replacing Monte Carlo estimates with quasi-Monte Carlo ones, which have advantages over the former, such as better replicability and the availability of deterministic bounds on their performance, as discussed, e.g., by Bacigalupo et al. (2019) and by Gnecco et al. (2012), respectively. It is worth mentioning that recent research (Mitchell et al., 2022) has shown how quasi-Monte Carlo methods could be applied to estimate Shapley values.

A further research direction – still related to computational savings – consists in using suitable machine-learning techniques for the prediction of Shapley values, as outlined by Passacantando et al. (2021b) in a different context (in which Shapley values were used to address the importance of arcs in graphs). In this way, proxies of the Shapley values could be found, making it possible to extend the present analysis to larger datasets in future extensions. This goal could be achieved also by using weighted linear regression-based estimates of the Shapley values, like SHapley Additive exPlanations (SHAP) and especially the Kernel SHAP, introduced by Lundberg and Lee (2017). The latter is based on a Weighted Least Squares (WLS) estimate (obtained starting from a dataset of coalitions and of their associated utilities) of the vector of coefficients of a linear model, in which the coefficients are just the Shapley values, and a suitable kernel function is used to give different importance to distinct observations. Differently from the approach considered in this work, Kernel SHAP estimates all the Shapley values simultaneously, instead of individually. Hence, when Kernel SHAP is used, one may reasonably expect a smaller number of observations needed to achieve a desired precision of the estimate of each Shapley value, with respect to the case in which the

Monte Carlo approach is exploited. Nevertheless, an advantage of the alternative Monte Carlo approach is that, by construction, the summation of its estimated Shapley values equals the value of the grand coalition (which is a property of the Shapley values: see, e.g., Maschler et al. (2013)). Moreover, as we have verified preliminarily, the estimates of the Shapley values obtained by Kernel SHAP with a small number of observations are not guaranteed to be always non-negative, even in case in which the average marginal contributions of the players (averaged over the available permutations) turn out to be non-negative. Hence, for a small number of observations, the Monte Carlo approach may be still preferable.

One way to overcome the current limitation on the maximum number of players allowed in the analysis (for reasons of computational feasibility) would be to perform such an analysis in two or more steps at different levels of aggregation. For example, the analysis could first be performed at a higher level by having each player correspond to a cluster of countries. Then, at a lower level, one could consider a second game of transferable utility for each cluster, where the players are only the countries included in that cluster.

Appendix A: Additional results (2004 expansion of the European Union)

In this appendix, we repeat the main steps of the analysis reported in Section 3 by choosing as players the 15 countries belonging to the European Union (EU) before its expansion in 2004, and by considering two different years for the analysis, respectively before and after 2004. In order to increase the likelihood of seeing differences in the results of the respective analyses, the two years selected are 2002 and 2006 (and not, say, 2003 and 2005, which may appear to be too close to 2004). Further, in order to avoid computational issues related to the approximate evaluation of the Shapley values (which would arise in the case of a too-large set of players), the 10 countries that joined the EU in 2004 are not included in the set of players. It is worth remarking that, in this way, it is still possible to take into account the exogenous changes due to the 2004 expansion of the EU, by comparing the same set of 15 players before and after that event (instead of considering two different sets of players for the two years, which would make the results of the two analyses related to the years 2002 and 2006 less easily comparable).

The specific 15 countries selected as players are: Austria (AUT), Belgium (BEL), Germany (DEU), Denmark (DNK), Spain (ESP), Finland (FIN), France (FRA), United Kingdom (GBR), Greece (GRC), Ireland (IRL), Italy (ITA), Luxembourg (LUX), the Netherlands (NLD), Portugal (PRT), and Sweden (SWE).

For completeness, we remind the reader that the 10 countries that joined the EU in 2004, and that are not included as players in the present analysis due to the computational reasons reported above, are: Cyprus (CYP), Czech Republic (CZE), Estonia (EST), Hungary (HUN), Latvia (LVA), Lithuania (LTU), Malta (MLT), Poland (POL), Slovakia (SVK), and Slovenia (SVN).

Figures A.1 (a) and (b) show a colored visualization of the average ranking in 2002 and 2006, respectively, of each of the 15 countries involved in the analysis made in this appendix, based on the approximate Shapley values.

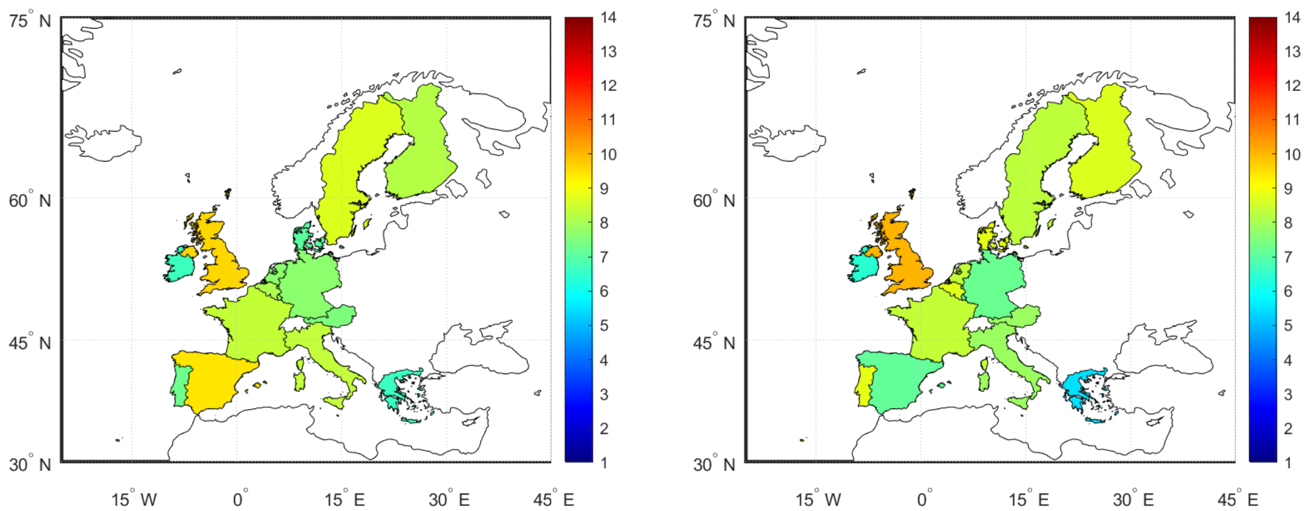

(a) Average rankings induced by the approximate Shapley values in 2002.

(b) Average rankings induced by the approximate Shapley values in 2006.

Figure A.1: Visual representations of the average rankings (from 1st to 14th) induced by the approximate Shapley values in 2002 and in 2006, respectively, for the 15 countries belonging to the European Union before its expansion in 2004.

Similarly, Figures A.2 (a) and (b) still refer to the case of the average ranking in 2002 and 2006, this time based on the cosine similarities.

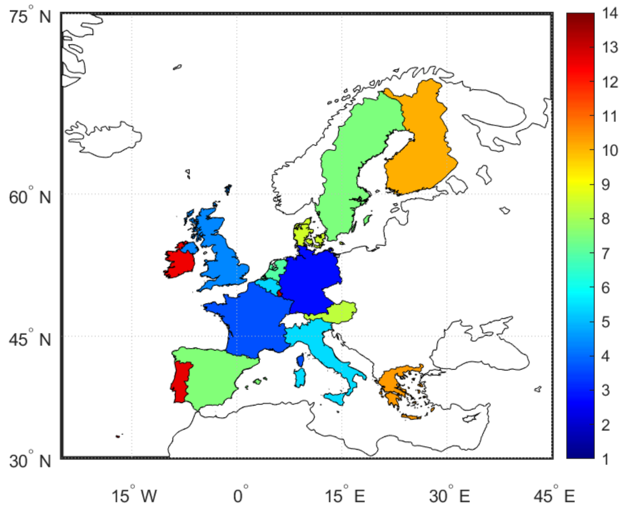

(a) Average rankings induced by the cosine similarities in 2002.

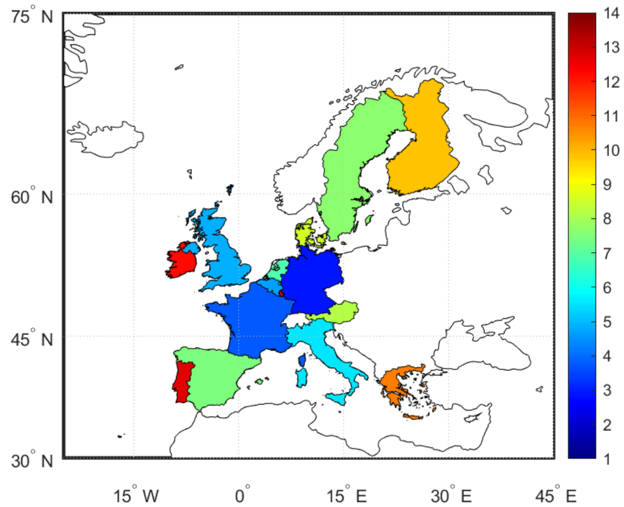

(b) Average rankings induced by the cosine similarities in 2006.

Figure A.2: Visual representations of the average rankings (from 1^{st} to 14^{th}) induced by the cosine similarities in 2002 and in 2006, respectively, for the 15 countries belonging to the European Union before its expansion in 2004.

Remarkably, a comparison of Figures A.1 and A.2 shows that, when moving from 2002 to 2006, the average rankings induced by the approximate Shapley value changed more than the average rankings induced by the cosine similarities, possibly suggesting a better capability of the approximate Shapley values in capturing changes (e.g., at the trade level) of the relationships among the 15 countries considered in the analysis, motivated by the successive entrance of the other 10 countries in the EU. In particular, when moving from 2002 to 2006, Germany and Spain experienced major improvements (i.e., lowerings) of their average rankings induced by the approximate Shapley values. This might be explained, respectively, by the following reasons:

– the geographical proximity of Germany with most of the newly-entered countries, and particularly, the presence of a common boundary of Germany with the Czech Republic, and one with Poland;

– the large immigration shock experienced from East-European countries to Spain after the 2004 EU expansion (de la Rica, 2009).

Acknowledgements

The work was partially supported by the “Contributi Liberali 2022” project “Application of Matrix Completion Techniques to the Definition of Economic and Financial Recommender Systems”, granted by Banca d’Italia. The authors wish to thank two anonymous Reviewers for the constructive feedback they provided on an earlier version of the manuscript.

References

Adigwe, E. O. (2022). A comparative analysis of competitive trade in a cluster market of the european union: the Revealed Comparative Advantage (RCA) index. *Naše Gospodarstvo/Our Economy 68*(1), pp.14–24.

Albora, G., L. Pietronero, A. Tacchella, and A. Zaccaria (2023). Product progression: a machine learning approach to forecasting industrial upgrading. *Scientific Reports 13*(1). Article no. 1481.

Amdahl, G. M. (1967). Validity of the single-processor approach to achieving large scale computing capabilities. In *Proceedings of the Conference of the American Federation of Information Processing Societies (AFIPS '67)*, Volume 30, pp. 483–485.

Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association 116*(536), 1716–1730.

Athey, S. and G. Imbens (2019). Machine learning methods that economists should know about. *Annual Review of Economics 11*, 685–725.

Aumann, R. J. J. (1994). Economic applications of the Shapley value. In J. F. Mertens and S. Sorin (Eds.), *Game-Theoretic Methods in General Equilibrium Analysis*, pp. 121–133.

Bacigalupo, A., M. Lepidi, G. Gnecco, F. Vadalà, and L. Gambarotta (2019). Optimal design of the band structure for beam lattice metamaterials. *Frontiers in Materials 6.* Article no. 2.

Belhumeur, P. N., J. P. Hespanha, and D. J. Kriegman (1997). Eigenfaces vs. Fisherfaces: recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence 19*, 711–720.

Bhatt, U., P. Ravikumar, and J. M. F. Moura (2019). Building human-machine trust via interpretability. pp. 9919–9920.

Bojnec, S. and I. Ferto (2016). Export competitiveness of the European Union in fruit and vegetable products in the global markets. *Agricultural Economics 62*(7), 299–310.

Castro, J., D. Gómez, and J. Tejada (2009). Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research 36*, 1726–1730.

de la Rica, S. (2009). The effect of the 2004 and 2007 EU enlargement on the Spanish labour market. IZA Discussion Papers, no. 4104.

Estmann, C., B. B. Sørensen, B. Ndulu, and J. Rand (2022). Merchandise export diversification strategy for Tanzania: promoting inclusive growth, economic complexity and structural change. *The World Economy 45*, 2649–2695.

Gnangnon, S. K. (2019). Aid for trade and export diversification in recipient-countries. *The World Economy 42*(2), 396–418.

Gnecco, G., Y. Hadas, and M. Sanguineti (2021). Public transport transfers assessment via transferable utility games and Shapley value approximation. *Transportmetrica A: Transport Science 17*, 540–565.

Gnecco, G., F. Nutarelli, and M. Riccaboni (2022). A machine learning approach to economic complexity based on matrix completion. *Scientific Reports 12.* Article no. 9639.

Gnecco, G., M. Sanguineti, and M. Gaggero (2012). Suboptimal solutions to team optimization problems with stochastic information structure. *SIAM Journal on Optimization 21*(1), 212–243.

Hadas, Y., G. Gnecco, and M. Sanguineti (2017). An approach to transportation network analysis via transferable utility games. *Transportation Research Part B – Methodological 105*, 120–143.

Hastie, T., R. Tibshirani, and M. Wainwright (2015). *Statistical learning with sparsity: the LASSO and generalizations.* CRC Press.

Hidalgo, C. (2021). Economic complexity theory and applications. *Nature Reviews Physics 3*, 92–113.

Hidalgo, C. A. and R. Hausmann (2009). The building blocks of economic complexity. *Proceedings of the National Academy of Sciences 106*(26), 10570–10575.

Kannen, P. (2020). Does foreign direct investment expand the capability set in the host economy? A sectoral analysis. *The World Economy 43*, 428–457.

Kuźnar, A. (2016). Poland's trade in services with Germany-EU membership experience. *Scientific Journal Warsaw University of Life Sciences 16*(4), 18–32.

Liao, H., X.-M. Huang, X.-T. Wu, M.-K. Liu, A. Vidmer, M.-Y. Zhou, and Y.-C. Zhang (2018). Enhancing countries' fitness with recommender systems on the international trade network. *Complexity 2018.* Article no. 5806827.

Lundberg, S. M. and S.-I. Lee (2017). A unified approach to interpreting model predictions. In *Proceedings of the Thirty-First International Conference on Neural Information Processing Systems (NIPS 2017)*, pp. pp. 768–4777.

Mann, I. and L. S. Shapley (1960). Values of large games, IV: evaluating the electoral college by Montecarlo techniques. The Rand Corporation.

Maschler, M., E. Solan, and S. Zamir (2013). *Game theory.* Cambridge University Press.

Mazumder, R., T. Hastie, and R. Tibshirani (2010). Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research 11*, 2287–2322.

Metulini, R., G. Gnecco, F. Biancalani, and M. Riccaboni (2022). Hierarchical clustering and matrix completion for the reconstruction of input-output tables. *ASta Advances in Statistical Analysis*.

Mitchell, R., J. Cooper, E. Frank, and G. Holmes (2022). Sampling permutations for Shapley value estimation. *Journal of Machine Learning Research 23*, 1–46.

Molnar, C. (2022). *Interpretable machine learning: a guide for making black box models explainable*. Independently published.

Morrison, G., S. V. Buldyrev, M. Imbruno, O. A. Doria Arrieta, A. Rungi, M. Riccaboni, and F. Pammolli (2017). On economic complexity and the fitness of nations. *Scientific Reports 7*(1). Article no. 15332.

Narayanam, R. and Y. Narahari (2010). A Shapley value-based approach to discover influential nodes in social networks. *IEEE Transactions on Automation Science and Engineering 8*(1), 130–147.

Passacantando, M., G. Gnecco, Y. Hadas, and M. Sanguineti (2021a). Braess' paradox: a cooperative game-theoretic point of view. *Networks 78*, 264–283.

Passacantando, M., G. Gnecco, Y. Hadas, and M. Sanguineti (2021b). On Braess' paradox and average quality of service in transportation network cooperative games. In *Special Volume of the AIRO Springer Series related to the International Conference Optimization and Decision Science (ODS) 2020*, pp. 27–37. Springer.

Sciarra, C., G. Chiarotti, G. Ridolfi, and F. Laio (2020). Reconciling contrasting views on economic complexity. *Nature Communications 11*. Article no. 3352.

Shapley, L. S. (1953). A value for n-person games. *Contributions to the Theory of Games 2*, 307–317.

Tacchella, A., A. Zaccaria, M. Miccheli, and L. Pietronero (2021). Relatedness in the era of machine learning. ArXiv preprint arXiv:2103.06017.

Štrumbelj, E. and I. Kononenko (2014). Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems 41*, 647–665.

Yeung, D. W. K., L. A. Petrosyan, and Y. Zhang (2021). Trade with technology spillover: a dynamic network game analysis. *International Game Theory Review 23*(1). Article no. 2050011.

Zhu, Z., G. Morrison, M. Puliga, A. Chessa, and A. Petersen (2018). The similarity of global value chains: A network-based measure. *Network Science 6*(4), 607–632.