Emerging Methodologies in Economics and Finance

Department of Economics
Democritus University of Thrace
Komotini, Greece

Emerging Methodologies in Economics and Finance

1. SUMMARY

THE TEAM. Our research team is operates within the Department of Economics, Democritus University of Thrace, Greece. Our research efforts are funded by a Research Grant from the European Union (Research Funding Program THALES) under the title "Study and Forecasting of Economic Data Using Machine Learning", MIS 380292.

MEMBERS. The research team is led by associate professors Periklis Gogas an economist (B.A., M.A., Ph.D.) and Theophilos Papadimitriou a mathematician (B.A.) and electrical engineer (M.Sc., Ph.D). Seven Ph.D. candidates and nine Master's students are actively working for the team.

RESEARCH INTERESTS. Our team's interests include both classic and emerging methodologies of Econometrics as they are applied to Economics and Finance. We currently work with: a) Machine Learning: Support Vector Machines for Classification and Regression and Deep Learning Architectures and b) Complex Networks: Threshold – Minimum Dominating Set, Weighted Dominating Set, and Multivariate Networks.

COMPLEX NETWORKS. Since 2012 we explore the field of Economic and Financial Networks using Graph Theory. Two PhD students and four graduate students work exclusively on this project. Our main goal in these research studies is to obtain a reduced version of the initial network which: a) contains an adequate amount of the total information for each goal (the amount of needed information depends on the application) and b) is more efficient easier to control and analyze. We refer to this as the *representation goal*. This task has been undertaken in the literature so far mainly through the MST approach in a variety of economic systems including stock markets (Bonanno et al, 2004; Tse et al, 2010), banking networks (Papadimitriou et al, 2013), macroeconomic networks (Hill, 1999; Gilmore et al, 2010; Dias, 2013) and others. In our work, we show that the Minimum Spanning Tree has inherent limitations when applied to economics and financial networks that may lead to inefficient and/or misleading results. These limitations stem from MST's algorithmic calculation process. To overcome these limitations we proposed a more efficient methodology based on a modified Minimum Dominating Set (MDS) algorithm, a methodology mainly used in the routing of computer and communications networks. The original MDS has its own disadvantages as in the optimization procedure the information included in the edges is not exploited. The new method we propose overcomes this problem and is called the Threshold-Minimum Dominating Set (T-MDS). The proposed T-MDS can: a) identify the smallest subset of nodes which b) represents the whole network more efficiently than the two other alternatives.

MACHINE LEARNING. Our initial attempt was to forecast financial time series using Machine Learning techniques and specifically Support Vector Machines for Classification and Regression. Support Vector Machines (SVM) is a binary supervised machine learning classifier. Proposed by

Cortes and Vapnik (1995) the basic notion behind the method is to find a linear separator (a hyperplane) of the two classes in the so called *feature space*. With the use of a kernel function the original *data space* is projected to a higher dimensional space called the feature space. Solving a convex minimization problem the algorithm converges to a linear separator that has the largest margin between the two classes. The margin between the classes is defined by a set of data points called *Support Vectors*. We find that the Supports Vector Machines methodology when used in macroeconomics and financial data provides superior out-of-sample forecasting accuracy as compared to classic econometrics techniques. Recently we started employing in these data Deep Learning Machines that is considered the current state-of-the-art.

2. Identity

Our research team operates within the Department of Economics, Democritus University of Thrace, Komotini, Greece. The members of the team are:

Researchers

Dr. Periklis Gogas, Associate Professor

He was born in Thessaloniki, Greece in 1969. He received his Ph.D. degree from the Department of Economics of the University of Calgary and his Master's degree from the University of Saskatchewan. His B.A. in Economics is from the University of Macedonia, Greece. Recently, a Visiting Scholar in Finance at the Ross School of Business of the University of Michigan. He also taught in the past at Plovdiv University, and the vocational center of the Athens Stock Exchange. His research interests include macroeconomics, financial economics, chaotic and non-linear dynamics graph theory and machine learning applied to macro and finance. He served in the past at the position of the Financial Director of a large Greek multinational corporation. He authored more than 40 articles in journals such as Journal of Money Credit and Banking, Journal of Banking and Finance, Economic Modeling, Journal of Forecasting, International Finance, Computational Economics, Open Economies Review, etc.

Dr. Theophilos Papadimitriou, Associate Professor

He was born in Thessaloniki, Greece, in 1972. He received the Diploma degree (B.Sc.) from the Department of Mathematics, Aristotle University of Thessaloniki, Greece, and the D.E.A. A.R.A.V.I.S (Automatique, Robotique, Algorithmique, Vision, Image, Signale) degree (M.Sc.) from the Department of Computer Science, University of Nice-Sophia Antipolis, France, in 1996 and the Ph.D. degree from the School of Engineering, Aristotle University of Thessaloniki, Greece in 2000. In 2001, he joined the Department of Economics of the Democritus University of Thrace in Komotini, Greece, where he served as a lecturer (2002-2008), assistant professor (2008-2013). Currently he holds the position of Associate Professor in the same department. Dr. Papadimitriou co-authored more than 80 journal papers, conference papers and book chapters combined. He served as a reviewer for various publications and as a member to scientific committees for Conferences and Workshops. Theophilos Papadimitriou current research interests include complex network, machine learning, and data analysis.

Ph.D. Students

- Vasileios Plakandaras
- Eftymia Chrysanthidou
- Efthimios Stathakis
- Georgios-Antonios Sarantitis
- Maria-Artemis Matthaiou
- Anna Agrapetidou
- Efi Parharidou

Master's Students

- Anna Amvrosiadou
- Pavlina Georgiadou
- Fotini Gkika
- Georgios Zikidis
- Dimitrios Karagkiozis
- Athina Kyprianidou
- Sophia Maltepioti
- Gerasimos Panagiotidis
- Konstantina Papathanasiou

Funding

The research team is funded by a research grant received by Prof. Papadimitriou from the European Union and financed through the European Social Fund (ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) Research Funding Program THALES under the title "Study and Forecasting of Economic Data Using Machine Learning" (MIS 380292).

Research Interests

Our team is interested in both classic and emerging methodologies of Econometrics as they are applied to Economics and Finance. Currently, we work with: a) Machine Learning (Support Vector Machines for Classification and Regression, Deep Learning Architectures) and b) Complex Networks (Threshold – Minimum Dominating Set, Weighted Dominating Set, and Multivariate Networks).

Selected Publications (click on a paper to open it)

1. P. Gogas, T. Papadimitriou, E. Chrysanthidou, "Yield Curve Point Triplets in Recession Forecasting", accepted for publication in *International Finance*.

2. V. Plakandaras, T. Papadimitriou, P. Gogas, "Forecasting daily and monthly exchange rates with machine learning techniques", accepted for publication in *Journal of Forecasting*.
3. V. Plakandaras, P. Gogas, R. Gupta, T. Papadimitriou, "US inflation dynamics on long range data", accepted for publication in *Applied Economics*.
4. I. Pragidis, P. Gogas, V. Plakandaras, T. Papadimitriou, "Fiscal shocks and asymmetric effects: A comparative analysis", accepted for publication in *Journal of Economic Asymmetries*.
5. V. Plakandaras, R. Gupta, P. Gogas, T. Papadimitriou, "Forecasting the US real house index", accepted for publication on *Economic Modelling.*
6. T. Papadimitriou, P. Gogas, G. Sarantitis, "Convergence of European Business Cycles: A Complex Networks Approach", *Computational Economics*, November 2014.
7. V. Plakandaras, T. Papadimitriou, P. Gogas, K.I. Diamantaras "Market Sentiment and Exchange Rate Directional Forecasting", accepted for publication on Algorithmic Finance.
8. P. Gogas, T. Papadimitriou, A. Agrapetidou, "Forecasting Bank Credit Ratings", *Journal of Risk Finance*, Vol. 15, No. 2, 2014, pp. 195-209.
9. T. Papadimitriou, P. Gogas, E. Stathakis, "Forecasting Energy Markets using Support Vector Machines", *Energy Economics*, Vol. 44, pp. 135-142, July 2014.
10. P. Gogas, T. Papadimitriou, M. Matthaiou, E. Chrysanthidou, "Yield Curve and Recession Forecasting in a Machine Learning Framework", *Computational Economics*, Vol. 45, no. 4, pp 635-645, April 2015.
11. P. Gogas, V. Plakandaras, T. Papadimitriou, "Public Debt and Private Consumption in OECD countries", *Journal of Economic Asymmetries*, Vol 11, no. 1, June, 2014, pp. 1-7.
12. P. Gogas, T. Papadimitriou, G. Sarantitis, "Testing Purchasing Power Parity theory in a DFA Framework and Rolling Hurst Exponent: the case of 23 OECD Countries", *Applied Financial Economics*, Vol. 23, no 17, pp. 1399-1406, 2013.
13. T. Papadimitriou, P. Gogas, B.M. Tabak, "Complex Networks and Banking Systems Supervision", *Physica A*, Vol. 392, pp. 4429-4434, Oct. 2013.
14. P. Gogas, T. Papadimitriou and E. Takli, (2013) ''Comparison of simple sum and Divisia monetary aggregates in GDP forecasting: a support vector machines approach'', *Economics Bulletin*, Vol. 33, No. 2, pp. 1101-1115, 2013.

3. Methodologies Employed

3.1 Complex Networks

In 2012 we started exploring the field of Economic and Financial Networks through the prism of Graph Theory. We are still highly committed in the field with two PhD students and four graduate students working exclusively on this project. Our main goal in these research studies is to obtain a reduced version of the initial network which: a) contains an adequate amount of the total information for each goal (the amount of needed information depends on the application) and b) is more efficient easier to control and analyze. We refer to this as the *representation goal*. This task has been undertaken in the literature mainly through the MST approach in a variety of economic systems including stock markets (Bonanno et al, 2004; Tse et al, 2010), banking networks (Papadimitriou et al, 2013), macroeconomic networks (Hill, 1999; Gilmore et al, 2010; Dias, 2013) and others. Nonetheless, the Minimum Spanning Tree has inherent limitations that stem from its algorithmic calculation process. We investigated these limitations and proposed a more efficient methodology based on a modified Minimum Dominating Set (MDS) algorithm, a methodology mainly used in the routing of computer and communications networks. The new method is called Threshold-Minimum Dominating Set (T-MDS). By comparing these methodologies (MST, MDS and T-MDS) we can allege that the T-MDS can: a) identify the smallest subset of nodes which b) represents the whole network more efficiently than the two other alternatives.

Consider the network depicted in Figure 1. In this network, the values of the edges correspond to distances and thus nodes with more similar behavior are connected by a low distance edge and vice-versa. The similarity can be calculated using correlation based metrics, concordance based metrics or even a z-score multivariate distance.

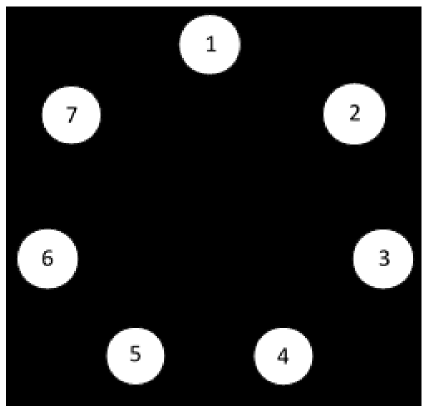

Figure 1. The nodes, edges and distances of the initial network

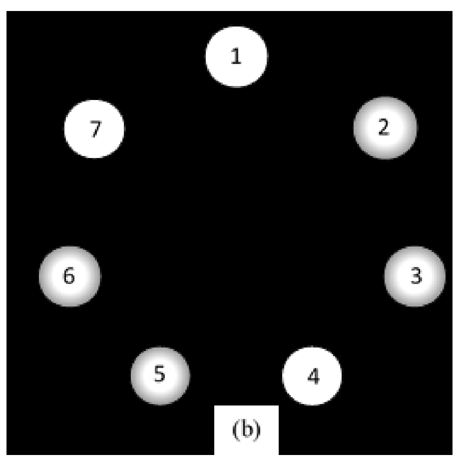

Figure 2. The core nodes of the MST based method.

The MST-based approach identifies a set of four core nodes, namely nodes 2, 3, 5 and 6, to describe the behavior of the entire network. They can be used as gauges for the behavior of their neighborhood. Whenever a shift in the behavior of a core node occurs, it reflects a change in the behavior of its neighborhood. Consequently, the state of each neighborhood (and eventually of the entire network) can be represented by the core nodes, which is the basic concept of the methodology in Papadimitriou *et al*. (2013).

However, by closely examining the topological features of the core nodes, we observe that:

- The MST procedure eliminates the edge e_{25} to avoid the creation of a loop in the path, although the corresponding distance d_{25} is one of the shortest in the network. This means that even though a) the behavior of nodes 2 and 5 is very similar and b) 2 is a core node, the behavior of node 5 is not represented by core node 2. In fact, this limitation of the MST algorithm leads to the identification of node 5 as a core node as well.
- On the other hand, the edge e_{34} is included in the MST, although the corresponding distance d_{34} is the maximum in the network. This means that core node 3 describes the behavior of node 4, even though the two nodes are very dissimilar.

Hence, the MST-based approach "misses" some essential relations (e.g. the e_{25} edge), and at the same time maintains some edges that connect nodes with low similarity in the process of identifying the core nodes (e.g. the e_{34} edge). These are important drawbacks that could affect the efficiency of the MST-based solution to the representation goal.

Next, we apply the simple MDS methodology in the theoretical seven nodes' network of Figure 1.

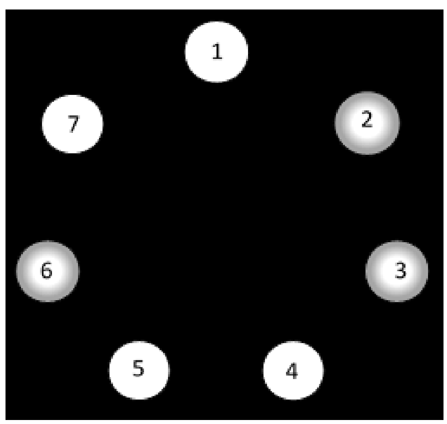

Figure 3. Minimum Dominating Set (gray nodes)

The absence of the no-loop restriction, in the MDS algorithm, increases the number of the considered edges in the network, which in turn is manifested through larger neighborhoods (since more nodes remain interconnected). As the network becomes denser, fewer MDS nodes are needed to reach the representation goal. Consequently, it is more probable that the cardinality of the MDS nodes' set will be smaller than the corresponding MST one. Indeed, by applying the MDS technique, we are able to represent the behavior of the whole network using just three nodes (one less than the MST methodology). This fact is noteworthy but not critical to our cause. We do search for a small subset of nodes, but most importantly we search for a set that best represents the entire network. The MDS nodes' set, as opposed to the MST-based one, is a better representation of the network: all the small distance edges are included in the representative nodes' neighborhoods. This implies that all node pairs with similar behavior are included in the neighborhoods. On the other hand though, the same is true for all the large distance edges (see in Figure 3 the edges e_{34}, e_{35}, e_{46} and e_{67}). These edges connect nodes with dissimilar behavior and their inclusion in the MDS neighborhoods may jeopardize the representation goal. Node 4 in the example belongs to the neighborhoods of nodes 3 and 6. In both cases, the edges linking node 4 with the two representative nodes describe long distances, meaning that the node in question does not have similar behavior with the nodes that represent it. Consequently, every decision that we make for node 4 based on the behavior of nodes 3 and 6 is potentially unreliable.

Threshold-Minimum Dominating Set (T-MDS)

We suggest that the limitations of the MST-based approach and the classical MDS technique can be effectively addressed with the inclusion of one extra step prior to the identification of the MDS: the imposition of a threshold on the edges' distances in the initial network. By doing so, we ensure that only the edges with low distance (i.e. the edges which connect nodes with high similarity) survive, while all the other edges are eliminated.

Definition: *We call a* ***Threshold – Minimum Dominating Set*** *the two step methodology for the identification of the most representative nodes of a network, defined as:*

Step 1. *A thresholding on the edges' distances leading to the elimination of all edges that correspond to large distances.*

Step 2. *The identification of the MDS nodes on the remaining network.*

The suggested process is depicted in Figure 4. In the left part of the Figure, in Step 1 we impose a threshold of $p = 4$ to remove any edges (depicted with the dashed lines) with distance greater than 4. Thus, we eliminate the edges that connect nodes with dissimilar behavior; this results in the isolation of nodes 3 and 6 from the rest of the network. Then, the MDS is identified on the right part of Figure 4..

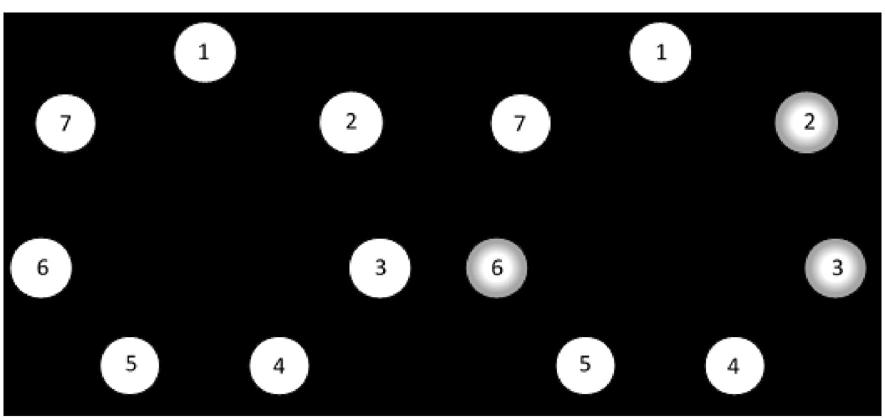

Figure 4. Threshold-Minimum Dominating Set (colored nodes)

Apparently, the Minimum Dominating Set consists of only three nodes: one dominant node (node 2) and two isolated ones (nodes 3 and 6). The dominant node of the network describes the behavior of the interconnected nodes 1, 4, 5, and 7. The edges connecting each of these nodes with the dominant one represent a short distance and thus this node can successfully act as a representative agent for the neighborhood.

The Threshold-Minimum Dominating Set (T-MDS) retains the advantages of the Minimum Spanning Tree and the classical Minimum Dominating Set while overcoming their inherent limitations: it is a compact and reliable subset able to describe efficiently the topology of the entire network. Among the three presented methodologies, the T-MDS based approach yields the best solution to the representation goal. The described methodology was successfully applied using a simple computer to real world networks that include more than 4000 nodes.

Application Range

Banking networks

- "Complex Networks and Banking System Supervision", *Physica A*, 392, pp. 4429-4434.
- "Bank Supervision using the Threshold-Minimum Dominating Set", *Physica A*, under review.

We used the T-MDS methodology to create an auxiliary supervision/monitoring mechanism that can be used by the banking system supervising authority, usually the central bank. This system is both efficient with respect to the required resources needed and can also promptly identify a set of banks that are in distress so that immediate and appropriate action can be taken by the supervising authority. The T-MDS is used to identify the smallest and most efficient subset of banks that can be used as a) sensors of distress of a manifesting banking crisis and b) provide a path of possible contagion. We propose the use of this method as a supplementary monitoring tool in the arsenal of a Central Bank. One dataset we used includes the 122 largest American banks in terms of their interbank loans. At threshold level of t = 0.8 our method identifies 47 T-MDS nodes, with 21 of them being isolated ones. According to this, the interconnected nodes in the thresholded network are 101, and they can be monitored by the 26 dominant nodes.

U.S. Gross State Product Synchronization

In this study we created the U.S. states' network, based on their Gross State Product (GSP). We use it to explore the inter-relations of the GSP growth rates and the degree of their business cycle synchronization. Moreover, this approach allowed us to identify a subset of states that can efficiently describe the collective behavior of the entire network. The policy implications are obvious: the monitoring of these key states from the federal fiscal and monetary authorities may provide an alternative and efficient tool in the designing and implementation of economic policy. To do so, we used four alternative measures of similarity: correlation, weighted correlation, the sign concordance index (SCI), and the weighted sign concordance index. Both types of measures (correlations and SCI) have indicated an increased similarity between the U.S. states GSP in the more recent years, evidenced both by the T-MDS methodological analysis and the simple network metrics. In both cases the weighted versions (where the more recent GSP values received a higher weight) corresponded to denser networks in which the overall correlations are higher. Thus, this can be viewed as evidence of an increased business cycles synchronization

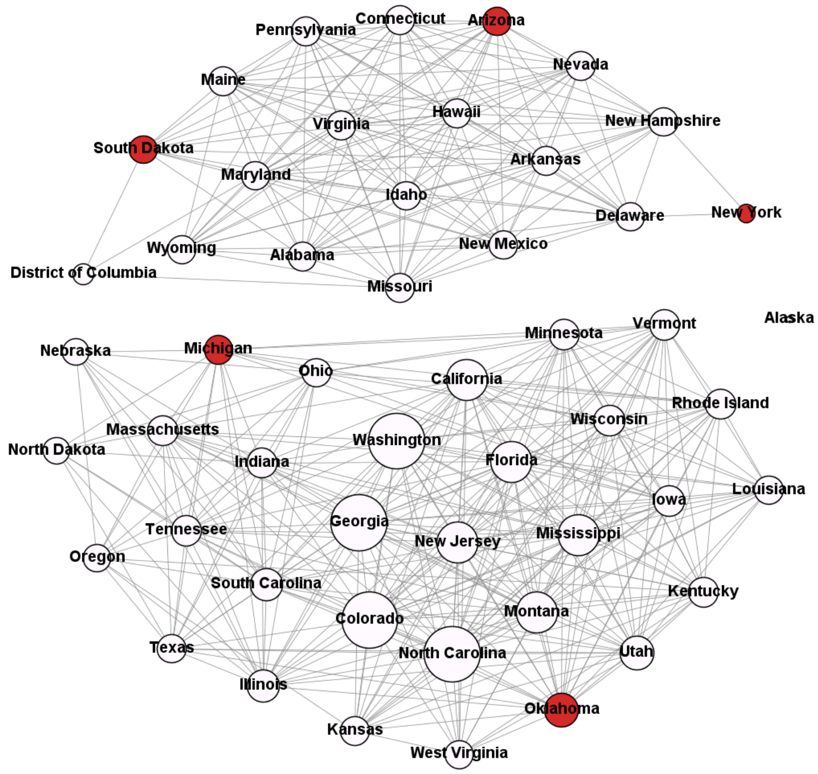

Figure 5. Network topology in the case of a weighted Sign Concordance Index and threshold level t=0.75 (nodes sized according to their node degree)

Another interesting result is identification and examination of crisis dispersion paths between the U.S. states. The imposition of a threshold on the network's edges removes low strength edges and allows only for the existence of the highly important and informative ones. This procedure creates a map of GSP inter-relations that could be used as a guide for contagion in the case of a crisis manifestation in one or more U.S. states. By studying this map, policy authorities could intervene timely and effectively to contain the crisis and/or curtail its effects and further contagion.

The U.K. Consumer Price Index

During the summer of 2014, our PhD student Georgios Sarantitis was employed as a research associate at the Bank of England. His job description was exclusively to apply the T-MDS methodology on the calculation of the United Kingdom's Consumer Price Index (CPI). He modeled the United Kingdom's Consumer Price Index as a complex network and he applied clustering and optimization techniques to study the network's evolution through time. By doing this, he provided a dynamic, multi-level analysis of the mechanism that drives inflation in the U.K. The analysis that we performed on our team revealed that the CPI-classes' network exhibits an evolving topology through time which depends substantially on the prevailing economic conditions in the U.K. We identified non-overlapping communities of these U.K. CPI classes and we observed that they do not correspond to the actual categories they officially belong to; a finding suggesting that diverse forces are driving the inter-relations of the CPI classes which are stronger for classes in different categories rather than for the classes within them. Finally, we created and presented a reduced version of the U.K. CPI that fulfils all the core inflation measure criteria and can possibly be used as an appropriate measure of the underlying inflation in the U.K. Since this new measure uses only 14 out of the 85 U.K. CPI classes, it can be employed to complement the Bank of England's arsenal of core inflation measures without the need for further resource allocation.

Income Inequality

- "Income Inequality: A State-by-State Complex Network Analysis", *Social Indicators Research*, under review.

We investigated the developing patterns of inequality in the U.S. using complex network analysis and the Threshold-Minimum Dominating Set (T-MDS). We used two alternative measures of income inequality, the top 1% share of income and the Gini coefficient. We performed a dynamic analysis over four consecutive periods running from 1916 to 2012.

Table 1. T-MDS metrics for the top 1% inequality measure

	1916-1929	1930-1944	1945-1979	1980-2012
T-MDS cardinality	22	28	15	3
Isolated nodes	14	22	10	0
Dominant nodes	8	6	5	3

Our findings reveal a heterogeneous pattern of income inequality and economic integration of the U.S. states according to each focal period. Furthermore, the empirical findings differentiate slightly with respect to the two inequality measures employed.

Business Cycle Synchronization

"International Business Cycle Synchronization since the 1870s: Evidence from a Novel Network Approach", *Physica A*, Under review.

We examined the issue of business cycle synchronization from a historical perspective in 27 developed and developing countries from the 1870s to the 2010s. Based on the T-MDS, our results reveal heterogeneous patterns of international business cycle synchronization during fundamental globalization periods since the 1870s. In particular, the proposed methodology reveals that worldwide business cycles decoupled during the Gold Standard, though they were synchronized during the Great Depression. The Bretton Woods era was associated with a lower degree of synchronization as compared to that during the Great Depression, while worldwide business cycle synchronization increased to unprecedented levels during the latest period of floating exchange rates and the Great Recession. This provides empirical evidence in support of lower business cycle synchronization within periods of fixed exchange rates and lack of an independent monetary policy.

3.2 Machine Learning

The spark, in 2010, that led to the creation of our team one year later, was the common interest of Dr. Gogas (economist) and Dr. Papadimitriou (mathematician/electrical engineer) to fuse methodologies from other fields to pure economic applications. Our initial attempt was to forecast financial time series using Machine Learning techniques and specifically Support Vector Machines for Classification and Regression.

Support Vector Machines (SVM) is a binary supervised machine learning classifier. Proposed by Cortes and Vapnik (1995) the basic notion behind the method is to find a linear separator (a hyperplane) of the two classes in the so called *feature space*. With the use of a kernel function the original *data space* is projected to a higher dimensional space called the feature space. Solving a convex minimization problem the algorithm converges to a linear separator that has the largest margin between the two classes. The margin between the classes is defined by a set of data points called *Support Vectors*. An example of an SVM classification employing the RBF kernel is depicted in Figure 2.

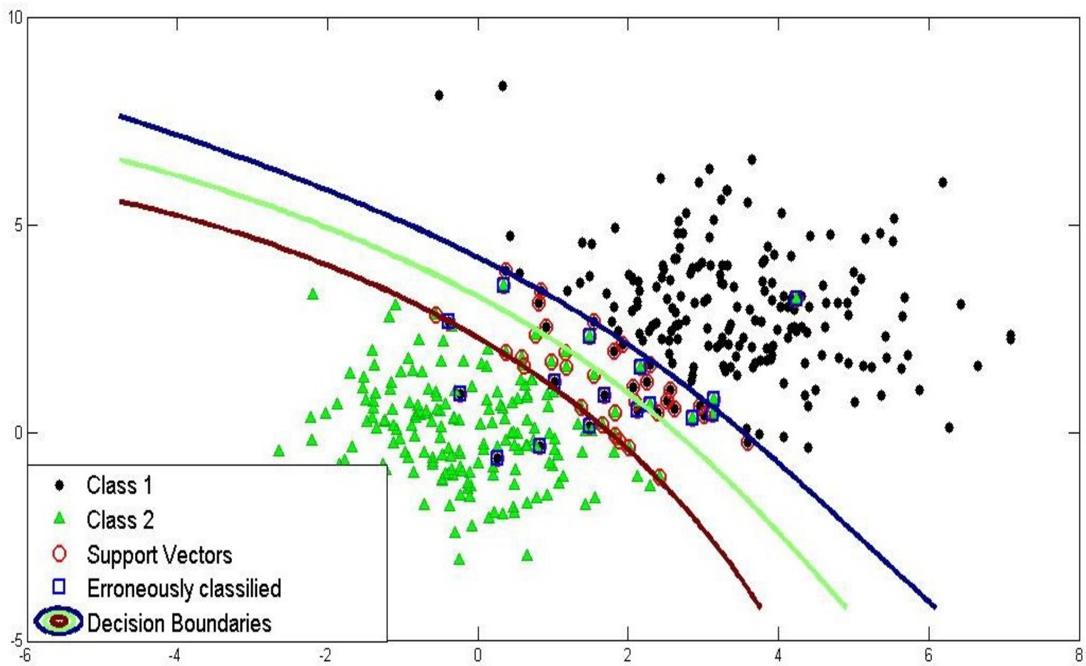

Figure 6: An example of an SVM classification using the RBF kernel. The two classes are separated with a linear separator at a higher dimensional space called the feature space. The separator in the feature space when it is re-projected back into the original dimensions becomes a non-linear function of the initial data space. The encircled observations are the Support Vectors defining the decision boundary and the observations with a square rounding are misclassified observations.

Application Range

Forecasting Bank Failures

We explored the problem of forecasting bank failures using the SVM methodology with great success. We used a dataset that includes 1443 U.S. banks and consists of 962 solvent banks and 481 banks that failed during the period 2003-2013. We start with an initial set of 144 variables that come from the publicly available financial statements of the banks. Next, we employ the variable selection methodology called *Local Learning* and we select only 5 variables out of the original 144. These 5 variables[1] are then used as the explanatory variables to train the optimum SVM model. This produced an out-of-sample forecasting accuracy of 97.67% in a one-period-ahead (one year) forecasting window. With a forecasting window of two and three years ahead the model's accuracy fell to 83.48% and 64.35% respectively. The out-of-sample forecasting accuracy for the failed banks was 97.39% and for the solvent banks 97.81%. These results are summarized in Table 2.

Table 2. The foresting accuracy on each class

	True Solvent	True Insolvent
Estimated Solvent	223	3
Estimated Insolvent	5	112

As a follow up, we investigated the three missed insolvent cases: a) In the case of The First National bank of Davis, Davis, Oklahoma, during the 2011 examination, significant unrecognized losses were discovered that exceeded the bank's capital and allowance for loan and lease losses, b) The First State Bank received an enforcement action on 21st of December 2010 and in a very short time, on the 28th of January 2011 it was closed and c) The Community's Bank was established in 2001 and in 2013 it failed. The same investigation for the five misclassified solvent case revealed that the financial institutions received an enforcement action from FDIC or financial help from the U.S. Treasury.

Stress Testing Tool

Moreover, we propose the use of the previously defined hyperplane separating the solvent from the insolvent banks as a new form of **bank stress testing**. Solvent banks that are close to the separating hyperplane are more vulnerable to class changing (become insolvent). Therefore, a solvent bank's distance from the separating hyperplane may be used as a measure of its robustness. By doing this, multi-level stress tests can be performed: a) system-wide, where alternative scenarios on various macroeconomic and/or financial indicators are reflected on the five explanatory variables of the banks' financial statements that are used for the classification and b) institute specific based on their financial condition. Moreover, for a bank that that is

[1] 1. Tier 1 (core) risk-based capital/total assets (t-1), 2. Provision for loan and lease losses/total interest income (t-1), 3. Loan loss allowance/total assets (t-1),4. Total interest expense/total interest income (t-1), 5. Equity capital to assets (t-1).

forecasted to fail within the next year, the distance from the hyperplane may be exploited and render it solvent. This distance can be used **in time** in order to prescribe a set of measures that will render the bank solvent.

Our proposition is also equipped with a sensitivity analysis of each bank in distress to the five variables that define its position in the data space.

Forecasting in Energy Markets

"Forecasting Energy Markets using Support Vector Machines *", Energy Economics*, 44, pp. 135-142, July 2014.

The efficiency of an SVM-based forecasting model for the next-day directional change of electricity prices that are very volatile was also investigated. We first fitted the best autoregressive SVM model and then augmented it with various relevant variables. We tested this model on the daily Phelix index of the German and Austrian control area of the European Energy Exchange (EEX) wholesale electricity market. The out-of-sample forecasting accuracy we achieved is 76.12% over a 200 day period.

We further explored the efficiency of the SVM methodology in the electricity market by forecasting "price spikes”. In electricity markets, the comparatively high upward or downward movements of prices within a short period are called spikes. In effect, when we are trying to forecast price spikes in electricity, we perform outlier detection. The problem was attacked using Multiclass SVM and Deep Learning architectures. The forecasting accuracy of the optimal model is depicted in Table 3.

Table 3. Spikes forecasting

	SVM	Deep Learning
Normal Cases	76%	100%
Positive Spikes	70%	85%
Negative Spikes	56%	98%

Forecasting Exchange Rates

"Forecasting Daily and Monthly Exchange Rates with Machine Learning Techniques*", Journal of Forecasting*, forthcoming.

We tested and compared the forecasting ability of various Machine Learning, Neural Networks, and Econometrics architectures on monthly and daily spot prices of five selected exchange rates: EUR/USD, JPY/USD, NOK/AUD, NZD/BRL and PHP/ZAR. In doing so, we combine a novel smoothing technique (initially applied in signal processing) with a variable selection methodology and the regressors are used in the regression estimation methodologies. After the decomposition of the original exchange rate series using the Ensemble Empirical Mode Decomposition (EEMD) method into a smoothed and a fluctuating component, Multivariate Adaptive Regression Splines (MARS) are used to select the most appropriate variable set from a very large set of explanatory variables that we collected. The selected variables are then fed into the forecasting models that produce one-period-ahead forecasts for the two components: smoothed and fluctuating.

We implemented two versions of this hybrid methodological setup; an autoregressive and a structural one. The autoregressive model consistently outperforms all alternative models in forecasting out-of-sample in monthly frequency. Using daily data, the structural EEMD-MARS-SVR model is superior for three out of the five exchange rates in our sample. A close alternative is the structural EEMD-MARS-NN (Neural Networks) model that outperforms in forecasting the other two exchange rates. Nevertheless, only the structural EEMD-MARS-SVR methodology consistently outperforms the benchmark random walk model in out-of-sample forecasting. The above findings corroborate with the microstructural aspect of the exchange rate market on the short run and the importance of macroeconomic dynamics (such as PPP, UIP, etc.) on the long run, validating exchange rate theory. Overall, our models seem to capture the different data generating processes between short and long horizons. The weak form of exchange rate market efficiency is rejected for both sampling frequencies due to the fact that the autoregressive model outperforms the RW benchmark.

Yield Curve and Recession Forecasting

"Yield Curve Point Triplets and Recession Forecasting*", International Finance*, forthcoming.

Several studies have highlighted the yield curve's ability to forecast economic activity. These studies use the information provided by the slope of the yield curve—i.e., pairs of short- and long-term interest rates. We constructed three models to forecast the positive and negative deviations of real U.S. GDP from its long-run trend over the period from 1976Q3 to 2011Q4: one that uses only pairs of interest rates and two that draw on more than two points from the yield curve. We employ two alternative forecasting methodologies: the probit model, which is commonly used in this line of literature, and the support vector machines (SVM) approach from the area of machine learning. Our empirical results show that the SVM model with the RBF kernel and three interest rates as input variables achieves a 100% out-of-sample forecasting accuracy for recessions (unemployment gaps) and the best overall accuracy is 80%. Thus, it appears that correct identification of upcoming unemployment gaps may be achieved at the cost of reduced accuracy of forecasting inflationary gaps or in other words at the cost of some extra inflation. The interest rates that produce these results are the 3-month T-bill rate and the 2- and 3-year government bond rates. Long-term rates do not appear to complement our models' forecasting ability. Our interpretation of this finding is that, since the U.S. is a developed industrialized country, the United States is considered to have a stable long-term economic outlook that is not affected by short-term dynamics and fluctuations and instead adheres to its long-run potential output. Thus, agents' views of future economic activity are not significantly affected by short-term events or fluctuations, thus rendering the long-term rates uninformative in terms of short-term recession forecasting.

4. References

G. Bonanno, G. Caldarelli, F. Lillo, S. Micciche, N. Vandewalle, and R. N. Mantegna (2004), "*Networks of equities in financial markets*", European Physical Journal B, Vol. 38, pp. 363-371.

J. Dias, (2013). Spanning trees and the eurozone crisis. Physica A: Statistical Mechanics and its Applications, 392, 5974-5984.

C.G. Gilmore, Lucey, B. M. & Boscia, M. W. (2010). Comovements in government bond markets: A minimum spanning tree analysis. Physica A: Statistical Mechanics and its Applications, 389(21), 4875-4886.

R.J. Hill, (1999). Comparing price levels across countries using minimum-spanning trees. Review of Economics and Statistics, 81(1), 135-142.

T. Papadimitriou, P. Gogas, B.M. Tabak (2013), *"Complex Networks and Banking Systems Supervision"*, *Physica A: Statistical Mechanics and its Applications,* Vol. 392, no. 19, pp. 4429 – 4434

C.K. Tse, Liu, J. & Lau, F. (2010). A network perspective of the stock market. Journal of Empirical Finance, vol. 17(4), pp. 659-667.