

Understanding Systemic Risk in Global Sovereign Credit Markets: A Network Topology Approach to Time-Series Analysis with Financial Applications

Aakash Pattabi

Dept. of Economics
{apattabi@stanford.edu}

Eric Gilliam

Dept. of Political Science
{egillia3@stanford.edu}

Ashwin Sreenivas

Dept. of Computer Science
{ashwinsr@stanford.edu}

December 10, 2018

Github: <https://github.com/ashwinsr/cs224w>

I. Introduction

Sovereign default – the acknowledged failure of a national government to pay back an owed debt – has accrued significance in recent years as debt crises wracked Europe and South America. Current macroeconomic risk profiling techniques leave much to be desired, and a holistic understanding of systemic risk in the global macroeconomy is notably absent from most policy discussions. Contemporary approaches to modelling financial risk compute various pairwise covariance measurements which in isolation fails to capture the “domino” effect of financial crisis spillover observed in the wake of the 2008 global recession. Thus, in this paper, we use network techniques to better profile risk and interconnectedness in the global macroeconomy. Network modelling is particularly important in this context as nations that owe each other debts bet on the likelihood that these debts will be repaid in Credit Default Swap (CDS) markets.

We construct and analyze a weighted, directed influence network using time series data on credit default swap prices for the sovereign debts of the world’s major nations. We find that during the observation period, the countries least exposed to spillover risk from other nations are either financial bastions (such as the United States), or those that suffer from high endogenous default risks independently of global trends. This empirically demonstrates that the majority of price movements and volatility in credit markets for nations with unstable economies are mostly due to changes in risk related to the economy of the country as a whole rather than risk spillover from other countries. Financial bastions, especially the US given the “exorbitant privilege” inherent in being the world’s dominant currency, seem less exposed to risk spillovers, but not entirely safe. Furthermore, our results seem to indicate that many “emerging market” countries whose economies are not at all exposed to one another effect one another and cause debt spillovers. This empirically demonstrates an effect of credit market bettors (such as banks) using “rules of thumb” to monitor risk in their portfolios, as we see largely independent countries without much obvious systemic connectedness being highly connected in the volatility network. This may be due to financial managers maintaining these investments within common “Emerging Markets Portfolios,” treating countries as similar for reasons beyond shared economic fundamentals. This is consistent with our expectation that bettors will de-risk from entire portfolios if they perceive one investment in that portfolio class to perform poorly, as such behavior could induce

financial risk connectedness where it ought not obviously exist.

II. Related Literature

We build on a small body of pre-existing work on financial networks. The sparsity of the literature is primarily due to the difficulty of private groups in obtaining direct transaction data of financial actors within national or global economic networks. In fact, there is no single reporting body which provides data on the direct relationships between actors in financial networks. [3] constructs a method to approximate the desired financial relationships given the lack of direct, explicit relationships in public data. The authors construct A , a 13 x 13 weighted adjacency matrix of interconnectedness between the largest 13 financial institutions in the United States, with element A_{ij} in this matrix describing the effect that a shock to bank i will have on bank j 's share price volatility. This was done using the forecast error variance decomposition method for time series analysis detailed below.

In a follow up paper in 2017, [2] expands upon [3] by introducing LASSO regularization to the network construction process. By using the LASSO to regularize the adjacency matrix, the authors impose a structural assumption of independence between many entities in the network to make the problem of estimating the edge weights from data numerically tractable. This method allows larger financial networks to be built and analyzed with fewer dimensionality and data issues. As [2] works with private actor data, they face more immediate dimensionality issues that demand regularization, compared to previous work with data on public bonds or smaller networks.

In terms of measuring risk in financial networks, the authors of [1] recognized that interconnected financial networks often did not have a principled approach to identifying systemically important nodes. To identify key financial institutions that were exposed to the most risk, they analyzed a network of US FED emergency loans: loans that large financial institutions made to each other. They constructed a network in which nodes were large financial institutions and edges were the amount of outstanding debt between them. The authors then applied a modified version of the PageRank algorithm on the resultant graph to identify nodes that were recursively vulnerable (i.e. held large amounts of debt from institutions that themselves held large amounts of debt). This allowed them to better capture the recursive domino effect of debt defaults in a more comprehensive way, than the traditional method of simple threshold values. As such, [1] motivates our own use of the PageRank algorithm in this research.

While [3] and [2] provide us with foundational work on converting interdependent financial time-series data into a well-modelled network, there remains significant work to be done in exploring the structure and properties of the resultant networks. In particular, [3] “eyeballs” the most prominent nodes in the network without a clear analytical technique; notes that they see some clustering behavior; and calculates basic eigenvector centrality coefficients for the nodes. We carry out a more principled approach. By integrating more sophisticated algorithms to analyze networks - such as PageRank and community detection algorithms - we can make more powerful claims about the resulting graph.

III. Network Construction

Data

We acquire data from Bloomberg Professional Services’s data subscription service. We are primarily concerned with the daily reported market price for the Credit Default Swaps for all the countries for which data is available. We acquired daily price data for the 65 countries tracked in Bloomberg’s index, including major developed economies such as the United States, the United Kingdom, and Germany; intermediately sized economies such as Turkey and France; and developing economies such as Malaysia. Our data span the world – from eastern Europe to southeast Asia. The records span the trading days

between January 1, 2006 and January 1, 2018. In this window, we capture major sovereign credit events including the 2008 financial crisis; the Greek debt crisis and its ensuing spillover effects into Eurozone countries; and the ongoing sovereign credit crisis in Venezuela.

One challenge in working with financial data – especially financial data that is available for free or through an academic subscription service – is data incompleteness. We found that of the 65 nations for which data were available, many nations’ time series were incomplete, leading to an imbalanced panel of data un-conducive to network construction our algorithm (detailed below). Furthermore, it is unfortunately impossible to differentiate between panel imbalance due to poor data collection as opposed to swaps for some countries simply not being traded over the observation period. Our analysis must assume the latter is true for any missing data; if the former were true, traditional data recovery techniques like mean imputation would skew the volatility measures from which we construct risk [4].

We applied three corrections to the data to enable our application of vector autoregression. First, we interpolated missing daily observations under a hypothesis of “observations are missing due to sparse trading.” [4]. Second, for each time series, we calculated the intra-week standard deviation in swap price for each week between the data collection start and end dates, a common, heuristic for asset volatility in financial time series analysis. This yielded a panel of 65 countries by 626 observations. Finally, we retained only the 30 countries with the most available data in the original data set to make parameter estimation for our autoregression tractable.

Vector Autoregressive Process

The foundational method we use to construct a financial dependency network is the vector autoregression (VAR), a multivariate time-series method that works particularly well in identifying relationships between endogenous regressors. VAR models are multivariate generalizations of the univariate autoregressive moving average (ARMA) process. A univariate $ARMA(p)$ model is of the form:

$$y_t = \theta_0 + \theta_1 y_{t-1} + \theta_2 y_{t-2} \dots \theta_p y_{t-p} + \epsilon_t = \theta_0 + \sum_{i=1}^p \theta_i y_{t-i} + \epsilon_t$$

A multivariate model $VAR(p)$ extends the univariate model by parametrizing the response y_t in terms of the p lagged observations for the remaining variables in the model. The error term in each linear equation in a multivariate VAR model ϵ_t is sometimes referred to as an “innovation” and captures the *forecast error* in the target contingent upon observing its past states as well as the past states of other variables [10]. The model parameter p is the backwards-looking horizon, which encodes the econometrician’s assumption about the duration of the dependence between y and all other endogenous variables z .

Econometricians often use structural methods to model the innovations, occasionally including exogenous variables in the model at the cost of requiring the imposition of identifying restrictions to enable each coefficient to be estimated. We do not need to apply such restrictions, as we are only concerned with the forecast error at each time period – our problem in constructing the network is a prediction problem with the implicit assumption that the sovereign debt and credit system is a closed system (else, it could not be modelled as a network).

Finally, we consider explicitly the forecast error at time t over horizon h . This error is due to compounding in the innovations across time horizons between t and $t + h$. The forecast error variance is simply the forecast error at the horizon squared. Consequently, a *forecast error variance decomposition matrix* (FEVD) is a matrix that quantifies how important each innovation in each covariate is in explaining the forecast error out h periods from t . This can be computed recursively.

Namely, for a time series of N variables, the forecast error variance decomposition outputs a $N \times N$ matrix V where V_{ij} is the fraction of forecast error variance of variable i due to innovations in covariate j at time horizon h . We construct our graphs where edge $i \rightarrow j$ is a weighted directional edge of either V_{ij} or V_{ji} depending on the application. We use a generalized $VAR(p)$ model proposed by [6] that is insensitive to the ordering of variables in the time series. As such, we construct our network using one of either a row-stochastic or column-stochastic matrix depending on the application. In the former case, A_{ij} of the weighted, fully connected adjacency matrix A denotes the percentage of volatility in i 's swap prices explained by volatility in j 's over the horizon of p weeks. In the latter, this same quantity is denoted by A_{ji} .

Finally, we tune the autoregressive horizon parameter p . Many credit default swaps, especially for large, stable economies, are traded infrequently, as investors are unlikely to make money going short or long on these assets. Setting p too small risks estimating the autoregressive parameters from moving windows in the data with no or insignificant movement in some assets, thereby zeroing out the estimated effects of these nations' volatility on the others. So, we set p as large as can to tractably estimate all parameters in the model, corresponding to $p = 8$ weeks, or two thirds of a financial quarter.

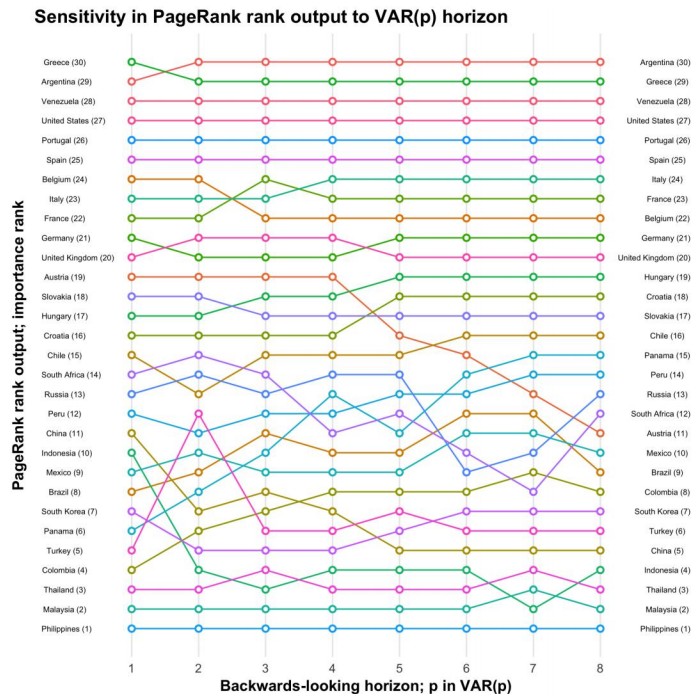


Figure 1: PageRank scores are highly sensitive to p for low values of p , but mostly stabilize as p -horizon increases.

We evaluated this choice experimentally, comparing PageRank scores for each nation in the network for each network constructed with a different value of p . High variance in the PageRank rankings indicates that nations' own-influences – the diagonal elements in A – fluctuate strongly in response to the choice of p , which may be due to time-invariant data as described above. For large p , the rankings stabilize (with some exceptions, e.g. Austria), which is desirable as it indicates that we are more closely estimating the “true” historical form of the network, especially for nations whose swaps are exchanged rarely.

IV. Network Analysis

Constructing the Network

We immediately re-purpose the coefficients obtained from the generalized FEVD, for use as a weighted adjacency matrix. Since we are concerned with volatility in one nation that results from shocks to other nations, we also eliminate all self edges, as these quantify how each nation’s volatility changes in its own time series, endogenously.

Visualizing the Network

Given the large number of nodes and edges in this network, the most revealing visualization comes from interpreting edge weights as a heat-map:

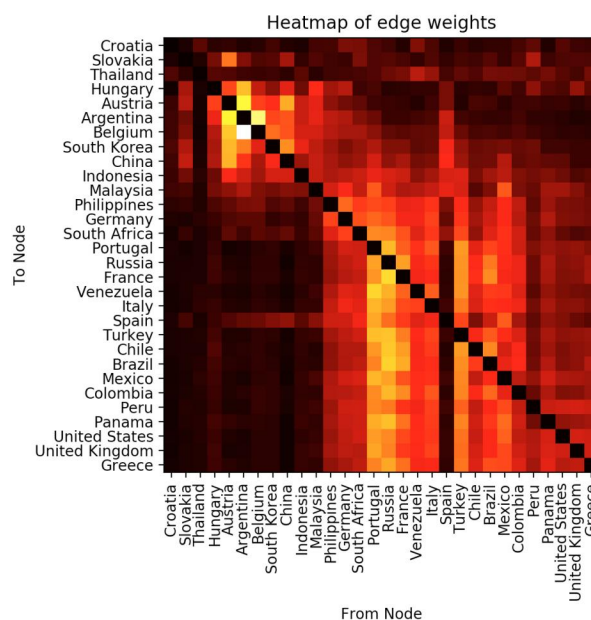


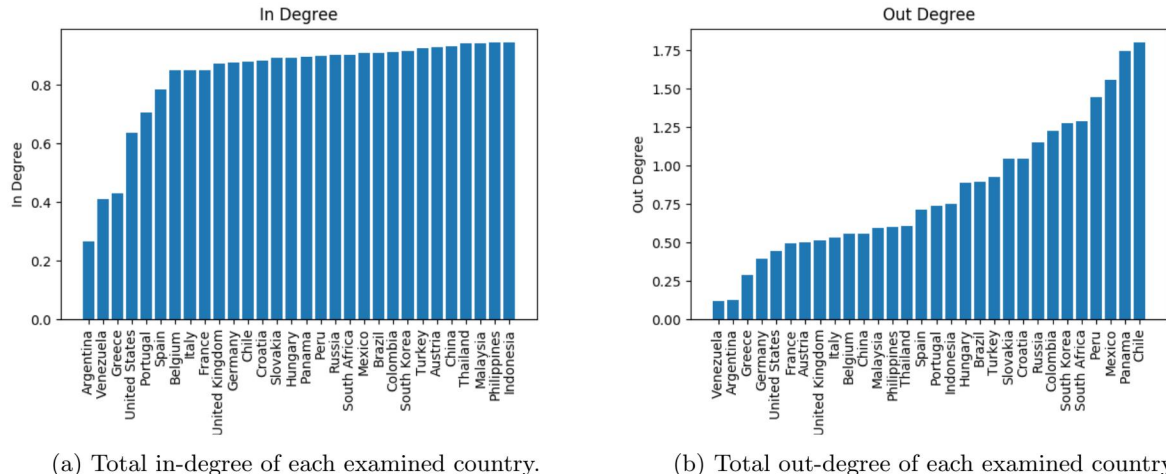
Figure 2: A heat-map of edge weights in the adjacency matrix immediately reveals community structure.

A clear community structure appears to exist within this network. Countries appear divided into two communities with strong edges within each community (but weak edges between the two groups). We will delve further into community structures, in the *Community Detection* section.

Basic Network Analysis

Prior to interpreting results, we posit two potential null models. The first null model is an influence graph with random edge weights. This suggests random relationships of pairwise influence for each pair of nations in the graph. The second null model assign each node’s influence on another node as a function of some shared economic fundamentals. For example, one might hypothesize that high-Gross Domestic Product (GDP) countries will be highly influential. Alternatively, one might hypothesize that pairwise edge weights ought to be proportionate to the current accounts deficit and credit between the pair of countries. The results of such a null model would simply sort countries by influence according to their national wealth.

Now, we examine basic properties of the graph structure. The in-degree of a node represents how “influenced” a node is — a larger in-degree means that a greater fraction of the variance of the node’s innovations come from external shocks. Conversely the out-degree of a node represents how “influential” a node is — a larger out-degree implies that a node is responsible for a larger share of variance in the innovations of other nodes.



(a) Total in-degree of each examined country.

(b) Total out-degree of each examined country.

Figure 3: In- and out-degrees of national volatility in the directed influence network.

Consider the outlier countries in Figure 3(a) and 3(b). Argentina, Venezuela, Greece, and the United States appear particularly isolated with low in-degrees and low out-degrees. This grouping seems surprising at first, but actually illustrates a foundational concept in economic risk analysis and asset pricing: that assets can be priced similarly because either investors place similar risk premiums on them, or because their economic fundamentals are the same (or both) [4]. The degree distributions above measure the changes in market prices, not fundamental values, of these countries’ debts. The United States’ independence arises from its powerful position as the world’s reserve currency and largest economy, isolating it from shocks around the world. Conversely, countries like Argentina and Venezuela, wracked by financial crises, have been intentionally isolated by shrewd financial managers, seeking to restrict the influence of their shocks on the rest of the global economy. We discuss this result further in the *PageRank* section, and in Section V.

PageRank

While countries with high in-degrees are susceptible shocks, countries that are influenced by countries with high in-degrees are potentially even more susceptible to ripple effects of financial crises. This model of influence is easily captured by the PageRank algorithm, which accurately models the “influenced by the influenced” view of networks. PageRank quantifies node importance with the assumption that nodes that are pointed to by influential nodes are also influential. Namely, given a vector of influence ranks r , the influence of node j is:

$$r_j = \sum_{i \rightarrow j} \frac{r_i}{d_i} \implies r = Mr$$

The recursive definition of PageRank on the left is identical to defining r as the greatest (in absolute value) eigenvector of the column-stochastic transition matrix of the directed graph, M . In our application, nodes with the highest PageRank scores are the nations that are most susceptible to credit volatility due to volatility in other countries.

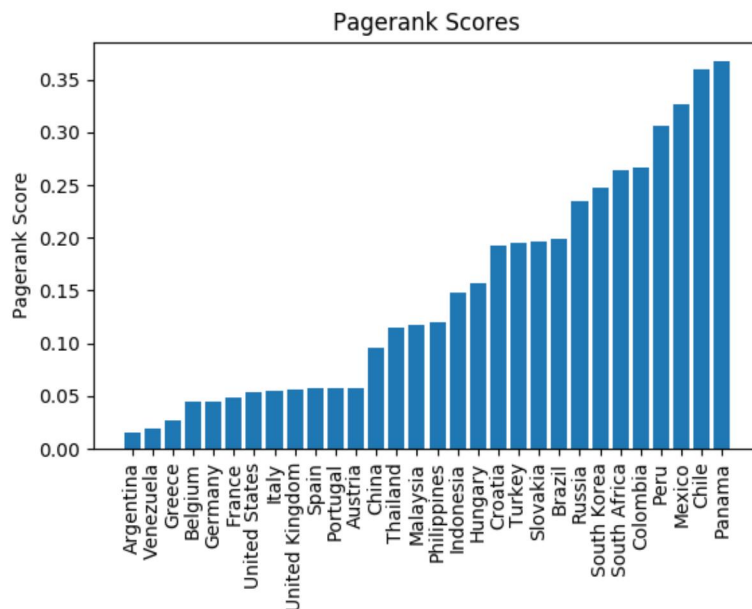


Figure 4: Countries with historically dubious economic fundamentals (such as Greece and Venezuela) are relatively un-influenced by other nations according to PageRank.

Surprisingly, we see countries such as Panama with extremely high PageRank scores. The reason for this rise becomes clear once one considers Figure 4 in concert with Figure 3: Panama has many heavy edges coming in from other heavily influenced countries such as Russia, South Korea, and South Africa, making it a “sink” of economic influence from other highly susceptible economies. The continued independence – evidenced by low PageRank scores – of countries such as Argentina, Venezuela, Greece and the United States further substantiates the discussion above.

Community Detection

Finally, we explore deeper community structure in the adjacency matrix. We hypothesized that since countries that are geographically proximate often engage in trade and are significantly exposed to each other’s economies, we might see communities form in geographic clusters. In fact, two community detection algorithms detected geographic trends in the data absent any prior geographic information. With the goal of constructing clusters based on two nodes’ total effects on one another we averaged the two directed edges between each node in the graph. We then use clustering techniques for undirected graphs to find the clusters with these new combined edge weights in the graph. First, we applied Kernighan-Lin bisection [5] to detect communities. Since the graph is fully connected, there are always two directed edges that reduce to one undirected edge.

Algorithm 1 Kernighan-Lin Bisection

```
1: function KL-BISECTION( $G$ , a graph)
2:   initialize: two communities,  $A$ , and  $B$  from the nodes  $V \in G$ 
3:   initialize:  $i \leftarrow 0$ 
4:   while  $t$  do here are vertices remaining
5:     Select  $a_i \in A, b_i \in B$  such that the reduction in modularity of  $A, B$  is maximized
6:     Swap  $a_i \rightarrow B, b_i \rightarrow A$ 
7:      $C_i \leftarrow$  modularity of these new  $A', B'$ 
8:      $i \leftarrow i + 1$ 
9:   return  $A', B'$  with minimum  $C_i$ 
```

Second, we used spectral community detection with the Fiedler vector on the same undirected weighted graph [7].



Figure 5: Community detection algorithms identify emerging markets and stable, primarily Western economies.

The Kernighan-Lin algorithm yields the two communities in the figure above. As is clear from the two figures, the two communities detected form what many in the financial sector refer to as the “developed” countries of Europe and the United States in community A, and the “developing” countries in community B. Interestingly, Argentina and Venezuela have crept into the “developed” countries in community A - however, this can be explained by the fact that both those countries appear to *behave* like developed countries - strongly independent from other nations in terms of credit default risk. The Fiedler vector implementation of community detection detects this difference and forms three communities - largely the same two as above, with the exception that Argentina and Venezuela form their own community.

V. Discussion

Understood within the context of international finance, these results are remarkably insightful. The in-degree and out-degree distributions in Figure 3 indicate the proportion of each country’s volatility due to price shocks to other nodes in the network. The four countries whose debt volatilities were least affected by shocks throughout the network according to the in-degree distribution are Argentina, Venezuela, Greece, and the United States in increasing order. In the PageRank chart the United States was similarly ranked near the bottom. Likewise, Argentina, Venezuela, and Greece ranked the lowest according to the PageRank metric, which ordered countries in ascending order of how influenced they were by the other nations in the network in the time series observation period.

To provide context: Argentina, Venezuela, and Greece all either defaulted, threatened to default, and/or received IMF bailouts to avoid defaulting in the given time period. In 2014, Argentina opted to

“selectively default” on over \$20 billion in debt due to economic troubles. Venezuela, amidst political and economic turmoil, has steadily halted payments over the last several years on several billion dollars worth of debt; while they have not technically defaulted, many refer to this as a “limbo default.” Finally, in 2010 the European Central Bank and the International Monetary Fund were forced to bail out Greece, which was suffering from a severe debt crisis. Greece received a loan of \$110 billion to rescue the country, which was in a severe recession, from imminent default on bond debts. In addition, Portugal, Spain, and Italy also found themselves in severe recessions and raised doubts as to the fortitudes of their central banks. Notice that Portugal, Spain, and Italy all rank low on the PageRank chart in Figure 4 as well.

Given this context, it seems surprising that community and influence detection algorithms group the United States with these insolvent or borderline-insolvent nations. The United States – while suffering from a financial crisis during the observation period– was never at risk of defaulting on its debts. This is due to the reputation of the United States as a financial bastion and the “exorbitant privilege” that the United States enjoys due to the dollar being the international reserve currency. Argentina, as it suffered from economic decline and irresponsible spending, elected to default on its debt because paying the debt would have put the country and its citizens into further economic distress. Why did they not simply print more Argentine Pesos? Even if the country was willing to deal with hyperinflation, this was not an option for most of their debts because most large debt contracts internationally are priced in terms of the currency of some stable, economic powerhouse. Such stable nations like the United States are perceived as being extremely unlikely to suffer from currency shocks, meaning that the monetary value of these contracts will fluctuate little over time. In some sense, the perceived stability of the United States and other strong Western states perpetuates future stability.

By far, the most common currency that denominates contracts in world markets is the United States dollar. Most of the world’s \$10 trillion in United States currency is used to settle contracts with no Americans involved. The United States benefits from this so called “exorbitant privilege.” Thus, the location of the US in the in-degree graph is precisely where experts would predict it to be, despite the fact that the United States is massively financially exposed to dozens of foreign economies. It is also likely that Belgium, Germany, France, the United Kingdom, and Austria have low PageRank scores as a result of relative confidence in the Euro as a store of value and their reputations as financial powerhouses with relatively stable financial systems.

We posit that Argentina and Venezuela grouped with the Western countries during clustering for one of two reasons. First, within the observation period, the majority of the southern European nations (Greece, Portugal, Spain, and Italy) experienced major debt crises. Intuitively, these nations would likely group with Europe given their severe exposure to each other due to sharing common leadership organizations, regulations, and even a common currency. In addition, it seems reasonable that Argentina and Venezuela would cluster with the economically struggling southern European nations due to similar economic situations.

Alternatively, Argentina and Venezuela’s debt default probabilities as predicted by the markets may have been less susceptible to volatility in other “emerging markets”. In particular, recall from [4] that volatility may stem from either systemic risk due to interconnectedness or from poor economic fundamentals, which Argentina and Venezuela certainly still have to this day. This would cause the majority of price movements in these two nations’ CDSs to result from changes in the fundamentals of the investment itself (that is, in the central banks and political situations of the countries). Contrast this behavior with other emerging markets, where price changes would likely stem from changing risk in the portfolio as a whole more often than changes in fundamental default rates of countries themselves. As such, we expect them to cluster away from other emerging markets, which are highly inter-correlated with each other but on their own have noisy economic fundamentals. Stated differently, a clustering in which countries whose credit fortunes are primarily self-determined – whether through extreme economic strength or extreme economic weakness – is highly intuitive. Furthermore, that most emerging

economies cluster together can be easily explained by portfolio management theory. Suppose an investor maintains a portfolio of emerging market swaps, one of which experiences significant volatility over some observation period. Because of common investing assumptions that emerging markets behave similarly, the investor might demand a higher risk premium on that more volatile asset as well as all others in the portfolio, effectively signaling a lack of confidence in those nations’ sovereign credit. And, the relationship between investor signaling and real financial outcomes is well documented – investor skepticism influences emerging markets’ sovereign credit by making it more difficult for nations to seek loans at reasonable rates in credit markets, effectively creating a self-fulfilling prophecy of credit troubles and even insolvency. Succinctly, one might imagine that emerging markets cluster together because a “bank run” on one of them (due to some idiosyncratic variation in volatility) could easily produce a “bank run” on all [9], [8].¹

VI. Conclusion

One major limitation on the applicability of network analysis is the small amount of real world data which can easily be translated into network form. As such, much of network analysis has, up to this point, focused mostly on social networks, collaboration networks, biological data, and other data-rich areas where the relationships between nodes in the network are readily available and clear. In this paper we statistically infer weighted, directed edges between each of the nodes in a graph of credit default swaps by borrowing methods from the field of econometrics. This kind of application may greatly increase the relevance of network analysis techniques to other fields.

Using network analysis methods on this rigorously constructed graph, we explored the interconnectedness of the world’s most consequential sovereign actors and applied knowledge of political economics, financial management theory, and current events to contextualize the degree distribution, PageRank, and community detection results of this graph. Basic network analysis techniques such as degree distributions are able to quickly yield financial insights about the CDS markets – countries like Argentina, Venezuela, Greece and the United States are far more independent than others from exposure to external financial shocks. More sophisticated network analysis using PageRank shows that countries such as Panama are most likely to suffer from financial ripple effects since they act as “shock sinks.” Finally, algorithms such as Kernighan-Lin graph bisection and spectral community detection infer geographic and financial relatedness – dividing the globe into the common “developing” and “developed” worlds without the need to encode any prior knowledge of these distinctions.

Absent the rigorous econometric work of [3] and [2], the construction of the graph would not have been possible. Network analysis confirmed a longstanding challenge of global financial macroeconomics – that differentiating between the effects of economic fundamentals and risk appetite in understanding asset prices and volatility is challenging. In this case, the United States, a country with strong economic fundamentals, clustered together with many near-insolvent nations whose high debt prices and volatilities stemmed from skeptical investors asking high risk premia to subsidize these central banks. Applying network analysis techniques alongside domain knowledge to this new data set yielded insights on systemic credit default risk, specifically on the remarkable amount of debt risk exposure that otherwise-unrelated “emerging market” economies have on one another. These results might be a remarkable indictment of the real world effects that financial rules of thumb may have on billion dollar economies and billions of lives in the developing world, who have no control over which Goldman Sachs portfolio they are held in.

¹The concept of Western investors treating emerging market economies as if they were a cohesive asset class whose values are interconnected in spite of having quite separate, unrelated national economies is expanded upon by Ruchir Sharma in *The Rise and Fall of Nations* and *Breakout Nations*

References

- [1] Stefano Battiston, Michelangelo Puliga, and Rahul Kaushik. “DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk”. In: *Nature* (2012), n. pag. (Web).
- [2] Mert Demirer et al. “Estimating Global Bank Network Connectedness”. In: *Journal of Applied Econometrics* (2018), pp. 1–15.
- [3] Francis X. Diebold and Kamil Yilmaz. “On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms”. In: *TUSIAD Economic Research Forum Working Paper Series* (2011).
- [4] Benjamin Hebert (Stanford Graduate School of Business). Personal communication.
- [5] Brian W. Kernighan and Shen Lin. “An efficient heuristic procedure for partitioning graphs”. In: *The Bell System Technical Journal* (1970), pp. 291–307.
- [6] Gary Koop, M. Hashem Pesaran, and Simon M. Potter. “Impulse response analysis in nonlinear multivariate models”. In: *Journal of Econometrics* (1996), pp. 119–147.
- [7] Fragkiskos Malliaros and Michalis Vazirgiannis. “Clustering and Community Detection in Directed Networks: A Survey”. In: *Physics Reports* (2013).
- [8] Ruchir Sharma. *Breakout Nations*. 2012.
- [9] Ruchir Sharma. *The Rise and Fall of Nations: Forces of Change in the Post-Crisis World*. 2016.
- [10] Eric Sims. “Graduate Macro Theory II: Notes on Time Series”. In: (2011), n. pag. (Web).

Contributions

Note: All three authors equally contributed to Sections I and II

Aakash Pattabi: Aakash was in charge of reading and understanding the econometric literature which informed our graph construction. In addition, he was chiefly responsible for obtaining and structuring the credit default data on which the network would be built. Also, he was responsible for using this data in constructing the financial network in a similar fashion to the graphs constructed by Diebold et al. in their work. He also contributed to the overarching planning of the project and helped interpret and debug all results throughout the process of writing the paper. Aakash wrote the paper’s network construction section and contributed to the graph analysis section.

Eric Gilliam: Eric’s reading was the inspiration for the idea behind the project; Eric was instrumental in defining the problem and reviewing the relevant literature; contextualizing the findings with knowledge of macroeconomic theory, financial management, and political economics; and evaluating the results of the PageRank and community detection algorithms. Eric additionally wrote the paper’s discussion and conclusion sections.

Ashwin Sreenivas: Ashwin helped the group understand the inner workings of the Vector autoregressive method which was used to help construct edge weights. He was chiefly responsible for carrying out the network network analysis methods and visualizations in the paper. These included performing the analysis for and subsequently visualizing the heat map of edge weights, in degree distribution, out degree distribution, PageRank scoring, and clustering which were all seen in Section IV of the paper. Ashwin wrote the paper’s graph analysis section.