



# FiNANCE FOR ENERGY MARKET RESEARCH CENTRE



## Untangling systemic risk in financialized Commodity markets

Julien Ling

Working Paper  
RR-FiME-19-03

January 2019

# Untangling systemic risk in financialized commodity markets

Julien Ling

EDF Energy - Finance for Energy Market Research Centre

January 28, 2019

Preliminary document

## **Abstract**

Systemic risk is a multifaceted concept that is of crucial importance for regulators. In order to ensure financial stability, they need to properly assess this risk, preventing financial shocks from affecting the real economy. In this study, we evaluate the extent to which the financialization of commodity markets contributes to systemic risk. We consider a system consisting of both commodity futures and financial markets in a sparse Vector AutoRegression (VAR) framework. It allows to distinguish two temporalities of systemic risk: we can assess "systematic" risk (integration) and propagation risk. In particular, we can identify which markets are influential in systemic risk and thus conduct a more in-depth investigation if necessary. Since we assume sparsity in the parameter matrices, we can rely on an algorithm using LASSO. In a static analysis, in the spatial dimension, we find that sectors are separated, except for metals and finance. Including the maturity dimensions proves necessary, as they connect all the sectors and thus cause the integration of the whole system. In our dynamic analysis, we focus on major financial events. We find that integration has been building up, was prominent and very high around each of these events between commodities and financial assets and among commodities, making common shocks a realistic possibility.

# 1 Introduction

Since the 2007-2008 financial and economic crises, regulators have tried to improve their understanding and monitoring of financial markets. The notion of systemic risk has gained interest worldwide and became their focus, considering how difficult it is to assess. The BIS (2001) defines systemic risk as "the risk that an event will trigger a loss of economic value or confidence in, and attendant increases in uncertainty about, a substantial portion of the financial system that is serious enough to quite probably have significant adverse effects on the real economy." From this definition, one can retain that the risk comes from the financial system and that it has a high probability of affecting the real economy, which is separated from the financial system. This system can be split into several layers, which are intra- and interconnected: institutions, instruments and markets.

Since the main events that are considered as the start of the 2008 crisis were the default of major financial institutions, the majority of the research focused on defaults (and their propagation) and their contribution to systemic risk. Please refer to Benoit, Colliard et al. (2017) for a review on this topic. More specifically, regulators and other international institutions have come up with a set of basic indicators to assess the "systemicity" of an institution. The Bank for International Settlements (BIS) published a methodology in BCBS (2018b) using the following criteria: size, cross-jurisdictional activity, complexity of the activity, nonsubstitutability of services and interconnectedness of the institution. Each criterion consists of several indicators. For example, the size score is computed as the "total exposures as defined for use in the Basel III leverage ratio", which is explained in BCBS (2018a). These exposures include derivatives, which are basically market valued and, if not netted, are inflated by a coefficient alpha, currently set to 1.4. Hence, derivative positions of financial institutions may pose even more regulatory pressure if their underlying assets are risky.

In particular, it is well known that individual commodity prices experience many jumps and thus their return distributions have fat tails, see e.g. Deaton and Laroque (1992). This does not prevent financial institutions to be involved in commodity spot and derivative markets, as shown by Benoit, Hurlin and Perignon (2015). Their increasing presence and its consequences have raised concerns regarding the influence of financial markets on commodity markets, which may be referred to as the financialization of commodity markets (see Cheng and Xiong (2014) for a review of this strand of literature). These two sectors actually interact in several ways. Since commodities appear in portfolios for the reasons mentioned above and in economic activ-

ities, they may in turn influence many financial assets, like bank or corporate stocks. Commodities are also influenced by financial institutions using them: Büyüksahin and Robe (2014) show that the presence of speculators on both equity and commodity markets generates correlation between their price returns. In addition, many large funds are dedicated to commodity investing and their interventions impact prices, for example when large synthetic Exchange Traded Funds (ETFs) or indices rebalance their portfolios or roll their futures contracts.

A strand of literature shows that futures markets, and later ETFs, have taken over the price discovery function from their spot counterpart. For example, Garbade and Silber (1983) show that spot markets play a little role compared to futures markets for some agricultural and metals markets. This move towards futures contracts and ETFs can also be witnessed for financial assets, as evidenced by Hasbrouck (2003) for the S&P500.

In fact, futures and spot markets are intertwined (for example because of arbitrage), making the propagation of information between the two as fast as market trading allows. As such, commodity futures present the potential to transmit financial shocks to commodity spot prices, which may affect economic activity and inflation, hence the real economy. Indeed, financial agents rebalancing their portfolios because of a shock on financial asset prices may generate large variations of commodity futures prices, which would subsequently translate into spot price movements. Due to the market trading speed, this potential influence generates systemic risk at a daily or even intra-daily frequency, as P. K. Jain, P. Jain and McNish (2016) find for Japanese equities.

Assessing systemic risk requires taking into account many interconnections between agents and markets. Regulating agents should work in the long-run. But since trading can occur at milliseconds or even faster, it seems more relevant to intervene on markets instead of agents (who need to disclose their positions, which may take time) to avoid flash crashes. Circuit breakers have been set if they were not already before the crises, but due to the potential leverage of financial institutions and to those interconnections, even small variations may lead to large losses, which may in turn affect other agents and markets.

If we need to account for the real economy, we need to consider many other factors than just financial markets. Commodities indeed have many economic and financial links, to other commodities, weather, financial assets, etc. Nevertheless, many papers have focused on specific links. For instance, Ewing, Malik and Ozfidan (2002) examined volatility spillovers between crude oil

and natural gas markets. Silvennoinen and Thorp (2016) study the correlation between crude oil and agricultural markets, which may have changed dramatically due to the development of biofuel.

All in all, it is important to incorporate as many markets as possible into a single system, so that one avoids missing relevant variables in the information transmission mechanism, and also to conduct a dynamic analysis to understand how the system evolves.

This work aims at providing tools for regulators to better understand and monitor a large system of markets from different sectors. We revisit the following relationships in a high dimensional setting: the integration of financial and commodity markets, the potential propagation of information and shocks between them and among commodities and the relationships between futures and spot markets. We use a sparse Vector AutoRegression (VAR) framework on daily data. Integration is measured in the partial correlation (PC) matrix, which gives the correlation coefficients of each pair of variables, conditional on all the other variables in the system. It thus filters out their influence and provides a more accurate measure of the dependence, hence of integration. The use of partial correlations here improves on unconditional correlations (see Appendix A for an explanation and example). While propagation may be more prominent at the intraday level, daily returns should contain more economic information, such as the transmission of information from one time zone (closing) to another (opening), hence daily propagation would dismiss microstructure noise. Propagation is measured in the Granger causality structure of the system, given by the AutoRegression (AR) matrix or matrices.

Such an analysis has both policy and industry implications. For the former, monitoring the evolution of the system can help prevent the build-up and occurrence of systemic crises by avoiding too much market integration and/or propagation. Disentangling the economic links from financial ones – by having or not economic foundations for them – would also help understand where the risks lie. Identifying those (ephemeral) financial links can also trigger intervention on these markets in order to break them and avoid unnecessary risk transmission.

For the latter, assessing integration would allow to allocate capital in an optimal way, so that the portfolio would benefit from diversification. Unless the manager can trade at very high frequency – and exit the markets before he suffers large losses –, investing in markets that have economic links rather than evanescent financial links may prove more robust and stable. For example, Lohre, Papenbrock and Poonia (2014) have recourse to graph theory

on equity sectors correlation matrices and can derive meaningful risk measures from the graphs. Moreover, they find that the location of the asset in the graph matters in terms of performance and market timing.

In Section 2, we review some of the literature related to our work here. In Section 3 we briefly introduce our database. In Section 4, we present the methodology we use and how it is relevant to study propagation and integration. In Section 5, we analyse our whole database (2000 to 2014) to provide long-term reference graphs to compare with and explain the measures and results we derive. In Section 6, we then conduct a dynamic analysis (with a rolling window), in which the problem becomes highly dimensional, thus the relevance of using LASSO. More specifically, we focus on periods around major financial events that may have affected commodity markets. In chronological order, first, we look at the day BNP Paribas froze the redemption of some of its investment funds because it could not value some of its structured products (August 9, 2007, one of the dates that are defined as the beginning of the subprime crisis). Second, we examine integration and propagation around the default of Lehman Brothers (September 15, 2008), which is often assumed to be the beginning of the economic crisis. Third and fourth, we assess whether the minimum and maximum returns on the S&P500 index in our sample, respectively October 9, 2008 and October 28, 2008, propagated to commodity markets, since we observed in the data that many of them experienced their extremum return on the following day. Fifth, we look into what may have happened around the day of the May 6, 2010 Flash Crash, since it has affected some contracts in our system, namely the S&P500 e-mini futures.

## 2 Related literature

As stated before, we revisit the following relationships in a high dimensional setting: the integration of financial and commodity markets, the potential propagation of information and shocks between them and among commodities and the relationships between futures and spot markets. We thus follow these three strands of literature, plus that of systemic risk.

Due to the spectacular (cascades of) defaults occurring since the beginning of the subprime crisis in August 2007, regulators and researchers have focused on studying the default of institutions and how the default can propagate and affect the whole financial system. For example, Acharya et al. (2017) developed the Systemic Expected Shortfall (SES), which measures

the propensity of an institution "to be undercapitalized when the system as a whole is undercapitalized". They thus focus on the increased risk faced by an institution due to an undercapitalized system (crisis).

On the contrary, Adrian and Brunnermeier (2016)'s CoVaR measures "the change in the value at risk of the financial system conditional on an institution being under distress relative to its median state", so the institution contributes to the risk of the financial system.

Other measures have been created, in particular regarding (the formation of) financial networks. Hautsch, Schaumburg and Schienle (2015) refines their analyses by modelling networks of institutions and computing the contribution of each to the Value-at-Risk (VaR) of the others, based on various variables such as leverage, size or macroeconomic context. Capponi and Larsson (2015) provide a model of interconnected agents and markets through their full balance sheets and can thus derive measures of "systemicness" of a shock on an institution or on a market by accounting for all the mechanisms and dynamics related to that shock.

Though very insightful and beautiful, these models and applications rely on information that may not be available at a relevant frequency. Benoit, Hurlin and Perignon (2015) infer the exposures of institutions from the disclosure of their VaR per sector, which may become quickly irrelevant as trading may occur much faster than the disclosure. Regulating or monitoring markets instead seems more efficient, but requires being able to process the influx of data as quickly as possible.

While some actions were taken by exchanges for preventing flash crashes (e.g. breakers), due to the potential leverage of financial institutions and to those interconnections, even small variations may lead to large losses. These losses may translate into defaults and further shocks for the financial system and the real economy. It is thus crucial, as has been identified by regulators (e.g. in BIS, FSB and IMF (2009)), to be capable of understand and examine all the links at play on markets. This is where our study contributes: we analyse a potentially highly interconnected system of markets and unveil economic and financial links from their complexity and the amount of data. For example, Lautier and Raynaud (2012) studied the integration of a very large system of futures contracts (220) and filtered their correlation matrix in order to retain only the most important connections. Our methodology differ and, in terms of systemic risk, our aim here is to distinguish integration and propagation, which represent different temporalities of dependence. It should allow to better understand the links between markets and monitor their risks more accurately.

Benoit, Hurlin and Perignon (2015) found that financial institutions have exposures to commodities. Gorton and Rouwenhorst (2006) indeed show that before 2004, a commodity futures index displayed interesting diversification benefits, with returns about the same as the S&P500 on average and zero or even negative correlation with it. Nonetheless, since these have been considered individually risky due to jumps and risk premia, they could very well generate losses for agents using them. In addition, commodities have many economic and financial links, to other commodities, weather, financial assets, etc. largely increasing the sources of risk, even more so if we consider commodity indices. In addition, Tang and Xiong (2012) show that commodities belonging to an index have larger correlations than commodities not included in it, creating more (financial) links between them. The literature on the financialization of commodity markets also unveiled potential ephemeral links between the two sectors and hence the influence of the financial system on commodity (futures) markets.

While there is a need for comprehensiveness, most papers have only focused on specific links, probably due to the lack of proper methodology. For instance, Zhang et al. (2008) looked at the spillover effects between US dollar exchange rate and oil prices. Hammoudeh and Yuan (2008) considered spillover between metals by taking into account interest rates and oil price shocks, while Park and Ratti (2008) assess how oil prices affect equity indices in European countries and the US. Ewing, Malik and Ozfidan (2002) examined volatility spillovers between crude oil and natural gas markets. Since the 2005 Energy Act, biofuel started being used and may have changed the links of energy markets and agricultural markets; hence the work of Silvennoinen and Thorp (2016).

We thus also contribute to the literature on the financialisation of commodity markets by providing tools to analyse a large system of interconnected markets, from both commodities and financial assets. Those tools should help unveil economic and financial (ephemeral, as shown by Büyüksahin and Robe (2014)) links. Focusing on the links between the two sectors (any pair of markets, one from each sector) should provide insight about which markets are financialised and how. Also, links that do not have reasonable economic foundations may also be qualified as financial, requiring deeper investigation and closer monitoring.

We also contribute to the literature on the propagation of information and price discovery function of spot and futures markets. While futures are derivatives – hence their price should be "derived" from the price of their underlying asset, as in the theory of storage by Kaldor (1939), Working (1949), Brennan (1958) and Telser (1958) –, it sometimes seems to be the



contrary. As we mentioned earlier, Garbade and Silber (1983), Schwarz and Szakmary (1994) and Hasbrouck (2003) show that spot markets play a little role compared to futures markets for some commodities and financial assets. The relationship is actually bi-directional, spot influencing futures and vice versa.

In addition, according to the World Federation of Exchanges 2017 report, equity and equity indices futures were the first class of futures in terms of volume (number of contracts). They are closely followed by commodity futures, representing at least 38% of all futures contracts, and their volume is following a positive trend. The importance of these two sectors and their connections thus represent significant risk. Indeed, since commodity futures are easier to trade than physical commodities, financial agents focus mainly on futures for their activity. Since they are the primary source of information on prices and since spot and futures markets are tightly connected, financial agents can quickly make commodity futures prices move, hence spot prices and the real economy.

Our contribution here lies in differentiating propagation and integration, which takes into account different temporalities of information transmission. In addition, we can assess which of the spot or futures market is dominant in terms of both propagation and integration, hence which one contributes the most to the price discovery process for all the markets in the system.

All in all, it is important to incorporate as many markets as possible into a single system, so that one avoids missing relevant variables in the information transmission mechanism, and also to conduct a dynamic analysis to understand how the system evolves.

We use a sparse Vector AutoRegression (VAR) framework on daily data to distinguish two components of systemic risk: we assess the integration of markets at the link level (pairwise), but also at a more global level (sector, system), and the propagation of returns between markets.

Coming back to the focus of this work – systemic risk –, the larger the system is, the more difficult the analysis can be. Having many sources of information in a single study requires to identify the relevant ones.

For example, Lautier and Raynaud (2012) also use graph theory to filter a very large correlation matrix and analyse the dynamics of the corresponding graph over time. It provides a visual representation that enables quick inspections and displays a meaningful structure, as markets from a same sector cluster and sectors connect through relevant markets. They then derive several measures from the graph and find an increasing integration (a condensing graph) in their system of commodities and financial assets, which reaches its maximum at the end of 2008, after the beginning of the crisis.

Using a VAR framework allows to distinguish between increasing integration and propagation. Working on a system of different markets (51 time series from 17 markets, 4 sectors), we rely on the seemingly reasonable assumption that not every market/contract is linked to every other. For example, why would the 12-month eurodollar contract be linked, contemporaneously or with a lag, to the 3-month natural gas, 12-month copper, 1-month soy, 2-month soy, 3-month soy, 4-month soy, etc. contracts? Based also on the results of Lautier and Raynaud (2012) that the contemporaneous dependence matrix can be filtered but still provides meaningful information, we use a calibration algorithm with LASSO penalisation from Barigozzi and Brownlees (2017). We choose LASSO for several reasons (see Appendix B for more details), but the most important ones are that it fits our assumption (sparse matrices), which can be represented as graphs, and it allows to work in high dimension (so we can increase the size of our system as much as we want or need to). Their algorithm calibrates a VAR model (AR and PC matrices) in a single step, which has several advantages, e.g. in terms of convergence properties of the parameters. It also avoids forcing propagation on the data by first calibrating the AR matrices as is common in other algorithms.

Barigozzi and Brownlees (2017) also provide an application of their algorithm to volatility spillovers between S&P100 stocks. While our interest lies in the very short term and the dynamics of the system, they focus on longer term horizons by combining the AR and PC matrices into a single dependence matrix.

This work also differs from Diebold, Yilmaz and Liu (2017), who also use a sparse VAR framework, but they use an adaptive elastic net (a mix of LASSO and Ridge penalisations) and focus on commodities only. Also, as in Barigozzi and Brownlees (2017), they summarise both the contemporaneous and lagged influences into a single matrix. They then derive variance decompositions and aggregate them into node-level and system-level directional connectedness to provide a global measure of systemic risk. Here, as stated above, we work on partial correlations (from the concentration matrix), which gives us the conditional (in)dependence structure of the system, compared to the variance decompositions. We also keep the contemporaneous and the lagged influences separated in order to disentangle system integration (in the former) from propagation (from the lagged links). It thus allows to identify which markets are actually involved in each component and monitor and intervene on them if necessary. Other works have also used sparse frameworks, but have only focused on one matrix: a sparse VAR with only the AR matrix being sparse or only the concentration/covariance matrix being sparse (see Barigozzi and Brownlees (2017) for some examples).

The sparsity of the matrices enables their representation as graphs: the partial correlation graph and the Granger causality graph. The nodes of the graph will represent the time series of our system and the edge between a pair of nodes will represent, in the former, the PC (undirected edge) and in the latter, their Granger causality (directed edge). In the integration (PC) graph, if some nodes are connected (i.e. belong to the component in the graph), they may be subject to a common factor, driving all their partial correlations. Even though Barigozzi and Brownlees (2017) recommend controlling for those common factors to get sparser matrices, we want to have the possibility to visualise whether financial and commodity markets are influenced by a common factor or not.

Based on these graphs, we derive graph theoretic measures that can help monitor potential propagation and integration and help identify which markets should be investigated in more depth. More particularly, the sparsity of the matrices allows to assess the importance of propagation and integration in terms of number of connections involved (degree centrality). In addition, we use the total communicability centrality measure by Benzi and Klymko (2013) to identify important contracts and also to assess how easy information can flow in the network (which is also a measure of integration when applied to the PC matrix). We adapt this measure to our application by using a weighted matrix instead of an adjacency matrix. It thus accounts for the individual weights (partial correlations) instead of using a single weight parameter, as we will explain below. This work should thus provide means for regulators to understand and monitor the short-term dynamics of markets and eventually be used to prevent the occurrence of crises by taking action on the identified markets if propagation risk or integration are too high.

### 3 Data

We collected futures prices for 17 different underlying assets from 4 different sectors (energy, finance, metals and agriculture) from Datastream, constructed continuous time series with constant maturity and computed daily returns.

We have 208 time series, with many maturities for some markets, but will only keep 3 for each market (short-, medium- and long-term contracts) for several reasons. First, not all maturities of each market are relevant. Second,

Table 1: Data summary

Market	Exchange	Node labels (maturity in months)
WTI	CME-US	WTI1 (1), WTI2 (3), WTI3 (12)
Brent	ICE-EU	Brent1 (1), Brent2 (3), Brent3 (12)
Heating Oil	CME-US	H.O.1 (1), H.O.2 (3), H.O.3 (12)
Gasoil	ICE-US	Gasoil1 (1), Gasoil2 (3), Gasoil3 (12)
US Nat. Gas	CME-US	USNat.Gas1 (1), USNat.Gas2 (3), USNat.Gas3 (12)
UK Nat. Gas	ICE-EU	UKNat.Gas1 (1), UKNat.Gas2 (3), UKNat.Gas3 (9)
Wheat	CME-US	Wheat1 (3), Wheat2 (5), Wheat3 (12)
Soy Bean	CME-US	Soybean1 (2), Soybean2 (4), Soybean3 (12)
Soy Oil	CME-US	Soyoil1 (1), Soyoil2 (3), Soyoil3 (12)
Corn	CME-US	Corn1 (3), Corn2 (5), Corn3 (12)
Eurodollar	CME-US	IR1 (1), IR2 (3), IR3 (12)
USD/EUR Fx Rate	CME-US	FX1 (3), FX2 (6), FX3 (12)
S&P500	CME-US	SP5001 (Spot), SP5002 (3), SP5003 (6)
Gold	CME-US	Gold1 (1), Gold2 (4), Gold3 (12)
Silver	CME-US	Silver1 (1), Silver2 (3), Silver3 (12)
US Copper	CME-US	USCu1 (1), USCu2 (3), USCu3 (12)
UK Copper	LME-EU	UKCu1 (1), UKCu2 (3), UKCu3 (12)

we want to have an overall balanced representation for each market. Third, working with too many time series, even in a high dimension framework with variable selection and filtering, can still lead to results that are difficult to visualise and interpret.

Table 1 details the underlying assets (Market column) we retained, the exchange on which they are traded (Exchange column) and the maturity of the contracts (Node labels and maturity) we kept. We thus chose three maturities for each market: the front-month contract (or actual spot), representing the spot value (short term); the 3-month maturity (or closest greater than 3), representing the medium term; and the 12-month maturity (or largest available if less than 12), representing the long term. We will use front-month or spot interchangeably.

We thus end up with 51 nodes/variables, with a total of 2889 daily observations of return for each, from 2000-01-24 to 2014-02-14, after removing missing data.

## 4 Methodology

An obvious way to think about propagation is to consider the effect of one (or several) market(s) on others. This is exactly what a Vector AutoRegression (VAR) is doing: it assumes that previous observations of a vector of random variables influence the current observation. More formally, if  $\mathbf{Y}_t$  is our vector of random variables  $Y_{i,t}$ ,  $i = 1, \dots, N$ , we have that:

$$\mathbf{Y}_t = \sum_{k=1}^p \mathbf{A}_k \mathbf{Y}_{t-k} + \mathbf{u}_t \quad (1)$$

with  $p$  being the order of the VAR and  $\mathbf{u}_t \sim \mathcal{N}(0, \Sigma_u)$ .

Table 2: Calibrated AR matrix of a VAR(1) model on metals markets in the spatial dimension.


	<b>UKCu1</b>	<b>USCu1</b>	<b>Silver1</b>	<b>Gold1</b>
<b>const</b>	0.0005	0.0004	0.0004	0.0003**
<b>L1.UKC<u>u</u>1</b>	-0.2432***	0.0362	0.0072	0.0163
<b>L1.USCu1</b>	0.2235***	-0.1005**	-0.0193	-0.0366
<b>L1.Silver1</b>	-0.0561**	-0.0459*	-0.0598**	-0.0140
<b>L1.Gold1</b>	0.0476	0.0522	0.0680	0.0213

\* = statistically significant at 10%, \*\* = statistically significant at 5%, \*\*\* = statistically significant at 1%. L1 means lagged value of 1 period.

Applying this framework on small systems gives interesting results. For example, if we consider only metals spot markets (here, two coppers, silver and gold) and calibrate a VAR(1) using OLS, we find the parameter estimates of Table 2. Keeping only the statistically significant coefficients (i.e. considering the others are 0), which correspond to Granger causality relationships, we can summarise these parameters as links in a graph, as shown in Figure 1. The nodes represent the time series of returns of each futures

contract. The directed links represent the Granger causality relationships. For example, the link going from USC<sub>Cu</sub> to UKC<sub>Cu</sub> means that the return on USC<sub>Cu</sub> at a date  $t$  will affect the return on UKC<sub>Cu</sub> at date  $t + 1$  (which is given by the corresponding entry in the VAR(1)).

Figure 1: Granger causality directed graph from VAR(1) on metals spot markets



The nodes represent the time series for each spot market. The directed links (arrows) represent the Granger causality between the different markets: for example, the arrow from USC<sub>Cu</sub> to UKC<sub>Cu</sub> means that the return on the copper traded in the US will Granger cause (influence) the return on the copper traded in London on the following day. The number on each edge represents the autoregression (AR) coefficients from Table 2.

The directionality of these links give interesting insights. There is some autocorrelation for both coppers and for silver and it is negative for all of them. This is consistent with the common view that commodity markets exhibit a mean-reverting behaviour (see Lutz (2010) for a review of explanations and tests). The link from USC<sub>Cu</sub> to UKC<sub>Cu</sub> could reflect the time difference between the markets: the information from the closing of the Chicago market would be incorporated the following day for the London market, for about 22%. Silver (Chicago) is influencing the two copper markets but is contributing only little to these markets. What is also interesting is that the gold market is not affected or affecting the others (at a statistical significance of 1%), but only with a lag 1. There is thus a clear separation between the reserve of value of the gold and industrial metals. This may be explained by a different temporality of the dependence: some cross-correlation at longer lags or maybe even in the contemporaneous observations.

We thus also study the contemporaneous correlation matrix. More precisely, we turn to the partial correlation (PC) matrix  $\mathbf{C}$ , because it encodes the conditional dependence structure of the time series. Indeed, if  $c_{ij} = 0$ , then  $Y_i$  and  $Y_j$  are independent conditional on the other variables

$(Y_k, k \neq i, j)$ . It thus filters out the influence of the other variables, which may have resulted in an exaggerated correlation if the two variables  $i$  and  $j$  are not actually correlated but are both correlated with another one for example (see Appendix A for more details).

If we turn to the PC matrix from the VAR(1) calibration, we find the


Table 3: Calibrated partial correlation matrix of a VAR(1) model on metals markets in the spatial dimension.

	UKCu1	USCu1	Silver1	Gold1
UKCu1	-1.0000	0.8516***	0.0169	0.0061
USCu1	0.8516***	-1.0000	0.1307***	0.0149
Silver1	0.0169	0.1307***	-1.0000	0.7321***
Gold1	0.0061	0.0149	0.7321***	-1.0000

\* = statistically significant at 10%, \*\* = statistically significant at 5%, \*\*\* = statistically significant at 1%.

one displayed in Table 3. We can first see that all off-diagonal coefficients are positive, indicating a positive dependence between all metals. There are only two high PC coefficients: the one between US and UK coppers (0.8516) and the one between silver and gold (0.7321). This results confirm the view that industrial metals are to be considered different of precious metals. The other coefficients being relatively small (between 0.1307 and 0.0061), one may consider that they are not relevant, but statistical tests show that the PC between silver and US copper is also significant at 1%.

Figure 2: Partial correlation undirected graph from VAR(1) on metals spot markets



The nodes represent the time series for each spot market. The undirected links represent the statistically significant partial correlation (PC) between the different markets. For example, the edge between UKCu and USCu means that, filtering out the influence of silver and gold, the two coppers are contemporaneously dependent. The number on each edge represents the PC coefficients from Table 3.

Keeping again only the statistically significant PC coefficients, we can represent them as a graph, shown in Figure 2. We find that the graph corresponding to the statistically significant PC matrix is linear. Though there

is no direct link between UKCu and Silver (their PC is not statistically significant, so assumed to be 0), there is a path connecting the two, through USC<sub>u</sub>. This means that even though they are not directly dependent, they will display direct dependence if we look at the unconditional correlation coefficient (see A for a detailed explanation), since it would be nonzero. Working on PC graphs thus gives a meaningful visualisation: if we see that two markets/sectors are connected, it means that they may be subject to common shocks, that they are integrated.

Working on a larger system would allow to filter out the influence of more variables and uncover a more accurate dependence between any pair of variables. Moreover, after filtering this influence, we may end up with many PC coefficients close to 0, potentially statistically insignificant. The more heterogeneous the system is, the more PC coefficients should be equal to 0 (since there would be no actual direct dependence between two unrelated assets). In addition, as per the results of Lautier and Raynaud (2012), filtering the unconditional correlation matrix still gives a meaningful dependence structure, with linear graphs in maturity dimensions and sectorisation. Diebold, Yilmaz and Liu (2017) also find clustering in commodities by sector, which also comforts us in using PC instead of unconditional correlations. We can thus assume that there may be sparsity in the dependence structure (every one of our contracts should be independent of many of the others).

Based on this assumption, the Least Absolute Shrinkage and Selection Operator (LASSO) regularisation is particularly suited for calibrating sparse parameter sets, by setting the irrelevant ones to 0 and taking these 0s into account when estimating the others. This regularisation thus allows to select the relevant parameters without giving any prior about them.

In addition, instead of successively estimating the AutoRegression (AR) matrix and then the partial correlation matrix of the residuals, we will estimate them simultaneously. Indeed, having a two-step estimation procedure gives nontrivial properties of the estimators and convergence, as investigated by Barigozzi and Brownlees (2017). Nevertheless, having to estimate both the AR matrices ( $N \times N$  for each lag) and the contemporaneous PC matrix ( $\frac{N \times (N-1)}{2}$ ) of the residuals makes this problem high-dimensional even for relatively small systems. For example, if we take our 17 different assets and keep only 3 contracts for each, we have a system of  $N = 51$  nodes/variables. In a VAR(1) model, this would mean  $51 * 51 = 2,601$  AR coefficients, plus  $51 * 50/2 = 1,275$  PC coefficients for a total of 3,876 parameters. The need for a large amount of data prevents one from working on high-dimensional



systems. We want our parameters to reflect both the short- and long-term dynamics of the markets. We thus rely on machine learning to provide robust parameter estimates in spite of the relatively low number of observations, by selecting only the relevant parameters. For example, Diebold, Yilmaz and Liu (2017), though also using machine learning to estimate the parameters of their VAR, use only 19 variables (sub-indices). We rely on LASSO instead of the adaptive elastic net they use, because LASSO has a stronger shrinking effect, giving more sparsity in the parameters than ridge or elastic net, which should be more suited to our study.

In the end, we will calibrate a high-dimensional VAR model on our data, with sparse AR and PC matrices. We use an algorithm developed by Barigozzi and Brownlees (2017), called "nets algorithm", in order to simultaneously estimate the sparse AR and sparse PC matrices thanks to a LASSO regularization. These two sparse elements are then represented as two graphs, as we saw before. The directed graph for the Granger causality links from the AR matrix will inform us about potential propagation of information and shocks, while the undirected graph for the contemporaneous PC will rather inform us about the integration of the system, and hence, the potential for systematic shocks.

This framework and the sparsity will allow us to assess systemic risk in different ways. We can, for instance, visualise the propagation links, which markets and sectors are involved, assess integration at a glance by looking at the PC graph and the number of components it has, how the markets or sectors are connected/integrated. More quantitatively, we compute the average or range of propagation coefficients, of integration coefficients, and derive some graph theoretic measures such as centrality in the graphs, to identify the one(s) that is (are) the most important market(s), contributing the most to shocks.

More specifically, we will briefly consider degree centrality (the number of links/neighbours of a node) and compare the results with total communicability centrality, developed by Benzi and Klymko (2013). This measure allows to take into account not only direct neighbours (degree), but also indirect neighbours (even infinitely far ones). More formally, if  $\mathbf{A}$  is the adjacency matrix of the graph, its powers ( $\mathbf{A}^k$ ) provides the number of paths – directed or not, depending on the graph – of length  $k$  between each pair of nodes. Hence, summing the powers of this matrix gives the total number of paths between each pair of nodes. But to dampen the effect of longer paths, it is common to add weights to the powers of the adjacency matrix. Here, the weights will be the  $\beta^k/k!$ , which allows the sum (power series) to converge

to the exponential of the weighted adjacency matrix as follows:

$$\sum_{k=0}^{\infty} \frac{\beta^k}{k!} \mathbf{A}^k = e^{\beta \mathbf{A}} \quad (2)$$

Each coefficient  $[e^{\beta \mathbf{A}}]_{i,j}$  thus gives the communicability between nodes  $i$  and  $j$ . To find the total communicability centrality of each node, we just sum each row.

In this paper, we will slightly change it: instead of taking  $\beta^k \mathbf{A}^k$ , we will directly take the sparse partial correlation matrix (with its diagonal set to 0 to avoid self-loops): it allows us to take into account the specific partial correlation of each link instead of a general  $\beta$  coefficient.

In addition, taking the sum of the centralities gives the total communicability of the network, allowing to have an overall measure of the ability for information to flow in the system (and see how it evolves over time for example) and even compare different network structures. We also use this as a global measure of integration of the system when we consider the PC graph.

## 5 The reference graphs (static, full sample)

We will first study the case where we calibrate a sparse VAR(1) using all our observations, which will provide us a reference case for future comparisons. We first calibrate on a subset of our variables (spatial dimension: only the front month contract for each of the 17 assets) and then calibrate on our 51 variables (3D). We can thus analyse the two graphs that emerge from this calibration: the graph of pairwise partial correlations of the residuals and the graph of Granger causality.

### 5.1 The partial correlation graph

Since the partial correlation (PC) between two variables measures their dependence conditional on the observation of all the other variables, it filters out their influence. If despite this filtering, some variables are still dependent, it means that they are directly and actually dependent on each other (instead of being correlated because they would both be correlated with another variable).

As explained before, we can assume that most of our markets are not conditionally dependent of many others, that there should be some clustering

(into sectors for example). If we instead observe that markets that should be unrelated are connected in the graph, it may mean that they are influenced by a common factor, and that it may represent a(n) (ephemeral) financial link. Hence, the number of components observed in the estimated graphs will tell us whether the markets under study are integrated (if the markets are connected) or if they are still subject to different fundamentals. In particular, if all the nodes belong to a single component, it means that the system is completely integrated, prone to systematic shocks.

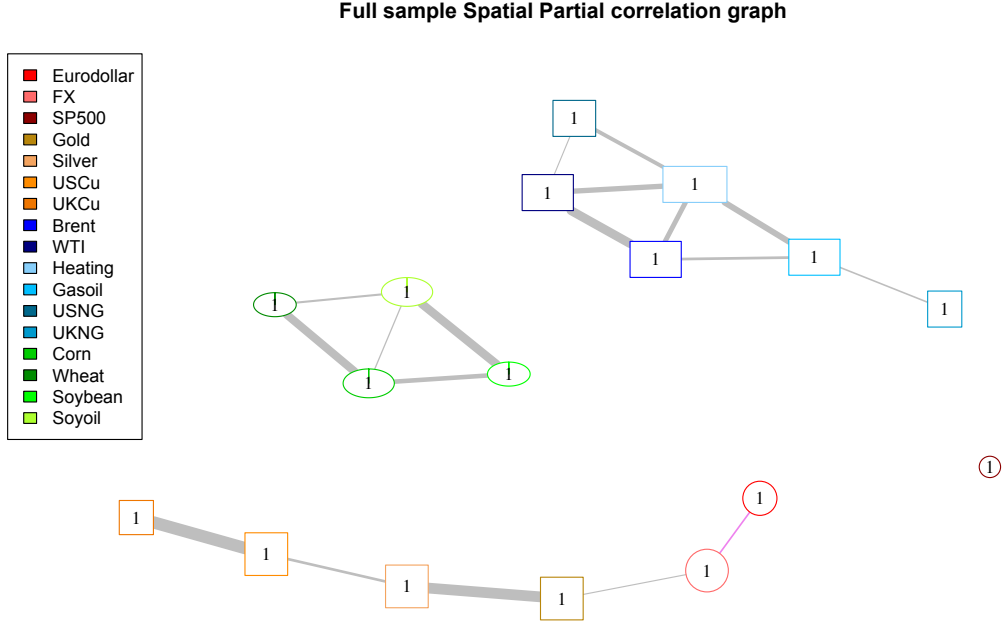
We can refine this analysis by looking at the values of the partial correlations. The average PC will tell whether this integration is strong or not. We can also check the minimum and maximum values of these partial correlations, in order to check their amplitude in the system. We will also assess which markets are the most central in terms of dependence with the others by looking at the total communicability centrality of the nodes.

### 5.1.1 In the spatial dimension

Figure 3 represents the calibrated partial correlation graph in the spatial dimension (only the front-month or spot). Nodes are coloured according to their sector (red for finance, blue for energy, green for agriculture and orange for metals), edges are coloured according to the sign of the partial correlation (grey for positive, violet for negative) and their width represents the absolute value of the coefficient. The PC range from -0.04 to 0.84, with an average of 0.28 (only one link has negative PC). Unconditional correlations range from -0.04 to 0.85, but are not as sparse (74% sparse vs. around 87% for partial correlation matrix, 18 links among the 136 possible ones). We remind that if there exist a path between two nodes (if they belong to the same component) in the PC graph, their unconditional correlation coefficient will be nonzero. We can notice several interesting insights from this graph. The first one is that it consists of three components, while we theoretically have 4 sectors: one cluster for energy markets, one for agricultural markets and one for metals and financial markets.

Lautier and Raynaud (2012) find similar results based on the filtering of the correlation matrix using Minimum Spanning Trees (MSTs), which constrains the graph to be connected (single component) but did not have metals markets. Here, allowing the existence of several components helps us assess whether the integration is systemwide or still "sectorwide". In addition, we allow for cycles here (while the MSTs do not), which tell us precisely which markets are connected with which others, instead of imposing to go through

Figure 3: Partial correlation matrix graph in the spatial dimension



The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details (colours) of each market in the legend. The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represents their degree: the bigger the node, the higher its degree. The width of the edges represents the absolute value of their weight/dependence: the wider the edge, the higher the dependence (in absolute value).

a certain node. The financial and metals sectors are connected through USD/EUR exchange rate and gold market, which can be considered a reserve of value. This component is moreover organised linearly, meaning that there is no "direct" dependence between all these markets, but will seem so in the unconditional correlation matrix, due to the other connections in the component. On the contrary, the other two components are quite connected and reveal actual integration in these sectors.

Finally, we can note that the S&P500 is not connected to others in this graph, meaning that its return innovations are independent of the others. In terms of portfolio management, it means that, as found by Gorton and Rouwenhorst (2006), if we focus only on the front-month contracts, an index made of commodity futures would be uncorrelated to this equity index, and would thus provide diversification benefits.

Table 4: Total communicability centrality of markets (from high to low) in the reference graph of partial correlations (PC) in the spatial dimension

	Centrality
Heating Oil 1	2.75711552
Brent 1	2.64689923
Light Crude 1	2.55987776
US Copper 1	2.52091147
UK Copper 1	2.39170158
Silver 1	2.27630636
SoyBean 1	2.18569361
Corn 1	2.14969241
Gold 1	2.13325709
SoyOil 1	1.97115158
Wheat 1	1.90001059
Gasoil 1	1.88942044
US Natural Gas 1	1.3720924
UK Natural Gas 1	1.04974776
SP500 1	1
USD EUR Fx Rate 1	0.9630333
Eurodollar 1	0.96096818

The degree of the nodes counts the number of direct neighbours that they have. But this is not enough, since indirect connections also matter, as a path in the PC graph means there is a nonzero unconditional correlation between the two extremities of the path. Hence, the total communicability centrality measure accounts for (infinitely) further "neighbours", in the sense that it accounts for every possible path between any two nodes, which are weighted (to dampen the influence of further neighbours). Therefore, if the graph consists of three components, the centrality of the nodes will represent their centrality only in their component.


Table 4 displays the total communicability centrality measure computed on this graph. We note that the heating oil has the largest degree (4) and is also the most central market, followed by the brent and WTI, which have a strong connection. Hence, these markets are the most susceptible to influence the others, but since the sectors are separated, the influence would be limited to energy markets.

In the agricultural sector, while corn has a larger degree (3 vs. 2), soybean is the most central, closely followed by corn. This result shows the importance of accounting for indirect neighbours.

Finally, in the metals/financial component, the copper markets have the most potential for information propagation, which can be explained by their strong link and their further position to the link of negative PC between eurodollar and USD/EUR exchange rate. This link affects their centrality, making them less central than even the S&P500, which does not have neighbours, meaning that they probably receive the shocks rather than drive them.

### 5.1.2 In 3D

Figure 4: Partial correlation graph in 3D



The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details (colours) of each market in the legend. The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represents their degree: the bigger the node, the higher its degree. The width of the edges represents the absolute value of their weight/dependence: the wider the edge, the higher the dependence (in absolute value).

Figure 4 presents the estimated partial correlation graph on the whole sample, in three dimensions. Adding the maturity dimensions, what is striking is that the graph has a single component (all nodes are present) instead

of several ones as in the spatial dimension only. Hence, working on commodities, it seems crucial to include the maturity dimensions, as they are the actual connecting variables. This graph means that all those markets are integrated (prone to systematic shocks) and that diversification benefits are actually much lower... But this integration is quite diverse: PCs range from -0.24 to 0.75 with an average of 0.16 and most of them (72%) are positive. The nets algorithm filtered 88% of the PC coefficients. This PC matrix translates into a full unconditional correlation matrix, with coefficients ranging from -0.82 to 0.99 with an average of 0.10 (only 60% are positive). We see here the effect of other variables on the unconditional correlation coefficients, which have a much larger amplitude than the PCs.

Markets are again clustered into sectors and here, finance and metals are in the middle of energy and agriculture. Their position may make them primary sources of systemic risk as they may drive the returns of the whole system. An exception is the US natural gas markets, which seems to connect only with metals instead of other energy markets. There are some natural connections between natural gas and copper, for example the usage of copper pipes to transport natural gas. There are also some economic ones: construction may require natural gas to heat housing, offices, etc., economic activity requires electricity, which may be produced with natural gas. Nevertheless, considering the seasonality of the natural gas prices, such connections may not hold at a daily frequency, as we can see from the coefficients, which are close to 0 (less than 0.03 in absolute value). It may mean that we need to filter more aggressively.

Table 5 presents the 10 most and least central markets according to the total communicability centrality measure. The most striking result is that the 12-month S&P futures contract (not the spot) is the most central one, confirming the observation of its position in the graph. This result naturally raises concerns of potential shocks led by financial markets in this system; we can thus look into it more particularly. Its maximum PC is around 0.28, so its influence is not over its neighbours but further down the paths. Nevertheless, the differences in centrality values are not very large until we reach the bottom of the ranking (so the system is highly integrated); the influence of the S&P500 futures may not be that much larger than others. In particular, we see that all 4 sectors are represented at the top, with four ranks taken by energy markets, three by metals, two by agricultural markets and only one for financial markets.

Medium- and long-term futures seem to be at the core of the system, conveying and receiving the most information contemporaneously, compared to

Table 5: Communicability centrality of the 10 most (left) and least (right) central markets (from high to low) in the reference graph of partial correlations in 3D

	Centrality		Centrality
SP500 3	2.84667941	Wheat 3	2.4987086
Gasoil 2	2.81147241	Soy Bean 1	2.46805167
Soy Oil 2	2.8034704	US Natural Gas 1	2.46415903
WTI 2	2.77946501	Gasoil 1	2.46101264
Silver 3	2.77499497	US Natural Gas 3	2.16668475
US Copper 2	2.77100391	Eurodollar 1	2.07323956
Brent 2	2.76630783	Eurodollar 3	2.01416571
Soy Bean 2	2.76122455	UK Natural Gas 1	1.98864143
Heating Oil 2	2.76078463	UK Natural Gas 2	1.86925835
Silver 1	2.7540656	UK Natural Gas 3	1.13756682

spot markets. This would mean that futures markets indeed are dominant in terms of price discovery. Only one spot market is in the top 10: silver. Then, investigating the maximum PC of silver contracts, we find that they are close to 0.5, larger than those of the S&P500 futures, but still not the largest ones in the system. In conjunction with degree centrality, the influence of that front-month contract (silver) is not only direct, it can also easily reach infinitely further nodes. Moreover, silver also has another contract in the top, its 12-month one, making it very influential: shocks from either its spot or futures markets could drive other markets and lead to systematic shocks. This central position in the system is economically reasonable and meaningful. Silver is very close to gold as they have similar uses (jewelry, electronics, reserve of value, etc.) and also to industrial metals, as we found in Section 4. Its price is thus related to business cycles, as are those of energy markets for example, making the connection with the other sectors.

If we now look at the least central markets, we find many energy markets, more particularly, most of the natural gas contracts. These natural gas markets are particular as their storage is difficult or costly, not allowing to dampen their seasonality. The links in their close neighbourhood appear thin on the graph, meaning very low PCs, which may explain their low centrality, even though the UK natural gas contracts appear around the "center" of the graph. The same holds for the 12-month eurodollar contract, which even has a quite high degree (13, second after the front-month silver, which has 15) and its PCs range from -0.08 to 0.52 (with 4 negative coefficients). If the focus is only on the transmission of information, it could be more rel-



evant to conduct this analysis with the absolute value of the coefficients to assess the amplitude of transmission instead of allowing for negative coefficients mitigating the influence of positive ones. It is not the purpose of this paper, but needs to be investigated in the future.

Partial correlation graphs are thus very rich in information, both visually and quantitatively. They allow to identify integration/diversification in terms of number of components, number of edges and weights of these edges. Identifying the most central nodes tells us which contracts/markets are the most influential, influenced or informative in the system and comparing the centrality or PCs of contracts from a same market can also hint at the dominant one in terms of price discovery.

Another way to find evidence of information/shock transmission is to study the Granger causality links, indicating a temporal transmission instead of a contemporaneous one.

## 5.2 The Granger causality graph

Since the AR coefficients are directly related to the notion of Granger causality, they first tell us whether some returns are Granger-caused by others (nonzero AR coefficients). They also inform us on the amplitude of this influence (autocorrelation is also allowed), we can see this as a kind of daily propagation of shocks and information.

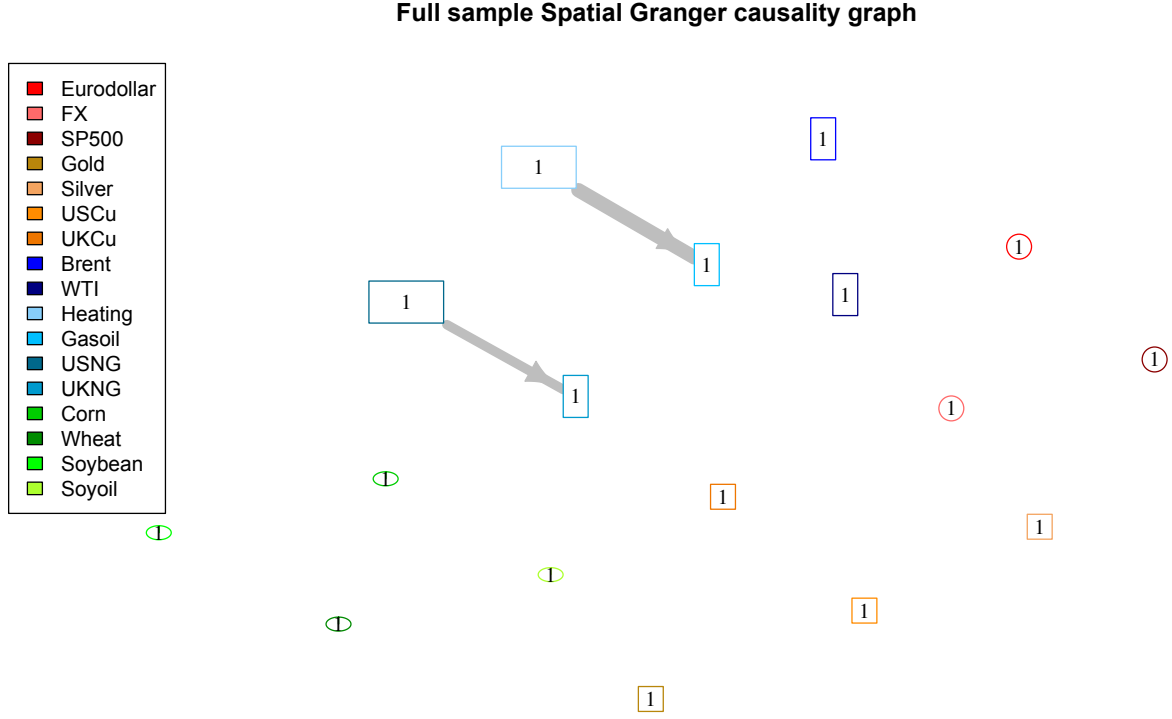
Several measures derived from the graph can be useful in assessing systemic risk (propagation risk). The number of clusters, of nodes and of links will tell us the range of the propagation, whether it is widespread ("marketwide", "sectorwide", systemwide) or if it is contained within a subset of variables. The average AR, minimum AR and maximum AR will tell us the possible amplitude of this propagation, whether it is positive or negative.

In addition, we can study the centrality of the nodes, to quantitatively assess which ones will be propagating information the most (which may need monitoring), and which ones the least, but considering that there are few links, it may not be necessary.

### 5.2.1 In the spatial dimension

It seems like there is not much Granger causality in the spatial dimension. Figure 5 shows the estimated Granger causality graph (from the AR matrix). Only two links, both positive, seem relevant when using our whole database:

Figure 5: Granger causality graph in the spatial dimension



The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details (colours) of each market in the legend. The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represents their degree: the bigger the node, the higher its degree. The width of the edges represents the absolute value of their weight/dependence: the wider the edge, the higher the dependence (in absolute value).

one from heating oil to gasoil (for 0.18) and one from US natural gas to UK natural gas (0.11).

The latter link is natural, as it would reflect the difference of trading hours (the US market closing after the UK market, information of the day in the US market is incorporated the following day in the UK market). The inverse information transmission is not seen, as it should appear in the contemporaneous correlations instead of here. The former link would be reflecting the same phenomenon (gasoil being quoted in EU). What is strange is that it does not show for copper (seen in Section 4) or oil markets. For the latter though, there has been a decoupling in the recent years, which may explain the absence of propagation.

Naturally, if we look at the centrality of this graph, the nodes influencing the

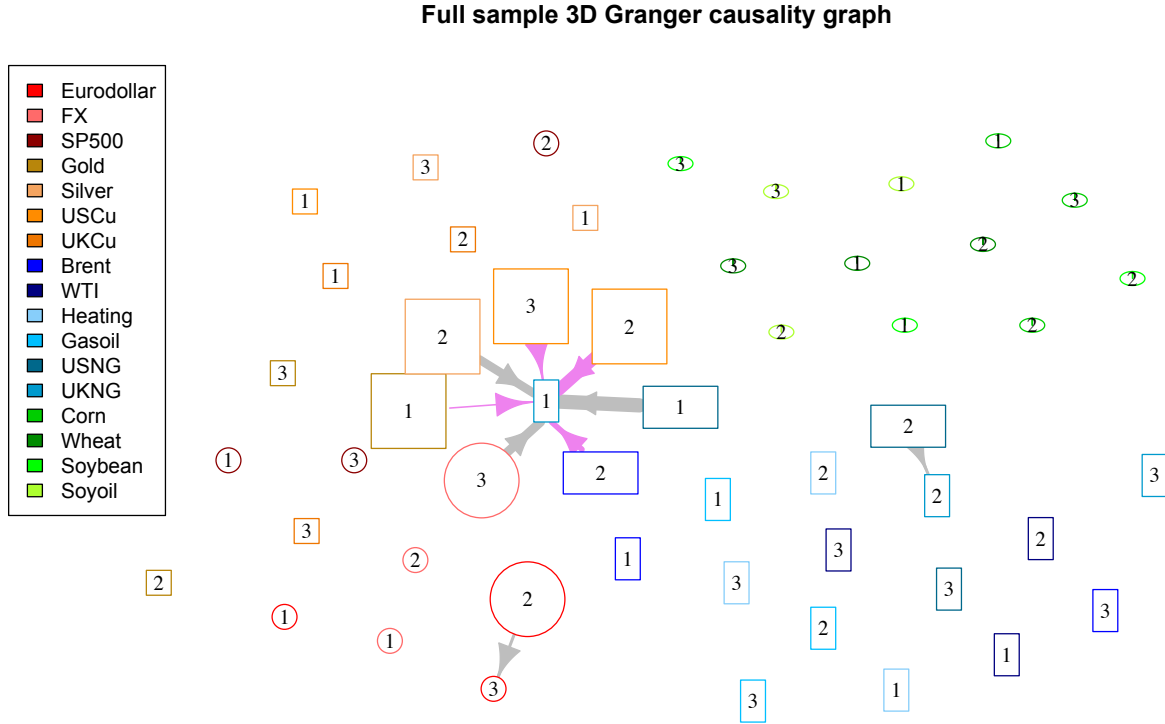
others are the most central ones, i.e. heating oil and US natural gas here, though the influence is limited.

Thus, there is not much propagation risk in the spatial dimension at a daily frequency, meaning that instead most of the dependence occurs contemporaneously, as integration, or maybe that it occurs at longer lags.

Let us check what happens when we incorporate the maturity dimensions too.

### 5.2.2 In 3D

Figure 6: Granger causality graph in 3D



The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details (colours) of each market in the legend. The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represents their degree: the bigger the node, the higher its degree. The width of the edges represents the absolute value of their weight/dependence: the wider the edge, the higher the dependence (in absolute value).

When including maturities in the system, the resulting Granger causality graph in Figure 6 remains quite simple but adds interesting features to

the spatial dimension. From that dimension, only the link from US natural gas to UK natural gas remains. Several other contracts join in and influence this spot UK natural gas. Its neighbours are: spot US natural gas (0.1210), spot gold (-0.0023), 3-month US copper (-0.1004), 3-month silver (0.0686), 3-month brent (-0.0404), 12-month US copper (-0.0077) and 12-month USD/EUR exchange rate (0.0724). Though their influence is not homogeneous (some have a negative influence, others have a positive one), it is relatively low, the most important one in absolute value still being from the US natural gas (0.1210). So here again, propagation risk is limited too. But why would metals futures markets influence natural gas spot markets the following day? There does not seem to be a reasonable explanation for these daily propagation links, except for US natural gas to UK natural gas, though some of these metals are used in natural gas transportation, hence information of future pipe or building constructions may be incorporated into spot natural gas prices...

We also find more economically sound links, from the 3-month US natural gas to 3-month UK natural gas and from 3-month eurodollar to 12-month eurodollar, but both influences are very limited (0.0155 and 0.0158 respectively).

Finally, in terms of centrality, as in the spatial dimension, the nodes influencing the others are the most central ones and here the UK natural gas is obviously the one receiving the most influence.

We will then analyse these different measures dynamically, around several events of interest, but will also briefly look at the overall picture (their evolution over the years).

## 6 Shocks and propagation, dynamic analysis

To conduct a dynamic analysis, our only option is to have recourse to rolling windows. We thus use a window of 252 observations, corresponding to approximately 1 trading year. We computed all the above mentioned measures, rolling over our whole sample and provide the most interesting ones in Appendix C, giving the bigger picture for analysing where the events lie. In addition to the previous measures, in a dynamic setting, we can also assess the stability of the graphs by looking at the survival ratio of the graph, compared to the reference graphs or the graphs of the previous day for example (see Figure 13 in Appendix C).

To give an overview, before 2004, there was basically no propagation in the system. In terms of integration, the number of components varied, but was over 4 (the number of sectors) and as high as 12 (having 17 markets). This would indicate that most markets were not even clustered into sectors, components being potentially futures with the same underlying or very similar ones (e.g. brent and WTI).

After the beginning of the financialisation, we find increasing propagation and integration, with more links and larger averages in both cases. There also seems to be a structural break around the beginning of the subprime crisis, leading to an accelerated integration, which peaked shortly at the end of 2008. The single component in the PC graph indicates that the system is fully integrated and markets could have potentially experienced simultaneous shocks...

We retain several dates to analyse: from the data, we retain the minimum and maximum observed return for the spot S&P500 index, respectively 2008-10-09 and 2008-10-28 (because on the day right after each observation, many commodities experienced extremum return too). We will also study what happens around the beginning of the subprime crisis (2007-08-09), the default of Lehman Brothers (2008-09-15) and the Flash Crash (2010-05-06) since they are financial events that may have affected the system.

Please note that for the animations to play, you need to use Adobe Acrobat Reader or similar software (they do not work with Preview on Mac OS).

### 6.1 The beginning of the subprime crisis (2007-08-09)

One of the dates that have been identified as the beginning of the subprime crisis is August 9, 2007. On that day, BNP Paribas stopped the valuation

(and subscriptions and redemptions) of three of its funds due to "the complete evaporation of liquidity in certain market segments of the U.S. securitization market" (NYT 2007-08-09)<sup>1</sup>. We will thus look at how our graphs (Figures are available upon request, due to their size) and measures (Appendix C) behave around that date.

As in the Granger reference graphs, the UK natural gas receives information from metals markets (which is strange) and from the US natural gas, though there are only 4 of the 7 (reference) relevant links to UK natural gas. The front-month soybean contract oddly joins on 2007-08-13 and the silver contract goes away on 2007-08-21, after central banks have increased their support and lowered their rates. Again, most of the Granger causality links are weak and do not have reasonable economic foundation, hence may not be relevant, except for regulation.

In the PC graph, markets are highly integrated, forming almost a single component. They are also clustered into sectors overall, with metals markets in the middle and with the exception of the gasoil, US natural gas and eurodollar (the S&P500 also has an unstable connection). The US natural gas is a peculiar market, as evidenced by Lautier and Raynaud (2012). With the addition of metals markets compared to their work, these natural gas contracts behave erratically and seem to separate this market from other energy markets. That market connects rather weakly to silver, gold or copper markets, which may be natural due to the use of these metals in some natural gas transportation systems. Interestingly, the gasoil contracts followed the same behaviour and even became temporarily segregated from the large component.

Another particularity is that the front-month eurodollar contract is not part of this graph around this event and other eurodollar contracts were even separated from all other contracts on 2007-07-25. Around that time, the short-term rate was constant, explaining the conditional independence with all the other variables in the system. The range of partial correlations fluctuates quite a lot until one week before the event. Afterwards, there is a relatively large drop (-0.2 to -0.3) in the minimum partial correlation from 2007-08-20 to 2007-08-22, while the gasoil market is living on its own and central banks were supporting financial markets.

Around the beginning of the subprime crisis, however, the gold contracts have taken the role of most connected nodes (instead of silver). Since they

---

<sup>1</sup><http://www.nytimes.com/2007/08/09/business/worldbusiness/09iht-09bnp.7054054.html>

play a role of reserve of value, we naturally find that 60-75% of their links are negative partial correlations. Their PCs are in general very low in absolute value, except for those with other metals. Silver markets are not as important in terms of reserve of value, but still have many negative partial correlations with other markets (81% and 83% of their links for the 3-month and 12-month respectively). Nevertheless, looking at the rankings provided by the total communicability centrality, the 12-month S&P500, which seems peripheral, was the second most important contract on 2007-07-25 and became first the following day. Other important contracts are naturally mainly metals, at the center of the graph, of the economy (industrial metals) and due to their role as reserve of value. In terms of total communicability, the trend is negative around the beginning of the subprime crisis, but looking at the bigger picture in Appendix 12, it explodes very early after. At that time, markets were highly complex, with intricate instruments, connections, etc. Few people anticipated how the markets would move and it took time for agents to figure out what was happening and it may explain the lag between the event and the subsequent burst in integration.

Overall, we do not find evidence of substantial systemic risk. Nevertheless, some indicators point to potential vulnerability to systematic shocks coming from stock markets: the presence of a large component in the PC graph, meaning that the system is integrated, partial correlations ranging from -0.3 to more than 0.9 and the S&P500 being the most central contract in the system...

## **6.2 The default of Lehman Brothers (2008-09-15)**

Lehman Brothers, one of the most important derivative dealers at that time, faced tremendous difficulties, until it had to default on 2008-09-15. This event has triggered the global financial crisis, which has spread to economies worldwide and thus, becoming a systemic event. Many interconnections were neglected, leading to largely unexpected losses for many entities in different sectors and huge commodity price drops. We thus want to analyse what happened around that date and other subsequent events.

We find overall two components in the Granger causality graph: one revolving around the UK natural gas, as usual, and one around an agricultural market (wheat or soybean). The latter component is dominated by the 12-month gold contract, which positively influences the 12-month wheat until Lehman Brothers defaults on 2008-09-15, to then turn to the front-

Figure 7: Granger causality (above) and partial correlation (below) **animated** graphs on the day of default of Lehman Brothers

The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details of each market on the figure (different colours). The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represent their degree: the bigger the node, the higher its degree. The width of the edges represent the absolute value of their dependence: the wider the edge, the higher the dependence (in absolute value).



month soybean contract (negative influence). There is no explanation for this Granger causality, which thus may be due to financial activity. From 2008-09-19 on, the 12-month eurodollar becomes positively influenced by this same gold contract. Financial agents may have been looking for safety in the longer contracts at that time, investing in both reserves of value. The former component sees the WTI, Brent and/or silver (various maturities) influence the UK natural gas. The links to the WTI and silver are natural due to time difference and their economic relationships. The connection to the Brent can also be related to their economy, but the time lag (one day) and maturities (12-month Brent and front-month UK natural gas) do not really make sense and may be caused by financial interferences.

The PC graph still has a single huge component, but sectors are all scattered and with a slightly smaller range of partial correlations (from -0.3 to less than 0.9). The maximum partial correlation seemed to decrease steadily, but skyrockets back on 2008-09-23, just one day after the G7 commits to protect the financial system. This announcement may have reassured financial agents that the crisis may then be contained and may have sustained their risky activities.

In terms of centrality, gold, silver, WTI and USD/EUR exchange rate have the highest degrees and are at the center of the graph. The total communicability centrality gives overall the same results, but also allows to see the substantial change that occurred on 2008-09-19. On that day, the 3-month maturity eurodollar contract went from 26<sup>th</sup> to 2<sup>nd</sup> most important contract (after the 12-month S&P500) and total communicability of the network spiked, so information could flow more easily. By acting on this contract, regulators may have been able to mitigate the effect of the default of Lehman Brothers and its subsequent cascade effects. In addition, the front-month USD/EUR exchange rate contract also became quite important in the system (from 28<sup>th</sup> on 2008-09-11 to 9<sup>th</sup> on 2008-09-19), showing that exchanges between countries (capital reallocations) may have driven the returns in the system and may even have propagated the crisis. Regulators could also have intervened on these currency markets to try to avoid too much reallocation. Other contracts got to lower ranks accordingly, but the two shortest maturities of gold contracts in particular became less important (even if the second one has the highest degree around the end of the study period). This may not be that counterintuitive, as gold would be anticorrelated to procyclical assets and relevant as an alternative mostly for financial investors.

All in all, markets are highly integrated (they form a single component in the partial correlation graph). There may have been capital flows between

countries and/or assets at play around the demise of Lehman Brothers, which may have spread the shock. Nevertheless, there is no significant sign of potential for daily propagation, since the number of Granger causality links and their coefficients are low.

### **6.3 The lowest return on S&P500 (2008-10-09) (overlapping with Lehman)**

We chose to study this event, because we noticed in the data that all the S&P500 contracts experienced their minimum return (going down by at least -7.23%) on that day and that some commodities experienced their minimum (or very infrequent) return the following day, on 2008-10-10. 33 of our 51 contracts had 1% negative shocks (observed), among which 10 are their minimum in our sample: USC<sub>u1</sub>, USC<sub>u3</sub>, UKC<sub>u2</sub>, UKC<sub>u3</sub>, H.O.1, Brent1, Brent2, WTI2, Soyoil1 and Soybean3. We thus wanted to check whether this was due to propagation from the S&P500 or not.

Unfortunately, the S&P500 does not appear in the Granger causality graphs around that day, so it does not explain what was observed. Nevertheless, there are more propagation links on 2008-10-09 and 2008-10-10 than before. The link from 12-month gold to 12-month eurodollar is still here and its coefficient increased to 0.20 (4 times larger than the maximum on the day before...). In addition, the 3-month UK copper joins to influence the 12-month eurodollar contract negatively (-0.11) and the 3-month eurodollar one positively (0.02). These low coefficients and their later decrease in absolute value are again pointing to potentially financial links.

The system still forms a single component in the PC graph and sectors are still scattered, but the USD/EUR exchange rate has made its way to the center of the graph (and the front-month made it to the top of the centrality ranking), with silver and gold. As seen for Lehman Brothers, this may mean that capital flows have driven most of the returns, with these two precious metals being reserves of value. Again, the 12-month S&P500 was also among the most important contracts, but not the spot. They had a 0.6076 PC though, so they still could have been the source of the shocks on that day. The following day, 2008-10-10, it was slightly lower, at 0.6010, but can still explain why many commodities experienced their minimum return, as the S&P500 crashed again. Coefficients still range from about -0.2 to 0.9, but with a temporary drop in the minimum starting on the day of the event

Figure 8: Granger causality (above) and partial correlation (below) **animated** graphs on the day of lowest return of the S&P500

The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details of each market on the figure (different colours). The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represent their degree: the bigger the node, the higher its degree. The width of the edges represent the absolute value of their dependence: the wider the edge, the higher the dependence (in absolute value).

(2008-10-09). Finally, total communicability dropped on 2008-10-09 and recovered slightly the following day.

In the end, we find that in those troubled times, markets have been largely changing (as seen with the centrality measure), but the lowest return observed on the S&P500 on 2008-10-09 did not affect those of commodities the following day. The explanation/hypothesis may be instead that the following crash of the S&P500 on 2008-10-10 has driven the returns in the system, as the 3-month contract on S&P500 seems to be an important source of information in the system (very central).

#### **6.4 The largest return on S&P500 (2008-10-28) (overlapping with minimum S&P500)**

The S&P500 maximum returns occurred (on 2008-10-28, up by more than 10.24%) in expectation of imminent rate cuts by central banks. On the following day (2008-10-29), 35 contracts experienced 1% positive shocks (observed), among which USC<sub>1</sub>, USC<sub>2</sub>, UKC<sub>1</sub>, UKC<sub>2</sub>, UKC<sub>3</sub>, Silver<sub>1</sub>, Silver<sub>2</sub>, Silver<sub>3</sub>, Soy<sub>oil1</sub>, Soy<sub>oil2</sub>, Soy<sub>oil3</sub>, Wheat<sub>1</sub>, Soybean<sub>2</sub> and Corn<sub>3</sub> experienced their maximum return over our sample. Almost all markets experienced important returns (positive or negative), except for the UK natural gas. In particular, the returns on S&P500 were this time (2008-10-29) in the 10% bottom quantile, making the hypothesis of it driving the others not possible in this case. In addition, returns on the eurodollar and gold contracts were in the 5% quantile, respectively bottom and top, which also diverges from their positive relationship in the previous case. We thus wanted to assess whether the shock financial equities propagated this time to those commodities and other assets.

Figures are available upon request, due to the large size of this file. Again, in terms of propagation, the S&P500 does not appear in the Granger causality graph. And around this event, the graph is rather stable, with a negative link from front-month silver to 12-month eurodollar appearing and disappearing. The average coefficient was slightly positive but became negative after 2008-10-23 (never more than 10% in absolute value), with very little change. The propagation links do not have any meaningful economic explanation and would thus again be financial links.

As for the PC graphs, markets are still integrated (single component) and seem to move around in the graph, with a widely spread energy sector (crude oils stand around the middle and meddle with other sectors). For example, on the day of the event (2008-10-28), all sectors but energy seem well separated. While the S&P500 nodes seem at the periphery of the graph, its 12-month contract is still among the most central (though it is only connected to S&P500 contracts and with PCs between 0.5 and 0.56), with silver, gold and the exchange rate. But on the following day, when some commodities experience their maximum return, all sectors are intertwined and identifying the center of the graph is more difficult. Looking at the total communicability centrality measure, we find that precious metals and exchange rate markets have been very important on that day and experienced large returns, which, instead of equity markets, may have driven the returns on other commodities. The communicability in the graph indicates that integration has been increasing in trend over this period, probably due to the support announcements by regulators and capital flows to safe assets.

To summarise, the highest return observed in our sample for the S&P500 does not seem to contribute to the highest return observed for several commodities on the following day through propagation. Instead, we explain the positive shocks on 2008-10-29 on commodity prices by the returns on precious metals and exchange rate and their role in seeking safety during crises.

## 6.5 The Flash Crash (2010-05-06)

The Flash Crash corresponds to a financial event occurring on the e-mini S&P500 futures (that we have in our sample), at that time for maturity June 2010 (our SP500 2). There has been a sell order for 75,000 contracts, corresponding to around 4 billion dollars, which has triggered many reactions from other trading algorithms. Most of the losses were recovered quickly, but some equities were still impacted even after the end of the day. We thus examine whether this few-hour financial "shock" could have had some impact on the real economy, through commodity futures.

On the day of the Flash Crash, the Granger causality graph has many more links than during our other events of interest and looks quite different from our reference graph. The S&P500 is still not present on the following day, so it may not have triggered propagation to commodities. Though this event occurred on the front-month E-mini S&P500 futures contract, its daily return on that day was -3.63%, while those of the front-month and 3-month

Figure 9: Granger causality (above) and partial correlation (below) **animated** graphs on the day of the Flash Crash

The shape of the nodes represents their sector: circles for Finance, rectangles for Energy, ellipses for Agriculture and squares for Metals, with details of each market on the figure (different colours). The colour of the edges represents the sign of the weight/dependence: grey for positive, violet for negative. The size of the nodes represent their degree: the bigger the node, the higher its degree. The width of the edges represent the absolute value of their dependence: the wider the edge, the higher the dependence (in absolute value).

eurodollar contract were more than 30% (the two interest rates increased from 5.25 bp to 7.10 bp). These returns do not seem to come from propagation, since other markets contributed only +0.009% and +0.039%. The 12-month eurodollar (which receives the most from others) seems to behave differently as we observe a 1.81% return, while the observations of the previous day contributed -0.86%. Propagation risk has thus increased substantially on the day of the Flash Crash, but to a long-term eurodollar contract, which may not present too much direct importance in terms of systemic risk.

In terms of partial correlations, the corresponding graph still consists of a single component, but there are many more links than the other events, as can be seen in Figure 12. The metals and financial sectors are at the center of the graph and intertwined, with gold, silver and exchange rate being the most connected markets. Energy markets are also scattered and among the most central nodes (still with the 12-month S&P500 at the top of the ranking), while the agricultural sector is still clustered. Before the day of the event, the number of links in the graph increases, being mostly negative partial correlation links. It hence makes the average partial correlation decrease, but we also find a large increase in the maximum partial correlation on the eve of the Flash Crash (from 0.81 for Brent1-Brent2 to 0.93 for WTI1-WTI2), which could raise concerns about simultaneous shocks again since they are among the most central nodes.

Around the Flash Crash, markets seem to have integrated a lot as sectors are not as clearly separated as before in the PC graph. This goes with a large increase of the range of dependence, which in turn increases the risk of systematic shock. In addition, relatively many propagation links exist at that time, also largely increasing their influence in absolute value, which means that propagation risk is also larger, but did not seem to present a substantial risk.

## 7 Conclusion

The evolution of commodity prices since the early 2000s, raised concerns, since they seem closely related to equity prices or indices. The consequences of the financialisation of commodity markets could be an increased influence of financial agents on commodity futures prices, which may in turn affect commodity spot prices. Since these (futures and spot) commodities are used by individuals in their daily life, by firms in their economic activity or by

financial agents in their portfolios, the risk of affecting their prices would affect the real economy, hence falling under the scope of systemic risk.

The literature on the financialisation of commodities has largely studied the (contemporaneous) correlations between commodity returns and financial asset returns. However, there may be two temporalities for the influence/dependence: contemporaneous (integration) or over time (propagation). For example, we observe that some commodity markets experienced their extremum return on the day after the S&P500 experienced its extremum over our sample (which occurred in October 2008, after the default of Lehman Brothers).

A VAR framework is thus intuitive and suited for distinguishing propagation and integration. We thus rely on the algorithm of Barigozzi and Brownlees (2017), which estimates the sparse AutoRegression (AR, Granger causality) and Partial Correlation (PC) matrices in a single step, providing substantial advantages. The former matrix gives the propagation structure in the system while the second gives the integration structure. Both can be represented as meaningful graphs and can be quick risk inspection tools.

We claim that the maturity dimension (using futures of different maturities) is necessary for assessing systemic risk as futures markets are a place for price discovery, according to the literature. We indeed find in our spatial reference graph that when the maturity dimension is excluded, the different sectors are not connected, i.e. they do not exhibit common factors (in particular, financial markets and all commodity markets). On the contrary, once we include the 3-month and 12-month contracts in our 3D reference graph, all the sectors connect: they are highly integrated and thus present higher risk. To study propagation, in particular the aforementioned dates of extrema, we conduct a dynamic analysis using a one-trading-year rolling window. We find that there was little propagation at play and that the S&P500 contracts do not even appear in the Granger causality graph. The extrema observed for some commodities the day after the extrema of the S&P500 are thus not due to propagation. Instead, integration was dominant, but the S&P500 is only responsible for the negative shocks as its contracts were among the most central ones and they experienced negative shocks on day after their extremum. On the contrary, we find that after the day of their maximum return, they did not experience large shock, but a small negative shock. This, in addition to the centrality and positive shocks on precious metals and exchange rate, leads to the conclusion that the search for safety drove those maximum returns.



## References

- Acharya, Viral V. et al. (2017). “Measuring Systemic Risk”. In: *The Review of Financial Studies* 30(1), pp. 2–47.
- Adrian, Tobias and Markus K. Brunnermeier (2016). “CoVaR”. In: *American Economic Review* 106(7), pp. 1705–41.
- Barigozzi, Matteo and Christian Brownlees (2017). “Nets: Network Estimation for Time Series”. In: *Working Paper*.
- BCBS (2018a). *Basel III: Finalising post-crisis reforms*. Tech. rep. Basel Committee on Banking Supervision, Bank for International Settlements.
- BCBS (2018b). *Global systemically important banks: revised assessment methodology and the higher loss absorbency requirement*. Tech. rep. Basel Committee on Banking Supervision, Bank for International Settlements.
- Benoit, Sylvain, Jean-Edouard Colliard et al. (2017). “Where the Risks Lie: A Survey on Systemic Risk”. In: *Review of Finance* 21(1), pp. 109–152.
- Benoit, Sylvain, Christophe Hurlin and Christophe Perignon (2015). “Implied Risk Exposures”. In: *Review of Finance* 19(6), pp. 2183–2222.
- Benzi, Michele and Christine Klymko (2013). “Total communicability as a centrality measure”. In: *Journal of Complex Networks* 1(2), pp. 124–149.
- BIS (2001). *Report on Consolidation in the Financial Sector*. Tech. rep. Bank for International Settlements.
- BIS, FSB and IMF (2009). *Guidance to Assess the Systemic Importance of Financial Institutions, Markets and Instruments: Initial Considerations*. Tech. rep. Bank for International Settlements, Financial Stability Board and International Monetary Fund.
- Brennan, Michael J (1958). “The supply of storage”. In: *The American Economic Review* 48(1), pp. 50–72.
- Büyüksahin, Bahattin and Michel A. Robe (2014). “Speculators, commodities and cross-market linkages”. In: *Journal of International Money and Finance* 42, pp. 38–70.
- Capponi, Agostino and Martin Larsson (2015). “Price Contagion through Balance Sheet Linkages”. In: *The Review of Asset Pricing Studies* 5(2), pp. 227–253. DOI: 10.1093/rapstu/rav006. eprint: /oup/backfile/content\_public/journal/raps/5/2/10.1093/rapstu\_rav006/2/rav006.pdf. URL: <http://dx.doi.org/10.1093/rapstu/rav006>.
- Cheng, Ing-Haw and Wei Xiong (2014). “Financialization of commodity markets”. In: *Annu. Rev. Financ. Econ.* 6(1), pp. 419–441.
- Deaton, Angus and Guy Laroque (1992). “On the behaviour of commodity prices”. In: *The Review of Economic Studies* 59(1), pp. 1–23.
- Diebold, Francis X., Kamil Yilmaz and Laura Liu (2017). “Commodity Connectedness”. In: *NBER, Working Paper No. 23685*.

- Ewing, Bradley T, Farooq Malik and Ozkan Ozfidan (2002). “Volatility transmission in the oil and natural gas markets”. In: *Energy Economics* 24(6), pp. 525–538.
- Garbade, Kenneth D. and William L. Silber (1983). “Price Movements and Price Discovery in Futures and Cash Markets”. In: *The Review of Economics and Statistics* 65(2), pp. 289–297.
- Gorton, Gary B and K. Geert Rouwenhorst (2006). “Facts and Fantasies about Commodity Futures”. In: *Financial Analysts Journal* 62(2), pp. 47–68.
- Hammoudeh, Shawkat and Yuan Yuan (2008). “Metal volatility in presence of oil and interest rate shocks”. In: *Energy Economics* 30(2), pp. 606–620.
- Hasbrouck, Joel (2003). “Intraday Price Formation in U.S. Equity Index Markets”. In: *The Journal of Finance* 58(6), pp. 2375–2400.
- Hautsch, Nikolaus, Julia Schaumburg and Melanie Schienle (2015). “Financial Network Systemic Risk Contributions”. In: *Review of Finance* 19(2), pp. 685–738.
- Jain, Pankaj K., Pawan Jain and Thomas H. McInish (2016). “Does high-frequency trading increase systemic risk?” In: *Journal of Financial Markets* 31, pp. 1–24.
- Kaldor, Nicholas (1939). “Speculation and Economic Stability”. In: *The Review of Economic Studies* 7(1), pp. 1–27.
- Lautier, Delphine and Franck Raynaud (2012). “Systemic Risk in Energy Derivative Markets: A Graph-Theory Analysis”. In: *The Energy Journal* 33(3), pp. 215–239.
- Lohre, Harald, Jochen Papenbrock and Muddit Poonia (2014). “The use of correlation networks in parametric portfolio policies”. In: *Working paper*.
- Lutz, Björn (2010). *Pricing of Derivatives on Mean-Reverting Assets*. Vol. 630. Lecture Notes in Economics and Mathematical Systems. Springer-Verlag Berlin Heidelberg.
- Park, Jungwook and Ronald A. Ratti (2008). “Oil price shocks and stock markets in the U.S. and 13 European countries”. In: *Energy Economics* 30(5), pp. 2587–2608.
- Schwarz, Thomas V. and Andrew C. Szakmary (1994). “Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis”. In: *Journal of Futures Markets* 14(2), pp. 147–167.
- Silvennoinen, Annastiina and Susan Thorp (2016). “Crude Oil and Agricultural Futures: An Analysis of Correlation Dynamics”. In: *Journal of Futures Markets* 36(6), pp. 522–544.
- Tang, Ke and Wei Xiong (2012). “Index Investing and the Financialization of Commodities”. In: *Financial Analysts Journal* 68(6), pp. 54–74.

- Telser, Lester G (1958). “Futures trading and the storage of cotton and wheat”. In: *Journal of Political Economy* 66(3), pp. 233–255.
- Working, Holbrook (1949). “The Theory of Price of Storage”. In: *The American Economic Review* 39(6), pp. 1254–1262.
- Zhang, Yue-Jun et al. (2008). “Spillover effect of US dollar exchange rate on oil prices”. In: *Journal of Policy Modeling* 30(6), pp. 973–991.

# Appendices

## A Correlations vs. partial correlations

When we want to assess the dependence between two random variables, the linear correlation coefficient is commonly used, even though it has flaws. One of them is that it may be inflated if we do not take into account variables that may have an effect on both variables under consideration. More formally, if we take a set of three random variables  $X, Y, Z$ , we can compute their unconditional correlation coefficients, in their unconditional correlation matrix  $\rho$ :

$$\rho = \begin{pmatrix} 1 & \rho_{X,Y} & \rho_{X,Z} \\ \rho_{X,Y} & 1 & \rho_{Y,Z} \\ \rho_{X,Z} & \rho_{Y,Z} & 1 \end{pmatrix} \quad (3)$$

If we want to assess the dependence between  $X$  and  $Y$  and just look at  $\rho_{X,Y}$ , we may get a biased result. Indeed, what if  $Z$  is a common factor for these two variables or an intermediary variable between the two? It would create a "artificial" unconditional correlation  $\rho_{X,Y}$ . Hence, we need to filter out the influence of the variable  $Z$  from the dependence between  $X$  and  $Y$ . The partial correlation  $C_{X,Y}$  between the variables  $X$  and  $Y$  precisely serves this purpose: it corresponds to the linear correlation between these two variables, conditional on the other variables, here  $Z$ . We would thus get the following partial correlation matrix  $\mathbf{C}$ :

$$\mathbf{C} = \begin{pmatrix} 1 & C_{X,Y} & C_{X,Z} \\ C_{X,Y} & 1 & C_{Y,Z} \\ C_{X,Z} & C_{Y,Z} & 1 \end{pmatrix} = \begin{pmatrix} 1 & \rho_{(X,Y)|Z} & \rho_{(X,Z)|Y} \\ \rho_{(X,Y)|Z} & 1 & \rho_{(Y,Z)|X} \\ \rho_{(X,Z)|Y} & \rho_{(Y,Z)|X} & 1 \end{pmatrix} \quad (4)$$

where  $|A$  denotes the conditionality on variable  $A \in \{X, Y, Z\}$ . This partial correlation matrix can be easily obtained from the unconditional correlation matrix: we take the inverse of the unconditional correlation matrix  $\rho$  (or covariance matrix), then normalise it (by dividing each term by the corresponding diagonal terms) and finally multiply by -1 as follows:

$$\mathbf{C} = - \begin{pmatrix} \frac{D_{X,X}}{D_{X,X}^{1/2} D_{X,X}^{1/2}} & \frac{D_{X,Y}}{D_{X,X}^{1/2} D_{Y,Y}^{1/2}} & \frac{D_{X,Z}}{D_{X,X}^{1/2} D_{Z,Z}^{1/2}} \\ \frac{D_{X,Y}}{D_{X,X}^{1/2} D_{Y,Y}^{1/2}} & \frac{D_{Y,Y}}{D_{Y,Y}^{1/2} D_{Y,Y}^{1/2}} & \frac{D_{Y,Z}}{D_{Y,Y}^{1/2} D_{Z,Z}^{1/2}} \\ \frac{D_{X,Z}}{D_{X,X}^{1/2} D_{Z,Z}^{1/2}} & \frac{D_{Y,Z}}{D_{Y,Y}^{1/2} D_{Z,Z}^{1/2}} & \frac{D_{Z,Z}}{D_{Z,Z}^{1/2} D_{Z,Z}^{1/2}} \end{pmatrix} \quad (5)$$

$$= \begin{pmatrix} -1 & -\frac{D_{X,Y}}{D_{X,X}^{1/2} D_{Y,Y}^{1/2}} & -\frac{D_{X,Z}}{D_{X,X}^{1/2} D_{Z,Z}^{1/2}} \\ -\frac{D_{X,Y}}{D_{X,X}^{1/2} D_{Y,Y}^{1/2}} & -1 & -\frac{D_{Y,Z}}{D_{Y,Y}^{1/2} D_{Z,Z}^{1/2}} \\ -\frac{D_{X,Z}}{D_{X,X}^{1/2} D_{Z,Z}^{1/2}} & -\frac{D_{Y,Z}}{D_{Y,Y}^{1/2} D_{Z,Z}^{1/2}} & -1 \end{pmatrix} \quad (6)$$

where  $\mathbf{D} = \rho^{-1}$  (or the inverse of covariance matrix). The partial correlation matrix is to the concentration matrix (inverse of the covariance matrix) what the correlation matrix is to the covariance matrix. We can also obtain it by regressing each variable on the others and after some computations on the regression coefficients.

Let us take the three eurodollar contracts from our 3D reference graph (Figure 4) as an example. If we compute the correlation matrix  $\rho$  for these three time series (IR1, IR2, IR3), we get very high correlation coefficients:

$$\rho = \begin{pmatrix} 1 & 0.9994 & 0.9424 \\ 0.9994 & 1 & 0.9533 \\ 0.9424 & 0.9533 & 1 \end{pmatrix} \quad \mathbf{\Gamma} = \begin{pmatrix} 0.0003615 & 0.0003557 & 0.0002504 \\ 0.0003557 & 0.0005651 & 0.0004958 \\ 0.0002504 & 0.0004958 & 0.0011149 \end{pmatrix} \quad (7)$$

We note that the highest coefficients are, in descending order, the one between IR1 and IR2 (0.9994), the one between IR2 and IR3 (0.9533) and finally the one between IR1 and IR3 (0.9424). These values may be inflated by the influence of the variables not involved in each pair. Let us denote  $\mathbf{F}$  the concentration matrix (inverse of the covariance matrix  $\mathbf{\Gamma}$ , also called information or precision matrix). If we look at the partial correlations by normalising the concentration matrix, we find the following:

$$\mathbf{C} = \begin{pmatrix} -1 & 0.7403 & -0.1716 \\ 0.7403 & -1 & 0.5172 \\ -0.1716 & 0.5172 & -1 \end{pmatrix} \quad \mathbf{F} = \mathbf{\Gamma}^{-1} = \begin{pmatrix} 7838.30 & -5507.77 & 632.03 \\ -5507.77 & 7061.22 & -1807.80 \\ 632.03 & -1807.80 & 1730.46 \end{pmatrix} \quad (8)$$

We see that after conditioning on all the other variables in our system (of 51 variables), the linear dependence between IR1 and IR2 is only  $C_{IR1,IR2} = 0.7403$  now instead of  $\rho_{IR1,IR2} = 0.9994$ . Hence, the difference of 0.2591 was

due to the influence of other variables. The same applies to the other coefficients, but we interestingly find that the partial correlation between IR1 and IR3 is actually negative,  $C_{IR1,IR3} = -0.1716$ , while the unconditional one was  $\rho_{IR1,IR3} = 0.9424$ . Therefore, if we do not control for the influence of other variables that may alter the unconditional correlation matrix, we may get different results.

In addition, the concentration matrix  $\mathbf{F}$  encodes the conditional dependence between pairs of variables in a multivariate normal distribution. Hence, an entry  $C_{i,j}$  of the partial correlation matrix (which has the same 0s as the information matrix) is equal to 0 if and only if  $X_i$  and  $X_j$  are independent, conditional on  $\{X_k\}_{k \neq i,j}$ . Though the multivariate Gaussian distribution is a strong assumption to make in empirical studies, it is a common one. We can thus build a graph based on those coefficients, with a link when the partial correlation is nonzero as we do in this paper. Nevertheless, even if some variables are conditionally independent, they can still be unconditionally dependent. Let us consider three variables  $X, Y, Z$  again. If their partial correlation graph is

$$X - -Y - -Z$$

then the partial correlation between  $X$  and  $Z$  is 0, but their unconditional correlation is not (see Eq. 12 and 13 for a numerical example).

This feature of paths of partial correlations can highlight the presence of potential common factors when some nodes in the graph form a separate component, as in the reference partial correlation graph in the spatial dimension, in Figure 3, which we use as a way to visualise and measure the integration in our system (if different components/sectors start connecting, they become integrated).

All in all, we choose to work on partial correlations instead of unconditional correlations for their many benefits, but still come and go between the two, since unconditional correlations are also useful.

## B Explanation of LASSO penalization

Let us consider a simple linear model of the form:

$$Y = \beta_0 + \beta_1 X_1 + \dots + \beta_p X_p + \epsilon \quad (9)$$

where  $Y$ , the dependent variable, is explained by a constant and  $p$  explanatory variables  $\{X_1, \dots, X_p\}$ . If we have a set of data of length  $n$  (we have

$n$  observations of  $Y$ ,  $(y_1, \dots, y_n)$ , and the corresponding  $n$  observations of  $\{X_1, \dots, X_p\}$ ,  $(x_{i,1}, \dots, x_{i,p})_{i=1..n}$ , we will try to find the parameters  $(\beta_0, \dots, \beta_p)$  that will allow the estimated values  $\hat{Y}$  (on the regression line) to be the closest to the observed  $Y$ . Formally, in the case of an Ordinary Least Squares (where we minimise the sum of squared errors) estimation, the objective function would be:

$$\min_{\beta_0, \beta_1, \dots, \beta_p} \frac{1}{2} \sum_{i=1}^n \left( y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 \quad (10)$$

where  $j$  stands for variable  $X_j$  and  $i$  stands for observation  $i$  (of  $Y$  or  $X_j$ ). In traditional Econometrics works, we do not work in high dimension and have parameter estimates that can be tested. We rely on these statistical tests to identify the important explanatory variables of the regressions and analyse their parameters more particularly.

Penalizing with LASSO would slightly change that objective function by adding a term (which represents a constraint):


$$\min_{\beta_0, \beta_1, \dots, \beta_p} \frac{1}{2} \sum_{i=1}^n \left( y_i - \beta_0 - \sum_{j=1}^p \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^p |\beta_j| \quad (11)$$

with  $\lambda \geq 0$  the penalizing parameter. This parameter basically defines how much sparsity we want in the parameters: the greater this  $\lambda$ , the more coefficients will be set to 0. Indeed, if coefficients are very small, a large  $\lambda$  will increase the value of the objective function, so in order to minimise this function, we should set them to 0, so that they don't contribute to the minimisation problem anymore. Meanwhile, the nonzero parameters are calibrated at the same time, taking into account the 0s. We should thus obtain more realistic estimates of those nonzero parameters than if we calibrate the parameters first (if it is even possible, since we are in high dimension) and only keep the statistically significant ones. As we will see below, the estimates for the nonzero parameters will compensate for setting some parameters to 0.

Other regularizations are possible, but they may not set the small ones to 0. Another famous one is the Ridge regularization, which will provide many nonzero coefficients, due to the form of the constraint. Figure 10 shows how these two types of constraints translate into geometry. LASSO penalises the objective function with an L1-norm on the parameters,  $|\beta|$ , while ridge penalises it with an L2-norm,  $\beta^2$ , hence their shape in the figure (square vs. circle respectively, in blue). If we consider these regularizations as constraints in the optimization problem, we see in this figure that, on the one hand, the LASSO constraint will generally give an optimum in a corner of the square,

setting a parameter to 0 ( $\beta_1$  here). On the other hand, with the Ridge constraint, an optimum may be found with both parameters being different from 0.

Figure 10: Geometry point of view, difference between LASSO (left) and Ridge (right) prenalizations



The figure was taken from Wikipedia. The red ellipses represent different level ellipses of the objective function of a least square error minimisation problem. The blue areas represent the constraints added: square-shaped for LASSO (L1-norm, in the form of  $|\beta|$ , on the left) and circle-shaped for ridge regularization (L2-norm, in the form of  $\beta^2$ , on the right).

In particular, the partial correlation matrix (or the concentration matrix) could potentially be sparse, which comforts us in using this approach. It should at least be sparser than the unconditional correlation matrix, since it is "polluted" by the influences of all the variables. Applying this penalization to the calibration of the the concentration matrix  $\mathbf{F}$ , we would get many 0s, meaning conditional independence between many variables, after filtering out the influence of the others.

Let us come back to our example with the three eurodollar contracts (IR1, IR2 and IR3) from Appendix A. We may wonder whether the coefficient  $F_{IR1,IR3} = 632.0388$  (and maybe even  $F_{IR2,IR3} = -1807.8044$  and  $F_{IR3,IR3} = 1730.4653$ , considering other values are very large) is "normal" or if it is small enough compared to the others to be set to 0. If we look at their partial correlations from the whole  $51 \times 51$  matrix, we get the following PC matrix  $\mathbf{C}^{LASSO}$  and concentration matrix  $\mathbf{F}^{LASSO}$  (which was calibrated with



LASSO and then the PC was derived from it):

$$\mathbf{C}^{LASSO} = \begin{pmatrix} -1 & 0.6548 & 0 \\ 0.6548 & -1 & 0.5211 \\ 0 & 0.5211 & -1 \end{pmatrix} \quad \mathbf{F}^{LASSO} = \begin{pmatrix} 6999.69 & -4313.20 & 0.000 \\ -4313.20 & 6199.35 & -1632.13 \\ 0.000 & -1632.13 & 1582.37 \end{pmatrix} \quad (12)$$

Their Looking at the reference graph in Figure 4, we indeed see that there is no link between these two nodes. Comparing with the partial correlation matrix (and the corresponding concentration matrix) found in Eq. 8, the coefficient between IR2 and IR3 is slightly larger when using LASSO (0.5211 vs. 0.5172). The coefficient between IR1 and IR2 is lower when penalizing with LASSO (0.6548 vs. 0.7403). This may be explained by the 0 set for  $C_{IR1,IR3}^{LASSO}$ : while it was negative when not using LASSO, to compensate for this mitigating effect (of the negative coefficient), the dependence between IR1 and IR2 may have been decreased (or it could have affected other partial correlations too).

Looking at the corresponding unconditional correlations, we find the following matrix (extracted from the  $51 \times 51$  matrix):

$$\rho = \begin{pmatrix} 1 & 0.9994 & 0.9424 \\ 0.9994 & 1 & 0.9533 \\ 0.9424 & 0.9533 & 1 \end{pmatrix} \quad \rho^{LASSO} = \begin{pmatrix} 1 & 0.9169 & 0.8370 \\ 0.9169 & 1 & 0.9230 \\ 0.8370 & 0.9230 & 1 \end{pmatrix} \quad (13)$$


The coefficients are not exactly equal to the ones in the previous section, without using LASSO (and the ordering is not the same), but are still quite close. The difference may also be due to some noise in the unconstrained case.

Moreover, applying the *nets* algorithm of Barigozzi and Brownlees (2017) applies LASSO penalization to the estimation of both the AutoRegression and partial correlation matrices in one step. This allows the optimization to have better convergence properties and allows it to select the parameters among both AR coefficients and partial correlation coefficients, instead of selecting the former first and then the latter.

## C Measures in the dynamic analysis over the whole period

### C.1 Measures related to the Granger causality graphs


Figure 11: Measures related to the links of the Granger causality graphs



The vertical bars represent the events under consideration: red for the beginning of the subprime crisis, blue for the default of Lehman Brothers, green for the day of minimum return on the S&P500, orange for the day of its maximum return and violet for the Flash Crash. The blue line (#Granger links) represents the number of links of propagation. The yellow line (Average propagation) represents the average value of the propagation coefficients. Note that when there is no link, the average is NA, hence the discontinuity.


## C.2 Measures related to the partial correlation (Integration, PC) graphs

Figure 12: Measures related to integration of markets



The vertical bars represent the events under consideration: red for the beginning of the subprime crisis, blue for the default of Lehman Brothers, green for the day of minimum return on the S&P500, orange for the day of its maximum return and violet for the Flash Crash. The blue line (Total communicability) represents the total of the easiness of information flow in the system; the values have been normalized by the minimum value to fit on the same scale. The orange line (%links PC) represents the inverse of the sparsity in the graph; it gives the proportion of nonzero PC in terms of total links possible (full graph). The yellow line (#clusters PC) represents the number of factors (components, clusters) in the PC graph.

Figure 13: Measures related to the structure of the partial correlation graphs  
(3)



The vertical lines represent the events under consideration: red for the beginning of the subprime crisis, blue for the default of Lehman Brothers, green for the day of minimum return on the S&P500, purple for the day of its maximum return and orange for the Flash Crash. The blue line (Reference) represents the proportion of links that are common to the daily graphs and the reference graph. The red line (Previous graph) represents the proportion of links that are common to the graph on that day and the graph on the previous day.

## **Finance for Energy Market Research Centre**

Institut de Finance de Dauphine, Université Paris-Dauphine

1 place du Maréchal de Lattre de Tassigny

75775 PARIS Cedex 16

[www.fime-lab.org](http://www.fime-lab.org)