Financial Network Analysis

Causal Inference Approach

Doctoral Dissertation

presented by

Katerina Rigana

under the supervision of

Giovanni Barone-Adesi, Antonietta Mira, and Ernst Wit

A thesis submitted in fulfillment of the requirements for the degree of Philosophiae Doctor in Finance at the Swiss Finance Institute, Università della Svizzera italiana.

October 2023

Dissertation Committee

Prof. Giovanni Barone-Adesi	Swiss Finance Institute and Università della Svizzera italiana, Switzerland
Prof. Antonietta Mira	Università della Svizzera italiana, Switzerland
Prof. Ernst-Jan Camiel Wit	Università della Svizzera italiana, Switzerland
Prof. François Degeorge	Swiss Finance Institute and Università della Svizzera italiana, Switzerland
Prof. Erich Walter Farkas	Swiss Finance Institute, ETH Zurich and University of Zurich, Switzerland
Prof. Silvia Figini	Università di Pavia, Italy

When you can measure what you
are speaking about and express it
in numbers, you know something
about it.

Lord Kelvin

Abstract

This doctoral dissertation contributes to the field of financial network analysis by exploring contagion risk estimation through a causal inference framework, specifically within the Forex market. By integrating causal inference, network contagion analysis, and machine learning techniques, this research introduces innovative metrics and frameworks that surpass traditional risk management methodologies. It offers a nuanced understanding of contagion dynamics among individual currencies, detailing the pathways of contagion and evaluating the systemic risk impact of major events, such as the COVID-19 pandemic.

Building on traditional methods, this research introduces a distinctive measure of contagion in the Forex market through a causal network approach. This stands in contrast to the widely used instrumental variable and Granger causality methods. Employing causal inference theory – already validated in fields such as genetics, medicine, and climate science, to distinguish causality from correlation in observational data – contagion pathways are identified. The resulting Network Contagion (NECO) measure evaluates both the market as a whole and individual currencies in terms of diversification and exposure to systemic risk.

Further advancing the domain, the dissertation provides a unique framework for Value at Risk. This framework takes into account the intricate dynamics of contagion, improving its dependability in volatile market situations. This new method specifically addresses long-standing challenges related to non-normality, non-stationarity, and unforeseen market shocks. When compared with conventional VaR measures, the superior accuracy of this approach is evident.

Acknowledgements

Embarking on this academic journey has been an enlightening experience, and its successful culmination would not have been possible without the support and guidance of many.

First, I extend my deepest appreciation to my advisors, Professor Giovanni Barone-Adesi and Professor Antonietta Mira. Their insightful feedback, encouragement, and unwavering support have consistently inspired and guided me.

I am immensely grateful to Professor Ernst C. Wit, an expert in causal inference. His guidance, enthusiasm, and expertise have been pivotal in shaping this research. Our brainstorming sessions and discussions have not only influenced this work but have also enriched my perspective on academic research.

A special mention goes to Sam Cook. Her technical acumen, attention to detail, and valuable insights have been instrumental in our collaborative efforts. Our weekly meetings, always infused with her positive attitude, have been a consistent highlight.

To my colleagues and lab partners at the Institute of Finance, the Institute of Computing, and at FNA – our shared moments, fun discussions, and mutual learning have been both a pleasure and privilege.

Last but certainly not least, my heartfelt thanks go to my family and friends. Their faith in my abilities, continuous encouragement, and boundless love have served as my steadfast pillars. Their unwavering belief, even in the face of challenges, has propelled me forward.

In closing, I extend my appreciation to all who indirectly contributed to this research. A special note of recognition goes to the members of the dissertation committee. I look forward to their invaluable feedback and thank them in advance for their time and expertise in reviewing this doctoral dissertation.

Contents

Chapter 1

Introduction

This dissertation comprises three essays within the domain of financial network analysis, focusing on contagion risk estimation through a causal inference framework. It is submitted as a requisite for the attainment of a PhD in Economics from the Università della Svizzera italiana and of a PhD in Finance at the Swiss Finance Institute.

In this introductory chapter, we describe the main ideas and lay the theoretical foundation for our research. We start by introducing key financial concepts, specifically, contagion and its importance in risk management. The analytical framework used in this dissertation is based on network theory and causal graphical models. Afterward, we provide a summary of the following three chapters and highlight the main contributions of each of the three essays.

1.1 Contagion

Contagion plays a pivotal role in the spreading and subsequent assessment of risk associated with asset or financial instrument holdings. This has been increasingly recognised in the past 20 years within financial research [66, 159, 3, 82, 53, 135, 93, 177]. The precise delineation of contagion remains, however, a subject of ongoing debate within the financial literature. This is highlighted by the absence of a unanimous consensus on its exact definition [68, 42, 35]. Various interpretations and methods exist to quantify the spread of the infection, which adds to the complexity of the matter.

Drawing inspiration from epidemiology, which deals with the transmission of diseases between individuals or organisms, contagion can be seen as the diffusion of financial shocks among market participants [152, 91, 164]. The historical roots of contagion in the financial domain can be traced back to the 19th century, when it was frequently invoked to describe the rapid spread of panics and bank runs [121, 169, 92]. These early conceptualisations mainly revolved around what today would be categorised under the domain of behavioural finance, focussing on sentiments, emotions, behavioural biases, and crowd psychology [116, 101].

Prompted by significant events, such as the 1987 market crash and the 1997 Asian financial crisis, contagion emerged as a phenomenon characterised by unexplained shocks or correlations, setting it apart from anticipated changes called spillovers and interdependence [178]. The idea of “pure contagion" was developed, referring to fluctuations in asset prices that cannot be explained by individual trends, which are instead believed to be influenced by factors transmitted from one asset to another [214]. This definition highlights the causal aspirations of research on contagion: It is crucial to understand both the source and the target of contagion.

The global financial crisis of 2007-2010 marked a pivotal moment in contagion research, highlighting the need for robust estimation and modelling techniques to enhance preparedness for future crises [156, 71, 58, 82, 52, 195, 20]. This resurgence of interest in contagion stemmed from the recognition of its role in financial instability and systemic risk. Researchers focused mainly on identifying the sources of contagion and explaining the global financial crisis. As part of a comprehensive approach to measure and identify contagion, several methodological approaches were developed, including a copula approach [179, 218], vector moving average, and vector autoregression (VAR) models [45, 14, 15]. However, these approaches often lacked a causal interpretation and the ability to effectively handle system shocks.

For financial regulators, gaining a complete understanding of contagion mechanisms and their sources is crucial. This understanding is vital to assessing the resilience of financial networks, particularly in the context of systemic risk [202]. Regulators and central banks must closely scrutinise contagion dynamics to safeguard the stability of the financial system. In a broader context of risk management, understanding contagion remains paramount. It transcends the traditional focus on individual asset holdings, extending to a systemic perspective. This perspective recognises the intricate web of interconnections within financial markets that can lead to contagion effects [70]. Effective risk management should encompass both aspects: securing individual asset portfolios and mitigating systemic risks. Therefore, understanding the dynamics of contagion becomes essential not only for macroeconomic dynamics, but also for prudent risk management within the investment and banking sectors, ensuring financial stability and improving overall risk management strategies.

1.2 Risk Management and Contagion

Risk management has long been an essential pillar of finance and investment, guiding strategies to assess, control, and mitigate potential losses associated with financial decisions [109]. Central to the concept of risk management is the idea of understanding the vulnerabilities inherent in financial instruments and the broader market environment. With the evolution of global markets, the complexities surrounding risk management have increased, underpinning the need to address sources of risk beyond traditional methods [147, 40].

Value-at-Risk (VaR) is a prevalent risk metric used by banks and financial institutions to quantify the maximum potential loss of a portfolio over a specified time horizon for a given confidence level [115]. While VaR remains a cornerstone in risk management, its traditional implementations often make assumptions about the normality of return distributions and the linearity of asset relationships, which have been criticised as being simplistic and potentially misleading [16, 90, 157]. This can be particularly problematic when considering the potential of contagion effects, as contagion inherently challenges these assumptions [143, 23, 64].

Contagion introduces an additional layer of complexity to risk management by highlighting the need to account for the interdependencies among different financial instruments and markets. Traditional risk assessment methods, when relying solely on the historical performance of a single asset, may be blind to potential systemic shocks. On the other hand, incorporation of contagion into the analysis leverages the intricate web of relationships among assets, allowing the extraction of valuable insights from the behaviour of related assets that could affect the asset of interest [2].

Thus, it becomes evident that to enhance the accuracy of risk assessment, it is indispensable to incorporate contagion effects. Recognising the potential of one asset's behaviour to be influenced by others permits a more informed evaluation of potential risks. This approach allows for the incorporation of information from a spectrum of assets, rather than limiting the analysis to a single asset past performance.

So far, the discussion about contagion may be confused with inter-correlations. Although correlations have some predictive power, this is only true if the external circumstances do not change. Correlations are not robust against external shocks, which are common in financial markets. On the other hand, causal relationships persist in the presence of shock [167, 168, 97]. Therefore, our aim is to focus on a causal interpretation of contagion: *causal contagion*. Addressing the assumptions that underlie most risk assessment methods and integrating causal

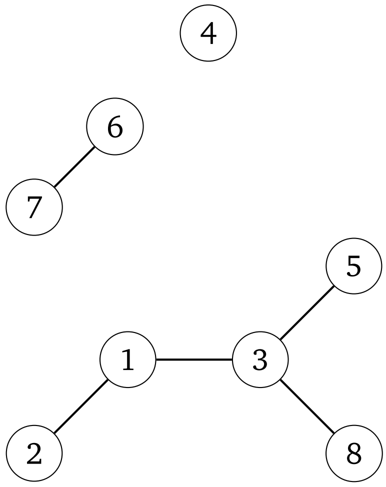

Figure 1.1: Example of a network

contagion provides a more robust framework to protect against unexpected financial shocks.

In conclusion, as markets become increasingly interconnected, the importance of considering contagion in risk management increases. Incorporating causal contagion aids to capture the multifaceted nature of risks associated with asset prices. By doing so, it paves the way for more resilient financial systems, capable of withstanding and mitigating the adverse impacts of unexpected shocks, ensuring a more stable and secure financial landscape for all market participants.

1.3 Networks Theory in Finance

This thesis relies on network theory to study contagion within financial systems. Network theory offers a structured approach to understanding the interrelations of entities within a system. Conceptually, a network consists of interconnected entities which could be items, individuals, or larger groups such as institutions or nations. Mathematically, these networks take the form of graphs. In these graphs, the entities, termed nodes or vertices, are connected by lines known as links or edges. A good overview of the theoretical background can be found in [155].

Figure 1.1 shows an example of a network. There are eight nodes that are varyingly connected by links. Not all nodes have to be connected; for example, node 4 is on its own and has no relationship, as defined by the model, with the other nodes. There are two subgroups within the network $\{6, 7\}$ and $\{1, 2, 3, 5, 8\}$

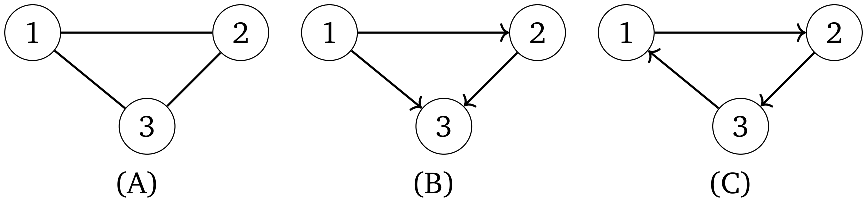

Figure 1.2: Simple examples of: an undirected network (A), a directed acyclic graph or DAG (B) and a directed cyclic graph (C).

— each are called *connected components*. The sequence of links $\{1-3, 3-8\}$ is called a *path* from node 1 to node 8. In network theory, a component of an undirected network is a subgraph, in which any two nodes are connected to each other by paths but that is not connected to the rest of the network.

Network theory is increasingly being used to analyse financial issues, particularly those related to contagion and systemic risk [23, 2, 196, 195, 59]. In finance, there are two main types of networks that can be analysed. One group analyses the interconnections between objects that have clearly defined links, like ownership relationships between companies, and so it is a question of data mining in order to recreate such a network. When the relationships are not observed, like how different assets and financial instruments affect each other, then the links themselves have to be estimated before being analysed.

Networks theory further distinguishes between directed and undirected networks. Undirected networks show relationships that are either symmetric by nature or defined in a way where the direction is irrelevant to the researcher. For instance, in social networks, if one individual is acquainted with another, the reverse is inherently true. An example of a undirected network in finance can be one where the links between investors show the presence of having common assets under management. Figure 1.2(A) shows an example of an undirected network. A directed network has arrows that go from one node to the other, as in Figure 1.2(B) and 1.2(C), and describe relationships that are not automatically reciprocal by nature. Examples of directed networks include payment networks, where transactions between entities have a clearly defined direction, and trade networks, where imports and exports between countries are distinctly directed. Finally a correlation network where nodes represent assets, and edges indicate a significant correlation in their price movements, are an example of undirected network. Whereas causal networks, used to depict causal relationships between various financial variables, are an example of a directed network.

Links can be further quantified by a weight. In the case of import-export networks for example, the dollar amount of the transactions would represent the

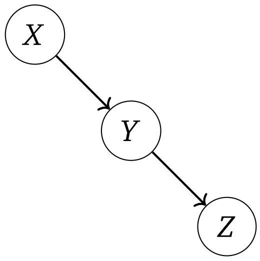

Figure 1.3: Example of a causal network

weight of the links between each pair of nodes, in this case countries. The density of a graph or network G is calculated as the ratio of the size of the network, that is, the number of links, relative to the maximum size, that is, the total number of possible links. The density for an undirected network G:

$$\text{density}(G_{undirected}) = \frac{\text{size of } G}{\text{maximum size of } G} = \frac{\#links}{N(N-1)/2} \tag{1.3.1}$$

where N is the number of nodes in the network.

For directed networks, there are twice as many possible links as for undirected networks, because each pair of nodes may be joined by up to two links.

$$\text{density}(G_{directed}) = \frac{\text{size of } G}{\text{maximum size of } G} = \frac{\#links}{N(N-1)} \tag{1.3.2}$$

Within directed networks, we distinguish between acyclic and cyclic networks. An acyclic network represents a network where no path can be found such that a node is connected to itself. Figure 1.2(C) shows an example of a cyclic network where there is a path for all nodes such that one goes back to itself. A network that is directed without such cycles is called a **directed acyclic graph (DAG)**.

We are going to represent our causal networks as DAGs. We will interpret Figure 1.3 as saying that X has a causal effect on the variable Y directly and on Z indirectly. In such a diagram, Y is a child of X and X is the parent of Y. Y and Z are descendants of X and X and Y are ancestors of Z. This allows us to look at the acyclicity condition of a DAG such that no variable represented as a node in our causal network can be an ancestor of itself.

1.4 Causal Inference in Finance

In finance, it is essential to understand complex systems and their multifaceted interactions and relationships. While traditional methods provide insight into correlations, the growing complexity of financial systems requires techniques that can elucidate causality. Granger causality, instrumental variables, and causal graphical models present three important methodologies in the field of financial causality research. Whereas the former two have a long history in economic and econometrical research, the application of the latter is more rare in this area. Causal Graphic Models, anchored in causal inference, offer a robust alternative framework to discern the underlying causal structures [83].

1.4.1 Granger Causality

Historically rooted in economics and econometrics, Granger Causality was introduced as a convenient technique within the broader category of transfer entropy-based methods for causal analysis [186]. Granger causality relies on temporal correlations within time series to establish causal relationships [88, 57]. Its core tenet suggests that if the past values of one variable can improve our understanding of the future trajectory of another variable, then a Granger-causal relationship is said to exist [87, 103].

The foundational idea behind Granger causality is that if a variable X provides unique information that enhances our prediction of a future value of another variable Y, then X is said to Granger-cause Y. This is underscored by two essential principles: the causal event precedes its effect, and the causal event possesses exclusive insights about the subsequent values of the effect.

Formally, the hypothesis for Granger-causality of X on Y can be denoted as:

$$\mathbb{P}[Y(t+1) \in A | \mathscr{I}(t)] \neq \mathbb{P}[Y(t+1) \in A | \mathscr{I}_{-X}(t)] \tag{1.4.1}$$

In this equation, $\mathbb{P}$ signifies probability, A denotes an arbitrary non-empty set, and $\mathscr{I}(t)$ and $\mathscr{I}_{-X}(t)$ represent the information sets available up to time t inclusive and exclusive of X, respectively. Acceptance of this hypothesis suggests that X Granger-causes Y. However, Granger causality does not imply actual cause-and-effect relations, but identifies potential predictive chronologically ordered associations.

The approach is, in fact, not without pitfalls. By leaning heavily on temporal correlations, Granger causality sometimes falls short in nonlinear systems where past patterns are not necessarily predictive for future behaviours [205, 102].

Additionally, inherent assumptions like linearity, the need for stationary data, and perfect-time measurements can prove restrictive, especially in the dynamic world of finance marked by abrupt market shifts [144, 193].

Furthermore, the methodology can sometimes skirt the edges of true causal relationships, potentially missing crucial variables or being ensnared by hidden confounders, leading to possibly misleading interpretations [203, 144, 193]. While Granger causality serves as an insightful tool for preliminary explorations of potential causal dynamics, it is imperative to approach its results with a critical lens, especially when other causal inference techniques are available.

1.4.2 Instrumental Variables

Instrumental Variables (IV) present an alternative paradigm to Granger causality, stemming from quasi-experimental designs. At the heart of the IV approach is the use of instruments, variables that exert influence on predictors, but remain uncorrelated with the outcome's error terms. This allows researchers to control for potential endogeneity, an issue in which predictors could be correlated with the error term, leading to biased estimates [8].

For a comprehensive grasp of IV, consider a typical econometric model,

$$y = X\beta + \epsilon \tag{1.4.2}$$

In this representation, y stands for a $T \times 1$ vector of dependent variables, X is a $T \times k$ matrix of independent variables, β is a $k \times 1$ vector of parameters we seek to estimate, and ϵ denotes a $k \times 1$ vector of error terms. Ordinary Least Squares (OLS) might be the immediate consideration for estimating β. However, consider scenarios where X might be correlated with ϵ. This correlation could lead to biased OLS estimators.

To address this, we introduce a matrix of instrumental variables, Z. This $T \times k$ matrix Z is correlated with X but is not correlated with ϵ. In other words, it is only related to y *through* X. With Z in place, we can construct a consistent IV estimator as:

$$\beta_{IV} = (Z'X)^{-1}Z'y \tag{1.4.3}$$

The key insight here is that the exogenous variables Z operate as instrumental variables. The instruments, represented by $(Z'Z)^{-1}Z'X$, offer estimates of the component of X that remains uncorrelated with ϵ. The two-stage least squares is a notable extension of this IV approach. A good overview can be found in [207].

This methodology situates itself within a broader framework of quasi-experimental methods. These techniques, including Difference-in-Differences, regression discontinuity designs, matching, propensity score designs, and comparative interrupted time series designs, aim to replicate the robustness of experimental designs in observational settings [190]. Each method offers unique ways to address common issues in causal inference, from selection bias to endogeneity, without the need for randomised control.

However, just as with Granger causality, instrumental variables are not without their own set of limitations. The foundational assumptions underpinning the IV method, especially the validity and relevance of the instruments, can be challenging to verify in real-world scenarios. Ensuring that the chosen instrument only affects the outcome through the predictor (the exclusion restriction) is crucial, and any violation could lead to misleading causal estimates [96]. Additionally, similar to Granger causality, the IV approach is contingent on specific model dependencies and, at times, leans on assumptions that may appear stringent or unrealistic in certain contexts [220, 162].

1.4.3 Causal Graphical Models

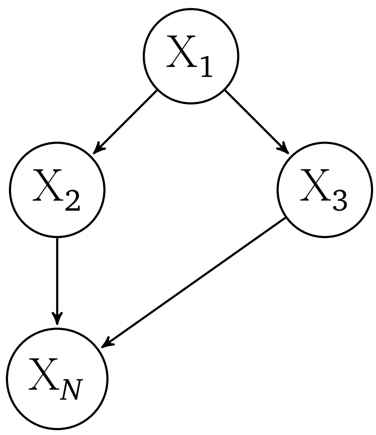

Figure 1.4: A Directed Acyclic Graph (DAG) G_X representing the relationships among financial assets. Nodes (e.g., X_1, X_2) represent distinct financial assets, while the arrows indicate conditional dependencies. For example, an arrow from X_1 to X_2 indicates that X_2 is conditionally dependent on X_1. In this DAG, X_2 and X_3 are conditionally independent, but both depend on X_1.

Directed Graphical Models Directed Graphical Models (GM) employ directed graphs to represent conditional dependencies between random variables. Their primary objective is to represent the joint probability distribution and to decompose it based on the graph structure. For example, in the context of finance where we consider N distinct financial assets, their logarithmic returns can be represented by the multivariate random variable $X = (X_1, \ldots, X_N)$. In this representation, each of the N assets corresponds to a node in our network. This setup allows us to associate X with a set of N nodes, integrated into a Directed Acyclic Graph (DAG), as illustrated in Figure 1.4. The absence of arrows within the DAG G_X defines conditional independence. Such a perspective enables us to transform the intricate multivariate distribution into a product of simpler univariate densities:

$$P_X(x) = \prod_{i=1}^{N} P_{X_i}\left(x_i \mid x_{pa(i)}\right) \tag{1.4.4}$$

where $x_{pa(i)}$ denotes the parent nodes of the i-th node.

In directed GMs, the directed edges primarily signify statistical or probabilistic dependencies. They do not always allude to causal relationships. If a directed edge comes from node X_1 and ends at node X_2, it suggests that X_2 is conditionally dependent on X_1, given the remaining nodes. However, it does not infer a causal relationship between X_1 and X_2. It is possible for a directed GM to represent contagion effects, but such representation might lack causal underpinning.

In the context of directed GMs, the global Markov property offers a way to describe conditional independence using the structure of the graph. It states that if there is no d-connecting path (also referred to as an active trail) between two sets of nodes, given a conditioning set within the graph, then the associated random variables for these nodes are conditionally independent based on the variables in the conditioning set.

This property is notable for its ability to simplify the understanding of conditional independence relationships in a directed GM. Rather than delving deeply into the joint distribution, one can infer these relationships directly from the graph. By ensuring that the graph accurately captures the conditional independencies of the distribution, the Global Markov property emphasises the utility and reliability of directed GMs in representing complex distributions and facilitating insights into relationships between variables.

Causal Graphical Models Causal Graphical Models are specialised versions of directed GMs, where the directed edges represent explicit causal relationships as opposed to mere probabilistic associations. To ensure that a directed GM can be interpreted causally, an additional *intervention stability* criterion is introduced. This criterion requires that

$$p\left(x \mid \mathrm{do}(x_1)\right) = \prod_{x_i \neq x_1} p\left(x_i \mid x_{\mathrm{pa}(i)}\right) \tag{1.4.5}$$

This equation delineates the essence of causality, differentiating between mere conditional probabilities and genuine causal interactions. The need for this distinction arises from the inherent inadequacy of traditional statistical and probabilistic frameworks in capturing causal relationships. The act of intervention, often denoted as $\mathrm{do}(x_1)$ or symbolised by $\|x_1$, essentially "forces a value" on X_1, making X_1 independent of its immediate precursors. This renders it possible to define the causal effect X_1 exerts on its dependents. Pearl's notation, $P(Y|do(T))$, conveys the probability distribution of an outcome Y when an intervention modifies the treatment T. Specifically, $P(Y|do(T=1))$ describes the outcome's distribution given the treatment, while $P(Y|do(T=0))$ reflects the outcome without the treatment. Challenges emerge in observational studies where interventions are not feasible. However, the right hand side of formula 1.4.5, which depends entirely on observables, suggests that it is still possible to calculate the causal effect.

Causal Inference If we could randomly intervene in any of the nodes within our causal GM and measure the consequences, we would have no problem discovering the causal effects within our systems. However, randomised experiments are not always feasible or ethical. For example, when determining whether smoking cigarettes causes cancer without randomly selecting individuals and forcing them to smoke. We cannot make the EUR/USD exchange rate fall 20%, just for our experiment. This is where causal inference theory such as formula 1.4.5 can be used and why it has received such attention in fields such as genetics, medicine, and climate change [163]. It offers a way to estimate causal effects without conducting experiments, using only correlations from observed data. Formally, the additional assumption needed is that of faithfulness, where it is assumed that no causal effect exactly cancels out downstream. In fact, two cancelling-out effects might not even be regarded as causal. If this is satisfied, algorithms like the PC algorithm can be used to interactively extract from correlation data the causal structure of a casual graphical model.

PC Algorithm The PC algorithm, named after its developers Peter Spirtes and Clark Glymour, is a widely adopted constraint-based method for the discovery of causal structures in directed graphs, specifically in the context of causal GMs. Given a data set, the algorithm starts by constructing an undirected graph in which all pairs of variables are connected. Subsequently, it uses conditional independence tests iteratively to prune edges that are statistically independent, thereby revealing the underlying graphical skeleton. In a subsequent step, the skeleton can be (partially) directed by identifying colliders and applying various logical rules, such as avoiding the introduction of cycles and additional colliders.

The PC-Stable algorithm is an extension of the PC algorithm, an approach in causal inference. It enhances the robustness of the original algorithm by accounting for sampling variability and noise in the data, providing a more reliable identification of causal relationships. By iteratively applying the PC algorithm on multiple bootstrapped samples of the data and aggregating the results, PC-Stable produces a more stable and accurate causal graph representation of the underlying relationships among variables. This is the algorithm that will be used as a basis of this dissertation research on causal continuity networks.

The primacy of causal inference is gaining momentum, particularly in arenas such as predictive modelling, stress testing scenarios (akin to central bank evaluations under hypothetical conditions such as a 20% plunge in the USD), and disciplines spanning machine learning to software engineering. Embracing causal inference accentuates models transcending mere predictive capabilities, illuminating the foundational causal architectures. Such insights are instrumental in fostering informed decision-making, especially within intricate scenarios.

1.5 Overview of the Dissertation

This section covers an overview of each of three essays. Chapters 2 and 3 can be found as co-authored papers with Ernst C. Wit and Sam Cook in [177] and [176].

1.5.1 First Essay: A New Way of Measuring Effects of Financial Crisis on Contagion in Currency Markets

Although various approaches to quantifying contagion have been proposed, many of them lack a causal interpretation. This lack of causality can result in unstable and non-robust predictions, especially during times of financial crises. We will present a new measure for contagion among individual currencies within the foreign exchange market and show how the paths of contagion work within the Forex market using causal inference. This approach will allow us to pinpoint sources of contagion and find which currencies offer good options for diversification and which are more susceptible to systemic risk, ultimately resulting in feedback on the level of global systemic risk. Special attention is paid to the effects of the global pandemic COVID-19.

We introduce a parsimonious causal model to describe market fluctuations in currency returns. From the causal interpretation of the model, we define a new measure of network contagion (NECO) that can be used to assess the market as a whole and analyse individual currencies. We examine the dynamic behaviour of the NECO measure in the Forex market between 2000 and 2021. Our findings indicate that financial contagion increases for most currencies during times of crisis, with the notable exception of safe haven currencies such as the Swiss Franc.

1.5.2 Second Essay: Causal Network Contagion Risk

Building upon the foundational concepts of contagion network analysis introduced in the first essay, the second essay delves deeper, exploring the intricacies of contagion and its implications for risk management. This chapter introduces an innovative approach to understanding risk in financial markets and a new way of calculating the Value at Risk (VaR): the Causal Network Contagion Value at Risk (Causal-NECO VaR). This approach quantifies VaR by especially considering the inherent risk dynamics caused by contagion. In doing so, it seeks to be a more reliable measure, particularly during unpredictable market conditions. The use of causal-based contagion makes the method less affected by spurious correlations and consequently more robust to change, such as sudden volatility clustering.

A major highlight of the chapter is the integration of causal inference-based networks and copula transformations in VaR computation. By taking this route, the authors effectively address challenges such as handling skewed distributions and adapting to external shocks in the market. The Causal-NECO VaR design is particularly attractive for its invariance properties, which can be vital during transitional financial periods like economic crises. This method does not merely present another alternative to VaR computation; it introduces a robust measure designed with the intricacies of real-world market disruptions in mind.

The most significant contribution of this essay lies in its potential to transform risk management decisions by eliminating spurious correlations. It successfully lays the foundation for incorporating causal inference in finance, showcasing the tangible benefits of deciphering the underlying causal system from existing financial data. In conclusion, this essay not only presents a novel method for VaR computation; it also paves the way for a more comprehensive and causal understanding of financial risk, setting the stage for more accurate risk management strategies in the face of financial unpredictability.

1.5.3 Third Essay: Conformal Updating and Causal Network Contagion

Expanding on the core insights of the first two essays, this third essay navigates the intricacies of risk estimation during unprecedented financial turbulence. Specifically, we address the enduring challenges associated with non-normality, non-stationarity, and unexpected market shocks. To address these challenges, we introduce the Conformal Causal Network Contagion Value at Risk (CCNC VaR), a novel framework that combines the resilience of causal network contagion analysis with an innovatively adapted form of conformal updating, tailored specifically for the VaR metric.

Conformal updating, an innovative technique primarily rooted in machine learning, is used to enhance the reliability of predictive models by fine-tuning predictions adaptively as new data becomes available, without altering the foundational prediction model. Recognising the method's potential, we have tailored it for our context by enabling it to adjust to uncertainties stemming from model assumptions and potential changes in the underlying causal structure. Rather than sticking to a static algorithm, our adaptation of conformal updating allows for slow adjustments to the targeted VaR levels (α_t), ensuring the model remains robust, all while accounting for potential misspecifications that might skew ex-post results away from the desired confidence level (α).

The integration of conformal updating, a technique borrowed from the realm of machine learning, ensures our framework's accuracy, especially in scenarios where market disruptions are abrupt and wholly unexpected. We validate the effectiveness of our framework by applying it to the Forex market, a domain that has recently faced increased volatility, structural shifts, and notable central bank interventions during the COVID-19 recession. A comparative analysis with standard VaR measures from the literature underscores our method's superior accuracy.

In conclusion, the CCNC VaR not only offers a refined theoretical foundation but also holds considerable practical merit, effectively handling non-normality and non-stationarity in financial data. This essay underscores the potential of an integrative approach to capture the complexities of risk in today's financial landscape.

Chapter 2

A New Way of Measuring Effects of Financial Crisis on Contagion in Currency Markets

The following chapter was published as:
Katerina Rigana, Ernst-Jan Camiel Wit, and Samantha Cook. A new way of measuring effects of financial crisis on contagion in currency markets. *International Review of Financial Analysis*, 90:102764, 2023. doi:10.1016/j.irfa.2023.102764.

2.1 Introduction

In order to see how resilient financial networks are to contagion, financial regulators need to understand how contagion propagates and where the sources of contagion reside [202]. There are various ways to define financial contagion.

The first descriptions of contagion are predominantly in terms of what today we would call behavioural finance — in terms of sentiments, emotions, behavioural biases and crowd effects [93]. Modern research into contagion within the academic literature began to appear in the wake of the 1987 market crash and exploded during the 1995 Mexican crisis, the 1997 Asian financial crisis, and the related 1998 Russian financial crisis that followed. The main goal was to explain how a series of potentially local issues could spread from country to country, creating a financial crisis with global repercussions [55, 35]. These events caught many economists off guard and so most papers during this period concentrate on explaining [55, 30], or explaining away [67, 36, 119], this new market behaviour. Most of the definitions of contagion during the period define it as shocks or correlations that are unexplained or unexpected or significantly higher than usual, in contrast to expected and explainable changes and correlations, called spillovers and interdependence [178].

The 2007-2010 global financial crisis rekindled interest in analysing contagion. This time around there has been less research into whether contagion exists and more into how to estimate and model it properly to be ready for future crises. During this period we see the appearance of network theory applications that set the tone for future research [156, 71, 58, 82, 195, 20].

There have been various approaches to contagion estimation: the copula approach [179, 218]; the vector moving average and variance decomposition, related to the Impulse Response modelling [45, 14, 15]; and new approaches of Vector Autoregression (VAR) estimation in structural models [39, 57, 79, 78, 11, 5, 80, 4]. Each of these methods has its own advantages, but they lack an intervention-based causality interpretation. Traditionally when it comes to causality economists have relied on the Granger causality concept [201], but this concept merely relies on temporal correlations rather than structural causation [205].

We will analyse contagion through changes in price due to factors spreading from currency to currency that cannot be explained by individual trends. This type of contagion can be interpreted as "pure contagion" defined by [214]. Combining causal networks with a structural VAR model, we are able to extract from the log-returns of the exchange rates which part of a change in the value of a currency is caused by the idiosyncratic characteristics of the currency and which part is caused by contemporaneous contagion effects from other currencies. The structural VAR part of our approach is an extension of [79]. Similar VAR approaches have been used to analyse systemic risk within the foreign exchange market; however, none look at the causal direction of such effects.

In order to analyse the contagion paths in the foreign exchange market, we will recreate directed partial correlation-based networks using the causality concepts of [161] extended into practical methods by [198] and [37].

The remainder of this paper is organised as follows. In Section 2.2, we first explain the basic concepts of network theory and causal inference and then describe how we estimate contagion using Causal Graphical models within a structural VAR model. In Section 2.3, we present the dataset and in Section 2.4 we analyse contagion within a subset of currencies on the Forex market during the years 2000-2021; and in Section 2.5 we conclude by summarising the main advantages of this approach to measuring contagion in finance, while noting the remaining challenges.

2.2 A Structural Equation Model for Contagion

Our measure of contagion within a network is based on the theory of causal graphical models. In this section, we will introduce the main underlying concepts and notation. A network is a collection or system of interconnected objects such as things, people or groups of people (like institutions or countries, for example). In mathematics, these interconnections are represented as graphs, defined as a set of nodes or vertices that are connected by links or edges.

In this work, we will analyse directed networks by estimating causal links that are not directly observed between financial instruments or assets. Such networks are called causal networks, a specific type of graphical model (also often called Bayesian or belief networks). These networks incorporate probabilistic relationships between nodes in the form of a Directed Acyclic Graph or DAG [108, 160].

Bayesian networks can be analysed using several different types of algorithm. There are the so-called constraint-based algorithms that look for conditional independence [37], and the main algorithm of this group, the PC algorithm [198], is the one we will use in Section 2.2.3. The other group of algorithms is the so-called score-based algorithms, which maximise a causal objective function. A comparison between these methods in terms of speed and accuracy can be found in [187]. Importantly, they find that "constraint-based algorithms are more accurate than score-based algorithms for small sample sizes."

2.2.1 Causal Graphical Models

Here we present the mathematics underlying the causal graphical models. We consider the N assets that we want to analyse and their logarithmic returns are represented as the multivariate random variable $X = (X_1, \ldots, X_N)$. Each of these N assets will represent a node in our network. We can analyse X as a directed graphical model (GM) represented by a causal DAG. The arrows in the network G_X are to be interpreted in terms of conditional independence (CI) with the additional causal interpretation.

A DAG G_X is causal for a probability distribution $f(x)$ if $f(x)$ recursively factorizes with respect to G_X and the following intervention formula holds for all subsets A in V, the set of all nodes:

$$\forall A \subseteq V \quad f(x \| x_A^*) = \prod_{\alpha \in V \setminus A} f(x_\alpha | X_{pa(\alpha)} = x_{pa(\alpha)})\Big|_{X_A = x_A^*} \tag{2.2.1}$$

where $f(x \| x_A^*) = f(x | \mathrm{do}\{X_A = x_A^*\})$ is the distribution of X under the intervention in which a subset A of the nodes V in the network have been imposed a value x_A^*. An example of this could be when a central bank decides to fix its exchange rate with respect to another currency and we want to see the impact on the other currencies in the market. The conditional distributions $f(x_\alpha|X_{pa(\alpha)})$ are assumed stable under interventions that do not involve x_α — hence the condition that α represent nodes in V but not in in the intervention subset A. This conditioning by intervention allows for a much more specific causal interpretation than the conditional distribution. The random variable of interest X is a causal GM if it is a directed GM, as described above, such that the intervention factorization in Equation 2.2.1 holds. This definition is often referred to as Lauritzen's causal graph with interventions by replacement [131] or as Pearl's do-intervention [161].

2.2.2 Defining Contagion

In this section, we will present a network-based measure for contagion. It is based on the causal graphical models presented in the previous section, combined with an autoregressive approach to contagion as described in [79]. [79] introduce a structural vector autoregression (VAR) to analyse the contagion impact on the cost of insuring public debt during the European sovereign debt crisis.

We propose a structural equation model for the evolution of the log returns X_{it} on a financial asset i at time t through a structural VAR consisting of an autoregressive part $AR(i,t)$ and a network contagion part that we will call NECO(i,t):

$$\underbrace{X_{i,t}}_{\text{LogReturn on Asset i}} \leftarrow \alpha_{0,i} + \overbrace{\sum_{\ell=1}^{L} \alpha_i^\ell X_{i,t-l}}^{\text{AR(i,t)}} + \overbrace{\sum_{j\in pa(i)} \beta_{ji} \underbrace{X_{j,t}}_{\text{Assets that affect asset i}}}^{\text{Contagion=NECO(i,t)}} + \underbrace{\varepsilon_{i,t}}_{\text{Noise}} \tag{2.2.2}$$

where ℓ are the number of considered lags for the autoregressive part, α_i^ℓ are the autoregressive coefficients at lag ℓ for asset i and β_{ji} the causal effects of asset j on asset i or X_j on X_i. In the causal literature, causal effects are defined as the partial derivative of the expected log return $\frac{d}{dx_j}E(X_{i,t}|X_{-i,t}, X_{i,t-\ell})$ which in our case is equal to β_{ji}. Just like we did for the autoregressive part, lags can be added to the NECO part. We leave this extension for future research. The arrow in Equation 2.2.2 is to be interpreted in a generative manner, such as in a structural equation model [25]. The right-hand side is the driving force behind the value of X_{it}.

Equation 2.2.2 can be summarised as:

$$X_{it} = \alpha_{0,i} + \mathrm{AR}_{it} + \mathrm{NECO}_{it} + \varepsilon_{it} \tag{2.2.3}$$

where the NECO_{it} measures the totality of the contagion effect on the market. As a measure of the impact of contagion on the price of an individual asset, we propose the Network Contagion Factor (NECOF), which is computed as follows:

$$\mathrm{NECOF}(i) = \frac{\sigma^2_{i,NC}}{\sigma^2_{i,AR} + \sigma^2_{i,NC} + \sigma^2_i} = 1 - \frac{\sigma^2_{i,AR} + \sigma^2_i}{\sigma^2_{i,AR} + \sigma^2_{i,NC} + \sigma^2_i} \tag{2.2.4}$$

where $\sigma^2_{i,AR} = V(AR_{it})$, $\sigma^2_{i,NC} = V(\mathrm{NECO}_{it})$ and $\sigma^2_i = V(\varepsilon_{it})$, under the assumption that AR_{it}, NECO_{it} and ε_{it} are independent. This independence assumption is realistic, given that the considered contagion is assumed to be instantaneous and therefore by definition isolated from the idiosyncratic effects AR_{it}.

The NECOF is expressed in percentages and shows the impact of contagion on the return of the asset i. A NECOF of 0% would mean that contagion has no impact on the considered asset. On the opposite end of the spectrum, a NECOF of 100% would indicate that the impact of contagion for the given asset is absolute. The NECOF measure alone is very useful to identify which assets are at higher risk of outside influence - information that can be useful for investment and diversification strategies alike. We call the NECO a network-based measure because it depends on the underlying causal graph.

2.2.3 Identifying Contagion Paths

This section shows how the causal networks and causal NECO coefficients β_{ji} in Equation 2.2.2 are estimated. Once we estimate the NECO coefficients and the NECOF we can recover the contagion factor for each financial instrument. In order to estimate the causal coefficients correctly, we first need to establish the causal structure, which means finding all the causal parents for every considered asset i in our graph.

We estimate the causal structure using a more robust version of the standard PC-algorithm from [198] called the PC-stable algorithm [37][1]. The PC-Algorithm uses two steps in order to find the sources of contagion, which are summarised in A.1. We add a third step to estimate the size of the contagion effects, β_{ji}, by performing a series of linear regressions on Equation 2.2.4. Given the set of parents for each asset, the non-zero β_{ji} are to be estimated from the obtained

[1]We perform the PC-Algorithm using the pcalg package in R as described in [118, 94]

DAG. If at the end of Step 2, we achieve a Completed Partially Directed Acyclic Graph (CPDAG), a subset of Markov equivalent DAGs that can explain our data, we will also find a multiset of possible NECO estimates $\hat{\beta}_{ij}$. We can combine these into a range estimator as in [137]. Once we have estimated the NECO we can estimate the NECOF of interest for each financial instrument in our model using the following equation:

$$\widehat{\text{NECOF}}_i = 1 - \frac{RSS_{NECO}(i)}{SS(i)} \tag{2.2.5}$$

that we obtain by applying the Type II[2] Sums of Squares to Equation 2.2.3 such that:

$$RSS_{NECO}(i) = \sum_t \left[X_{i,t} - \widehat{\text{NECO}}(i,t)\right]^2 \text{ and } SS(i) = \sum_t \left[X_{i,t} - \mu_{i,t}\right]^2 \tag{2.2.6}$$

2.2.4 Community Detection

Using the contagion paths established in the previous section we can identify communities of financial instruments. These communities are groups of nodes that are more connected among themselves within the group than with the other groups. These groups are called communities, clusters or modules. Communities can be seen as sub-graphs that have specific properties not shared by the whole network and this allows for a next-level analysis of the network, moving from a single node to a more meaningful structure. These communities can be also seen as meta-nodes when representing and analysing very large networks, where considering and plotting each singular node would not be practically feasible.

There are many different algorithms to identify communities within a network; for a comparative study see [130]. We will be using the Louvain algorithm from [24] to establish communities among the nodes, because it is a benchmark among the clustering algorithm thanks to its robust and efficient results, making the results easier to compare with other studies.

2.2.5 Creating Dynamic Contagion Maps

In the previous sections, we defined a measurement of the contagion and the sources of this contagion, assuming that the NECO and its coefficients would remain constant for the entire time period in consideration. In this section, we will add a dynamic component, and in doing so not only allow for the AR(i,t)

[2] Similar to Type I but not dependent on the order of entry of terms into the model [225].

and NECO(i,t) in Equation 2.2.2 to change with time but also for the whole causal structure of our contagion graph to change with time. A dynamic version of (2.2.2) is written as,

$$X_{i,t} \leftarrow \alpha_{0,i} + \sum_{l=1}^{L} \alpha_{l,i}^{t} X_{i,t-l} + \sum_{\forall j \in pa(i,t)} \beta_{j,i}^{t} X_{j,t} + \varepsilon_{i,t} \tag{2.2.7}$$

describing a dynamic causal graphical model. From an inferential point of view, we will estimate the coefficients in a piecewise-constant way. At each timepoint t we evaluate a new DAG and the associated $\widehat{\text{NECOF}}$ estimates, creating a sequence of contagion maps. The contagion effect is considered to be contemporaneous within the considered window of time $[t-1,t]$. The length of the window will vary with the use case and depends on the data being analysed, what kind of short or long-term trends are associated and the purpose of the study.

2.3 Description of Empirical Exchange Rates 2000-2021

The Forex market is an important financial market, trading \$6.6 trillion per day [221]. Given that the most traded exchange rates are those over the USD we consider the interaction of 23 exchange rates over the USD for the years 2000-2021 as published daily by the Federal Reserve of New York. This allows us to evaluate the networks among highly traded currencies by expressing their value in terms of the US Dollar, the most liquid of all currencies, based on reliable historical data. Alternative approaches include using exchange rates based on a special drawing right (SDR) as in [215] or a benchmark based on the average, or geometric average, of different exchange rates as in [107] and [81]. SDR reflects the price of a basket of five major currencies and is periodically rebalanced and published on a daily basis by the International Monetary Fund (IMF). Another option for the base currency is to choose a currency that is of lesser importance, but still not completely illiquid. An example of this approach can be found in [122] which uses the Turkish Lira as a base. A comparison of different currencies being used as the base currency can be found in [128]. One final approach taken by [63] is that of ignoring the base currency issue altogether and using each exchange rate as a separate financial asset.

We consider the log returns on the spot exchange rates. Table 2.1 shows summary statistics for the 23 currencies considered here. The distribution of log returns on the exchange rates is often assumed normal, but as with most financial assets, there is the presence of fat tails as can be seen in the last column

Code	Currency Name	Minimum	Median	Mean	Maximum	StDev	Skewness	Kurtosis	Jarque Bera
AUD	Australian Dollar	-0.0771	-0.0003	-0.0000	0.0822	0.0080	0.6261	11.8178	31335
BRL	Brazilian Real	-0.0967	0.0000	0.0002	0.0867	0.0104	-0.0052	8.3711	15548
CAD	Canadian Dollar	-0.0507	0.0000	-0.0000	0.0381	0.0056	-0.0714	5.6143	6998
CHF	Swiss Franc	-0.1302	0.0000	-0.0001	0.0889	0.0067	-1.1203	35.1797	275708
CNY	Chinese Yuan Renminbi	-0.0202	0.0000	-0.0000	0.0182	0.0015	0.1610	22.2054	109425
DKK	Danish Krone	-0.0580	-0.0000	-0.0000	0.0494	0.0061	-0.1555	4.5734	4662
EUR	Euro	-0.0463	0.0000	-0.0000	0.0300	0.0060	-0.0775	2.4789	1369
GBP	British Pound	-0.0443	0.0000	0.0000	0.0817	0.0060	0.6707	10.6119	25385
HKD	Hong Kong Dollar	-0.0045	0.0000	0.0000	0.0033	0.0003	-1.2511	26.3958	155978
INR	Indian Rupee	-0.0376	0.0000	0.0001	0.0394	0.0044	0.1993	10.2796	23481
JPY	Japanese Yen	-0.0522	0.0001	0.0000	0.0334	0.0062	-0.3187	4.3307	4251
KRW	South Korean Won	-0.1322	-0.0001	-0.0000	0.1014	0.0067	-0.5492	50.7331	571339
LKR	Sri Lankan Rupee	-0.0339	0.0000	0.0002	0.0641	0.0029	2.7858	75.1652	1260440
MYR	Malaysian Ringgit	-0.0366	0.0000	0.0000	0.0277	0.0035	-0.2993	8.5426	16271
NOK	Norwegian Krone	-0.0644	-0.0001	0.0000	0.0612	0.0077	0.2195	4.7907	5135
NZD	New Zealand Dollar	-0.0593	-0.0002	-0.0001	0.0618	0.0082	0.3886	4.7691	5180
SEK	Swedish Krona	-0.0530	-0.0000	0.0000	0.0547	0.0074	-0.0544	3.8772	3338
SGD	Singapore Dollar	-0.0238	-0.0001	-0.0000	0.0269	0.0033	0.0239	4.9815	5507
THB	Thai Baht	-0.0353	0.0000	-0.0000	0.0447	0.0037	0.1773	11.5344	29547
TWD	New Taiwan Dollar	-0.0342	0.0000	-0.0000	0.0320	0.0031	-0.3537	16.1074	57676
VEB	Venezuelan Bolivar	-11.5129	0.0000	0.0028	5.8126	0.1857	-37.8364	2974.1949	1963940216
ZAR	South African Rand	-0.0916	-0.0001	0.0002	0.0843	0.0108	0.2563	4.3936	4341

Table 2.1: Overview of the summary statistics for the dataset of log returns on individual 23 exchange rates over the USD, for the period January 2000 to April 2021. The higher the Jarque Bera Test the less normally distributed the data is (all of the statistics have a p-value < 0.0001 *). There are 5325 observations for each currency, with no missing values.*

of Table 2.1. [112] analyse this problem, without finding any better alternative that would hold for every currency and time frame. Some studies even find that trading strategies based on the assumption of log-normality do in fact maximise profit [184]. The presence of fat tails will cause the significance level for the individual conditional independence tests within the PC-Algorithm to be empirically slightly higher than the nominal value.

Figure 2.1 shows the series of 23 log returns for our time frame, from January 2000 until April 2021. As we can see the CNY, HKD, LKR, MYR and VEB present periods of very low variance. These currencies have been pegged to the USD, whereby the Central Banks keep their exchange rate within a clearly pre-defined band. Within these bands, the market is allowed to operate, and therefore some level of contagion can spread to and from these currencies. China switched from a fixed exchange rate to a less restricted regime in July 2005, with Malaysia following suit. This has resulted in the value of those currencies getting closer to their perceived market value — for the MYR an appreciation and for the CNY a depreciation. The Chinese government is only slowly allowing more flexibility of the exchange rate [174], but it remains a highly influential rate and is the first emerging market currency to be held as a reserve by the International Monetary Fund (IMF). The Venezuelan Bolívar (VEB) also shows an unusual his-

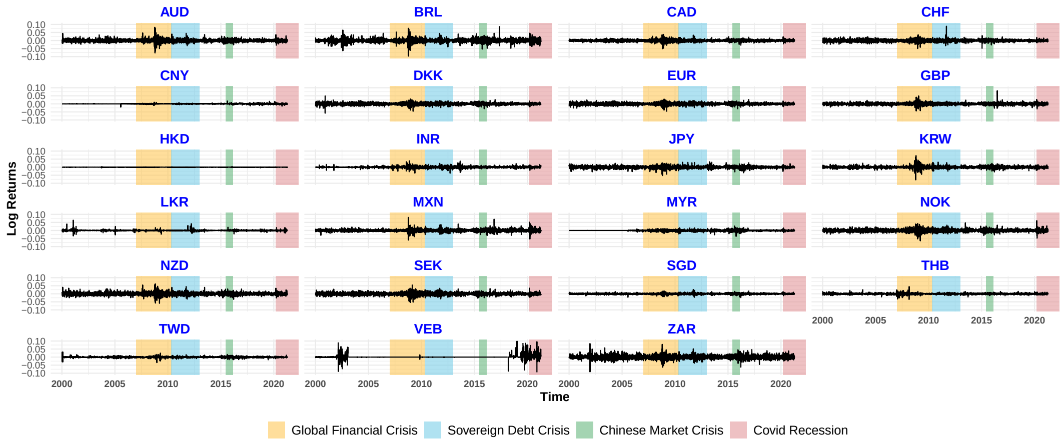

Figure 2.1: Log returns for each currency. The financial crises highlighted are the Global Financial Crisis (August 2007 - April 2010) with the following Sovereign Debt Crisis (May 2010 - December 2012), the Chinese Market Crisis (June 2015 - February 2016) and the COVID-19 Recession (starting March 2020).

tory of returns. This dynamic reflects periods of hyperinflation and the subsequent government-sanctioned devaluation of the currency. We will show that our model is capable of handling even these extreme types of behaviour.

In Figure 2.1 the highlighted areas show the financial crisis considered in the subsequent analysis. The financial crises highlighted are the Global Financial Crisis (August 2007 - April 2010) with the following Sovereign Debt Crisis (May 2010 - December 2012), the Chinese Market Crisis (June 2015 - February 2016), and the COVID-19 Recession (starting March 2020). All exchange rates have been impacted by the Global Financial Crisis at least in some way; even highly managed currencies like the Venezuelan Bolvar (VEB) or the Chinese Yuan Renminbi (CNY) show some turbulence during this period. In the following sections, we will analyse the impact of these crises on contagion in the Forex market.

2.4 Contagion in the Currency Markets

This section will present the results of applying the methods in Section 2.2 to the Forex returns data described in Section 2.3. There is a vast amount of research into the interdependencies within the Forex market. The scope of this paper is to show how this innovative causal approach can find these dependencies within a single model with an immediate and easy interpretability of the results and gain a better understanding of the underlying dynamics of the market.

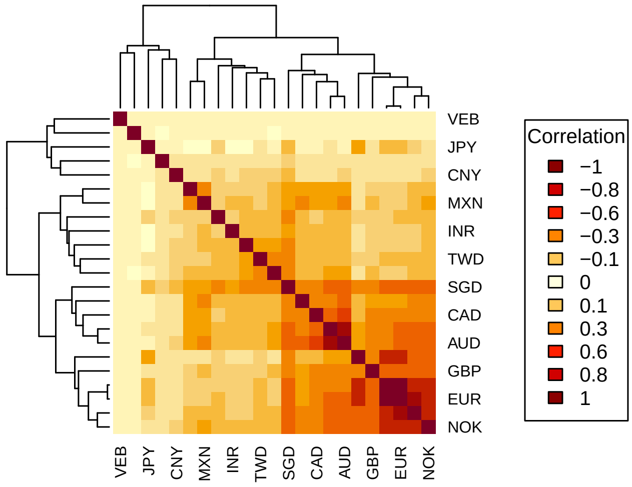

Figure 2.2: Correlation heatmap for the time period 2000-2021

What a static approach using the correlation networks can reveal is shown in Figure 2.2, which represents the correlation heatmap between the currencies considered. The heatmap can detect the outline of the main clusters present on the Forex market. The clusters include a European cluster (EUR, NOK, SEK, DKK, GDP, CHF), the Commonwealth cluster (AUD, NZD, CAD, ZAR, SGP), a small cluster of emerging economies (BRL and MXN) and then a somewhat sparse geographically based cluster of the Asian currencies. Applying the models from Section 2.2.1 we will be able to look more deeply and in more detail at the dynamics within the Forex of these currencies.

Figure 2.3 shows the causal network created from the complete set of returns data (2000 - 2021). The nodes are the 23 currencies in our dataset, and the links represent the causal effects of contagion from one currency to the other. The colour of the nodes shows the community classification, based on the Louvain clustering algorithm described in Section 2.2.4. The green arrows indicate a positive contagion coefficient, and the red arrows a negative coefficient of the corresponding causal effect. The width of the links reflects the strength of the causal effect.

The network shows some obvious connections, like the Euro (EUR) having an impact on the Danish Krone (DKK) exchange rate, and whole contagion paths, like the one starting from the Euro (EUR) to the British Pound (GBP), to the Canadian Dollar (CAD) and ending with an effect on the Mexican Peso (MXN). Most of the contagion coefficients are positive, apart from the effect of the Japanese Yen

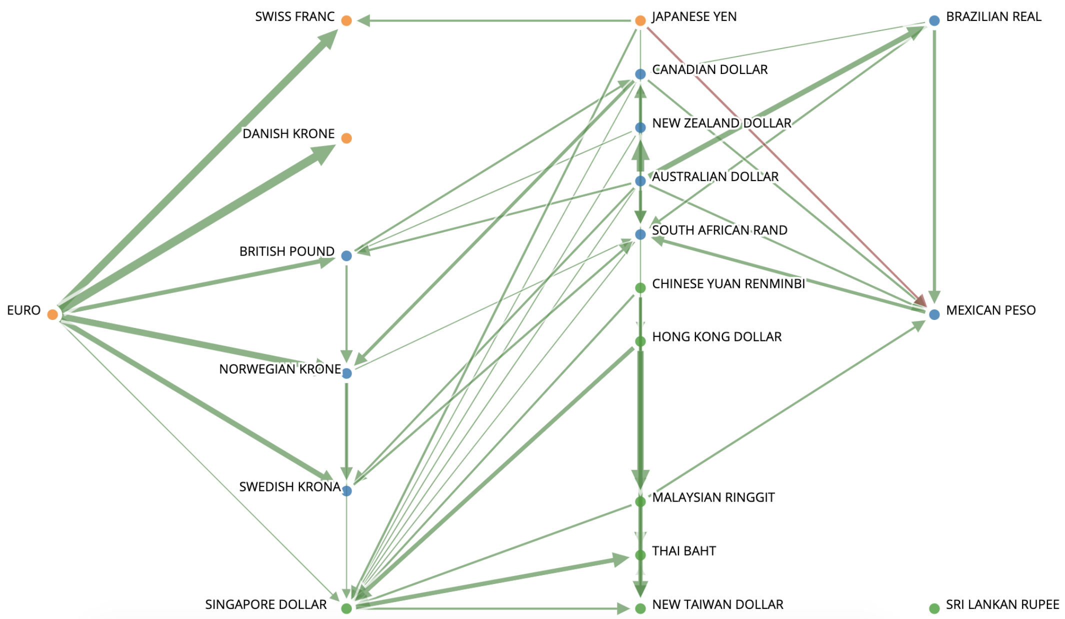

Figure 2.3: Causal network showing the interconnections and contagion paths within the 23 foreign exchange rates for years 2000-2021. Note that node colours represent Louvain clusters; green (red) arrows indicate a positive (negative) contagion coefficient of the corresponding causal effect; the width of the arrows reflects the strength of the causal effect.

(JPY) on the Mexican Peso (MXN), which is indicative of the carry trade activity between the pair [95, 165, 226]. With causal networks, we can establish where the currency of interest is positioned and, from a systemic risk point of view, where contagion could come from. Informed readers will be aware of some of these connections. Since causal inference has not been applied to the analysis of contagion within the Forex market before, we use them to validate our model. A key takeaway from this section is how much more informative Figure 2.3 is compared to Figure 2.2. Furthermore, our approach enables us to quantify these causal effects and compare them through time, as seen in the next Section.

2.4.1 Overall Development of Contagion on Forex

We estimate the model in Section 2.2.5 with a time window of one year (250 business days), rolled over every three months. This allows us to analyse the development of contagion on the Forex from a macroeconomic point of view. For each contagion map, we estimate the NECOF for each time period. We consider and analyse three contagion indices: Market NECOF, number of clusters and market density. We tested the reaction of these indices to major financial crises and reported detailed results in A.3.

Figure 2.4(a) shows the evolution of the market NECOF. The contagion increased significantly at the beginning of the 2000s from a NECOF of 17% to 37%, and then oscillated at a higher level, between 25% and 35%. For a more in-depth analysis of the change in contagion network statistics over time, refer to A.3. The rise at the beginning is caused by some of the currencies that were formerly pegged to the USD becoming freer, and hence more connected to the other currencies on the Forex. This initial rise is interrupted by a peak around the sovereign debt crisis of low- and middle-income countries that began in 2002 as defined by [129]. We tested the reaction of the three contagion indices to major financial crises and reported the results in A.3. After major events during financial crises, the market NECOF tends to increase (p-value < 0.0001). This is particularly evident in the case of the Global Financial Crisis (p-value $= 0.001$), the Sovereign Debt Crisis (p-value < 0.0001), and the COVID-19 Recession (p-value $= 0.036$). The NECOF values appear to be promising indicators of what is happening in the economy and how systemic risk evolves during periods of high uncertainty.

Figure 2.4(b) shows how clustering evolves over time and in response to economic events. The clustering effect is represented by counting the number of clusters present on the network using the Louvain algorithm: the lower the number of clusters, the higher the clustering of the network. The high number of clusters at the beginning of the considered period is driven by the many currencies that were pegged to the USD during the 2000s, whose exchange rates over the USD remained nearly constant. If a currency does not show any variation, it will by definition be counted as its own cluster, hence increasing the overall number of clusters detected. As currencies became more freely traded, similarly to the market NECOF rose, so did the number of clusters drop. This negative relationship is confirmed by a correlation between the market NECOF and the number of clusters of -0.6. Clustering is an important topic when analysing the Forex topology, especially during periods of financial crisis [215, 122, 128, 216]. There is no evidence that the number of clusters in the causal contagion network changes during periods of financial crisis (p-value $= 0.38$), although in the next section we will find that the structure of the clusters changes.

Figure 2.4 (c) shows that Forex is predominantly a sparse graph with a relatively low market network density overall. The effect of globalisation and increased interconnectedness is reflected in the positive trend in density development over the first decade of the 2000s. Once the Global Financial Crisis began in 2007, however, the density decreased slightly and then levelled off. What we can

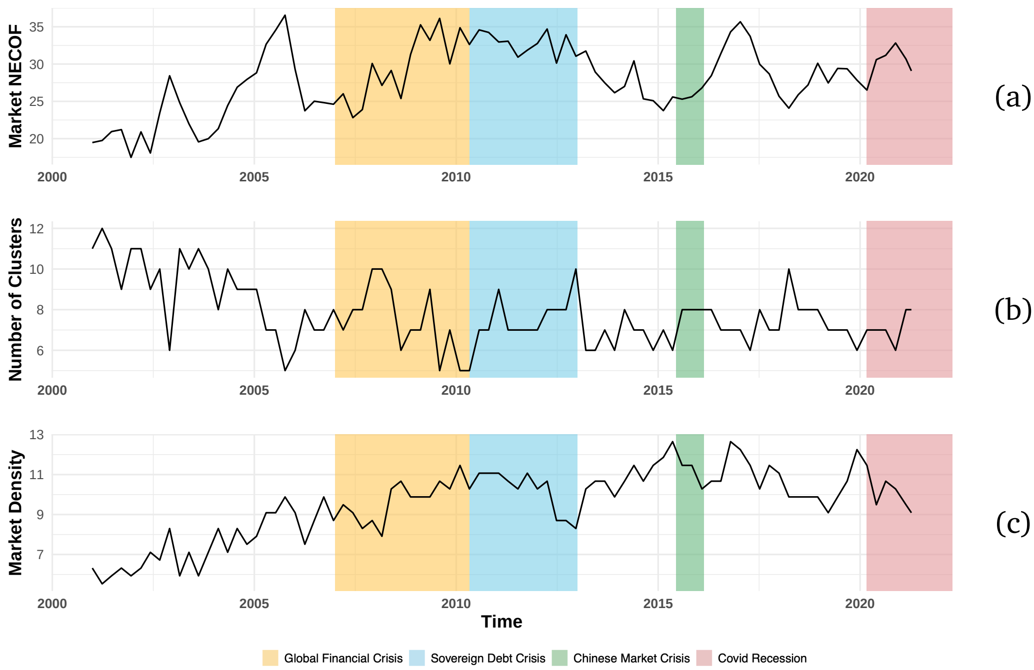

Figure 2.4: (a) The market NECOF as the average over all the currencies; (b) The number of Louvain clusters as a measure of clustering; (c) The market network density of the Causal Network, which is the percentage of links to all possible links. The financial crises highlighted are the Global Financial Crisis (August 2007 - April 2010) with the following Sovereign Debt Crisis (May 2010 - December 2012), the Chinese Market Crisis (June 2015 - February 2016) and the COVID-19 Recession (starting March 2020).

observe is that the contagion network tends to become denser during financial crises (p-value = 0.009).

2.4.2 Contagion Clustering in Forex

In Figure 2.5 we show the evolution of the clustering over 21 years of data. The structure of these cluster plots is similar to that of an adjacency matrix and shows the rows and columns labelled with the different currencies. However, instead of showing links or link weights, the matrix shows how often pairs of currencies belong to the same cluster. The more often two currencies are assigned to the same community by the Louvain algorithm during the specified time frame, the darker and larger the dot at the intersection of the pair becomes. When the dot linking two currencies is a solid dark blue, it indicates that the currency pair remained in the same cluster throughout the considered period. The Euro (EUR) and the Danish Krone (DKK) are the only pairs with a perfect 100% connection. Empty cells represent currencies that were never in the same cluster during the specified time frame.

The classically assumed geographically based clusters can be roughly identified within these cluster plots — but the structure is not so obvious and, more importantly, it changes in reaction to economic events. The main clusters can be roughly divided into a European cluster, a Commonwealth cluster, an Emerging Economies cluster and finally an Asian cluster. Within the European clusters, we have some clear oddities. The British Pound (GBP) switches between the European cluster and the Commonwealth cluster. The Japanese Yen (JPY) is often more connected to the European cluster and specifically the Swiss Franc (CHF), especially during a crisis. The behaviour of CHF and JPY during a financial crisis is very interesting and will be described in more depth in Section 2.4.3.

As stated in the previous section, we do not find a significant change in the number of clusters before and during a financial crisis. Where we find a difference however is in the stability of memberships within a cluster. Stability refers to whether a currency remains within the same cluster between two periods. In Figure 2.5 the increased stability is reflected in the prevalence of large, dark blue dots and fewer, smaller, light dots. By testing the stability of memberships within a cluster, we find that during the Sovereign Debt Crisis, the clustering stability was higher than that in the Global Financial Crisis (p.value=0.0139).

This effect is visible when we compare Figure 2.5(b) and Figure 2.5(c) – the clusters in Figure 2.5(c) (the sovereign debt crisis) are much more defined than the ones in Figure 2.5(b) (the Global Financial Crisis). The European cluster fully de-constructed during this period and does not really ever recover its original structure, unlike the other clusters. The European sovereign debt crisis seems to have caused structural changes to the contagion structure between the Euro and the rest of the European currencies, with effects still lasting to date. These changes cannot really be attributed to the growth and change within the Eurozone as such, because most of these changes took place well before the sovereign debt crisis hit.

During the turbulent times of the Global Financial Crisis, there was a notable shift in clustering. The Commonwealth cluster, traditionally comprised of nations like New Zealand, Australia, Canada, South Africa, and occasionally Britain, merged with the New Economies cluster, which includes countries such as Mexico and Brazil. The clustering effect of the Global Financial Crisis period can be seen much more clearly when considering changes in clustering during specific time periods than when considering clusters using the full 21-year time series.

During the Chinese Market Crash, the market NECOF only minimally went up and the global number of clusters did not change much during this crisis, but Figure 2.5(e) shows some signs of clustering nonetheless. The European, the Commonwealth and the Emerging Economies clusters are notably more defined. There are some changes within these communities: the Norwegian Krone (NEK) and Swedish Krona (SEK) join the Commonwealth cluster and the Indian Rupee (INR) and the South African Rand (ZAR) move from their respective clusters to the Emerging Economies cluster. It is notable that India, although a member of the Commonwealth, finds itself rarely if at all within the Commonwealth cluster.

Unsurprisingly, the Chinese Market Crash seems to have had the largest impact on the structure of the Asian cluster. Even though Figure 2.5(e) covers a relatively short period, the Asian currencies are spread all over the map and almost do not look like a clear cluster at all. During the crisis, several of the Asian currencies found more connection with currencies outside of the Asian cluster. The Chinese Yuan Renminbi (CNY), for example, interacts much more with other currencies than in the previous periods. Also, interesting is the behaviour of the Swiss Franc (CHF) during this crisis — it completely leaves the Euro-centred community and has only connections with Asian currencies, probably due to its traditional role of a safe-haven currency [98, 43, 111].

The period from March 2016 until February 2020 is a period of relative calm, and the clustering resembles the first plot of similar calm, apart of course from the European cluster. One other interesting mention is the Japanese Yen (JPY) which finally moves away from the Asian cluster almost completely. In March 2020 the COVID-19 Pandemic became a global crisis, and the COVID-19 Recession officially started. We see from Figure 2.4 that the market contagion expressed in the market NECOF increased immediately at the beginning of the pandemic, but the number of clusters did not change very much. What changed, as in the previous crisis, is the redistribution of the currencies within the clusters themselves — but in a different and more dramatic fashion. Figure 2.6 shows this redistribution. The network is much more compact and in fact, the density does sharply go down, i.e. the network has fewer links. This means we have new clustering with a lower number of links, but more significant links that lead to a much higher contagion on the markets. The CNY and the INR join the European cluster, whereas the GBP, SEK and NOK move together to join the Commonwealth cluster. That the impact of the COVID-19 pandemic on the contagion map and clustering within the Forex market would be somewhat different from the previous financial crisis was to be expected — even the markets reacted very

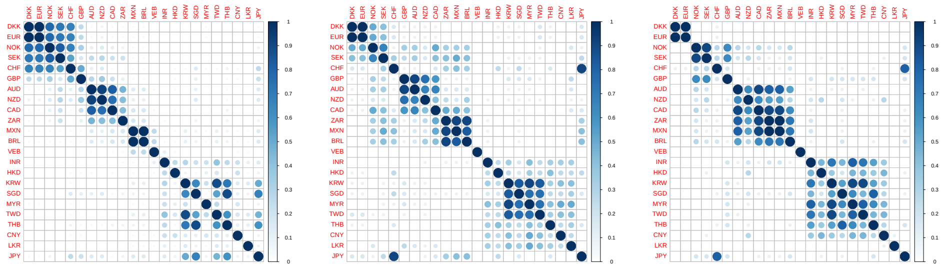

(a) Jan 2000 - Jul 2007 (b) Aug 2007 - Apr 2010 (c) May 2010 - Dec 2012

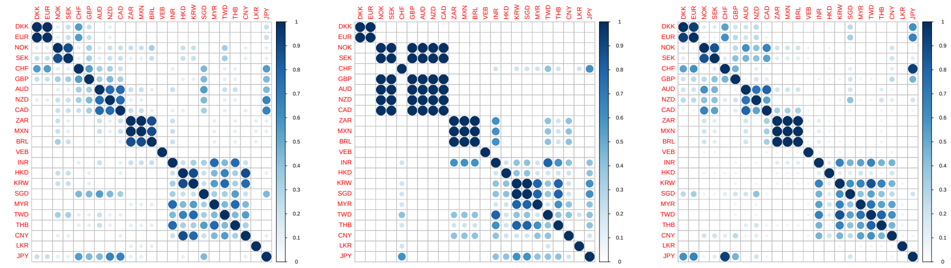

(d) Jan 2013 - May 2015 (e) Jun 2015 - Feb 2016 (f) Mar 2016 - Feb 2020

Figure 2.5: These figures show the evolution of the clustering over time and the impact of financial crises. The results are normalised to account for the different lengths of each period. We subdivide our dataset into the following periods: (a) Beginning Period: January 2000 - July 2007; (b) Global Financial Crisis: August 2007 - April 2010; (c) Sovereign Debt Crisis: May 2010 - December 2012; (d) Intermediate Period January 2013 - May 2015; (e) Chinese Stock Market Crash: June 2015 - February 2016; (f) Post-Crisis Period: March 2016 - February 2020.

differently. The stock markets fell faster than ever before[3] and the impact of the COVID-19 Recession seems to be having an impact on the global structure of the world economy [32, 206]. The cross-border financial interventions and foreign aid spending reached unprecedented levels during the COVID-19 pandemic [158]. All of these could also explain changes within the clustering.

2.4.3 Individual Network Contagion Dynamics

In this section, we will examine some interesting behaviours of individual currencies to illustrate the different types of scenarios that can be encountered on the

[3]For example the S&P 500 index fell by 34% between Feb. 19 and March 23, which constitutes the fastest fall in market in history, for further details see [181].

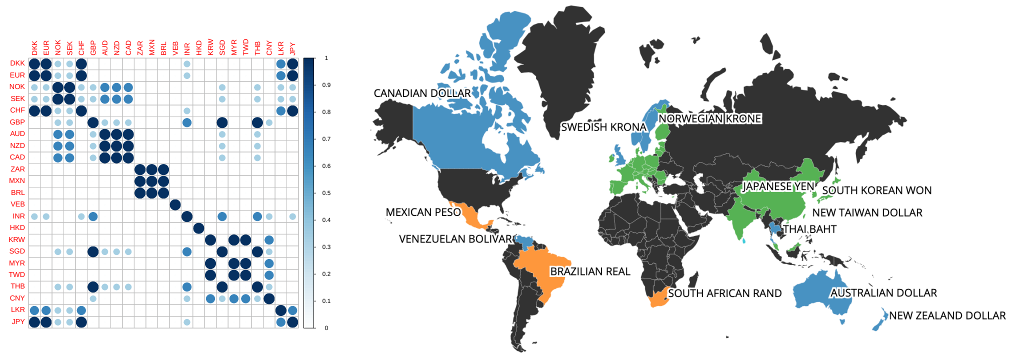

Mar 2020 - Apr 2021
(a) (b)

Figure 2.6: (a) Impact of the COVID-19 Recession on the contagion-based clustering, for the period starting in March 2020. (b) Louvain clustering of the contagion network during the COVID-19 Recession. Green Cluster: CHF, CNY, DKK, EUR, HKD, INR, JPY, KRW, MYR, SGD, TWD; Orange cluster: BRL, MXN, ZAR; Blue cluster: AUD, CAD, GBP, NOK, NZD, SEK, THB, VEB; Turquoise Cluster: LKR. The countries in black are not part of the analysis.

Forex market. Figure 2.7 shows how the individual network contagion factors, NECOFs, evolve for the chosen currencies: Swiss Franc (CHF), Danish Krone (DKK), Euro (EUR), British Pounds (GBP), Japanese Yen (JPY) and Malaysian Ringgit (MYR). EUR and DKK represent currencies with consistent high and low levels of contagion, respectively. As the membership of the cluster changes frequently in the GBP, it demonstrates the importance of dynamic analysis. In times of crisis, CHF and JPY are considered safe haven currencies. MYR was pegged to the USD until 2005 when it was abandoned and became a floating currency. We look at MYR in more detail to demonstrate how our model can deal with artificially controlled market movements.

We previously discussed the pair EUR and DKK and their very high correlation. As described in [210], this is a well-known special relationship and serves as a good example of how to read the causality network and NECOFs. Considering the causality links we are able to see immediately that it is the EUR that influences the DKK and not vice versa, as shown in Figure 2.8 and 2.2(b). The NECOF measure is much more meaningful than a correlation analysis in describing the risk of contagion of a financial asset. Two currencies can have a very high correlation and yet completely opposite NECOFs. In this example, the DKK shows a NECOF of almost 100% most of the time, with a median of 99.6%, while the NECOF of the EUR is usually at 0% — increasing significantly only once at

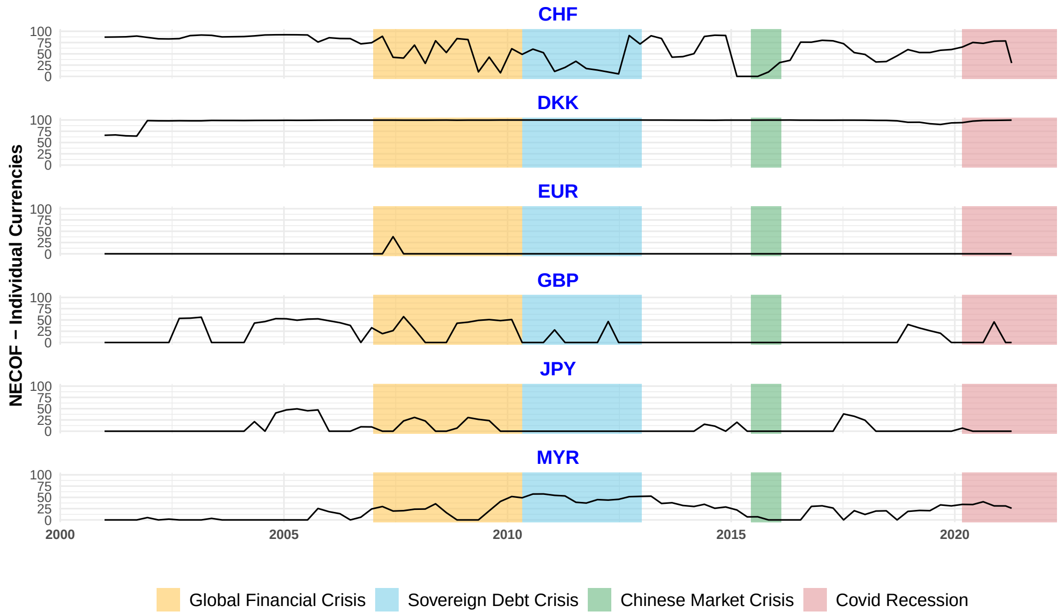

Figure 2.7: NECOF through time for the currencies CHF, DKK, EUR, GBP, JPY and MYR. The financial crises highlighted are the Global Financial Crisis (August 2007 - April 2010) with the following Sovereign Debt Crisis (May 2010 - December 2012), the Chinese Market Crisis (June 2015 - February 2016) and the COVID-19 Recession (starting March 2020).

the beginning of the Global Financial Crisis in 2007. It is evident that this causal relationship between the EUR and the DKK has been maintained throughout the 21-year period considered, demonstrating that the model can identify such a relationship simply by reviewing observational data on the prices, without requiring further analysis.

The next interesting example of a NECOF path is that of the GBP. In the previous section, we saw that the GBP has a tendency to switch clusters between Europe and the Commonwealth cluster. It usually is a currency that influences others and generally has a low NECOF around 0% — this is in line with what [79] find for the United Kingdom based on Corporate Default Swap spreads (CDS). In our study, however, we find that GBP's NECOF doesn't stay at zero throughout. This is especially true for the Global Financial Crisis (p-value = 0.0292). From a qualitative point of view, we see a NECOF for the GBP rise from 26% before the start of the Global Financial Crisis to 57% in the first months. Curiously, the GBP was hardly ever under the influence, from a contagion point of view, of the EUR — unless, again, there was a crisis. The GBP experienced contagion from the

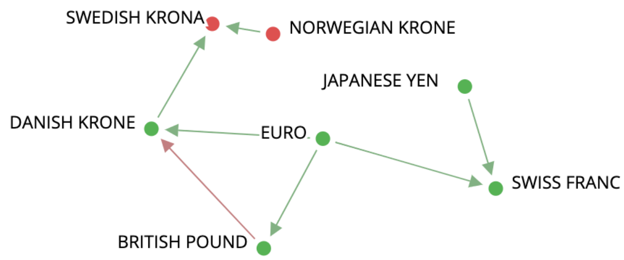

Figure 2.8: Subgraph of the Causal Network for the GBP in 2019.

EUR in 2007, during the Global Financial Crisis, and then during Brexit, around the time when the first draft withdrawal agreement was negotiated and endorsed by the EU members at the end of 2019, as seen in Figure 2.8.

The CHF is traditionally considered a safe-haven currency [98, 43, 111]. The CHF presents a relatively volatile and often high NECOF which is in contrast to the assumed safety of the CHF. CHF has a mean NECOF 62% and the highest volatility in terms of NECOF among all of the currencies considered, with a standard deviation of 28.5. What speaks for the safe haven status is CHF's behaviour during periods of financial crises: CHF NECOF reacts significantly to a period of financial crisis (p-value < 0.0001) and it goes down during the Global Financial Crisis (p-value 0.0016), the Sovereign Debt Crisis (p-value < 0.0001) and Chinese Market Crisis (p-value < 0.0001). The details of all the tests performed are in A.3. The low NECOFs in Figure 2.7 during the European Sovereign Debt Crisis are an example of this behaviour. The sharp fall of the NECOF at the end of 2014 was due to the intervention of the Swiss central bank, which tried to peg the currency to the EUR to prevent the increase in the value of CHF. This attempt was however scrapped in January 2015. And after the Chinese market crash passed, the NECOF shot up again.

In the past, the JPY was seen predominantly "as a low-interest rate or funding currency" [98], but the most recent studies classify it as a safe haven currency [27, 111]. As with the CHF, the JPY seems to appreciate in value during a crisis and during high volatility periods [173]. As Figure 2.7 shows, the JPY presents a very low NECOF most of the time. A median NECOF of 0% and a mean NECOF of 7% seem to validate the consideration of the JPY as a safe-haven currency. We further find a dependency of the CHF on the JPY, illustrated in Figure 2.8. The arrow goes from the JPY to the CHF in 75% of the networks (never in the opposite direction), which indicates that there is contagion that goes from the JPY to the CHF. The causal effect from the JPY to the CHF and the lower NECOF of the JPY in general would suggest that the JPY could be even considered a better safe haven than the CHF. This is exactly what [62] and [43] find when comparing these two currencies in terms of safe-haven characteristics.

The last currency we will consider in detail is the MYR. The MYR shows the NECOF evolution of an Asian currency that was pegged to the USD until the 2005. The NECOF fell in 2008 and after the 2015 Chinese stock market crash, mainly because of the intervention by the Bank Negara (the Central Bank of Malaysia) to prevent the exchange rate over the USD from plummeting. At times of important intervention by a central bank, the contagion from other currencies clearly decreases. The Malaysian economy and financial sector have grown during the 20 years considered here, and it is interesting how resilient to contagion it seems to be, unlike the other South Asian currencies maintaining an average NECOF of 20.5% with a maximum of 57% during the Sovereign Debt Crisis. Details for all NECOFs of all currencies can be found in the A.2.

2.5 Conclusion

Financial contagion measures have always had causal aspirations. In this work, we have unified this concept of causality with a defined measure that allows for a quantifiable causal interpretation of contagion relationships on the Forex market. We corroborate and extend results from different studies within one unifying framework and are able to answer very practical questions like how contagion spreads on the Forex market and which currencies are at the highest risk of contagion at any given time.

We have recreated a series of causal networks based on 23 exchange rates over the USD, spanning over 21 years, and present both the overall development of contagion on the Forex market as well as individual network contagion dynamics. We have shown how to read these causal networks as contagion maps to pinpoint sources of contagion and how the contagion paths on the Forex market evolve through time. We were able to identify a new promising group of financial indicators that take contagion and systemic risk directly into account. The newly identified measure of network contagion (NECO) seems to be of value for both a market level evaluation and the analysis of single currency alike. We discussed the network contagion factors' (NECOF) evolution for a subset of currencies, to demonstrate how this metric can be used to identify and evaluate a currency from an investment and hedging point of view.

In contrast to correlation networks, we obtain causal directions. Because they are inherently sparse, causal networks do not have to be filtered, e.g. via Minimum Spanning Tree (MST) methods [140], and can be easily analysed and evaluated. [120] compare the use of different centrality measures in constructing MSTs for the Forex market, (including degree, betweenness, closeness and eigenvector). Filtering is found to be important, but results are sensitive to the exact measure of correlation used as well as the distance measure. There are other filtering methods and so the choice of the method itself will also affect results [195, 213, 188, 141].

The application of causal graphical models to financial data is in its infancy and there are still interesting challenges. In our case, we assumed normality for log returns and found our sample to be stationary, but a model that could deal automatically with non-normality, fat-tails, heteroskedasticity and non-stationarity of the data would be advantageous for further applications in finance. The trading on the Forex market is active 24 hours a day [85] so we were able to use prices for all currencies at the same time instant. This is not the case for most other markets and assets being traded, and hence the impact of asynchronously observed returns on the analysis would have to be taken into account [28]. Whereas the contemporaneous contagion is the most significant in a fast moving and liquid market, it could also be of interest to consider economic cycles and longer time delays. Lastly, any measurable confounder can be added to our model, should we want to see how other variables, e.g. interest rates or inflation, impact the contagion. For unmeasured confounders and latent variables the Fast Causal Inference algorithm (FCI) [199, 200] can be used. Although our approach has been validated on the Forex market, it can be extended to markets of many other financial instruments. Looking at applications beyond financial contagion, our approach fits well the recent demand for explainability of machine learning and artificial intelligence methods [126].

Chapter 3

Causal Network Contagion Risk

3.1 Introduction

Risk is a crucial aspect of financial research, encompassing its definition, measurement, management, and pricing [182, 147]. Whatever its precise definition, there seems to be a consensus that contagion plays a role when evaluating the risk of holding an asset or financial instrument [66, 159, 3, 82, 53, 135, 93, 177].

Causal inference refers to strategies that allow one to draw causal conclusions based on data [161]. This is particularly challenging when experiments are not feasible and the researcher must rely solely on observational data [100, 217]. Separating correlation from causality has always been an important issue in any field of empirical research. The increased use of machine learning and AI methods has not only led to a major overhaul in the approach to handling data but a careful re-evaluation of causal algorithms for big data [175, 171, 209]. In medicine, this has already led to a renewed evaluation of causal inference methods [75, 192, 89, 125, 219]. As machine-learning and AI methods evolve, more and more research fields are adopting causal inference principles for actionable results, explainability, and safety in applications [7, 163, 194, 185].

The finance industry has been more reluctant to embrace novel developments in causal inference, partly as a result of the age-old conundrum of correlation versus causality [60, 105]. However, recently there have been some examples of using causal inference to improve stress testing [72], empirical research in the accounting field [86] and assessing causal factors determining the success of start-ups [74]. In finance, the two most common modeling approaches are Granger causality [87, 57] and Instrumental Variables [8]. Granger Causality can be seen as part of the transfer entropy based methods of causal analysis [186] and for nonlinear systems [102, 154]. The instrumental variable designs

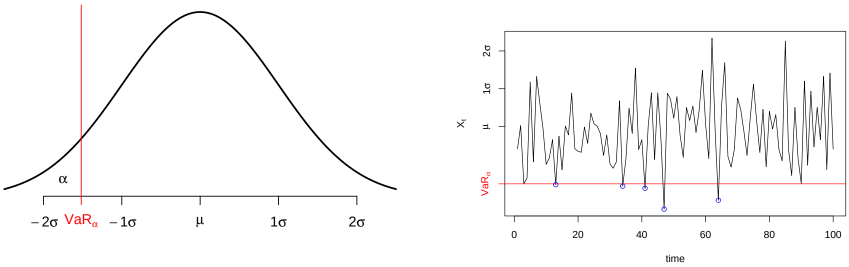

Figure 3.1: The Value at Risk, VaR$_\alpha$, is the α quantile of the distribution of a financial instrument. In the example above, the VaR for $\alpha = 5\%$ is depicted for a normally distributed instrument both cross-sectionally (left) and longitudinally (right).

are part of the quasi-experimental approaches [190].

Most modelling approaches of the effect of contagion on VaR do not have any explicit causal outset [64]. Most contagion-based VaR approaches, such as *CoVaR* [212], *SDSVaR* [3] and various other alternative attempts to integrate contagion within the risk management and the VaR measure [166, 151, 12], focus on correlations in systemic risk rather than causality.

In principle, the use of causal models offers the prospect of robust prediction [167, 168, 97]. If cause and effect are correctly identified, then the prediction will offer clear risk guarantees under any novel circumstances. There are some drawbacks as well. Any causal analysis is subject to model uncertainty, or it is possible that in stable financial circumstances, purely predictive methods perform better. However, as the financial system is subject to constant external shock, it is rare for the financial system to be in an ergodic equilibrium state. The financial system is subject to continuous outside influences, such as systemic, behavioural, and regulatory moves, ranging from wars to central bank interventions.

In this paper, we apply a causal network approach to risk management. In Section 3.2 we introduce a risk measure that relies on network contagion to capture volatility and spillover effects. When assessing the risk position, we consider price changes in related assets that have been identified from the data. Section 3.3 analyses the performance of this new approach in a simulation study. Section 3.4 applies the method to the Forex market and Section 3.5 concludes.

3.2 Value at Risk and Network Contagion

Value at Risk (VaR) is an established risk measurement used in finance to evaluate and compare the risk of holding a specific financial instrument or portfolio of different instruments. The VaR represents the α% quantile loss under normal market conditions for a specific holding time period [115, 114, 133, 48]. The future value of an asset can be considered as a random variable. Identification of the distribution of the value of this instrument X gives access to the value at risk at any level α,

$$\mathrm{VaR}_\alpha(X) = \inf\{x \leq \mathbb{R} : F_X(x) > \alpha\} \tag{3.2.1}$$

where F_X is the cumulative distribution function of X [10]. This corresponds to the α-quantile $q_\alpha(X)$ of the distribution, shown in Figure 3.1. In practical terms, the portfolio value or return will be smaller than the corresponding VaR over the holding period with a probability of at most α per cent.

Usually in finance, the Value at Risk is computed on the basis of the returns or log-returns on investment instead of the price for their statistical properties and computational ease. The returns-based VaR is then easily recomposed into the dollar loss VaR using the current price of such a portfolio. The Value at Risk measure has been developed within the investment banking sector to estimate market risk of their market positions, a task becoming more difficult especially due to the high volatility and interconnectedness of the markets and an increasing use of derivative instruments.

The base of the VaR measure, looking into establishing the probability of a loss, can be traced back to [54]. A more specific risk measure resembling VaR dates back to the 1920s, when the New York Stock Exchange imposed capital requirements on its member firms based on a similar metric [104]. The modern development of the VaR resembling today's approaches started in the 1980s [73]. It gained prominence in the 1990s when VaR was made the official risk measure for capital requirements for all financial institutions [19]. The success of the VaR measure was facilitated by the development of more accurate and rigorous estimation techniques. The most prominent example of such a method is the *RiskMetrics*^{TM} by J.P.Morgan [136]. For further analysis and more in-depth history see [50] and [115].

The VaR has some caveats. It does not represent the absolute worst-case scenario. Moreover, it has to be recomputed for each portfolio, as it is not an additive measure. Furthermore, it is not a coherent measure of risk, as the VaR of a portfolio can be larger than the VaR of the individual components of such a portfolio [10]. It is also important to choose the appropriate computational method for the particular scenario of interest. Nevertheless, the advantages of VaR are (i) it is an easily interpretable, standardized metric, (ii) it can be used to compare financial instruments and investments (iii) its probabilistic definition allows it to be tuned to the risk profile of the investor. An alternative to the value at risk is the expected shortfall (ES). The ES tends to need more data to be accurate [223], but it is a coherent measure of risk. There are ways to link the VaR to the ES [208].

3.2.1 Most Common Ways to Estimate Value at Risk

There are various methods for estimating the VaR. A parametric approach is the variance-covariance method, also called the delta-normal method, popularised by J.P. Morgan as the RiskMetrics method [136]. The parametric variance-covariance approach assumes a distribution for the returns on each instrument whose parameters are then estimated to find the VaR. The most common method is based on the multivariate normal distribution, which leads to an easy and quick estimation of the $\mathrm{VaR}_\alpha^{\mathrm{VC}}(X) = \mu_x - z_\alpha \cdot \sigma_x$, where μ_x is the mean and the σ_x the standard deviation of past realisations of the log-return X, and z_α is the Z-score of a standard Normal distribution at α%. The biggest drawbacks of this method are the, not always realistic, underlying assumptions of normally distributed log-returns, a constant volatility through time, and no correlations among the financial instruments.

To counteract the latter drawback, volatility can also be modelled using ARCH and GARCH models. When dealing with non-linearly priced instruments like options, the VaR is usually estimated using a Monte Carlo or stochastic simulation [133, 134]. The Monte Carlo VaR is computed as the quantile of the simulated returns X^{sim}, simulated from a chosen stochastic model for the behaviour of X – usually a GARCH process, $\mathrm{VaR}_\alpha^{\mathrm{GARCH}}(X) = \hat{q}_\alpha(X^{sim})$

In nonparametric VaR models, the joint risk factor distribution is constructed using historical data rather than assuming a specific functional form [29]. The easiest method would be to resample the past returns within the estimation window and pick the α-quantile for the VaR directly. The Historical Simulation (HS) develops this idea further by bootstrapping from the past returns, $\mathrm{VaR}_\alpha^{\mathrm{HS}} = \hat{q}_\alpha(X^{\mathrm{boot}})$, where X^{boot} represent the increased sample through bootstrapping. The HS is an easy and fast method. Even though it is nonparametric, it still relies on an assumption of stationarity in distribution and specifically volatility, often violated in practice due for example volatility clusterings [90].

Filtered historical simulation (FHS) combines the best of both parametric and non-parametric approaches [17, 18]. FHS runs an HS on volatility-rescaled past returns, thus maintaining the non-parametric nature of HS while allowing for varying volatility. Rescaling is done in two steps within a conditional volatility model (e.g., GARCH or AGARCH). Returns are first standardised by the estimated volatility of the day of the return and then rescaled by the forecasted volatility for the VaR holding period, $VaR^{\text{FHS}} = q_\alpha(X^{\text{boot}}_{re-scaled})$ This transformation reflects current market conditions and thus requires shorter estimation windows to simulate extreme events. In contrast to other Monte Carlo based approaches, the correlation matrix does not have to be estimated, as all rescaled returns at a specific event time are sampled together. Extreme observations or time-varying correlations may not be adequately considered by the FHS VaR [170]. Extensions of the FHS have been proposed by [110] as the Volatility-weighted HS and by [146], who combine FHS with extreme value theory. For comparisons, see [48], [142] and [170].

3.2.2 Definition of Causal Network Contagion Value at Risk

In this section, we propose a novel VaR procedure, based on causal network contagion. The method is based on the causal framework proposed by [177]. What separates this approach from standard dependencies analysis like [113], is the ability to identify the direction of the contagion: which assets influence, or export contagion risk to, each other and vice versa. The method does not assume any a priori causal model and considers only what can be gathered through observed data — through the external manifestation through the path of returns for each asset and their co-dependences.

Gaussian Causal VaR. In this paragraph, we first derive the form for the value at risk in case the causal contagion system can be described by Gaussian noise. We assume that the log returns for each instrument $X_{i,t}$ is described by a structural equation model (SEM), consisting of an autoregressive part and a contagion part,

$$\underbrace{X_{i,t}}_{\text{LogReturn of Asset i}} \leftarrow \alpha_{0,i} + \overbrace{\sum_{\ell=1}^{L} \alpha_i^{\ell} X_{i,t-l}}^{\text{Autoregressive}} + \overbrace{\sum_{j \in pa(i)} \beta_{ji} \underbrace{X_{j,t}}_{\text{Assets that affect asset i}}}^{\text{Contagion}} + \underbrace{\varepsilon_{i,t}}_{\text{Noise}} \tag{3.2.2}$$

where L is the number of lags considered for the autoregressive part, α_0 is the intercept, α_i^ℓ are the autoregressive coefficients of lag ℓ for asset i, $pa(i)$ are the causal parents of instrument i, β_{ji} are the causal effects of instrument j on instrument i, and $\epsilon_{i,t}$ are the error terms, assumed to be independent and normally distributed with variance σ^2. Contagion here is defined as the instantaneous causal structure of the financial system.

The SEM can be written more concisely in matrix form:

$$X_t = (\mathbb{1} - B)^{-1}\left(\alpha_0 + A \cdot X_{t-1:t-\ell} + \varepsilon_t\right) \tag{3.2.3}$$

where $\epsilon_t \sim N(0, \Sigma)$, where A is the matrix of the autoregressive coefficients and B the matrix of contagion effects. So conditional on the previous log returns, the log return is given as

$$X_t | X_{t-1:t-\ell} \sim N\left((\mathbb{1} - B)^{-1}\left(\alpha_0 + A \cdot X_{t-1:t-\ell}\right), (\mathbb{1} - B)^{-1}\Sigma(\mathbb{1} - B)^{-\top}\right) \tag{3.2.4}$$

Using the distribution in Equation 3.2.4, we can incorporate the causal network contagion (NECO) into the VaR definition. The Gaussian Causal VaR for each instrument i corresponding to a level of risk at α% is given as

$$\mathrm{VaR}_\alpha^{\mathrm{NECO}}(X_{(t,i)}) = \left[\ (\mathbb{1} - B)^{-1}\left(\alpha_0 + AX_{t-1:t-\ell}\right)\ \right]_i - z_\alpha \sqrt{[(\mathbb{1} - B)^{-1}\Sigma(\mathbb{1} - B)^{-\top}]_{ii}} \tag{3.2.5}$$

General Causal VaR. The definition of the Causal VaR above is based on the assumption of a multivariate normal distribution for the instruments considered for the causal networks. This assumption represents log-returns relatively well, but it is still often violated in reality where financial instruments are concerned. In this paragraph, we will relax the normality assumption by using the Gaussian copula transformation to capture the fat tails and other non-Gaussian behaviour typical for the returns on financial instruments [46, 1]. Copulas provide a flexible tool for understanding dependence between random variables, in particular for non-Gaussian multivariate data [153].

The log returns of the instrument i follow a marginal distribution F_i. We assume that there is a latent temporal process Z, which can be described by the mean of the causal SEM in Equation 3.2.2. This SEM describes the idealized, Gaussian version of the financial process, centred around zero with variance one. The connection between the idealized process Z and observable X is given through the copula transformation,

$$X_{t,i} = F_i^{-1}\left(\Phi\left(Z_{t,i}\right)\right), \tag{3.2.6}$$

where Φ is the CDF of a standard normal distribution. This gives us a direct way of defining a general causal VaR,

$$\mathrm{VaR}_\alpha^{\mathrm{NECO}}(X_{(t,i)}) = F_i^{-1}\left(\Phi\left(\left[(\mathbb{1} - B)^{-1}AZ_{t-1:t-\ell}\right]_i - z_\alpha \sqrt{[(\mathbb{1} - B)^{-1}(\mathbb{1} - B)^{-\top}]_{ii}}\right)\right), \tag{3.2.7}$$

where $Z_{t,i} = \Phi^{-1}(F_i(X_{t,i}))$. The copula transformation allows us to perform our analysis in the latent space without the need to assume lognormal distribution for the returns. The other advantage of using a Gaussian Copula is that, unlike most other copulas, it can handle high dimensionality, i.e., many interrelated financial instruments.

3.2.3 Estimation of Causal NECO Value at Risk

We assume that we observe multivariate time-series, $D_X = \{x_{t,i}\}_{it}$, of log-returns of p financial instruments instruments across N epochs. The estimation of the Causal NECO VaR is done in four steps. First, we estimate the underlying marginal distributions F_i for the instruments. Secondly, we estimate the causal structure on the transformed scale to see what financial instruments impact the risk of the instrument of interest. Thirdly, given the causal structure, we then estimate the contagion coefficients from Equation 3.2.2. Finally, the estimated marginal distributions $\widehat{F}_i$ and coefficients $\widehat{A}$ and $\widehat{B}$ are used to compute the VaR as in Equation 3.2.7.

Estimating the marginal distribution of financial instruments. For each instrument i, we estimate its marginal distribution F_i non-parametrically as the adjusted empirical distribution function,

$$\widehat{F}_i(x) = \frac{0.5 + \sum_{i=1}^{N} 1_{\{x_i \leq x\}}}{N+1}.$$

We then use this empirical distribution, to define a transformed dataset, D_Z, of normally distributed variables, $D_Z = \left\{z_t \in \mathbb{R}^p \mid z_{t,i} = \Phi^{-1}(\widehat{F}_i(x_{t,i})), t = 1, \ldots, N\right\}$. We use the adjusted empirical distributions, in order not to get degenerate values for z.

Discovery of Causal Structure We estimate the causal structure based on the transformed data D_Z in the form of a causal network as in [177], which is based on the PC-stable algorithm from [37]. The PC-stable algorithm is a more robust version of the original PC algorithm [198]. We implement the PC-stable algorithm with the R package `pcalg` [118, 94].

The causal connections in the structural equation model can be estimated on purely observational data. The algorithm is based on two main insights. The first insight is that if there exists a separating variable S that makes Z_1 and Z_2 conditionally independent, then Z_1 and Z_2 are not directly causally connected.

Secondly, if no such separating sets can be found between Z_1 and Z_2 as well as between Z_2 and Z_3, but Z_1 and Z_3 are separated, but not by Z_2, then it must be that Z_1 and Z_3 are direct causes of Z_2.

Estimation of Causal Effects Given the structure of the structural equation model in Equation 3.2.2, the coefficients A and B can be estimated via standard least squares. If there are no links present on the causal network, then the coefficient is put to zero. If the algorithm from the previous step is unable to deliver a fully directed network — where for some couples of instruments both directions of the contagion are just as likely given the structure of the causal network — we obtain a multiset of possible coefficients and combine these into a range estimator [137]. Given that the financial system is very big and interconnected, the chances of finding a fully directed network are very high.

Causal NECO Value at Risk With the estimated $\widehat{F}_i$, $\widehat{A}$ and $\widehat{B}$ we can then compute the VaR using Equation 3.2.7. We assume that any interim payments on the considered assets are either zero or reinvested continually in the asset itself, as is done in mutual funds. Furthermore, we take the time period for the VaR evaluation to be equal to the frequency of the recorded log returns — so if we consider daily log returns, we will compute the 1-day VaR_α.

3.3 Performance of NECO Value at Risk

In this section, we analyse the performance of the Causal NECO Value at Risk in predicting the correct levels of risk in various simulation studies. The performance of the Causal NECO VaR will be compared to the other four standard methods defined in Section 3.2.1, and analysed for different contagion structures. We consider the most common backtesting methods defined in Section 3.3.1, which are model-free and are good for comparing different methods.

3.3.1 Backtesting Measures

Backtesting allows for testing how well the VaR performed, using simulations or past historical data. An overview of the most common backtesting methods can be found in [31]. Part of the backtesting will be performed using the R package `MSGARCH` [9]. Below we describe the tests we will consider.

Average Exceedance Rate. The computationally easiest and most immediate test for the performance of the VaR is looking at the number of times the actual X_t log returns fell below the VaR, called violations or exceedances. Given the VaR definition from Equation 3.2.1, we expect X_t to be smaller than the VaR_α about α% times. The average exceedance rate $\hat{\alpha}$ should therefore be near the targeted α level of the VaR_α.

Actual over Expected Ratio. Actual over Expected Ratio (AE) measures whether the VaR computation method tends to have more or fewer violations than expected given the target α. A good method has the number of exceedances n_1 being αN. Defining $AE = \frac{n_1}{\alpha N}$, $AE > 1$ signifies that the method is not restrictive enough and underestimates the risk of the underlying investment. An $AE < 1$ shows the opposite — a method that is overly conservative and overestimates the risk. Both directions of this error can lead to costly mistakes.

LR Test of Unconditional Coverage. The coverage rate for the targeted rate α. The unconditional coverage test from [127] tests if the proportion of violations is significantly far away from the expected level or not, using the likelihood ratio test statistic $LR_{UC} = -2ln\left[\frac{(1-\alpha)^{n_0}\cdot\alpha^{n_1}}{(1-\hat{\alpha})^{n_0}\cdot\hat{\alpha}^{n_1}}\right]$ that asymptotically has a χ^2 distribution under the null hypothesis $H_0 : E\hat{\alpha} = \alpha$.

LR Test of Conditional Coverage. Conditional Coverage (CC) Test from [34] expands the UC test for the detection of violations clustering in time. No violation occurrence should be informative about the performance of the next-step VaR. CC checks that exceedance realisations $\{\mathbb{1}_{\{x_{1,i} \leq \text{VaR}_\alpha(1,i)\}}, \cdots, \mathbb{1}_{\{x_{T,i} \leq \text{VaR}_\alpha(T,i)\}}\}$, often called hit series, are distributed independently and identically. It uses a likelihood ratio test with 2 degrees of freedom.

Dynamic Quantile Test. The Dynamic Quantile (DQ) test described in [61] and [51] tests whether the violations, i.e. exceedance realisations, are not only uncorrelated among themselves but also with other lagged variables. The LR_{DQ} follows asymptotically a χ^2_p distribution with p degrees of freedom.

Absolute Deviation. The mean and max absolute deviation (AD) show the actual loss that would occur if an investor or bank had relied on the VaR prediction. As pointed out by [145] the AD measure is of great importance as large violations can lead to bank failures when capital requirements implied by the VaR threshold forecasts are not sufficient to protect against losses that are actually realized.

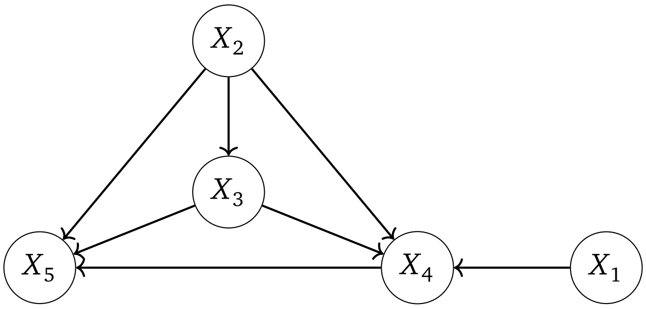

Simulated Network	
nr. links	7
density	0.7
market NECOF	47%

Figure 3.2: Simulated financial network with 5 instruments related to section 3.3.2

Average Quantile Loss. The Average Quantile Loss [84] is a weighted loss measure, defined as $QL_i(\alpha) = \frac{\sum_{t=1}^{N}\left(\alpha - \mathbb{1}_{\{x_{t,i} \leq \mathrm{VaR}_{\alpha}(t,i)\}}\right)\left(x_{i,t} - VaR_{\alpha}(t,i)\right)}{N}$. If for two methods we have $QL^1 < QL^2$, then method 1 is preferable over method 2.

3.3.2 Comparison of NECO VaR with Other Methods

We start by comparing the performance of our causal VaR method with the more traditional methods. We simulate data for the contagion network of 5 financial instruments, as shown in Figure 3.2. To test the ability of the various methods to deal with non-normality, we take exponentially distributed returns with an added shock to the system every 100 days. The procedures are trained on 250 time points, and tested for a further 100 days for their out-of-sample performance. This process is repeated 20 times.

Figure 3.3 shows both the temporal performance and the overall performance of the various VaR methods targeting $\alpha = 0.05$. It shows that the causal NECO VaR outperforms the other methods, both in its ability to deal with the non-normality and the external shocks to the system. Table 3.1 shows that the causal NECO method has good coverage for the three tests, LR_{UC}, LR_{CC} and LR_{DQ}, achieving 94%, 96% and 88%, respectively. Each simulated sample contains 100 out of sample VaR predictions each. From Figure 3.3 it is immediately clear that the two methods struggling the most are VarCovar and GARCH; these two methods rely heavily on the assumption of normality. The FHS-GARCH improves the GARCH method significantly, but as it is based on merging the HIST and the GARCH, it is swayed by the shocks as much as HIST in underestimating the risk.

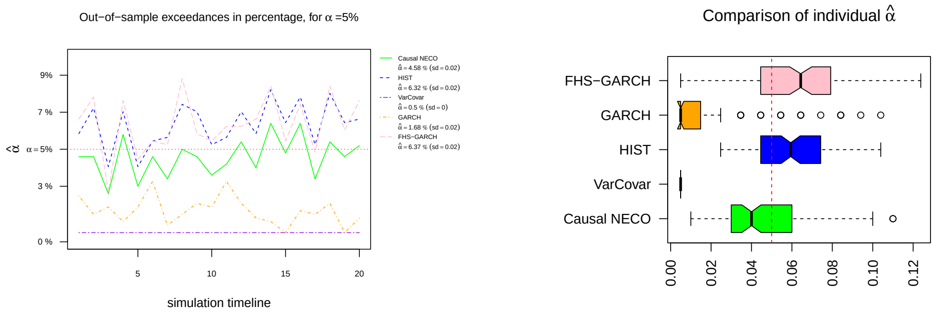

(a) (b)

Figure 3.3: Fraction of VaR exceedances within the out-of-sample window of 100 days for each Value-at-Risk model for all of the considered financial instruments. (a) shows how closely it follows the target of 5% *and (b) shows the overall precision through a boxplot.*

3.3.3 Effect of Training Window

The performance of any VaR method depends on the accuracy of the estimated model. In this simulation study, we vary the training window from $N = 50$ up to $N = 500$ observations for fitting the causal model. Then we apply the causal NECO VaR to an out-of-sample time series of 100 time points. Figure 3.4 shows the results. Although the standard deviation of the achieved VaR level does not depend much on the size of the training window N as seen in Figure 3.4 (b), the VaR tends to have a lower bias with the increase of the estimation window as seen in Figure 3.4 (a), (c), (d). This impact is especially significant when targeting an α below 5%.

3.3.4 Effect of Number of Variables

Next, we study the effect of the size of the financial system on the performance of the causal NECO VaR. We simulate log-return financial networks with $p = 5, 10, 20, 50$ instruments. The training window is set to $N = 250$. We fit the causal NECO on the training window, and results are collected on 100 out-of-sample data. The simulations are repeated 20 times. From Figure 3.5 we see the impact of the size of the network is negligent, especially for the more common levels of α below 10%.

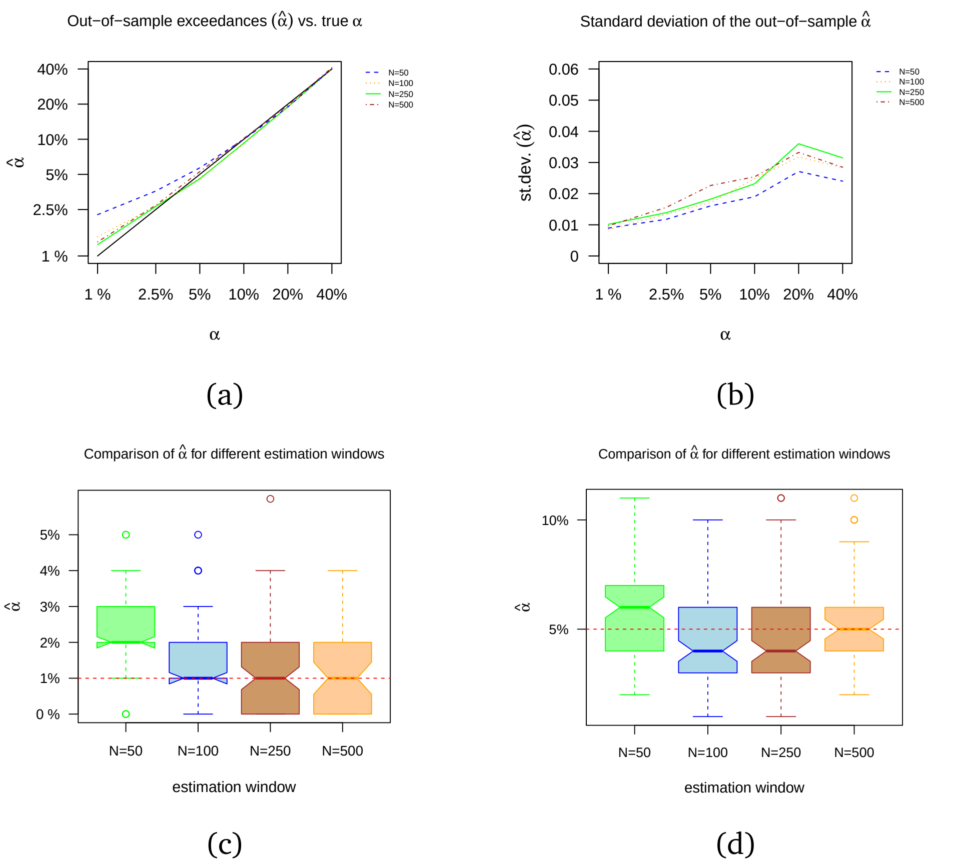

Figure 3.4: **Causal NECO VaR:** *(a) Comparison on the log scale between different target α levels and different lengths of the estimation window in terms of the considered number of observations (N). (b) Standard deviation for the individual $\hat{\alpha}$. (c) Comparison for target $\alpha = 1\%$ for the different N. (d) Comparison for target $\alpha = 5\%$ for the different N.*

	Causal-NECO	VarCovar	HIST	GARCH	FHS
mean($\hat{\alpha}$)	**0.0458**	0.0050	0.0632	0.0168	0.0637
st.dev($\hat{\alpha}$)	0.0203	**0**	0.0196	0.0236	0.0257
LRuc.accept	**0.9400**	0	**0.9400**	0.2100	0.8600
LRcc.accept	**0.9600**	0	0.9300	0.2900	0.8200
DQ.accept	**0.8800**	1.0000	0.7400	**0.8800**	0.6200
AE.mean	**0.9160**	0.0000	1.1760	0.2400	1.1860
AE.sd	0.4052	**0.0000**	0.3962	0.4774	0.5182
AD.mean	**0.0303**	NA	0.0375	0.1249	0.0634
AD.max	**0.1200**	-Inf	0.1300	1.9500	2.2600
CompareQL	1.0000	2.5811	**0.9997**	1.9885	1.0329

Table 3.1: Simulation Backtesting Results, for $\alpha = 5\%$

3.3.5 Effect of Market Contagion

We compare the impact of market contagion, expressed in market NECOF, on the achieved $\hat{\alpha}$ VaR level. The study is performed on a network of $p = 10$ instruments and an estimation window of $N = 250$. Market contagion is a function of a number of causal links in the financial network and of the size of the causal contagion effects. We simulate different network structures with varying effect sizes and express the market contagion in terms of network contagion factor (NECOF). Figure 3.6 (a) and Table 3.2 show that there is no discernible trend in the efficiency of the VaR estimation as the contagion levels on the market increase. For small α below the 5%, the method fares slightly better for lower contagion levels, which would suggest that the choice of the length N of the training window is especially crucial when targeting lower α levels in the case of extreme contagion levels. However, extreme contagion levels would suggest a dense financial contagion network, which is not a typical situation for the financial market.

3.3.6 Effect of Volatility

Volatility, captured by the term $\Sigma = \sigma^2 I$ in Equation 3.2.4, is the system's stochasticity. We simulate systems with $p = 5$ instruments, contagion level 47%. We fit each system using a training window of $N = 250$ time points. We consider five different volatility levels σ, with additional shocks three times the standard deviation. As Figure 3.6 (b) and Table 3.1 show, the volatility change appears to have a negligible effect on the performance of causal NECO VaR. Irrespective of volatility, the NECO VaR tends to be somewhat liberal at low VaR target values α. This is mainly due to the external shocks included in the simulation.

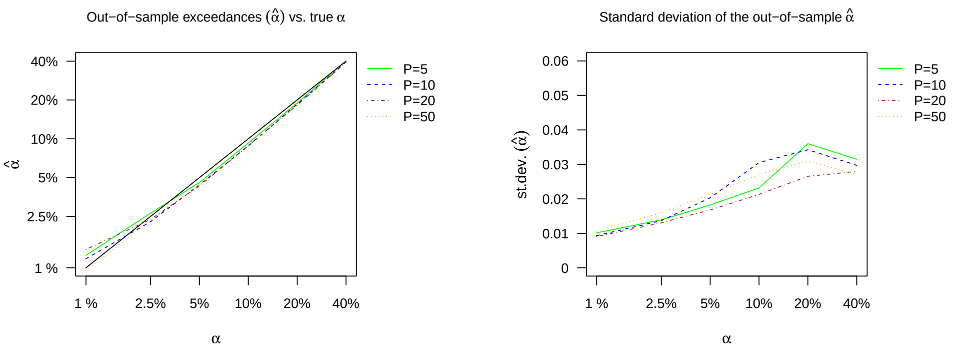

(a) (b)

Figure 3.5: **Causal NECO VaR:** *(a) Log-scale comparison between achieved $\hat{\alpha}$ and target α VaR levels for a different number of instruments (P); (b) Standard deviation of the achieved VaR levels for different values of the target α VaR level for different number of instruments (P).*

3.3.7 Computational Time

The underlying causal inference approach [177] is well adapted to sparse networks, that have relatively few causal links. This seems to be the case for financial networks, as there are usually clear pathways through which contagion flows. Figure 3.7 shows that for such sparse networks, computational time is low, even for a high number of financial instruments. It has to be said, that for very dense networks or for very high number of nodes, the PC-stable algorithm at the core of our method can have convergence issues [117].

3.4 Measuring Risk on Forex

The Forex is a very liquid and important financial market, trading $6.6 trillion per day [221]. Most traded exchange rates are those over the US Dollar (USD). In this study, we analyse the Forex market of 20 exchange rates over the USD, selected to provide a representative overview of the most commonly traded currencies, as well as some less common ones to showcase different challenges in risk estimation. Given that we consider liquid assets we will consider a VaR with a 1-day holding period.

Alternative approaches to considering the US dollar as base currency could consider exchange rates based on the special drawing right (SDR) as in [215]. SDR reflects the price of a basket of five major currencies and is periodically re-

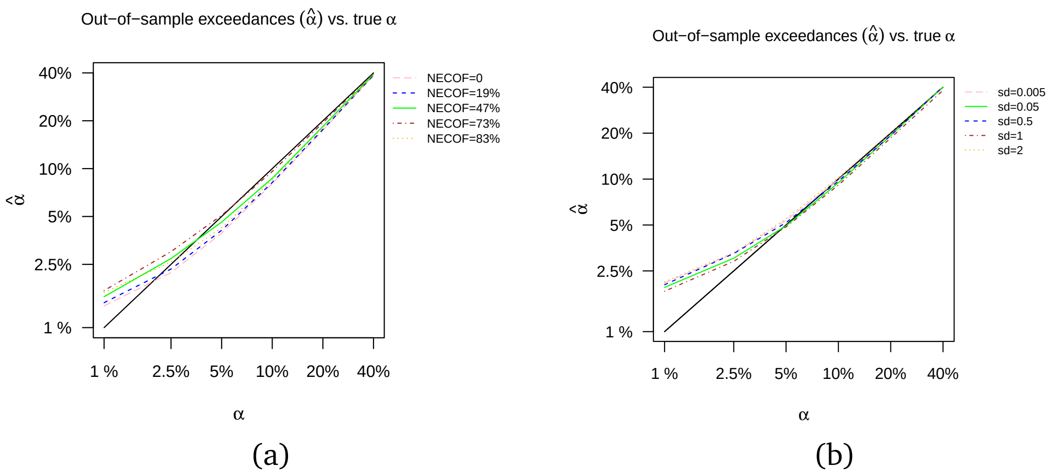

Figure 3.6: **Causal NECO VaR:** *log-scale comparison for the between target α and achieved $\hat{\alpha}$ VaR levels, as a function of (a) different levels of market contagion (NECOF) and (b) different levels of volatility (σ).*

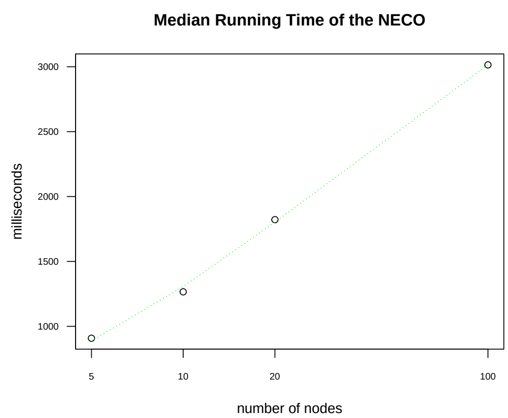

Figure 3.7: Average computation time (across 100 simulations) of the causal NECO network identification and estimation as a function of the number of financial instruments (nodes).

	Contagion (NECOF)					Volatility (sd)				
	0%	19%	47%	73%	83%	0.005	0.05	0.5	1	2
mean($\hat{\alpha}$)	0.039	0.041	0.046	0.051	0.042	0.054	0.049	0.052	0.049	0.050
sd($\hat{\alpha}$)	0.015	0.017	0.022	0.021	0.020	0.027	0.025	0.026	0.026	0.032
LRuc.accept	0.960	0.940	0.920	0.940	0.900	0.910	0.880	0.910	0.860	0.860
LRcc.accept	0.990	0.980	0.960	0.950	0.960	0.870	0.910	0.950	0.890	0.900
DQ.accept	0.880	0.860	0.820	0.810	0.880	0.850	0.890	0.830	0.860	0.840
AE.mean	0.786	0.819	0.922	1.017	0.849	1.074	0.988	1.036	0.972	1.098
AE.sd	0.306	0.329	0.436	0.427	0.399	0.538	0.495	0.531	0.525	0.643
AD.mean	0.024	0.026	0.029	0.035	0.035	0.000	0.000	0.000	0.000	0.000
AD.max	0.150	0.120	0.130	0.160	0.240	0.010	0.000	0.010	0.020	0.000
CompareQL	0.963	0.972	1.000	0.959	1.032	0.983	1.034	1.000	1.028	0.986

Table 3.2: Backtesting results for target $\alpha = 5\%$ at different levels of market contagion (left) and different levels of volatility (right).

balanced. It is published on a daily basis by the International Monetary Fund (IMF). Another option for the base currency is to choose a currency that is of lesser importance, but not completely illiquid. An example of this approach can be found in [122] which uses the Turkish Lira as a base. A comparison of the different currencies that are used as the base currency can be found in [128]. Finally, [63] ignores the issue of the base currency altogether and uses each exchange rate as a separate financial asset.

3.4.1 Foreign Exchange Rates 2000-2021

We consider the log returns on the spot exchange rates over the USD. Table 3.3 shows the summary statistics of the 20 considered exchange rates for the years 2000-2021. Most of the currencies have considerable fat tails in their log returns. The source of the exchange rates is the Federal Reserve of New York.

3.4.2 Fitting Causal Network Contagion on Forex

We consider causal networks with different values of the lag L in Equation 3.2.2. In order to choose the best number of lags, we make use of the Akaike information criterion, given as $\text{AIC}(\ell) = 2k(\ell) - 2l_\ell(D_X)$, where l is the log-likelihood, and $k(\ell)$ is the number of all non-zero autoregressive parameters A and contagion coefficients B. Figure 3.8 shows the value of the AIC for different lags. The minimum is obtained at lag $\ell = 1$, which is used in the estimated NECO model.

Given the presence of fat tails in the log return, we use a copula transformation, as described in section 3.2.2. We fit a marginal distribution F_i for the

	$\widehat{\text{NECOF}}$	Minimum	Median	Mean	Maximum	StDev	Skewness	Kurtosis	Jarque Bera
AUD	0%	-0.0771	-0.0003	-0.0001	0.0822	0.0080	0.6197	12.0299	31085
EUR	0%	-0.0463	0	-0.0001	0.0300	0.0059	-0.0557	2.5411	1375
NZD	63.9%	-0.0593	-0.0003	-0.0001	0.0618	0.0082	0.3770	4.8517	5124
GBP	20.2%	-0.0443	-0.0001	0	0.0817	0.0060	0.7047	10.8604	25491
BRL	5.8%	-0.0967	0	0.0002	0.0867	0.0105	-0.0040	8.1012	13949
CAD	37.6%	-0.0507	-0.0001	-0	0.0381	0.0057	-0.0669	5.4645	6350
DKK	98.5%	-0.0580	-0.0001	-0.0001	0.0494	0.0060	-0.1284	4.7956	4902
HKD	0.3%	-0.0045	0	-0	0.0033	0.0003	-1.2257	25.3347	137697
INR	1.3%	-0.0376	0	0.0001	0.0394	0.0045	0.1980	9.9731	21174
JPY	6.2%	-0.0522	0.0001	-0	0.0334	0.0062	-0.3144	4.4245	4245
KRW	13.3%	-0.1322	-0.0001	-0	0.1014	0.0068	-0.5511	49.7260	525804
MXN	37.8%	-0.0596	-0.0001	0.0001	0.0811	0.0072	0.7485	11.0121	26250
NOK	40.6%	-0.0644	-0.0002	-0	0.0612	0.0078	0.2380	4.8194	4985
SEK	63.3%	-0.0530	-0	-0	0.0547	0.0074	-0.0482	3.9502	3319
ZAR	42.4%	-0.0916	-0.0002	0.0001	0.0843	0.0109	0.2626	4.2637	3922
SGD	59.5%	-0.0238	-0.0001	-0.0001	0.0269	0.0033	0.0313	4.9083	5121
LKR	0.4%	-0.0339	0	0.0002	0.0641	0.0029	2.5329	76.6587	1254466
CHF	65.6%	-0.1302	0	-0.0001	0.0889	0.0067	-1.1545	36.8482	289720
TWD	39.5%	-0.0342	0	-0	0.0248	0.0030	-0.3858	9.8420	20714
THB	37.9%	-0.0353	0	-0.0001	0.0447	0.0037	0.1609	12.2860	32104

Table 3.3: Overview of the summary statistics for the dataset of log returns on individual 20 exchange rates over the USD, for the period January 2000 to April 2021. The higher the Jarque Bera test statistic the less likely the data are normally distributed — all of the statistics have a p-value of 0. Each sample is 5101 observations long with no values missing.

log returns for each of the 20 currencies. We fit the causal network to the transformed log returns as explained in Section 3.2.3. Finally, we obtain the 1-day ahead causal NECO VaR values via Equation 3.2.7.

3.4.3 Results and Backtesting

We will consider the Value at Risk for each individual currency rate at the α level of 5%. We consider a training window of $N = 250$ trading days. We will compare the performance of the Causal NECO approach to the more established methods HIST, VarCovar, GARCH and FHS-GARCH with the backtesting measures from Section 3.3.1 applied to 100 out-of-sample VaR predictions, for each of the 20 currencies during 20 non-overlapping periods between 2000 and 2021.

Table 3.4 shows that the causal NECO VaR beats all other methods in most of the categories. Given that at least during calm periods, the overall Forex volatility is quite low, the losses are not extreme for any of the methods. In fact, the maximum loss AD.max is the same for all, and most other comparisons are relatively close. The FHS seems to be the best competition for the causal NECO VaR on this level.

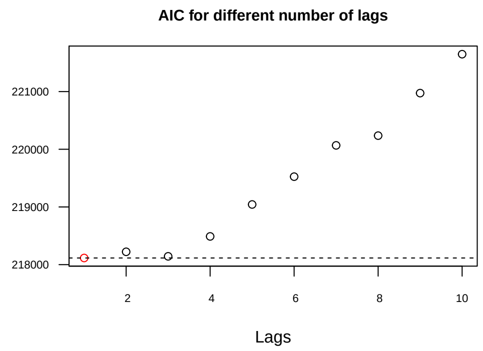

Figure 3.8: Overall AIC for each number of lags on the whole of the Forex dataset considered.

	Causal-NECO	VarCovar	HIST	GARCH	FHS
mean($\hat{\alpha}$)	**0.0511**	0.0597	0.0594	0.0549	0.0574
st.dev($\hat{\alpha}$)	**0.0202**	0.0603	0.0392	0.0338	0.0318
LRuc.accept	**0.9500**	0.8100	0.6700	0.7600	0.7800
LRcc.accept	**0.9600**	0.9000	0.8000	0.8700	0.8800
DQ.accept	**0.8500**	0.8000	0.7800	0.8200	**0.8500**
AE.mean	1.0220	1.1060	1.0995	**1.0085**	1.0595
AE.sd	**0.4046**	1.2171	0.7928	0.6818	0.6416
AD.mean	0.0032	0.0031	0.0032	**0.0030**	**0.0030**
AD.max	0.1200	0.1200	0.1200	0.1200	0.1200
CompareQL	1.0000	0.9804	1.0247	0.9934	**0.9801**

Table 3.4: Backtesting of the VaR out-of-sample predictions in the Forex market (2000-2021) for the different methods at target $\alpha = 5\%$. The best results are presented in bold.

Figure 3.9 (a) shows the distribution of actual exceedances for all currencies and periods. The Causal-NECO is the most centred around the target value of 5%. Figure 3.9 (b) shows how these exceedances change through time. The Causal-NECO method seems to be able to adapt properly to changing conditions of the underlying network. Unlike the other methods, it is barely affected by the various financial crises in this period and achieves the nominal 5% level throughout the evaluation period.

In addition, we consider each of the 20 currencies individually. Figure 3.10 shows the individual boxplots of the exceedance fractions for each of the given methods. The causal NECO performance is consistently less variable and more centred around the 5% level.

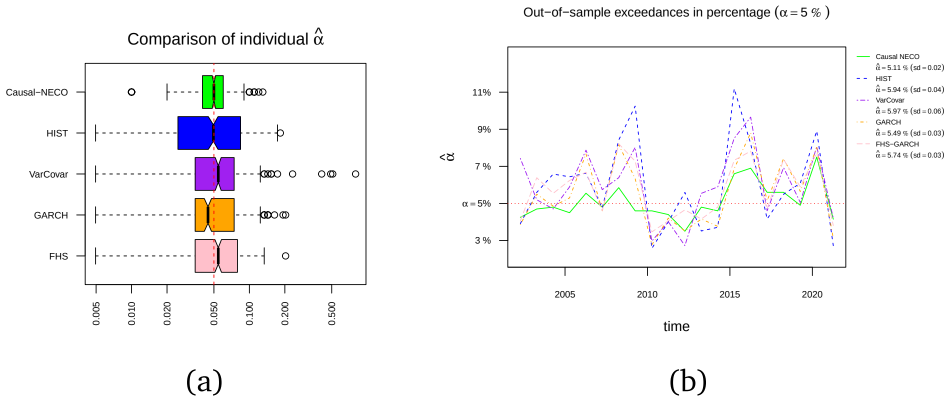

Figure 3.9: (a) Overall fraction of exceedances $\hat{\alpha}$ within the out-of-sample window of 100 days for each Value-at-Risk model across the Forex market (2000-2021). (b) Fraction of the exceedances over time.

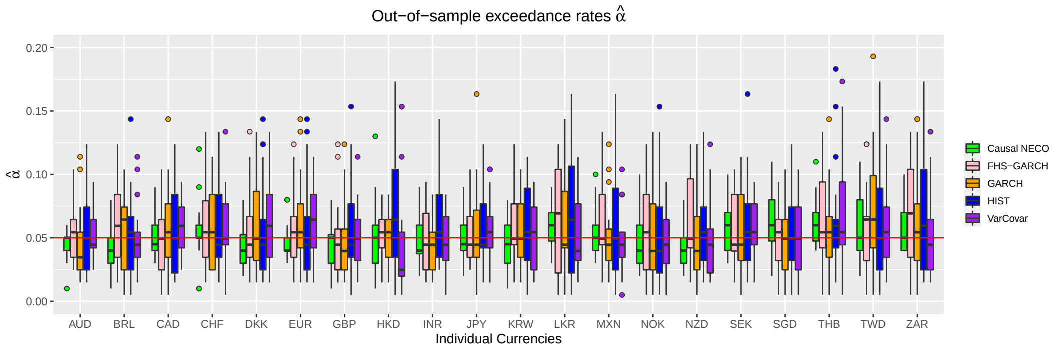

Figure 3.10: Number of actual exceedances within the out-of-sample window of 100 days for each Value-at-Risk model, expressed in terms of the $\hat{\alpha}$ computed for each period and shown for each currency separately.

3.5 Conclusion

This paper introduces an innovative way of eliminating spurious correlations from risk management decisions. It lays the groundwork for using causal inference in finance showing how practical advantages can be obtained by inferring the underlying causal system from available financial data.

The value at risk measure is calculated every day by financial managers at banks and investment firms alike. The causal NECO VaR could be an interesting addition, that is easy to compute and offers a robust and competitive addition to the standard approaches. The volatility is not modelled directly but is addressed through the contagion effect and in part through the copula transformation, offering an alternative view on modelling volatility.

As a way to deal with the non-normality of financial data, we use a Gaussian copula transformation. Gaussian copulas are not always appropriate for the analysis of financial data and have received quite a bad reputation after their apparent role in the 2008 financial crisis. [47] and [138] analyse what happened and show the consequences of the indiscriminate use of a Gaussian copula to estimate risk, without checking the appropriateness of the assumptions and as a stand-in to compensate for a lack of data and information on correlations. When used correctly the Gaussian copula can prove rather useful. [139] find for example that “most pairs of currencies and pairs of major stocks are compatible with the Gaussian copula”. In our case, it is also important to note that we do not treat the Gaussian copula as a way to model risk but to improve the estimation of our causal estimation of contagion.

Possible extensions could attempt to account for a wider class of distributions using a Student-t copula instead of a Gaussian copula and time-varying errors with an added GARCH component [113]. Both these extensions, however, are beyond the scope of this paper, and could in theory be incorporated without affecting the underlying causal structure of our model.

It is possible to add any measurable confounder to our model, in order to determine how other variables, like interest rates, inflation or any other stock market returns, might affect Forex contagion. When unmeasured confounders and latent variables are present, the Fast Causal Inference algorithm (FCI) [199, 200] is a viable option.

Chapter 4

Conformal Updating and Causal Network Contagion

4.1 Introduction

The foreign exchange market (Forex) is a critical component of international trade and finance, with a daily turnover of more than 7.5*trillionin*2022[49].*Thisdynamicmarketp*
f requencytraders, onlineretailinvestors, andsmallercommercialbanks, aregrowingtheirpres
33].*ThisunderscorestheneedfornewresearchonForexdynamics, providinganopportunitytode*

Value-at-risk (VaR) is a fundamental and pivotal risk measure in the contemporary financial industry, systematically quantifying the potential downside loss in the value of an asset or portfolio over a specified temporal horizon, all within a predetermined level of confidence [115]. Particularly within the intricate realm of the Forex, VaR provides invaluable information on potential risk exposure and equips decision makers with the analytical tools necessary to devise robust and comprehensive risk management strategies, especially when faced with the uncertainties of currency risk or when actively navigating the tumultuous seas of currency trading. To this day, it remains the most widely adopted risk measure and holds a central position within the international banking regulatory framework, as emphasised by its role in the Basel Accords conventions [124]. However, as the global financial markets evolve, conventional VaR models, which once held the limelight within the financial industry, are increasingly being viewed as inadequate in deciphering the multilayered complexities of today's Forex environment. These traditional approaches often anchor themselves on rather simplistic assumptions of normality and stationarity, often turning a blind eye to the intricate dynamics, unpredictable volatilities, and unanticipated structural shifts that are now hallmarks of financial data [147, 123]. Resorting to such oversimplified frameworks can culminate in misleading risk assessments, thereby posing amplified challenges not only to market participants, but also to the broader interconnected global economy.

Research into VaR methods has decreased after the Global Financial Crisis in 2008. New approaches introduced the Expected Shortfall (ES) measure, also known as Conditional Value at Risk (CVaR). ES complements VaR by providing insights into the expected losses beyond the VaR level [224]. It represents the average loss that can be expected if the portfolio's returns fall beyond the VaR threshold. ES offers a more comprehensive view of tail risk, which is crucial for risk management. However, the method is intrinsically connected to the calculation of VaR [208]. For this reason, we will focus on the VaR in this paper, leaving further extensions for future work.

We introduce an innovative approach to measuring VaR by leveraging a comprehensive framework that combines elements from causal analysis based on causal inference [161], network contagion [177] and conformal updating techniques [13]. This blended approach, referred to as Conformal Causal Network Contagion VaR (CCNC VaR), seeks to uncover the intricate web of causal relationships and interconnectedness that underlie the behaviour of currency exchange rates [178].

The first and possibly most innovative contribution in this paper is the development of a conformal updating for the network contagion approach to VaR estimations. Traditional VaR models often struggle with the non-normality and non-stationarity of financial data. Conformal updating, mainly used in machine learning, allows for frequent rechecking of whether the target quantile is being reached as model estimation continues. This information is used to improve estimations without any model assumptions. These models adapt to changing distributions and variability, offering a more accurate representation of risk [44]. Secondly, we apply a causal approach, based on causal inference [161]. These are especially good for the analysis during a crisis. Thirdly, the presence of contagion in the financial markets motivates the use of network modelling.

In this chapter, we will provide an in-depth analysis of each of these three components, highlighting their significance in the context of Forex risk management. By incorporating the CCNC VaR framework into their standard risk management tools, financial institutions can improve their risk assessment capabilities and make informed decisions in an ever-evolving Forex. This framework can be easily extended to other markets and applications. Section 4.2 provides an introduction to the Forex and the empirical challenges it presents for risk management. Section 4.3 introduces the CCNC approach to VaR estimation. Section 4.4 presents the performance of CCNC on the Forex dataset, with a specific focus on the impact of the COVID-19 recession. Finally, Section 4.5 provides concluding remarks.

4.2 Empirical Problem

4.2.1 Forex and Data Source

We examined the interactions between the 20 most liquid exchange rates over the US Dollar (USD) published daily by the Federal Reserve of New York for the years 2000-2021 for our dataset. We consider the following currencies: the Australian Dollar (AUD), Brazilian Real (BRL), Canadian Dollar (CAD), Danish Krone (DKK), Euro (EUR), British Pound (GBP), Hong Kong Dollar (HKD), Indian Rupee (INR), Japanese Yen (JPY), South Korean Won (KRW), Sri Lankan Rupee (LKR), Mexican Peso (MXN), New Zealand Dollar (NZD), Norwegian Krone (NOK), Singapore Dollar (SGD), Swiss Franc (CHF), Swedish Krona (SEK), South African Rand (ZAR), New Taiwan Dollar (TWD), and Thai Baht (THB).

We chose the US Dollar (USD) as the base currency for the exchange rates in the analysis, as it is the most liquid currency and holds the status of the dominant international currency. According to [49] “the USD was involved in nearly 90% of global FX transactions, making it the single most traded currency in the FX market”. The global significance and role of the USD is often challenged and discussed during financial crises [38, 204, 56, 22]. Despite the rise of potential currency competitors, for the period considered in this chapter, it remains the international currency of reference [49].

4.2.2 Forex Summary Statistics and Lognormality

For the selected 20 currencies, the evolution of their log-returns from January 2000 until April 2021 is depicted in Figure 4.1. The summary statistics of the considered time series are presented in Table 4.1. In addition to the typical summary statistics, we also incorporate a metric for network contagion - the Network Contagion Factor (NECOF) - for each currency. Understanding network contagion is a crucial element when analysing the behaviour of exchange rates, and relying solely on correlation and volatility levels may not provide reliable indications [81, 69, 177]. The higher the NECOF the higher the contagion for that particular currency.

As shown in Table 4.1, the data exhibits non-normality and fat tails with significant differences in contagion levels among individual currencies. Both HKD

Figure 4.1: Log returns for each currency. The financial crises highlighted are the Global Financial Crisis (August 2007 - April 2010) with the following Sovereign Debt Crisis (May 2010 - December 2012), the 2015/2016 Crises (March 2015 - June 2016) and the COVID-19 recession (starting March 2020).

and LKR show unusual periods of very low variability. This can be attributed to the fact that respective Central Banks pegged these currencies to the USD, keeping their exchange rate within a pre-defined band. However, some level of contagion can still spread to and from these currencies within these bands and is seen from the NECO not being zero for these currencies. AUD, EUR, INR and JPY have a NECO of 0 %, which shows that contagion from other currencies tends not to be of importance to estimate their price. The EUR is the only platykurtic currency, with a kurtosis below three. Currencies with log returns closest to that of a normal distribution and only slightly leptokurtic include DKK, JPY, NOK, NZD, SEK, SGD and ZAR. The least normally distributed currencies with very high kurtosis and Jarque-Bera values are the CHF, HKD, KRW, LKR and TWD.

The log transformation seems to make the log-returns more stationary when considering the whole period, as confirmed by performing the Augmented Dickey-Fuller test (ADF). Our approach can also handle non-stationarity. The overall volatility on the Forex has a tendency to be relatively low, which makes the spike in volatility and correlations during a crisis all the more difficult to handle by standard statistical approaches to risk management and forecasting. A deeper analysis of the dynamics of the volatility on the Forex can be found in [150].

	$\widehat{\text{NECOF}}$	Minimum	Median	Mean	Maximum	StDev	Skewness	Kurtosis	Jarque Bera
AUD	0%	-0.0771	-0.0003	0	0.0822	0.008	0.6261	11.8178	31335.1735
BRL	18.3%	-0.0967	0	0.0002	0.0867	0.0104	-0.0052	8.3711	15547.9966
CAD	43.8%	-0.0507	0	0	0.0381	0.0056	-0.0714	5.6143	6998.1322
CHF	64.4%	-0.1302	0	-0.0001	0.0889	0.0067	-1.1203	35.1797	275708.2298
DKK	96%	-0.058	0	0	0.0494	0.0061	-0.1555	4.5734	4662.0828
EUR	0%	-0.0463	0	0	0.03	0.006	-0.0775	2.4789	1368.7587
GBP	44.3%	-0.0443	0	0	0.0817	0.006	0.6707	10.6119	25385.1549
HKD	1.6%	-0.0045	0	0	0.0033	0.0003	-1.2511	26.3958	155978.276
INR	0%	-0.0376	0	0.0001	0.0394	0.0044	0.1993	10.2796	23480.9216
JPY	0%	-0.0522	0.0001	0	0.0334	0.0062	-0.3187	4.3307	4251.4023
KRW	21.6%	-0.1322	-0.0001	0	0.1014	0.0067	-0.5492	50.7331	571338.979
LKR	0.3%	-0.0339	0	0.0002	0.0641	0.0029	2.7858	75.1652	1260440.0855
MXN	40.2%	-0.0596	-0.0001	0.0001	0.0811	0.0071	0.7336	11.1455	28039.508
NOK	63.6%	-0.0644	-0.0001	0	0.0612	0.0077	0.2195	4.7907	5135.0209
NZD	68.1%	-0.0593	-0.0002	-0.0001	0.0618	0.0082	0.3886	4.7691	5180.4432
SEK	75.6%	-0.053	0	0	0.0547	0.0074	-0.0544	3.8772	3337.9567
SGD	64.6%	-0.0238	-0.0001	0	0.0269	0.0033	0.0239	4.9815	5506.5174
THB	24.3%	-0.0353	0	0	0.0447	0.0037	0.1773	11.5344	29546.8524
TWD	35.4%	-0.0342	0	0	0.032	0.0031	-0.3537	16.1074	57675.9338
ZAR	41.8%	-0.0916	-0.0001	0.0002	0.0843	0.0108	0.2563	4.3936	4341.3691

Table 4.1: Summary statistics for the log-returns of the considered 20 currencies from January 2000 to April 2021. High Jarque Bera Test results mean higher deviation from a normally distributed sample (all of the statistics have a p-value of 0).

4.2.3 Crises in the Currency Markets

Crises are periods of structural instability and tend to be times when traditional methods struggle. As our method is designed to be partially useful during these periods, we considered various crisis periods in the 2000-2021 range. The main four financial crises periods considered are the Global Financial Crisis (August 2007 - April 2010), the following Sovereign Debt Crisis (May 2010 - December 2012), the Crises of 2015/2016 (March 2015 - June 2016) and the COVID-19 recession (starting March 2020). The timing of the different crises was chosen according to the literature [149, 21] and the financial news, coordinated with the change in volatility on the market. There is a vast literature analysing the first two crises, [6] and [149] give a good overview.

The two-year period 2015/2016 presents several critical events for the Forex. In 2015 to the beginning of 2016, the global economy experienced a significant slowdown and, as the US was experiencing an earnings recession, the markets began to worry about a possible new recession coming [77]. On March 18, 2015, the USD experienced a flash crash when the USD fell more than 3% in under four minutes [172]. The Chinese Market Crash, often marketed as the period of March 2015 to June 201, had an impact not only on Asian economies, but also on other markets on a global scale [197, 227]. Finally, the Brexit vote that decided the UK to leave the European Union took place in June 2016. This event has had a primary impact on the volatility of the GBP. Other studies found changes in correlation and correlation-based contagion on the Forex [41]. We will refer to this partially overlapping sequence of events as the 2015/2016 Crises. The COVID-19 recession was the most recent crisis that had an impact on global financial markets. It differed from previous crises in that it was a real exogenous shock that did not originate within the financial system. The crisis had a global impact and resulted in long-term structural changes in multiple sectors, as well as unprecedented government interventions [26]. The Forex was also affected, with central banks implementing significant interventions. The crisis also had a high impact on volatility and contagion [21]. Network contagion grew rapidly at the beginning of the crisis, and the redistribution of clusters among currencies shifted to new patterns. While the number of causal links decreased, the contagion rates were much higher than before. Similar structural changes were observed in [32, 206]. We can see these changes in the dataset. Average volatility almost doubled and the mean correlation increased from 28.4% to 42.7% in the year before and after March 1, 2020. Table 4.2 shows that the 10-day rolling window volatility increases significantly for almost all currencies considered. Notable exemptions are the HKD, where volatility decreased, given central bank interventions on the peg to USD. Given all the changes in the Forex during the COVID-19 period, it is hypothesised that there may be significant disparities in VaR predictions between models based on observed correlation and those grounded in causal inference.

	$\widehat{sd}_{\text{before}}$	$\widehat{sd}_{\text{after}}$	p-value		$\widehat{sd}_{\text{before}}$	$\widehat{sd}_{\text{after}}$	p-value
AUD	0.0069	0.0072	**0.009**	KRW	0.0053	0.0047	0.339
BRL	0.0086	0.013	$<$ **0.001**	LKR	0.0018	0.0032	$<$ **0.001**
CAD	0.005	0.0049	0.734	MXN	0.0058	0.0102	$<$ **0.001**
CHF	0.006	0.0043	$<$ **0.001**	NOK	0.0069	0.0094	$<$ **0.001**
DKK	0.0056	0.0043	$<$ **0.001**	NZD	0.0074	0.0074	0.52
EUR	0.0056	0.0043	$<$ **0.001**	SEK	0.0068	0.0063	**0.003**
GBP	0.0054	0.0061	$<$ **0.001**	SGD	0.003	0.0028	**0.026**
HKD	2e-04	2e-04	$<$ **0.001**	THB	0.0031	0.003	**0.007**
INR	0.0036	0.0036	0.101	TWD	0.0026	0.0018	$<$ **0.001**
JPY	0.0057	0.004	$<$ **0.001**	ZAR	0.0098	0.0101	**0.005**

Table 4.2: Mean rolling 10-days volatility before and after COVID-19, and the p-value from the Wilcoxon Rank-Sum test, to see if there is a difference in the volatility before and after the start of COVID-19. In bold if p-value $>$ 0.05.

4.3 Conformal Causal Network Contagion Approach to Value at Risk

4.3.1 Modelling the Log-Returns

Causal Graphical Model Nothing happens in a vacuum in the real world and network theory is one way to operationalise complex interactions. For situations where randomized experiments cannot be used to study such interactions, causal inference is one successful approach [161]. The causal network contagion approach unifies these two concepts and strives to explain the interactions among the development of log-returns of different financial assets, showing how network contagion reacts in financial crises [176]. The assumed underlying structure for logarithmic returns $Z_{i,t}$ for asset i at time t is the result of the autoregressive part $\{Z_{i,t-1}, \cdots, Z_{i,t-l}\}$, the contemporaneous contagion effects from so-called causal parents within the estimated causal networks, and a noise component of independent and identically normally distributed $\varepsilon_{i,t}$:

$$Z_{i,t} \leftarrow \alpha_{0,i} + \sum_{\ell=1}^{L} \alpha_i^{\ell} Z_{i,t-l} + \sum_{j \in pa(i)} \beta_{ji} Z_{j,t} + \varepsilon_{i,t} \tag{4.3.1}$$

where $\alpha_{0,i}$ is the intercept, $pa(i)$ are the causal parents of instrument i and L is the maximal lag of the autoregressive components. Crucially, before being able to estimate the coefficients of this structural vector autoregressive model we need to first find the causal parents for each instrument, or, in other words, the causal network structure.

The PC-stable algorithm [37] is an extension of the PC algorithm [198]. The PC algorithm is a widely used causal inference method in observational data analysis. The aim is to uncover potential causal relationships among variables by identifying conditional independence relationships in the data. By iteratively testing for statistical independence and gradually building a causal graph structure, the PC algorithm infers plausible causal connections while accounting for confounding factors. By iteratively applying the PC algorithm to multiple bootstrapped samples of the data and aggregating the results, PC-Stable produces a more stable and accurate causal graph representation of the underlying relationships between variables and gives us the parents we need for Equation 4.3.1.

Gaussian Copula Transformation Many financial assets do not adhere to normality or lognormality assumptions, as seen in Section 4.2. To address this is-

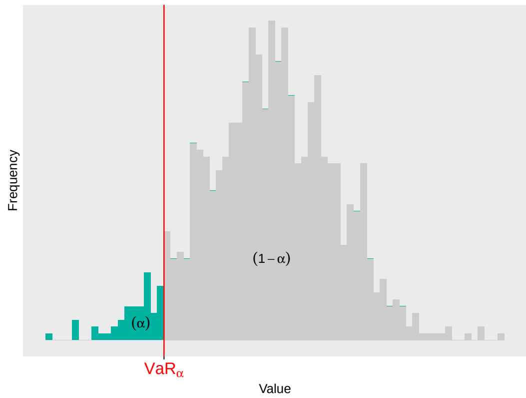

Figure 4.2: Value at Risk (VaR$_\alpha$) as the α quantile for the returns of a financial instrument.

sue, we consider a multivariate Gaussian copula as part of the underlying causal graphical model:

$$X_{t,i} = F_i^{-1}\left(\Phi\left(Z_{t,i}\right)\right) \tag{4.3.2}$$

where the F_i represents a non-parametric empirical cumulative distribution function. The variable Z represents the ideal data, which is assumed to follow a normal distribution, whereas X are the real-world observed log returns. This specific approach has been developed in [46] and [1]. In this context, the copula transformation is not utilised to examine the observed dependence structure among various financial assets [183]. Instead, it is part of the causal inference framework.

4.3.2 Conformal Causal Network Contagion Value at Risk (CCNC VaR)

Value at Risk Value at Risk (VaR) is a risk measure used in financial risk assessment and management. It serves as a tool to quantify and compare the potential risk associated with holding specific financial instruments or portfolios. VaR offers insights into the potential loss, at a given confidence level, over a defined holding time period under normal market conditions [115, 114, 133, 48].

The VaR at a level α is defined as the threshold below which the value or return of the future value of an asset or portfolio may potentially fall with a probability of α (Figure 4.2). For a financial instrument X, the VaR can be mathematically expressed as:

$$\mathrm{VaR}_{\alpha}(X) = \inf\{x \leq \mathbb{R} : F_X(x) > \alpha\} \tag{4.3.3}$$

where $F_X(x)$ represents the cumulative distribution function (CDF) of the log-return X.

Conformal updating Conformal updating is a technique used in machine learning and data analysis to improve the reliability of predictive models by incorporating adaptive tuning into their predictions [191]. This approach adapts the prediction regions as new data points are observed, without changing the underlying prediction model. Currently, conformal predictions are the most widely used prediction framework in deep learning models [65, 222]. This approach has found applications in various fields such as medicine, drug discovery, computer vision, and anomaly detection, where accurate and well-calibrated uncertainty estimates play a crucial role in decision-making [13].

It works by adapting to the uncertainty that arises from the assumptions of the model and potential yet undiscovered changes in the underlying causal structure. This is done by allowing slow adjustments to the targeted levels of α_t. This approach maintains the underlying model unchanged while countering potential misspecifications that could deviate the ex-post results from the required level α.

Given a step size parameter $\gamma > 0$, we consider the simple online update of the operational VaR level α_t, in order to achieve the VaR target level α:

$$\alpha_{t+1} := \alpha_t \times e^{\{\gamma(\alpha - \hat{\alpha}_{\Delta t})\}} \tag{4.3.4}$$

where the α is the desired target quantile, $\hat{\alpha}_{\Delta t}$ is the realised exceedance rate over the past Δt period computed as follows:

$$\hat{\alpha}_t = \frac{\sum_{j=1}^{N} \sum_{\delta=0}^{\Delta t} \mathbb{1}_{\mathrm{VaR}_{(t-\delta),j} > x_{(t-\delta),j}}}{N \cdot \Delta t} \tag{4.3.5}$$

This conformal updating approach is based on theoretical work by [180], [132], [211], [76] and [189]. They present the conformal prediction in the context of quantile regressions and extend the underlying theoretical assumptions. We have confirmed that the choice of the γ has minimal impact. Whereas larger step sizes get fast adaptation, closer to the target α, it can result in less smooth predictions. While this may have practical implications when deciding on hedging positions based on the predictions, it has less impact on the precision evaluation, which is our focus in this paper.

Conformal Causal Network Contagion VaR The copula transformation from Equation 4.3.2 says Z_{it} follows a normal distribution, which when combined with conformal updating, enables us to compute our Conformal Causal Network Contagion VaR (CCNC VaR) measure on the transformed scale as follows:

$$\mathrm{VaR}_{\alpha}(Z_{i,t}) = \mu_{Z_{i,t}} - z_{\alpha_t} \cdot \sigma_{Z_{i,t}} \tag{4.3.6}$$

where the α_t is updated each Δt steps as in Equation 4.3.4. Using the copula transformation we can then define the CCNC VaR on the original scale as:

$$\mathrm{VaR}_{\alpha}(X_{i,t}) = F_{X_i}^{-1}\left(\Phi_{0,1}\left(\mathrm{VaR}_{\alpha}\left(Z_{i,t}\right)\right)\right) \tag{4.3.7}$$

For a more comprehensive understanding of the CCNC computation, please refer to the pseudo-code provided in Appendix B.1.

4.4 CCNC VaR Performance on Forex

We will analyse the performance of the Conformal Causal Network Contagion approach in forecasting VaR on the dataset described in Section 4.2. We compare the results to four representative alternative methods. The direct competitor and the underlying method for the Conformal NECO is the parametric variance-covariance (VarCovar) method. The difference in performance between the Conformal NECO method and the VarCovar can be directly attributed to the conformal updating and the underlying causal approach. Another parametric approach is the GARCH method [133, 134]. The GARCH method simulates possible future returns with Monte Carlo assuming an underlying GARCH process and taking the appropriate α quantile. GARCH is an improvement over the VarCovar method because it allows for time-varying volatility and correlation, while still assuming normally distributed standardized residuals. However, GARCH struggles to accurately model asymmetric volatility clusters, as these are more prevalent during financial crises. A detailed analysis can be found, for example, in [157]. A non-parametric alternative is the historical simulation (HIST) method [29]. Historical simulation computes the quantile for the VaR based on bootstrapped, sampled with replacement, and past returns. Using realised returns eliminates the need to create models that perfectly describe the reality, but makes the model highly sensitive to what is included in the considered estimation window and hence its length [99, 106]. This method still depends on the stationarity in distribution, which is often violated during a financial crisis, when markets experience volatility clusterings [90]. The fourth and final method considered is the filtered historical simulation (FHS) that unites both the GARCH and HIST methods [17, 18]. FHS improves both the GARCH and HIST methods, performing better than these two separately. [170] shows that extreme observations or time-varying correlations degrade the performance of FHS.

This study compares the performance of VaR prediction techniques using eight backtesting measures presented in Table 4.3. Backtesting is a systematic and quantitative approach to evaluating and comparing VaR prediction techniques based on out-of-sample realizations. These measures evaluate different aspects of VaR predictions, including the extent to which they capture actual losses (coverage), whether violations of VaR predictions are independent, and the size of the prediction errors.

4.4.1 Global Backtesting Results

Table 4.4 displays the backtesting results of predicting 100 out-of-sample 1-day ahead VaR predictions over 50 rounds, using an estimation window of 250 throughout the dataset, for all of the 20 currencies considered. Results for $\alpha = 1\%$ and $\alpha = 10\%$ are shown in Appendix B.2. For the CCNC approach, we evaluate the underlying causal network structure only once and apply it separately to all individual currencies. The performance of all methods is relatively good on average over the whole period. The CCNC shows very good and most importantly consistent results, given the low standard deviations of the average exceedance rate and actual over-expected ratio. When compared to the theoretical counterpart VarCovar method, we see a huge improvement thanks to the causal network contagion and conformal updating. CCNC demonstrates remarkable stability with low variability in its out-of-sample α estimates. The α values consistently hover closely around the α target, reflecting a high degree of precision and reliability in its predictions. This can be seen in the boxplots in Figure 4.3. Figure 4.4 presents the actual quantile realisation out of the sample, with the target set at 5% and an expected exceedance rate close to this desired level. The performance is depicted over the years and four critical periods, as described in Section 4.4.2.

It is worth noting that the HIST and VarCovar approaches display high volatility and significant deviation, particularly during periods of crises. These methods exhibit a high percentage of failure in predicting losses accurately when market conditions change. On the other hand, they tend to overcompensate and fail to identify the return to a new normal by returning very restrictive predictions after

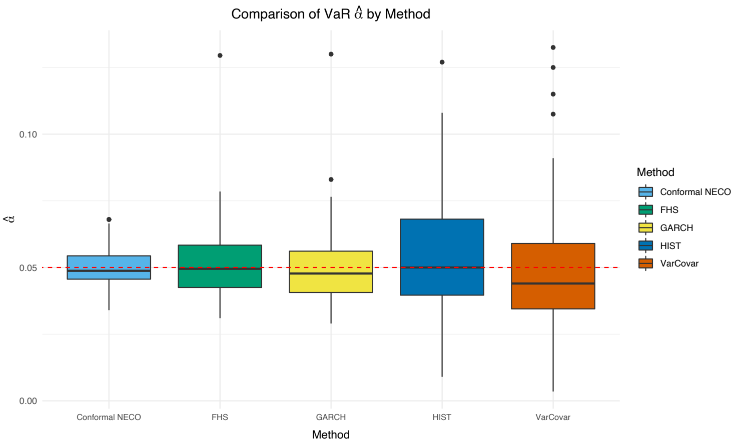

Figure 4.3: Boxplots illustrating the distribution of out-of-sample exceedance rates ($\hat{\alpha}$) for 20 currencies, allowing us to assess how closely each model aligns with the targeted $\alpha = 5\%$ throughout the entire period.

a sharp overcorrection.

Compared to these methods, FHS and GARCH tend to remain closer to the 5% target, especially FHS. However, a complete failure in prediction was observed during the first period of the Global Financial Crisis and the COVID-19 recession.

The tight out-of-sample predictions of the CCNC throughout the entire 21-year period are also confirmed by Figure 4.3. This figure shows the boxplots of all estimated exceedance rates, with CCNC being centred at the target and presenting a very narrow range. In comparison, the method VarCovar is wide and slightly off-centre. This tight band around the target for CCNC is precisely why conformal updating is used.

AE ratios above or below 1, mean that the VaR model underestimates or overestimates respectively, the number of exceedances. Both GARCH and CCNC achieve a target performance. In terms of the DQ and QL, FHS and GARCH slightly outperform CCNC, but the differences are small.

4.4.2 COVID-19 Recession and Handling Exogenous Shocks

For the analysis of the COVID-19 recession period, our analysis focuses on the period from December 2019 to March 2021, with the crisis period identified as starting in March 2020. The results of our backtesting for this sub-sample are summarized in Table 4.5. Only the conformal CCNC approach remained efficient and achieved the α target during this challenging period for the Forex. The CCNC method outperformed all other backtesting measures considered. The

Figure 4.4: The average out-of-sample exceedance rates $\hat{\alpha}$ over the 20 currencies. The CCNC tends to track the target $\alpha = 5\%$ more accurately during periods of financial crises than other methods.

Dynamic Quantile Test (DQ) results are particularly interesting, as only the CCNC approach maintained a high acceptance (non-rejection) rate. It is worth noting that extreme outliers or rare events can significantly impact the results of this test. Moreover, as the only method, CCNC keeps the AE close to 1 consistently (low standard deviation).

Figure 4.5 shows that the CCNC method has an out-of-sample exceedance rate that is close to the target α, while the other methods take significantly longer after the start of the COVID-19 crisis to get back to the target level, and often exceed 10% exceedance rates. HIST and VarCovar have the tendency to overcompensate when high volatility is inside the estimation window, making the length of the estimation window an important parameter for these methods during crisis periods, as confirmed in other studies as well.

Both the FHS and GARCH models demonstrated relatively strong performance when assessing their overall performance over a 20-year data period, as illustrated in Section 4.4.1. However, their performance notably deteriorates when analysing the COVID-19 period, and they even underperform some of the less complex methods like HIST and VarCovar in certain backtesting measures. In particular, the exceedance rate $\hat{\alpha}$, the AE, the DQ, and QL measures show a significant decline in accuracy. It appears that both FHS and GARCH models are highly sensitive to their estimations when log-returns encounter sudden shocks, resulting in incorrect predictions for a considerable duration. These methods are particularly affected by the challenge of maintaining a balance between accom-

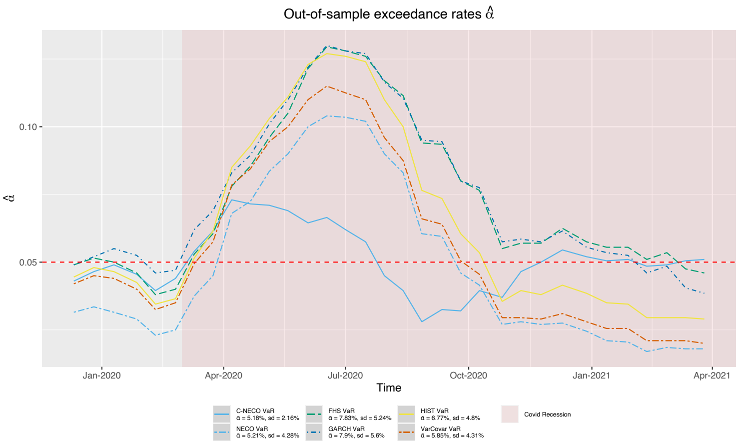

Figure 4.5: During the Covid Recession, the average out-of-sample exceedance rates $\hat{\alpha}$ of the CCNC more accurately tracks the target $\alpha = 5\%$ than other methods.

modating changes in the system and refraining from overweighting the temporary increases in volatility.

In order to evaluate the effect of conformal prediction, we include in Figure 4.5 and 4.6, and Table 4.5 an explicit comparison with CCNC non-conformal sister, NECO. Although the NECO method performs well compared to other methods, CCNC outperforms the NECO method across all backtesting statistics.

During the prediction process, conformal updating is applied when the model starts to fail. In the CCNC method, this might happen when there is minimal contagion within the considered causal network, and the parametric part fails if the autoregressive part lacks sufficient information. To provide a clearer illustration of this process, we examine four distinct currencies: AUD, NZD, HKD, and NOK.[1] We analyse the pair AUD and NZD, where AUD has little contagion, but has a significant impact on NZD, which shows high levels of contagion. HKD is an example of a currency that is highly managed by its central bank. It is also the only currency for which volatility decreased significantly during the COVID-19 crisis period. The final currency is NZD, which experienced a significant and sudden increase in volatility as the crisis unfolded, making risk estimation more challenging. The log-returns of the exchange rates for these four currencies are plotted against the different VaR predictions in Figure 4.7.

[1] For a comprehensive view of the performance across all 20 currencies, please refer to Appendix B.3 where the results for the remaining currencies are presented.

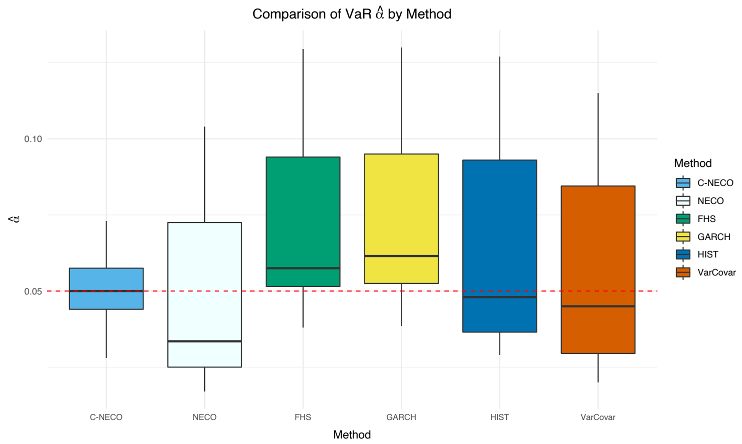

Figure 4.6: During the Covid Recession, all methods except for the CCNC experience bias and higher levels of variability in achieving the target VaR of $\alpha = 5\%$.

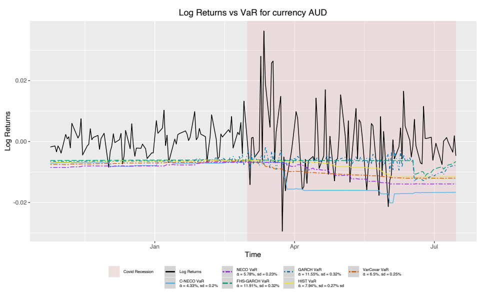

(a)

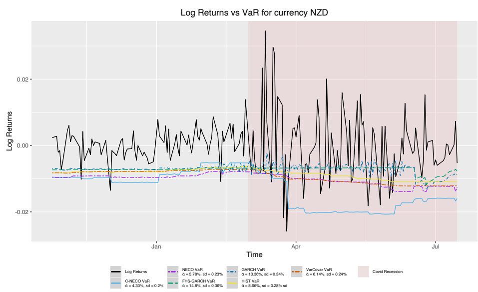

(b)

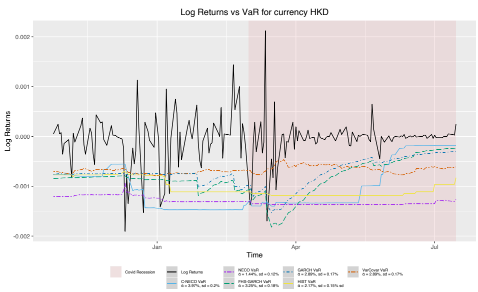

(c)

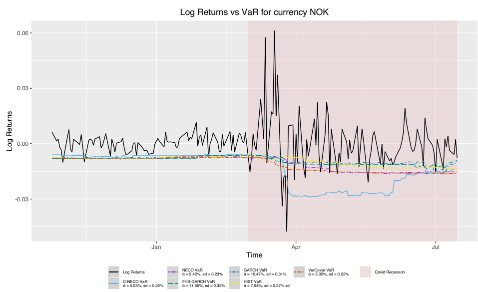

(d)

Figure 4.7: Out-of-sample predictions of VaR for six methods, including NECO without conformal updating. The predictions are plotted against the respective log-return realizations of the AUD, NZD, HKD, and NOK exchange rates over the USD during the Covid period. The VaR should closely follow the 5% quantile as possible. CCNC is the method that quickly identifies a sudden change in the volatility structure, both when volatility increases and when volatility goes down.

The performance of various parametric and non-parametric methods in handling the arrival of Covid for currencies such as AUD and NZD was found to be poor, with a significant amount of time passing before they returned to the targeted exceedance rate α. However, CCNC managed to follow the development of the log-return for both currencies. The next two currencies, HKD and NOK, highlight the importance of conformal updating and the difference between NECO, which does not use conformal updating, and CCNC. NOK experienced a significant jump in volatility after COVID-19, but the causal network contagion model was able to handle this change very well. Therefore, even though CCNC performed better, NECO was also able to perform well. When dealing with the HKD, which is a pegged currency, the central bank's interventions were significant and during the crisis, beyond the causal network measurements. As a result, NECO was unable to accurately predict the VaR for HKD, while CCNC was able to do so, due to the conformal updating that corrected the causal predictions.

4.5 Conclusion

The proposed Conformal Causal Network Contagion VaR (CCNC VaR) framework, which combines causal analysis, network contagion, and conformal updating techniques, represents a significant contribution to the field of risk management. This innovative approach offers a robust solution for addressing non-normality and non-stationarity in financial data. By incorporating conformal updating into the NECO approach, the CCNC VaR framework achieves improved accuracy and robustness, enhancing risk assessment capabilities.

The CCNC VaR framework has demonstrated its efficacy during times of crisis, proving robust measures even in the face of significant exogenous shocks to the system. By equipping financial institutions with a powerful tool for risk management, the CCNC VaR framework promises enhanced risk assessment and informed decision-making. Its versatility extends beyond the Forex, opening avenues for application in other markets and domains.

This paper also presented the theoretical background to the CCNC VaR framework and demonstrated its efficiency and robustness on the Forex through backtesting and comparisons with standard VaR measures. The COVID-19 recession served as a prime example of an exogenous shock to the system that highlights the potential of conformal updating for risk management, in general. We recommend further research to evaluate the performance of CCNC in other markets and domains.

	Measure	**Equation**	**Description**
$\hat{\alpha}$	Average Exceedance Rate	$\hat{\alpha} = \frac{\sum_{i=1}^{N} \mathbb{1}_{\alpha}}{N}$ $sd(\hat{\alpha}) = \sqrt{\frac{\sum_{i=1}^{N}(\mathbb{1}_{\alpha,i} - \hat{\alpha})^2}{N-1}}$ $\mathbb{1}_{\alpha} = \begin{cases} 0, & \text{if } x_i < VaR_i^{\alpha}, \\ 1, & \text{otherwise.} \end{cases}$	Computes the average exceedance rate, indicating how often the actual log returns fall below the VaR threshold for asset i.
AE	Actual over Expected Ratio	$AE = \frac{\sum_{i=1}^{N} \mathbb{1}_{\alpha}}{\alpha * N}$	Measures whether the VaR method underestimates or overestimates the number of exceedances compared to the expected level.
UC	Likelihood Test Unconditional Coverage	$LR_{UC} = -2\ln\left[\frac{(1-\alpha)^{n_0} \cdot \alpha^{n_1}}{(1-\hat{\alpha})^{n_0} \cdot \hat{\alpha}^{n_1}}\right]$ $n_1 = \sum_{i=1}^{N} \mathbb{1}_{\alpha}, \quad n_0 = N - n_1$ $\text{UC} = \%$ Not-rejected LR_{UC}	Tests if the proportion of violations significantly deviates from the expected level using a likelihood ratio test [127].
CC	Likelihood Test Conditional Coverage	$LR_{CC} = LR_{UC} + LR_{IND}$ $CC = \%$ Not-rejected LR_{CC}	Extends the UC test to account for violations clustering in time [34].
DQ	Dynamic Quantile Test	$LR_{DQ} = \text{Hit}^N Z(Z^N Z)^{-1} Z^N \frac{\text{Hit}}{\alpha(1-\alpha)}$ $\text{Hit} = \{\mathbb{1}_{\alpha,t} - \alpha, \cdots, \mathbb{1}_{\alpha,N} - \alpha\}$ $DQ = \%$ Not-rejected LR_{DQ}	Tests for dependence among violations and explanatory lagged variables Z, using a dynamic quantile approach [61, 51].
QL	Average Quantile Loss	$QL_i^{\alpha,B} = \frac{\sum_{t=1}^{N} \left(\alpha - \mathbb{1}_{\alpha,t}\right)\left(x_{i,t} - VaR_{i,t}^{\alpha,B}\right)}{N}$ $QL = \frac{\widehat{QL}^B}{\widehat{QL}^C}$	Computes the mean quantile loss penalising losses from VaR violations. When $QL < 1$ then the method B is more efficient than the benchmark method C [84].

Table 4.3: Backtesting measures to evaluate VaR performance

	CCNC	FHS	GARCH	HIST	VarCovar
$\hat{\alpha}$	**0.050**	0.052	**0.050**	0.054	0.052
$sd(\hat{\alpha})$	**0.019**	0.033	0.034	0.040	0.073
AE	**1.003**	1.048	**1.000**	1.086	1.032
$sd(AE)$	**0.389**	0.653	0.682	0.794	1.457
UC	**0.970**	0.820	0.810	0.700	0.690
CC	**0.970**	0.900	0.890	0.820	0.820
DQ	0.840	**0.880**	0.870	0.780	0.830
QL	1.000	**0.976**	0.984	1.025	1.016

Table 4.4: Backtesting study of different VaR models in the Forex for 20 currencies (2000-2021), using a 250-day estimation window for the target $\alpha = 5\%$. The best-performing results are in bold. If more than one result per row is in bold, no statistically significant difference was found (5% significance level).

	CCNC	NECO	FHS	GARCH	HIST	VarCovar
$\hat{\alpha}$	**0.0518**	0.0521	0.0783	0.0790	0.0677	0.0585
$sd(\hat{\alpha})$	**0.0216**	0.0428	0.0524	0.0560	0.0480	0.0431
AE	**1.0368**	1.0418	1.5668	1.5804	1.3539	1.1696
$sd(AE)$	**0.4310**	0.8557	1.0486	1.1197	0.9596	0.8624
UC	**0.9400**	0.6200	0.7100	0.6800	0.6800	0.6500
CC	**0.9600**	0.7400	0.7700	0.7200	0.7200	0.7400
DQ	**0.8000**	0.7500	0.6800	0.6600	0.6900	0.7200
QL	**1.0000**	1.0520	1.0604	1.0768	1.0687	1.0166

*Table 4.5: Backtesting study of different VaR models in the Forex for 20 currencies for the **Covid Recession** period, using a 250-day estimation window for the target $\alpha = 5\%$. The best-performing results are in bold. If more than one result per row is in bold, no statistically significant difference was found (5% significance level).*

Appendix A

Supplementary Material for Chapter 2

A.1 PC Stable Algorithm

STEP 1 - Finding the Skeleton We begin the procedure by defining all nodes of interest X and link all of them to create a complete graph, as in Figure A.1 (A).

Once we have the complete graph, we eliminate links between nodes X_i and X_j that are independent or conditionally independent, conditioning iteratively on a growing subset of the other nodes X_C in the graph. The algorithm considers all possible separating subsets X_C for each pair of nodes to test the hypothesis H_0 that nodes X_i and X_j are independent. If there is any subset X_C for which H_0 is not rejected at the specified significance level (often set at $\alpha = 0.05$), the link between the two nodes is removed from the network. Whatever is left at the end of the process is the so-called skeleton, an undirected graph G such that nodes (X_i, X_j) are connected with a link if and only if no set X_C can be found to make them conditionally independent. See an example of a skeleton in Figure A.1 (B).

STEP 2 - Applying causal orientation rules First, for each pair of nodes X_i and X_j that are not connected by a link, but both connecting to a common neighbour X_k, we check whether X_k belongs to the subset of links X_C from the previous step. Since we do not have a link between X_i and X_j, we know that a subset X_C exists such that these two are conditionally independent. If X_k does not belong to the subset X_C and yet we still have a link between these three nodes, we know that X_i and X_j influence X_k and not vice versa. This means we add directions $X_i \rightarrow X_k$ and $X_j \rightarrow X_k$ as seen in Figure A.2. This primary orientation is often referred to as a collider or inverted fork or V orientation. We can think of colliders as nodes that stop a path unless the analysis is conditioned on them — colliders

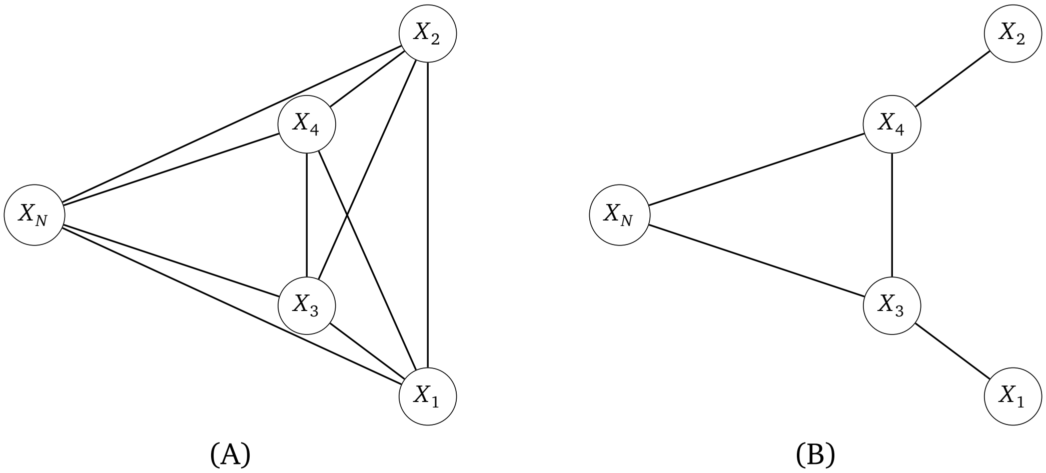

Figure A.1: (A) Complete causal graph for five financial instruments of interest. (B) The skeleton obtained from the complete graph using STEP 1.

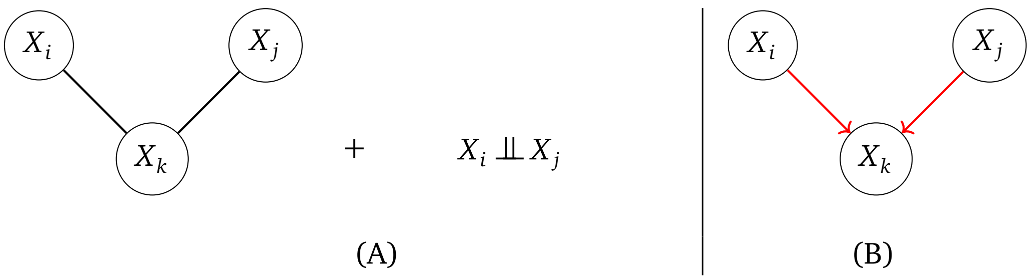

Figure A.2: Example of a collider.
Knowing that X_k is not a parent node for X_i and/or X_j as shown in (A), only one direction for the links between the three variables is logically possible – the red arrows shown in (B).

are random variables that appear on the left-hand side in a structural equation model with multiple variables on the right-hand side.

Once we have established all colliders within the graph, the PC algorithm tries to orient as many of the remaining links as possible by a set of consistency rules, as in [148]. The rules ensure that no newly directed link disrupts the previously established structure and no cycle is created. An example of such a rule is shown in Figure A.3.

Not all links can always be successfully oriented at the end of this step. If the graph we obtain after implementing the orientation rules is not fully directed, it is referred to as a Completed Partially Directed Acyclic Graph (CPDAG). In that case, a CPDAG is the best possible outcome we can obtain. It describes an equivalence class of DAGs that cannot be distinguished even with an infinite

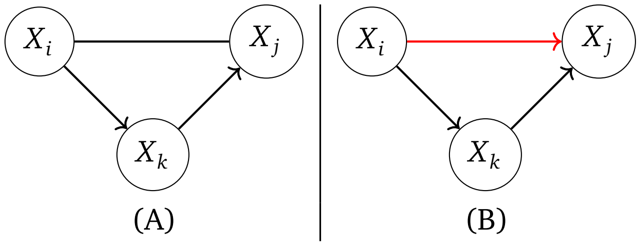

Figure A.3: Example of a consistency rule. The graph in (A) necessarily implies the graph in (B), because if the link went in the other direction the three nodes would form a cycle.

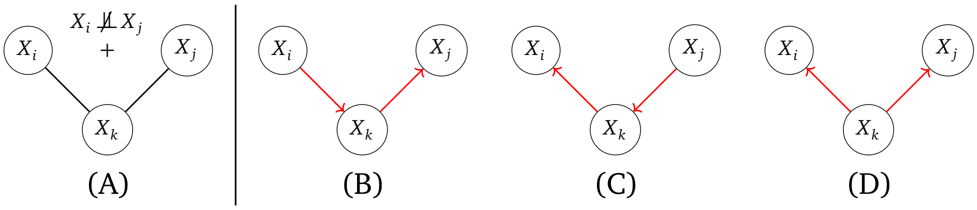

Figure A.4: Example of a CPDAG based on the Skeleton shown in (A). (B), (C) and (D) are Markov equivalent DAGs, or CPDAGs, under the assumption that X_k belongs to the subset of X_C that made X_i and X_j independent and so a collider solution as in Figure A.2 was not possible.

amount of data. An example of a CPDAG can be seen in Figure A.4. Unlike in Figure A.2, we cannot find a collider, and so all three DAGs (B),(C) and (D) can equally explain the skeleton in (A).

The more nodes or variables we have in a network, and therefore the more connectivity we have, the easier it is for the PC Algorithm to detect directions. This is in contrast to the curse of dimensionality that most models have. A trivial example would be a network with just two nodes: In such a situation no amount of observational data on these two nodes would allow the PC Algorithm to orient the graph and decide on the direction of the arrow between the two nodes.

A.2 Individual NECOFs

In this section, we present the individual NECOFs for all 23 currencies, according to formula 2.2.5. The values in Table A.1 are in percentages (0-100) and express the amount of contagion from other currencies in the overall period 2000-2021. Figure A.5 shows the development of the NECOF scores through time for each of the currencies.

Code	Currency Name	Minimum	Median	Mean	Maximum	StDev
AUD	AUSTRALIAN DOLLAR	0	0	0	0	0
BRL	BRAZILIAN REAL	0	0	5.68	52.1	14.73
CAD	CANADIAN DOLLAR	0	33.95	36.53	68.3	18.43
CHF	SWISS FRANC	0	72.9	61.73	92.3	28.48
CNY	CHINESE YUAN RENMINBI	0	0	0.08	7.2	0.78
DKK	DANISH KRONE	64.3	99.6	97.42	99.9	7.29
EUR	EURO	0	0	0.44	38	4.1
GBP	BRITISH POUND	0	0	15.51	57.4	21.76
HKD	HONG KONG DOLLAR	0	0	1.48	25	4.5
INR	INDIAN RUPEE	0	0	2.99	24.1	5.97
JPY	JAPANESE YEN	0	0	6.77	49.6	13.48
KRW	SOUTH KOREAN WON	0	6.4	11.42	47.5	13.68
LKR	SRI LANKAN RUPEE	0	0	1.28	13.9	3.07
MXN	MEXICAN PESO	0	37.1	36.48	73	20.13
MYR	MALAYSIAN RINGGIT	0	20.5	20.5	57.6	18.33
NOK	NORWEGIAN KRONE	0	40.2	38.5	81.4	26.76
NZD	NEW ZEALAND DOLLAR	34.7	64.55	63.89	81.9	12.3
SEK	SWEDISH KRONA	0	63.95	62.99	89.7	18.19
SGD	SINGAPORE DOLLAR	17.2	65.75	62.57	83.2	16.22
THB	THAI BAHT	0	39.9	38.27	62.6	15.58
TWD	NEW TAIWAN DOLLAR	2.9	39.6	39.1	75.8	18.74
VEB	VENEZUELAN BOLIVAR	0	0	0.57	10	1.68
ZAR	SOUTH AFRICAN RAND	0	48.45	41.43	75.1	20.64

Table A.1: Overview of the summary statistics for the Network Contagion Factors (NECOFs) on individual 23 exchange rates over the USD, for the period January 2000 to April 2021. The values are in percentages between 0 and 100.

A.3 Supplement for Section 2.4

In these supplementary materials, we provide the details of the various analyses performed in the manuscript.

A.3.1 Change in Contagion Network Statistics Over Time

In the paper, it was claimed that the various statistics describing the contagion in the Forex market change through time. Figures A.6-A.8 show the development of the Market NECOF, the number of clusters in the causal network and the market density, respectively. Table A.2 shows the significance level of the test for the need for a temporal spline for each of the three contagion indices. They show a highly significant result, meaning that the indices are not constant over time. From Figure A.6 it is clear that the market NECOF increases over time. Figure A.7 shows that the number of clusters in the contagion network decreases over time, whereas Figure A.8 illustrates the increase of the market density.

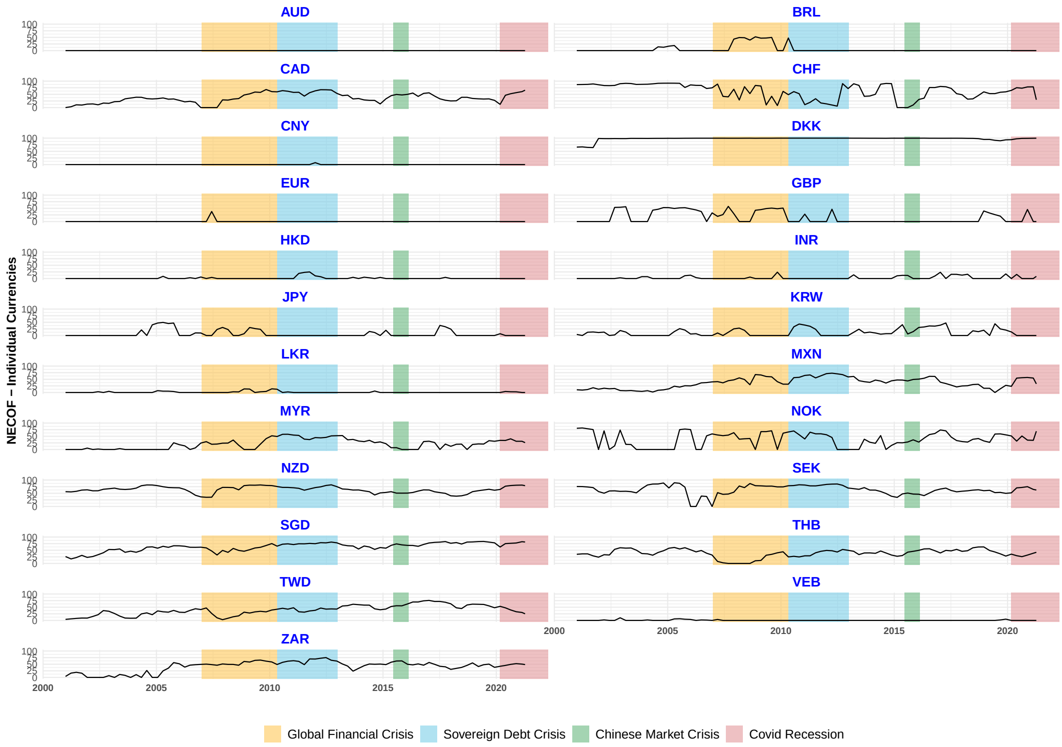

Figure A.5: NECOF through time for each currency. The financial crises highlighted are: the Global Financial Crisis (August 2007 - April 2010) with the following Sovereign Debt Crisis (May 2010 - December 2012), the Chinese Market Crisis (June 2015 - February 2016) and the COVID-19 Recession (starting March 2020).

A.3.2 Significance Test of the Reaction of Contagion Indices to Major Financial Crises

This section describes the formal tests for the various contagion indices to the various financial crises. We performed an Analysis of Variance and the results are shown in Table A.3. Both market NECOF and network density are significantly affected by the financial crises, whereas the number of clusters is not. Table A.4 and Table A.5 show the results of the linear regression for the market NECOF and the market density each. The market NECOF increased significantly during the Global Financial Crisis, the Sovereign Debt Crisis and the COVID-19 Recession, but stayed relatively flat during the Chinese Market Crisis. The market network density reacted mainly to the Sovereign Debt Crisis, the Chinese Market Crisis, and the COVID-19 Recession.

	Estimate	Std. Error	t value	Pr(>| t |)
Market NECOF	28.0709	0.3212	87.38	<2e-16 ***
Number of Clusters	7.8023	0.1329	58.7	<2e-16 ***
Market density	9.53663	0.08548	111.6	<2e-16 ***

Table A.2: Significance test for the temporal change in, respectively, market NECOF, number of clusters, and market density.

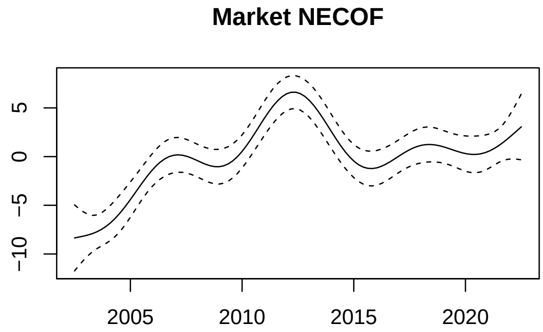

Figure A.6: Thin plate spline of the temporal change in market NECOF between 2000 and 2021. The upper and lower dotted lines are at 2 standard errors above and below the estimate.

A.3.3 Stability of Clustering Effects during the Global Financial Crisis and the Sovereign Debt Crisis

We tested cluster membership stability during the Global Financial Crisis and the Sovereign Debt Crisis. The results of the Welch two-sample t-test are shown in Table A.6. The test suggests greater stability during the Sovereign Debt Crisis than the Global Financial Crisis. Figure A.9 visually supports this finding.

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Market NECOF	4	511.5	127.88	7.916	1.95e-05 ***
Number of Clusters	4	10.77	2.692	1.064	0.38
Contagion Network Density	4	39.1	9.775	3.601	0.00941 **

Table A.3: Analysis of Variance of the reaction of market NECOF to financial crises.

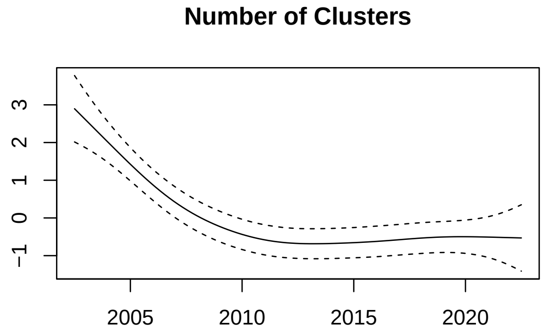

Figure A.7: Thin plate spline of the temporal development of the number of clusters between 2000 and 2021. The upper and lower dotted lines are at 2 standard errors above and below the estimate.

Parametric coefficients	Estimate	Std. Error	t-value	$\Pr(>\|t\|)$
(Intercept)	26.36189	0.56280	46.840	$< 2e-16$ ***
Global Financial Crisis	4.38811	1.28955	3.403	0.00104 **
Sovereign Debt Crisis	6.38356	1.33616	4.778	7.79e-06 ***
Chinese Market Crisis	-0.01233	1.88350	-0.007	0.99479
COVID-19 Recession	3.45177	1.62003	2.131	0.03615 *

Table A.4: Result of the linear regression showing the reaction of the market NECOF during the Global Financial Crisis, the Sovereign Debt Crisis, the Chinese Market Crisis and the COVID-19 Recession.

A.3.4 CHF NECOF Behaviour during a Financial Crisis

We claimed in Section 2.4.3 that NECOF of the Swiss Franc (CHF) has a tendency to go down in times of financial crises. Table A.7 presents the results of the Analysis of Variance on the NECOF of the CHF. The results show that financial crises have a notable and highly significant effect on the NECOF of the CHF during the four crisis periods considered.

Table A.8 shows the estimates, standard errors, t-values and p-values for the individual contagion effects during crises. The results indicate that the CHF NECOF is significantly affected by the Global Financial Crisis, Sovereign Debt Crisis, and Chinese Market Crisis with p-values less than 0.0001.

Parametric coefficients	Estimate	Std. Error	t-value	Pr(>| t |)
(Intercept)	9.0359	0.2307	39.163	$< 2e-16$ ***
Global Financial Crisis	0.8133	0.5287	1.538	0.12785
Sovereign Debt Crisis	1.1350	0.5478	2.072	0.04144 *
Chinese Market Crisis	2.2681	0.7722	2.937	0.00431 **
COVID-19 Recession	1.3541	0.6641	2.039	0.04472 *

Table A.5: Result of the linear regression showing the reaction of the market density during the Global Financial Crisis, the Sovereign Debt Crisis, the Chinese Market Crisis and the COVID-19 Recession.

Difference	GFC	SDC	df	t-value	p-value
Clustering stability	208.9	222.4	13.787	-2.818	0.01385

Table A.6: Welch Two Sample t-test result comparing the clustering stability between the Global Financial Crisis (GFC) and Sovereign Debt Crisis (SDC).

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Crises	4	28085	7021	13.91	1.1e-08 ***
Residuals	81	40871	505		

Table A.7: Analysis of Variance of the reaction of NECOF for the CHF to financial crises.

Parametric coefficients	Estimate	Std. Error	t-value	Pr(>| t |)
(Intercept)	74.137	3.145	23.570	< 2e-16 ***
Global Financial Crisis	-23.596	7.207	-3.274	0.00156 **
Sovereign Debt Crisis	-39.037	7.468	-5.228	1.31e-06 ***
Chinese Market Crisis	-58.937	10.527	-5.599	2.86e-07 ***
COVID-19 Recession	-8.523	9.054	-0.941	0.34933

Table A.8: Result of the linear regression showing the reaction of the CHF NECOF during the Global Financial Crisis, the Sovereign Debt Crisis, the Chinese Market Crisis and the COVID-19 Recession.

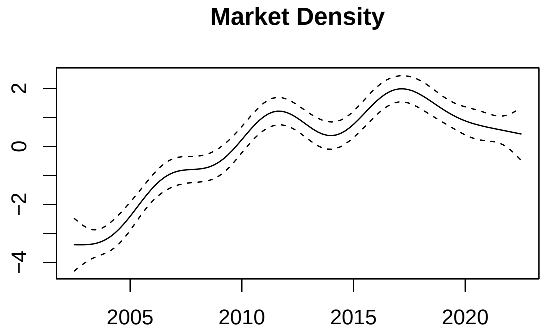

Figure A.8: Thin plate spline of the market density between 2000 and 2021. The upper and lower dotted lines are at 2 standard errors above and below the estimate.

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
Crises	4	5603	1400.6	3.274	0.0154 *
Residuals	81	34655	427.8		

Table A.9: Analysis of Variance of the reaction of NECOF for the GBP to financial crises.

GBP NECOF Behaviour during a Financial Crisis

We claimed in Section 2.4.3 that the NECOF of the British Pound (GBP) reacts to periods of financial crises. The Analysis of Variance results presented in Table A.9 indicate that being in a period of financial crisis has a statistically significant effect on the NECOF of the GBP. These results suggest that financial crises are an important factor in explaining the variation in the NECOF of the GBP, and should be carefully considered when analyzing the factors that affect this variable.

Table A.10 shows the individual contagion factors for the four financial crises. These estimates indicate the effect of each crisis on the NECOF of the GBP. The most significant effect observed is the increase in the GBP NECOF during the Global Financial Crisis.

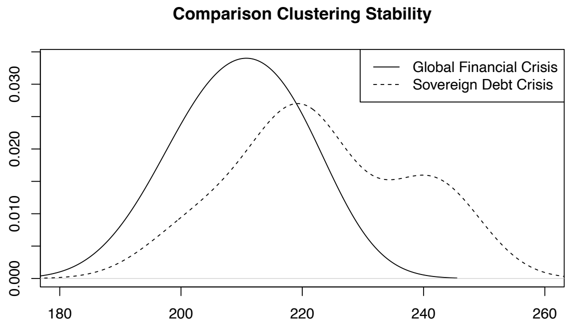

Figure A.9: Comparison of the estimated probability density function for the stability of memberships within a cluster during the Global Financial Crisis and the Sovereign Debt Crisis, where the right shift during the Sovereign Debt Crisis indicates higher stability in the clustering.

Parametric coefficients	Estimate	Std. Error	t-value	Pr(>| t |)
(Intercept)	16.461	2.896	5.683	2.01e-07 ***
Global Financial Crisis	14.731	6.636	2.220	0.0292 *
European Debt Crisis	-9.679	6.876	-1.408	0.1631
Chinese Market Crisis	-16.461	9.693	-1.698	0.0933
COVID-19 Recession	-9.932	8.337	-1.191	0.2370

Table A.10: Result of the linear regression showing the reaction of the GBP NECOF during the Global Financial Crisis, the Sovereign Debt Crisis, the Chinese Market Crisis and the COVID-19 Recession.

Appendix B

Supplementary Material for Chapter 3

B.1 Pseudo-Code for the CCNC VaR

Algorithm 1 provides the pseudo-code for the CCNC VaR estimation framework. This framework consists of a series of steps, including data transformation with copula, causal network analysis, estimation of the causal graphical model, VaR computation, and conformal updating. As conformal updating is designed for online updating, steps 4 and 5 must be iteratively executed, continuously adapting to new data as it becomes available.

B.2 Target α 1% and 10%

To show the consistency of the conformal NECO method we show the results of the backtesting for the $\alpha = 1\%$ in Table B.1 and $\alpha = 10\%$ in Table B.2. The boxplots can be seen in Figure B.1 (a) and Figure B.1 (b) respectively. There is minimal effect of the α level chosen on the performance.

Algorithm 1: CCNC VaR Estimation Pseudo Code

1 **Step 1: Data Transformation with Copula**
Data: Input data - log returns on forex exchange rates over USD
Result: Transformed data using copula
Fit a copula to the data: $copula \leftarrow fit\_copula(data)$
Transform the data using the fitted copula:
$transformed\_data \leftarrow transform\_with\_copula(data, copula)$

2 **Step 2: Analyse Causal Network**
Data: Transformed data
Result: Causal network estimation
Apply the PC algorithm to estimate the causal network:
$causal\_network \leftarrow pc\_algorithm(transformed\_data)$
Additional steps to validate and refine the causal network, if needed

3 **Step 3: Estimate the Causal Graphical Model**
Data: Transformed data, causal network
Result: Causal Graphical Model parameters
Knowing the causal structure, estimate : $causal\_\widehat{graphical}\_model \leftarrow$
$structural\_VAR\_estimation(transformed\_data, causal\_network)$

4 **Step 4: Value at Risk Computation**
Data: Transformed data, Causal Graphical Model parameters
Result: Value at Risk (VaR) estimation
Calculate VaR using the estimated causal network and other relevant parameters: $predict\_mean, predict\_sd \leftarrow$
$predict(transformed\_data, causal\_\widehat{graphical}\_model)$
$VaR\_transformed \leftarrow VaR\_compute(predict\_mean, predict\_sd, \alpha)$
$VaR \leftarrow inverse\_copula(VaR\_transformed, transformed\_data)$

5 **Step 5: Conformal Updating - Chaining the Alpha Target at each δt step**
Data: t realised VaR, t realised data, α_target, γ step size
Result: α
Determine direction and size of the error in prediction in the last t predictions:
$error \leftarrow$
$distance\_from\_target(previous\_VaR, previous\_data, alpha\_target)$
Update the α_target: $\alpha \leftarrow conformal\_update(error, \gamma)$

6 **Repeat Step 4 and 5 each Δt**

	CCNC	FHS	GARCH	HIST	VarCovar
$\hat{\alpha}$	**0.0126**	**0.0132**	0.0170	0.0153	0.0198
sd($\hat{\alpha}$)	**0.0129**	0.0167	0.0214	0.0177	0.0569
AE	**1.2560**	**1.3190**	1.6980	1.5340	1.9840
$sd(AE)$	**1.2947**	1.6703	2.1432	1.7731	5.6946
LR_{UC}	**0.9300**	0.9100	0.8800	0.8700	0.8500
LR_{CC}	**0.9500**	0.9400	0.9100	0.8900	0.8700
DQ	0.8200	**0.8900**	0.8300	0.7700	0.7800
QL	1.0000	**0.9404**	0.9530	1.0335	1.0171

Table B.1: Backtesting study of different VaR models in the Forex market for 20 currencies (2000-2021), using a 250-day estimation window for the target $\alpha = 1\%$. The best-performing results are in bold. If more than one result per row is in bold, no statistically significant difference was found (5% significance level).

	CCNC	FHS	GARCH	HIST	VarCovar
$\hat{\alpha}$	**0.1014**	**0.1013**	0.0907	**0.1032**	0.0893
sd($\hat{\alpha}$)	**0.0275**	0.0449	0.0466	0.0525	0.0832
AE	**1.0223**	**1.0126**	0.9069	**1.0317**	0.8934
$sd(AE)$	**0.3045**	0.4493	0.4664	0.5250	0.8317
LR_{UC}	**0.9600**	0.8500	0.8200	0.7700	0.7200
LR_{CC}	**0.9500**	0.8600	0.8500	0.8200	0.7600
DQ	0.8600	0.8900	**0.9000**	0.8200	0.8800
QL	1.0000	**0.9809**	0.9946	1.0090	1.0112

Table B.2: Backtesting study of different VaR models in the Forex market for 20 currencies (2000-2021), using a 250-day estimation window for the target $\alpha = 10\%$. The best-performing results are in bold. If more than one result per row is in bold, no statistically significant difference was found (5% significance level).

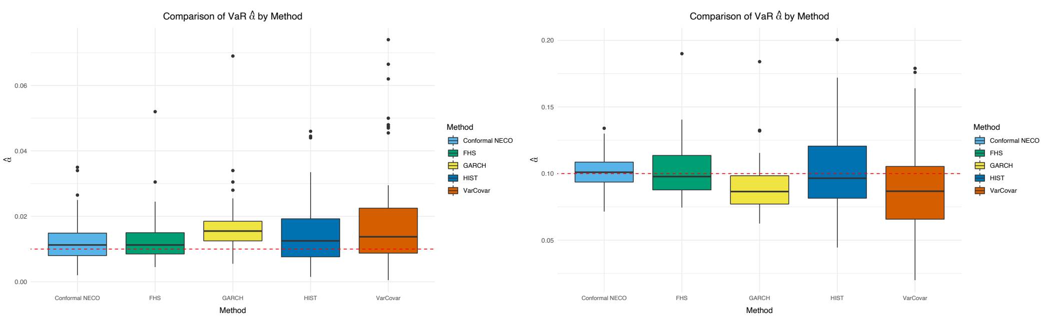

(a) (b)

Figure B.1: Boxplots illustrating the distribution of out-of-sample exceedance rates ($\hat{\alpha}$) for 20 currencies, allowing to assess how closely each model aligns with the targeted α throughout the entire period. (a) has the target $\alpha = 1\%$ and in (b) target $\alpha = 10\%$.

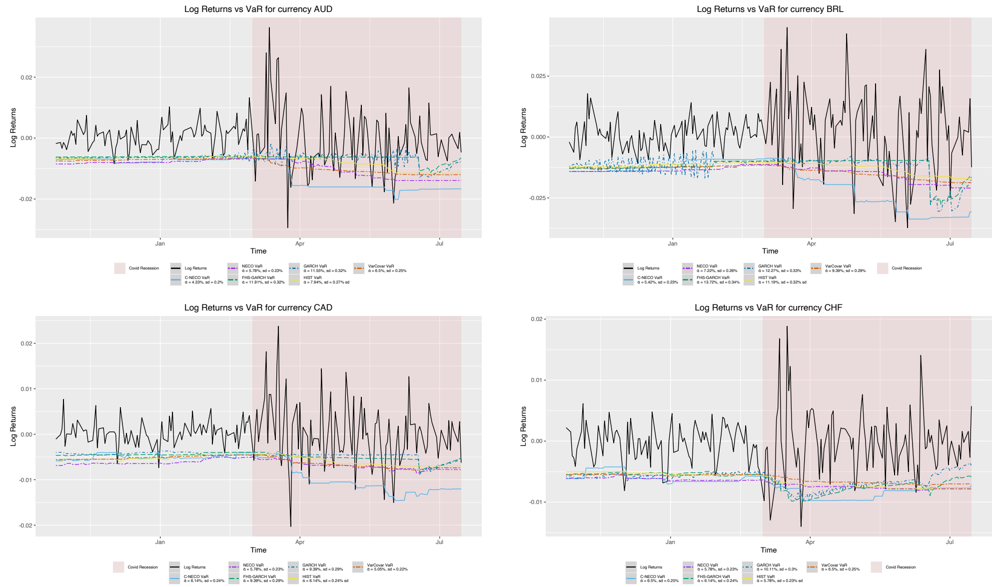

Figure B.2: Out-of-sample predictions of VaR for six methods, including NECO without conformal updating. The predictions are plotted against the respective log-return realizations of individual exchange rates over the USD during the COVID-19 period. The VaR are computed for all methods with an estimation window of 250 days, for 1-day ahead rolling predictions.

B.3 Individual Currencies during COVID-19

In this Section, we present the performance of each of the VaR estimations for all 20 individual currencies in our dataset. Out-of-sample predictions of VaR for the six methods, including NECO without conformal updating are presented in Figure B.2.

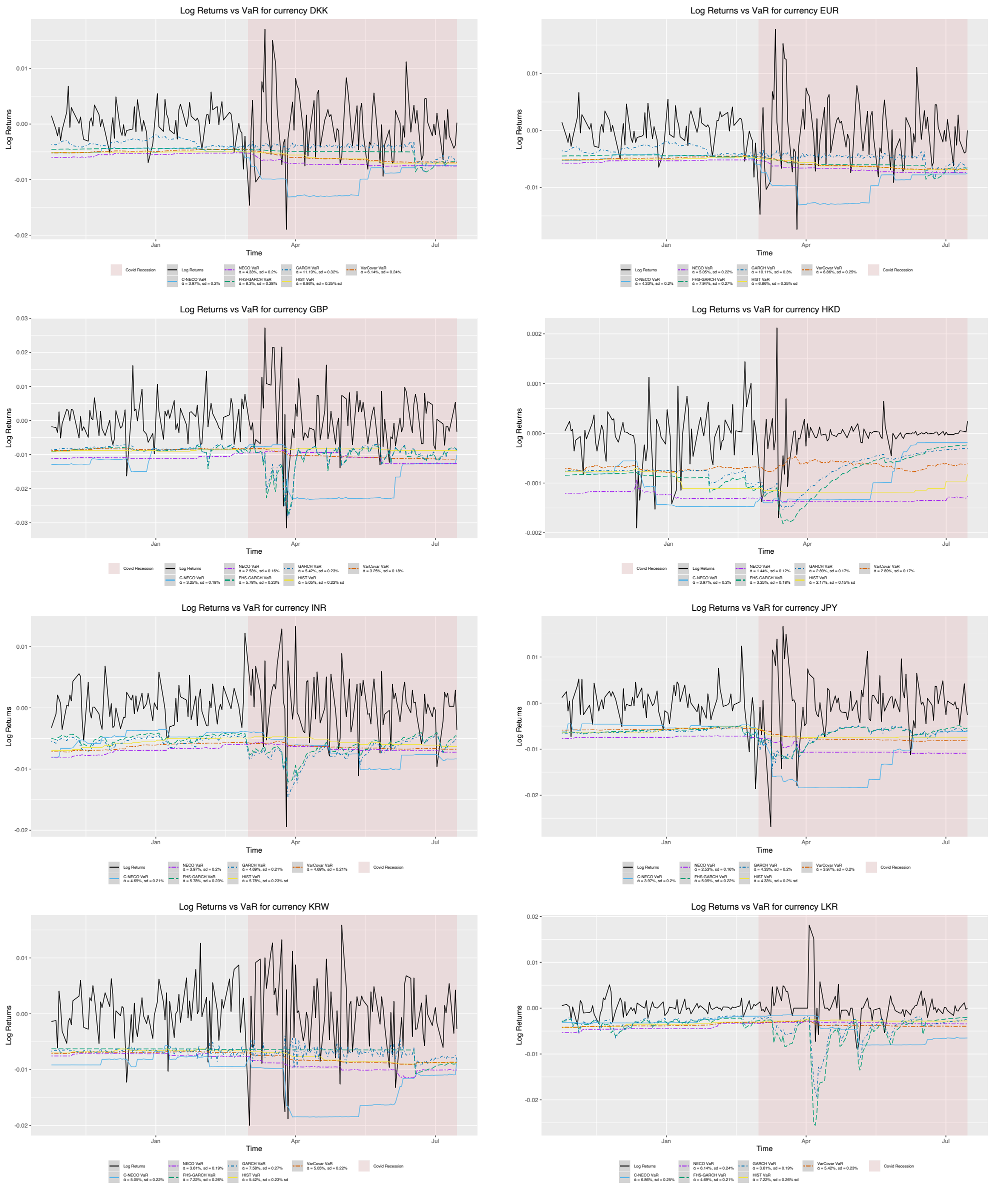

Figure B.2: Out-of-sample predictions of VaR for six methods, including NECO without conformal updating. The predictions are plotted against the respective log-return realizations of individual exchange rates over the USD during the COVID-19 period. The VaR are computed for all methods with an estimation window of 250 days, for 1-day ahead rolling predictions.

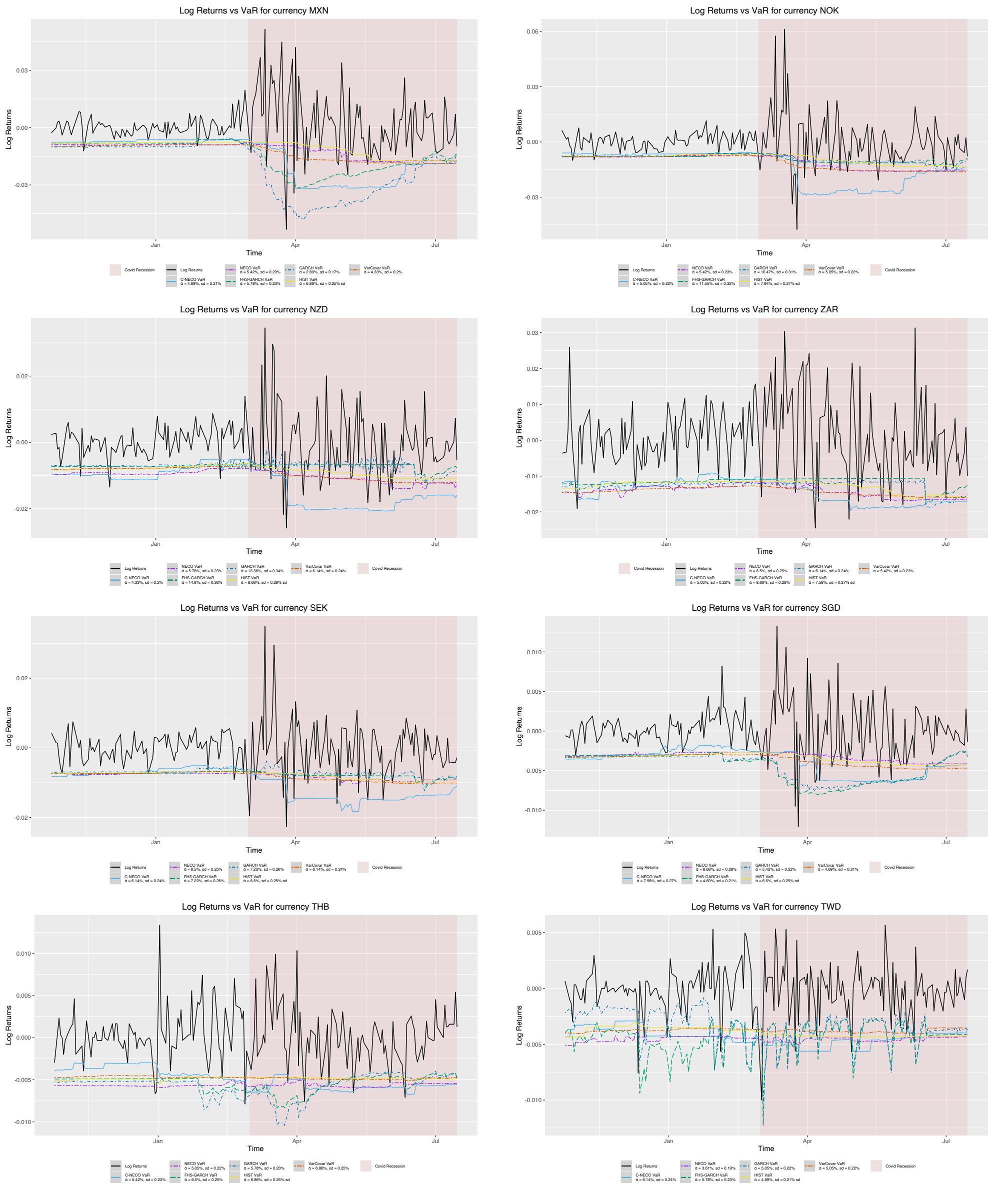

Figure B.2: Out-of-sample predictions of VaR for six methods, including NECO without conformal updating. The predictions are plotted against the respective log-return realizations of individual exchange rates over the USD during the COVID-19 period. The VaR are computed for all methods with an estimation window of 250 days, for 1-day ahead rolling predictions.

Bibliography

[1] Fentaw Abegaz and Ernst C. Wit. Copula gaussian graphical models with penalized ascent monte carlo em algorithm. *Statistica Neerlandica*, 69(4): 419–441, 2015.

[2] Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105 (2):564–608, 2015.

[3] Zeno Adams, Roland Füss, and Reint Gropp. Spillover effects among financial institutions: A state-dependent sensitivity value-at-risk approach. *Journal of Financial and Quantitative Analysis*, 49(3):575–598, 2014. doi: 10.1017/S0022109014000325.

[4] Arianna Agosto, Daniel Felix Ahelegbey, and Paolo Giudici. Tree networks to assess financial contagion. *Economic Modelling*, 85:349–366, 2020.

[5] Daniel Felix Ahelegbey, Monica Billio, and Roberto Casarin. Bayesian graphical models for structural vector autoregressive processes. *Journal of Applied Econometrics*, 31(2):357–386, 2016.

[6] Robert Z. Aliber and Gylfi Zoega. The 2008 global financial crisis in retrospect. *Cham, Switzerland: Palgrave Macmillan. Benediktsdóttir, Sigrídur, Jón Daníelsson, and Gylfi Zoega (2011).“Lessons from a Collapse of a Financial System.” Economic Policy*, 26:183–235, 2019.

[7] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul F. Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *CoRR*, abs/1606.06565, 2016. URL http://arxiv.org/abs/1606.06565.

[8] Joshua D. Angrist, Guido W. Imbens, and Donald B. Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.

[9] David Ardia, Kris Boudt, and Leopoldo Catania. Generalized autoregressive score models in R: The GAS package. *Journal of Statistical Software*, 88(6):1–28, 2019. doi: 10.18637/jss.v088.i06.

[10] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

[11] Stefan Avdjiev, Paolo Giudici, and Alessandro Spelta. Measuring contagion risk in international banking. *Journal of Financial Stability*, 42:36–51, 2019.

[12] Kee-Hong Bae, G. Andrew Karolyi, and Reneé M. Stulz. A new approach to measuring financial contagion. *The Review of Financial Studies*, 16(3): 717–763, 2003.

[13] Vineeth Balasubramanian, Shen-Shyang Ho, and Vladimir Vovk. *Conformal prediction for reliable machine learning: theory, adaptations and applications*. Newnes, 2014.

[14] Matteo Barigozzi and Marc Hallin. A network analysis of the volatility of high dimensional financial series. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 66(3):581–605, 2017.

[15] Matteo Barigozzi, Marc Hallin, Stefano Soccorsi, and Rainer von Sachs. Time-varying general dynamic factor models and the measurement of financial connectedness. *Journal of Econometrics*, 222(1):324–343, 2021.

[16] Giovanni Barone-Adesi and Kostas Giannopoulos. Non-parametric VaR techniques. myths and realities. *Economic Notes*, 30(2):167–181, 2001.

[17] Giovanni Barone-Adesi, Frederick Bourgoin, and Kostas Giannopoulos. Don't look back. *Risk*, 11:100–104, August 1998.

[18] Giovanni Barone-Adesi, Kostas Giannopoulos, and Les Vosper. VaR without correlations for portfolios of derivative securities. *Journal of Futures Markets*, 19(5):583–602, 1999.

[19] Basel Committee on Banking Supervision. Overview of the amendment to the capital accord to incorporate market risks. *BIS*, BCBS 23, January 1996.

[20] Stefano Battiston and Serafin Martinez-Jaramillo. Financial networks and stress testing: Challenges and new research avenues for systemic risk analysis and financial stability implications. *Journal of Financial Stability*, 35: 6–16, 2018.

[21] Walter Bazán-Palomino and Diego Winkelried. FX markets' reactions to COVID-19: are they different? *International Economics*, 167:50–58, 2021.

[22] Carol C. Bertaut, Bastian von Beschwitz, and Stephanie E. Curcuru. The international role of the u.s. dollar. *FEDS Notes*, 2021. URL https://doi.org/10.17016/2380-7172.2998.

[23] Monica Billio, Mila Getmansky, Andrew W. Lo, and Loriana Pelizzon. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559, 2012.

[24] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, oct 2008. doi: 10.1088/1742-5468/2008/10/P10008. URL https://dx.doi.org/10.1088/1742-5468/2008/10/P10008.

[25] Kenneth A. Bollen and Judea Pearl. Eight myths about causality and structural equation models. In *Handbook of Causal Analysis for Social Research*, pages 301–328. Springer, 2013.

[26] Claudio Borio. The COVID-19 economic crisis: Dangerously unique. *Business Economics*, 55(4):181–190, 2020.

[27] Dennis Botman, Irineu de Carvalho Filho, and W. Raphael Lam. The curious case of the yen as a safe haven currency: A forensic analysis. IMF Working Paper WP/13/228, International Monetary Fund, Asia and Pacific Department, November 2013.

[28] Patrick Burns, Robert Engle, and Joseph Mezrich. Correlations and volatilities of asynchronous data. *Journal of Derivatives*, 5(4):7, 1998.

[29] J. David Cabedo and Ismael Moya. Estimating oil price 'value at risk' using the historical simulation approach. *Energy Economics*, 25(3):239–253, 2003.

[30] Sara Calvo and Carmen M. Reinhart. Capital flows to latin america: Is there evidence of contagion effects. Policy, Research working paper WPS 1619, World Bank Group, Washington, D.C., 1996. URL http://documents.worldbank.org/curated/en/300461468753000150/Capital-flows-to-Latin-America-Is-there-evidence-of-contagion-effects.

[31] Sean D. Campbell. A review of backtesting and backtesting procedures. *Journal of Risk*, 9(2):1–17, 2007.

[32] Philipp Carlsson-Szlezak, Martin Reeves, and Paul Swartz. What coronavirus could mean for the global economy. *Harvard Business Review*, 3: 1–10, 2020.

[33] Alain Chaboud, Dagfinn Rime, and Vladyslav Sushko. The foreign exchange market. In *The Research Handbook of Financial Markets*, page 253. Edward Elgar Publishing, 2022.

[34] Peter F. Christoffersen. Evaluating interval forecasts. *International Economic Review*, 39(4):841–862, November 1998.

[35] Stijn Claessens and Kirsten Forbes. *International Financial Contagion*. Springer Science & Business Media, 2013.

[36] Daryl Collins and Nicholas Biekpe. Contagion: a fear for african equity markets? *Journal of Economics and Business*, 55(3):285–297, 2003.

[37] Diego Colombo and Marloes H. Maathuis. Order-independent constraint-based causal structure learning. *Journal of Machine Learning Research*, 15 (116):3921–3962, 2014. URL http://jmlr.org/papers/v15/colombo14a.html.

[38] Richard N. Cooper. The future of the dollar. Peterson Institute for International Economics Policy Brief 09-21, Peterson Institute for International Economics, Washington, DC, 2009.

[39] Rainer Dahlhaus and Michael Eichler. Causality and graphical models for time series. In Peter J Green, Nils Lid Hjort, and Sylvia Richardson, editors, *Highly structured stochastic systems*, pages 115–137. University Press, Oxford, 2003.

[40] Jon Danielsson, Robert Macrae, and Andreas Uthemann. Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140:106290, 2022.

[41] Thong M. Dao, Frank McGroarty, and Andrew Urquhart. The brexit vote and currency markets. *Journal of International Financial Markets, Institutions and Money*, 59:153–164, 2019. ISSN 1042-4431. doi: https://doi.org/10.1016/j.intfin.2018.11.004. URL https://www.sciencedirect.com/science/article/pii/S1042443118301367.

[42] Sharada N. Davidson. Interdependence or contagion: A model switching approach with a focus on Latin America. *Economic Modelling*, 85:166–197, 2020.

[43] Reinout De Bock and Irineu de Carvalho Filho. The behavior of currencies during risk-off episodes. *Journal of International Money and Finance*, 53: 218–234, 2015.

[44] Nicolas Dewolf, Bernard De Baets, and Willem Waegeman. Valid prediction intervals for regression problems. *Artificial Intelligence Review*, 56(1): 577–613, 2023.

[45] Francis X. Diebold and Kamil Yılmaz. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134, 2014.

[46] Adrian Dobra and Alex Lenkoski. Copula gaussian graphical models. Technical report, Department of Statistics, University of Washington, 2009.

[47] Catherine Donnelly and Paul Embrechts. The devil is in the tails: actuarial mathematics and the subprime mortgage crisis. *ASTIN Bulletin: The Journal of the IAA*, 40(1):1–33, 2010.

[48] Kevin Dowd. *Measuring market risk*. John Wiley & Sons, 2007.

[49] Mathias Drehmann and Vladyslav Sushko. The global foreign exchange market in a higher-volatility environment. *BIS Quarterly Review, December*, pages 33–48, dec 2022.

[50] Darrell Duffie and Jun Pan. An overview of value at risk. *Journal of Derivatives*, 4(3):7–49, 1997.

[51] Elena-Ivona Dumitrescu, Christophe Hurlin, and Vinson Pham. Backtesting value-at-risk: from dynamic quantile to dynamic binary tests. *Finance*, 33(1):79–112, 2012.

[52] Mardi Dungey and Dinesh Gajurel. Contagion and banking crisis – international evidence for 2007–2009. *Journal of Banking & Finance*, 60: 271–283, 2015. ISSN 0378-4266. doi: https://doi.org/10.1016/j.jbankfin.2015.08.007. URL https://www.sciencedirect.com/science/article/pii/S0378426615002186.

[53] Mardi Dungey and Eric Renault. Identifying contagion. *Journal of Applied Econometrics*, 33(2):227–250, 2018.

[54] Francis Y. Edgeworth. The mathematical theory of banking. *Journal of the Royal Statistical Society*, 51(1):113–127, 1888.

[55] Sebastian Edwards. Contagion. *The World Economy*, 23(7):873–900, July 2000. doi: 10.1111/1467-9701.00307. URL https://ideas.repec.org/a/bla/worlde/v23y2000i7p873-900.html.

[56] Barry Eichengreen. International currencies in the lens of history. *Handbook of the History of Money and Currency*, pages 335–359, 2020.

[57] Michael Eichler. Granger causality and path diagrams for multivariate time series. *Journal of Econometrics*, 137(2):334–353, 2007.

[58] Matthew Elliott, Benjamin Golub, and Matthew O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–53, 2014.

[59] Scott Ellis, Satish Sharma, and Janusz Brzeszczyński. Systemic risk measures and regulatory challenges. *Journal of Financial Stability*, 61:100960, 2022. ISSN 1572-3089. doi: https://doi.org/10.1016/j.jfs.2021.100960. URL https://www.sciencedirect.com/science/article/pii/S1572308921001194.

[60] Paul Embrechts, Alexander McNeil, and Daniel Straumann. Correlation and dependence in risk management: properties and pitfalls. In *Risk management: value at risk and beyond*, pages 176–223. Cambridge University Press, Cambridge, UK, 2002.

[61] Robert F. Engle and Simone Manganelli. CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4):367–381, 2004.

[62] Rasmus Fatum and Yohei Yamamoto. Intra-safe haven currency behavior during the global financial crisis. *Journal of International Money and Finance*, 66:49–64, 2016.

[63] Daniel J. Fenn, Mason A. Porter, Peter J. Mucha, Mark McDonald, Stacy Williams, Neil F. Johnson, and Nick S. Jones. Dynamical clustering of exchange rates. *Quantitative Finance*, 12(10):1493–1520, 2012.

[64] Jaime de Jesus Filho, Paulo Matos, and Ronald Fonseca. The role of contagion and integration in risk management measures. *Global Business Review*, 0(0), 2020. doi: 10.1177/0972150920933857. URL https://doi.org/10.1177/0972150920933857.

[65] Matteo Fontana, Gianluca Zeni, and Simone Vantini. Conformal prediction: a unified review of theory and new challenges. *Bernoulli*, 29(1): 1–23, 2023.

[66] Kristin J. Forbes and Roberto Rigobon. Measuring contagion: conceptual and empirical issues. In *International Financial Contagion*, pages 43–66. Springer, 2001.

[67] Kristin J. Forbes and Roberto Rigobon. No contagion, only interdependence: measuring stock market comovements. *The Journal of Finance*, 57 (5):2223–2261, 2002.

[68] Kristin J. Forbes, Roberto Rigobon, Graciela L. Kaminsky, and Andrea Repetto. Contagion in latin america: Definitions, measurement, and policy implications [with comments]. *Economía*, 1(2):1–46, 2001. ISSN 15297470. URL http://www.jstor.org/stable/20065404.

[69] Marcel Fratzscher. On currency crises and contagion. *International Journal of Finance & Economics*, 8(2):109–129, 2003.

[70] International Monetary Fund. Global financial stability report: Responding to the financial crisis and measuring systemic risks. Global Financial Stability Report 2008/001, International Monetary Fund, Washington, D.C., 2008.

[71] Prasanna Gai, Andrew Haldane, and Sujit Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470, 2011.

[72] Gelin Gao, Bud Mishra, and Daniele Ramazzotti. Causal data science for financial stress testing. *Journal of Computational Science*, 26:294–304, 2018.

[73] Ken Garbade. Assessing risk and capital adequacy for treasury securities. *Topics in money and Securities Markets*, 22, 1986.

[74] Mariia Garkavenko, Tatiana Beliaeva, Eric Gaussier, Hamid Mirisaee, Cédric Lagnier, and Agnès Guerraz. Assessing the factors related to a start-up's valuation using prediction and causal discovery. *Entrepreneurship Theory and Practice*, October 2022. URL https://journals.sagepub.com/doi/abs/10.1177/10422587221121291.

[75] Amy L. Gaskell and Jamie W. Sleigh. An introduction to causal diagrams for anesthesiology research. *Anesthesiology*, 132(5):951–967, 2020.

[76] Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34: 1660–1672, 2021.

[77] Patrick Gillespie. U.s. economy grinds to near halt at end of 2015. *CNN-Money*, January 2016. URL https://money.cnn.com/2016/01/29/news/economy/us-economy-gdp-fourth-quarter/index.html.

[78] Paolo Giudici and Iman Abu-Hashish. What determines bitcoin exchange prices? a network VAR approach. *Finance Research Letters*, 28:309–318, 2019.

[79] Paolo Giudici and Laura Parisi. CoRisk: Credit risk contagion with correlation network models. *Risks*, 6(3):95, 2018.

[80] Paolo Giudici, Peter Sarlin, and Alessandro Spelta. The interconnected nature of financial systems: Direct and common exposures. *Journal of Banking & Finance*, 112:105149, 2020.

[81] Paolo Giudici, Thomas Leach, and Paolo Pagnottoni. Libra or Librae? basket based stablecoins to mitigate foreign exchange volatility spillovers. *Finance Research Letters*, 44:102054, 2022.

[82] Paul Glasserman and H Peyton Young. How likely is contagion in financial networks? *Journal of Banking & Finance*, 50:383–399, 2015.

[83] Kun Glymour, Clark N.and Zhang and Peter Spirtes. Review of causal discovery methods based on graphical models. *FRONTIERS IN GENETICS*, 10, JUN 4 2019. doi: 10.3389/fgene.2019.00524.

[84] Gloria González-Rivera, Tae-Hwy Lee, and Santosh Mishra. Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk, and predictive likelihood. *International Journal of Forecasting*, 20 (4):629–645, 2004.

[85] Charles A.E. Goodhart and Thomas Hesse. Central bank forex internvention assessed in continous time. *Journal of International Money and Finance*, 12(4):368–389, 1993. ISSN 0261-5606. doi: https://doi.org/10.1016/0261-5606(93)90002-S. URL https://www.sciencedirect.com/science/article/pii/026156069390002S.

[86] Ian D. Gow, David F. Larcker, and Peter C. Reiss. Causal inference in accounting research. *Journal of Accounting Research*, 54(2):477–523, 2016.

[87] Clive W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1912791.

[88] Clive W. J. Granger. Some properties of time series data and their use in econometric model specification. *Journal of Econometrics*, 16(1):121–130, 1981. ISSN 0304-4076. doi: https://doi.org/10.1016/0304-4076(81)90079-8. URL https://www.sciencedirect.com/science/article/pii/0304407681900798.

[89] Sander Greenland. For and against methodologies: some perspectives on recent causal and statistical inference debates. *European Journal of Epidemiology*, 32(1):3–20, 2017.

[90] Pedro Gurrola-Perez and David Murphy. Filtered historical simulation value-at-risk models and their competitors. Bank of England working papers 525, Bank of England, 2015. URL https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2015/filtered-historical-simulation-value-at-risk-models-and-their-competitors.pdf.

[91] Andrew G. Haldane. Rethinking the financial network. In Stephan A. Jansen, Eckhard Schröter, and Nico Stehr, editors, *Fragile Stabilität – stabile Fragilität*, pages 243–278. Springer, 2013.

[92] Kristian Bondo Hansen. Contrarian investment philosophy in the american stock market: on investment advice and the crowd conundrum. *Economy and Society*, 44(4):616–638, 2015.

[93] Kristian Bondo Hansen. Financial contagion: problems of proximity and connectivity in financial markets. *Journal of Cultural Economy*, 14(4): 388–402, 2021. doi: 10.1080/17530350.2021.1879211. URL https://doi.org/10.1080/17530350.2021.1879211.

[94] Alain Hauser and Peter Bühlmann. Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *Journal of Machine Learning Research*, 13:2409–2464, 2012. URL https://jmlr.org/papers/v13/hauser12a.html.

[95] Alexandra Heath, Gabriele Galati, and Patrick McGuire. Evidence of carry trade activity. *BIS Quarterly Review*, September 2007.

[96] James Heckman and Salvador Navarro-Lozano. Using matching, instrumental variables, and control functions to estimate economic choice models. *Review of Economics and statistics*, 86(1):30–57, 2004.

[97] Christina Heinze-Deml, Jonas Peters, and Nicolai Meinshausen. Invariant causal prediction for nonlinear models. *Journal of Causal Inference*, 6(2): 20170016, 2018. doi: doi:10.1515/jci-2017-0016. URL https://doi.org/10.1515/jci-2017-0016.

[98] Callum Henderson. *Currency strategy: The practitioner's guide to currency investing, hedging and forecasting*. John Wiley & Sons, 2006.

[99] Darryll Hendricks. Evaluation of value-at-risk models using historical data. *Economic Policy Review*, 2(1):39–69, April 1996. URL https://dx.doi.org/10.2139/ssrn.1028807.

[100] Austin Bradford Hill. *The environment and disease: association or causation?* Sage Publications, 1965.

[101] David Hirshleifer and Siew Hong Teoh. Thought and behavior contagion in capital markets. In *Handbook of financial markets: Dynamics and evolution*, pages 1–56. Elsevier, 2009.

[102] Katerina Hlavácková-Schindler. Equivalence of granger causality and transfer entropy: A generalization. *Applied Mathematical Sciences*, 5(73): 3637–3648, 2011.

[103] Paul W. Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.

[104] Glyn A. Holton. History of value-at-risk: 1922-1998. Technical report, Contingency Analysis, Boston, United States, 2002.

[105] Kevin D Hoover. Causality in economics and econometrics. In *New Palgrave Dictionary of Economics*. Springer, 2006.

[106] Richard Hoppe. VaR and the unreal world: A value-at-risk calculation is only as good as the statistics that back it up-and they may not be as reliable as they seem. *RISK*, 11:45–50, 1998.

[107] Nikolai V. Hovanov, James W. Kolari, and Mikhail V. Sokolov. Computing currency invariant indices with an application to minimum variance currency baskets. *Journal of Economic Dynamics and Control*, 28(8):1481–1504, 2004.

[108] Ronald A. Howard and James E. Matheson. Influence diagrams. *Decision Analysis*, 2(3):127–143, 2005. URL `https://EconPapers.repec.org/RePEc:inm:ordeca:v:2:y:2005:i:3:p:127-143`.

[109] John C. Hull. *Risk management and financial institutions*. John Wiley & Sons, Hoboken, New Jersey, 2012.

[110] John C. Hull and Alan White. Incorporating volatility updating into the historical simulation method for value-at-risk. *Journal of Risk*, 1(1):5–19, 1998.

[111] Adrian Jäggi, Martin Schlegel, and Attilio Zanetti. Macroeconomic surprises, market environment, and safe-haven currencies. *Swiss Journal of Economics and Statistics*, 155(1):1–21, 2019.

[112] Ken Johnston and Elton Scott. The statistical distribution of daily exchange rate price changes: dependent vs independent models. *Journal of Financial and Strategic Decisions*, 12(2):39–49, 1999.

[113] Eric Jondeau and Michael Rockinger. The copula-garch model of conditional dependencies: An international stock market application. *Journal of International Money and Finance*, 25(5):827–853, 2006.

[114] Philippe Jorion. Risk2: Measuring the risk in value at risk. *Financial Analysts Journal*, 52(6):47–56, 1996.

[115] Philippe Jorion. *Value at risk: the new benchmark for managing financial risk*. The McGraw-Hill Companies, Inc., 2007.

[116] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

[117] Markus Kalisch and Peter Bühlmann. Estimating high-dimensional directed acyclic graphs with the pc-algorithm. *Journal of Machine Learning Research*, 8(22):613–636, 2007. URL http://jmlr.org/papers/v8/kalisch07a.html.

[118] Markus Kalisch, Martin Mächler, Diego Colombo, Marloes H. Maathuis, and Peter Bühlmann. Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47(11):1–26, 2012. doi: 10.18637/jss.v047.i11.

[119] G. Andrew Karolyi. Does international financial contagion really exist? *Journal of Applied Corporate Finance*, 16(2-3):136–146, 2004.

[120] Mansooreh Kazemi-Lari and Maman A. Djauhari. Analysis of a correlation network in world currency exchange market. *International Journal of Applied Mathematics and Statistics*, 44(14):202–209, 2013.

[121] Morgan Kelly and Cormac O'Grada. Market contagion: Evidence from the panics of 1854 and 1857. *American Economic Review*, 90(5):1110–1124, 2000.

[122] Mustafa Keskin, Bayram Deviren, and Yusuf Kocakaplan. Topology of the correlation networks among major currencies using hierarchical structure methods. *Physica A: Statistical Mechanics and its Applications*, 390(4): 719–730, 2011.

[123] Irina N. Khindanova and Svetlozar T. Rachev. Value at risk: Recent advances. *Handbook of Analytic Computational Methods in Applied Mathematics*, pages 801–858, 2019.

[124] Peter King and Heath Tarbert. Basel III: an overview. *Banking & financial services policy report*, 30(5):1–18, 2011.

[125] Candyce H. Kroenke, Romain Neugebauer, Jeffrey Meyerhardt, Carla M. Prado, Erin Weltzien, Marilyn L. Kwan, Jingjie Xiao, and Bette J. Caan. Analysis of body mass index and mortality in patients with colorectal cancer using causal diagrams. *JAMA Oncology*, 2(9):1137–1145, 2016.

[126] Kun Kuang, Lian Li, Zhi Geng, Lei Xu, Kun Zhang, Beishui Liao, Huaxin Huang, Peng Ding, Wang Miao, and Zhichao Jiang. Causal inference. *Engineering*, 6(3):253–263, 2020.

[127] Paul H. Kupiec. Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, 3(2):73–84, Winter 1995.

[128] Jarosław Kwapień, Sylwia Gworek, Stanisław Drożdż, and Andrzej Górski. Analysis of a network structure of the foreign currency exchange market. *Journal of Economic Interaction and Coordination*, 4(1):55, 2009.

[129] Luc Laeven and Fabian Valencia. Systemic banking crises revisited. IMF Working Paper 2018/206, International Monetary Fund, September 2018.

[130] Andrea Lancichinetti and Santo Fortunato. Community detection algorithms: a comparative analysis. *Physical Review E*, 80(5):056117, 2009.

[131] Steffen L. Lauritzen. Causal inference from graphical models. *Complex Stochastic Systems*, pages 63–107, 2001.

[132] Jing Lei and Larry Wasserman. Conformalized quantile regression. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1186–1194, 2014.

[133] Thomas J. Linsmeier and Neil D. Pearson. Risk measurement: an introduction to value at risk. ACE Reports 14796, University of Illinois at Urbana-Champaign, Department of Agricultural and Consumer Economics, 1996. URL https://EconPapers.repec.org/RePEc:ags:uiucar:14796.

[134] Thomas J. Linsmeier and Neil D. Pearson. Value at Risk. *Financial Analysts Journal*, 56(2):47–67, 2000.

[135] Gunter Löffler and Peter Raupach. Pitfalls in the use of systemic risk measures. *Journal of Financial and Quantitative Analysis*, 53(1):269–298, 2018. doi: 10.1017/S0022109017001041.

[136] Jacques Longerstaey, Peter Zangari, Martin Spencer, Christopher C. Finger, and Scott Howard. Riskmetrics™– technical document. Technical report, J.P.Morgan/Reuters, New York, NY, 1996. URL https://www.msci.com/documents/10199/5915b101-4206-4ba0-aee2-3449d5c7e95a.

[137] Marloes H. Maathuis, Markus Kalisch, and Peter Bühlmann. Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics*, 37(6A):3133 – 3164, 2009. doi: 10.1214/09-AOS685. URL https://doi.org/10.1214/09-AOS685.

[138] Donald MacKenzie and Taylor Spears. 'The formula that killed Wall Street': The gaussian copula and modelling practices in investment banking. *Social Studies of Science*, 44(3):393–417, 2014.

[139] Yannick Malevergne and Didier Sornette. Testing the gaussian copula hypothesis for financial assets dependences. *Quantitative finance*, 3(4):231, 2003.

[140] Rosario N. Mantegna. Hierarchical structure in financial markets. *The European Physical Journal B-Condensed Matter and Complex Systems*, 11 (1):193–197, 1999.

[141] Riccardo Marcaccioli and Giacomo Livan. A pólya urn approach to information filtering in complex networks. *Nature communications*, 10(1): 1–10, 2019.

[142] Carlos Martins-Filho, Feng Yao, and Maximo Torero. Nonparametric estimation of conditional value-at-risk and expected shortfall based on extreme value theory. *Econometric Theory*, 34(1):23–67, 2018.

[143] Roman Matkovskyy and Akanksha Jalan. From financial markets to bitcoin markets: A fresh look at the contagion effect. *Finance Research Letters*, 31:93–97, 2019. ISSN 1544-6123. doi: https://doi.org/10.1016/j.frl.2019.04.007. URL https://www.sciencedirect.com/science/article/pii/S1544612318308717.

[144] Mariusz Maziarz. A review of the granger-causality fallacy. *The Journal of Philosophical Economics: Reflections on Economic and Social Issues*, 8(2): 86–105, 2015.

[145] Michael McAleer and Bernardo Da Veiga. Forecasting value-at-risk with a parsimonious portfolio spillover garch (ps-garch) model. *Journal of Forecasting*, 27(1):1–19, 2008.

[146] Alexander J. McNeil and Rüdiger Frey. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3-4):271–300, 2000.

[147] Alexander J McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton University Press, 2015.

[148] Christopher Meek. Causal inference and causal explanation with background knowledge. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, UAI'95, page 403–410, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc. ISBN 1558603859.

[149] Michael Melvin and Mark P Taylor. The crisis in the foreign exchange market. *Journal of International Money and Finance*, 28(8):1317–1330, 2009.

[150] Lukas Menkhoff, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. Carry trades and global foreign exchange volatility. *The Journal of Finance*, 67(2):681–718, 2012.

[151] Norbert Metiu. Sovereign risk contagion in the eurozone. *Economics Letters*, 117(1):35–38, 2012.

[152] Peta Mitchell. *Contagious metaphor*. A&C Black, 2012.

[153] Abdolreza Mohammadi, Fentaw Abegaz, Edwin van den Heuvel, and Ernst C. Wit. Bayesian modelling of dupuytren disease by using gaussian copula graphical models. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 66(3):629–645, 2017.

[154] Alessandro Montalto, Luca Faes, and Daniele Marinazzo. Mute: a matlab toolbox to compare established and novel estimators of the multivariate transfer entropy. *PloS one*, 9(10):e109462, 2014.

[155] Mark Newman. *Networks*. Oxford university press, 2018.

[156] Erlend Nier, Jing Yang, Tanju Yorulmazer, and Amadeo Alentorn. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6):2033–2060, 2007.

[157] Cathy Ning, Dinghai Xu, and Tony S. Wirjanto. Is volatility clustering of asset returns asymmetric? *Journal of Banking & Finance*, 52:62–76, 2015. ISSN 0378-4266. doi: https://doi.org/10.1016/j.jbankfin.2014.11.016. URL https://www.sciencedirect.com/science/article/pii/S0378426614003781.

[158] OECD. COVID-19 spending helped to lift foreign aid to an all-time high in 2020 but more effort needed, 2021. URL https://www.oecd.org/newsroom/covid-19-spending-helped-to-lift-foreign-aid-to-an-all-time-high-in-2020-but-more-effort-needed.htm.

[159] Tatsuyoshi Okimoto. New evidence of asymmetric dependence structures in international equity markets. *Journal of Financial and Quantitative Analysis*, 43(3):787–815, 2008. doi: 10.1017/S0022109000004294.

[160] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco (CA), 1988.

[161] Judea Pearl. *Causality*. Cambridge University Press, 2009.

[162] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96 – 146, 2009. doi: 10.1214/09-SS057. URL https://doi.org/10.1214/09-SS057.

[163] Judea Pearl. The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3):54–60, 2019.

[164] Robert Peckham. Economies of contagion: financial crisis and pandemic. *Economy and Society*, 42(2):226–248, 2013.

[165] Mark Pengelly. Carry in favour? *Risk*, 22(7):40, 2009.

[166] M Hashem Pesaran and Andreas Pick. Econometric issues in the analysis of contagion. *Journal of Economic Dynamics and Control*, 31(4):1245–1277, 2007.

[167] Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):947–1012, 2016.

[168] Niklas Pfister, Peter Bühlmann, and Jonas Peters. Invariant causal prediction for sequential data. *Journal of the American Statistical Association*, 114(527):1264–1276, 2019.

[169] Mary Poovey. *Genres of the credit economy*. University of Chicago Press, 2008.

[170] Matthew Pritsker. The hidden dangers of historical simulation. *Journal of Banking & Finance*, 30(2):561–582, 2006.

[171] Yoshihiko Raita, Carlos A. Camargo Jr, Liming Liang, and Kohei Hasegawa. Big data, data science, and causal inference: A primer for clinicians. *Frontiers in Medicine*, 8:678047, 2021.

[172] James Ramage. Dollar rallies against euro, pound: U.S. currency rebounds after tumble it took wednesday following fed meeting. *The Wall Street Journal*, Mar 2015. URL https://www.wsj.com/articles/dollar-rallies-against-euro-pound-1426790239.

[173] Angelo Ranaldo and Paul Söderlind. Safe haven currencies. *Review of Rinance*, 14(3):385–407, 2010.

[174] Reuters. Timeline - china's reforms of the yuan exchange rate. *Thomson Reuters*, 2010. URL https://www.reuters.com/article/china-economy-yuan-idUSSGE65I02S20100619.

[175] Jonathan G Richens, Ciarán M Lee, and Saurabh Johri. Improving the accuracy of medical diagnosis with causal machine learning. *Nature Communications*, 11(1):1–9, 2020.

[176] Katerina Rigana, Ernst C. Wit, and Samantha Cook. Causal Network Contagion Value at Risk. *Available at SSRN*, 2022.

[177] Katerina Rigana, Ernst C. Wit, and Samantha Cook. A new way of measuring effects of financial crisis on contagion in currency markets. *International Review of Financial Analysis*, 90:102764, 2023. ISSN 1057-5219. doi: https://doi.org/10.1016/j.irfa.2023.102764. URL https://www.sciencedirect.com/science/article/pii/S1057521923002806.

[178] Roberto Rigobon. Contagion, spillover, and interdependence. *Economía*, 19(2):69–100, 2019.

[179] Juan Carlos Rodriguez. Measuring financial contagion: A copula approach. *Journal of Empirical Finance*, 14(3):401–423, 2007.

[180] Joseph P. Romano and Azeem M. Shaikh. Conformalized statistical inference for regression. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 33–41, 2013.

[181] Nouriel Roubini. Coronavirus pandemic has delivered the fastest, deepest economic shock in history. *The Guardian*, 25(March), 2020.

[182] Mark Rubinstein. Markowitz's "portfolio selection": A fifty-year retrospective. *The Journal of Finance*, 57(3):1041–1045, 2002. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2697771.

[183] Felix Salmon. Recipe for disaster: The formula that killed wall street. *Wired Magazine*, 17(3):17–03, 2009.

[184] Solomon Sarpong. Estimating the probability distribution of the exchange rate between ghana cedi and american dollar. *Journal of King Saud University-Science*, 31(2):177–183, 2019.

[185] Bernhard Schölkopf. Causality for machine learning. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 765–804. Association for Computing Machinery, New York, NY, United States, 2022.

[186] Thomas Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461, 2000.

[187] Marco Scutari, Catharina E. Graafland, and J. Manuel Gutiérrez. Who learns better bayesian network structures: Constraint-based, score-based or hybrid algorithms? In Václav Kratochvíl and Milan Studený, editors, *Proceedings of the Ninth International Conference on Probabilistic Graphical Models*, volume 72 of *Proceedings of Machine Learning Research*, pages 416–427. PMLR, 11–14 Sep 2018. URL https://proceedings.mlr.press/v72/scutari18a.html.

[188] M Ángeles Serrano, Marián Boguná, and Alessandro Vespignani. Extracting the multiscale backbone of complex weighted networks. *Proceedings of the National Academy of Sciences*, 106(16):6483–6488, 2009. doi: 10.1073/pnas.0808904106.

[189] Matteo Sesia and Emmanuel J Candès. A comparison of some conformal quantile regression methods. *Stat*, 9(1):e261, 2020.

[190] William R. Shadish, Thomas D. Cook, and Donald T. Campbell. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference.* Houghton, Mifflin and Company, 2002.

[191] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.

[192] Daniel D. Shapiro and Pavlos Msaouel. Causal diagram techniques for urologic oncology research. *Clinical Genitourinary Cancer*, 19(3):271–e1, 2021.

[193] Ali Shojaie and Emily B Fox. Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9:289–319, 2022.

[194] Julien Siebert. Applications of statistical causal inference in software engineering. *Information and Software Technology*, page 107198, 2023.

[195] Kimmo Soramäki and Samantha Cook. *Network theory and financial risk*. Risk Books, 2016.

[196] Kimmo Soramaki, SR Cook, and Alan Laubsch. A network-based method for visual identification of systemic risks. *Journal of Network Theory in Finance*, 2(1):67–101, 2016.

[197] Didier Sornette, Guilherme Demos, Qun Zhang, Qun Zhang, Peter Cauwels, Vladimir Filimonov, and Qunzh Zhang. Real-time prediction and post-mortem analysis of the shanghai 2015 stock market bubble and crash. *Swiss Finance Institute Research Paper*, 15(31), 2015. doi: 10.2139/ssrn.2693634.

[198] Peter Spirtes and Clark N. Glymour. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9(1):62–72, 1991.

[199] Peter Spirtes, Christopher Meek, and Thomas Richardson. Causal inference in the presence of latent variables and selection bias. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 499–506, 1995.

[200] Peter Spirtes, Clark N. Glymour, Richard Scheines, and David Heckerman. *Causation, prediction, and search*. MIT press, 2000.

[201] Stavros Stavroglou, Athanasios Pantelous, Kimmo Soramaki, and Konstantin Zuev. Causality networks of financial assets. *The Journal of Network Theory in Finance*, 3(2):17–67, 2017.

[202] Joseph E. Stiglitz. Contagion, liberalization, and the optimal structure of globalization. *Journal of Globalization and Development*, 1(2), 2010. doi: doi:10.2202/1948-1837.1149.

[203] Patrick A. Stokes and Patrick L. Purdon. A study of problems encountered in granger causality analysis from a neuroscience perspective. *Proceedings of the national academy of sciences*, 114(34):E7063–E7072, 2017.

[204] Paola Subacchi. *The people's money: How China is building a global currency*. Columbia University Press, 2016.

[205] George Sugihara, Robert May, Hao Ye, Chih-hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting causality in complex ecosystems. *science*, 338(6106):496–500, 2012. doi: 10.1126/science.1227079.

[206] Łukasz Sułkowski. COVID-19 pandemic; recession, virtual revolution leading to de-globalization? *Journal of Intercultural Management*, 12(1): 1–11, 2020.

[207] Rui Sun. Reevaluating the effect of tax and expenditure limitations: An instrumental variable approach. *Public Finance Review*, 42(1):92–116, 2014.

[208] James W. Taylor. Forecasting value at risk and expected shortfall using a model with a dynamic omega ratio. *Journal of Banking & Finance*, 140: 106519, 2022. ISSN 0378-4266. doi: https://doi.org/10.1016/j.jbankfin.2022.106519. URL https://www.sciencedirect.com/science/article/pii/S0378426622001133.

[209] Eric J. Tchetgen Tchetgen and Tyler J. VanderWeele. On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21 (1):55–75, 2012.

[210] Jens Thomsen. Denmark and the euro: A special relationship. *BIS Review*, 22, 2003.

[211] Ryan J. Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32, 2019.

[212] Adrian Tobias and Markus K. Brunnermeier. CoVaR. *The American Economic Review*, 106(7):1705, 2016.

[213] Michele Tumminello, Tomaso Aste, Tiziana Di Matteo, and Rosario N Mantegna. A tool for filtering information in complex systems. *Proceedings of the National Academy of Sciences*, 102(30):10421–10426, 2005.

[214] Caroline Van Rijckeghem and Beatrice Weder. Sources of contagion: is it finance or trade? *Journal of International Economics*, 54(2):293–308, 2001.

[215] Gang-Jin Wang, Chi Xie, Feng Han, and Bo Sun. Similarity measure and topology evolution of foreign exchange markets using dynamic time warping method: Evidence from minimal spanning tree. *Physica A: Statistical Mechanics and its Applications*, 391(16):4136–4146, 2012.

[216] Gang-Jin Wang, Chi Xie, Peng Zhang, Feng Han, and Shou Chen. Dynamics of foreign exchange networks: a time-varying copula approach. *Discrete Dynamics in Nature and Society*, 2014, 2014.

[217] Graham W. Warren, Anthony J. Alberg, Andrew S. Kraft, and K. Michael Cummings. The 2014 surgeon general's report: "the health consequences of smoking–50 years of progress": A paradigm shift in cancer care. *Cancer*, 120(13):1914–1916, 2014. doi: https://doi.org/10.1002/cncr.28695. URL https://acsjournals.onlinelibrary.wiley.com/doi/abs/10.1002/cncr.28695.

[218] Xiaoqian Wen, Yu Wei, and Dengshi Huang. Measuring contagion between energy market and stock market during financial crisis: A copula approach. *Energy Economics*, 34(5):1435–1446, 2012.

[219] Elizabeth J. Williamson, Zoe Aitken, Jock Lawrie, Shyamali C. Dharmage, John A Burgess, and Andrew B Forbes. Introduction to causal diagrams for confounder selection. *Respirology*, 19(3):303–311, 2014.

[220] Jeffrey M Wooldridge. Should instrumental variables be used as matching variables? *Research in Economics*, 70(2):232–237, 2016.

[221] Philip D Wooldridge. FX and OTC derivatives markets through the lens of the triennial survey. *BIS Quarterly Review, December*, pages 15–19, 2019.

[222] Chen Xu and Yao Xie. Conformal prediction set for time-series. *arXiv preprint arXiv:2206.07851*, 2022.

[223] Yasuhiro Yamai and Toshinao Yoshiba. Value-at-risk versus expected shortfall: A practical perspective. *Journal of Banking & Finance*, 29(4):997–1015, 2005.

[224] Yasuhiro Yamai and Toshinao Yoshiba. Value-at-risk versus expected shortfall: A practical perspective. *Journal of Banking & Finance*, 29(4):997–1015, 2005. ISSN 0378-4266. doi: https://doi.org/10.1016/j.jbankfin.2004.08.010. URL https://www.sciencedirect.com/science/article/pii/S0378426604001499. Risk Measurement.

[225] Brian S Yandell. *Practical data analysis for designed experiments*. Routledge, 1997.

[226] Ziyun Zhang, Su Chen, and Bo Li. Does previous carry trade position affect following investors' decision-making and carry returns? *International Review of Financial Analysis*, 80:102056, 2022.

[227] Zhao Zhao, Huwei Wen, and Ke Li. Identifying bubbles and the contagion effect between oil and stock markets: New evidence from china. *Economic Modelling*, 94:780–788, 2021. ISSN 0264-9993. doi: https://doi.org/10.1016/j.econmod.2020.02.018. URL https://www.sciencedirect.com/science/article/pii/S026499931931819X.