

The Evolving Causal Structure of Equity Risk Factors

Gabriele D’Acunto

gabriele.dacunto@isi.it

Sapienza University of Rome, Italy

ISI Foundation, Turin, Italy

Francesco Bonchi

francesco.bonchi@isi.it

ISI Foundation, Turin, Italy

Eurecat, Barcelona, Spain

Paolo Bajardi

paolo.bajardi@isi.it

ISI Foundation

Turin, Italy

Gianmarco De Francisci Morales

gdfm@isi.it

ISI Foundation

Turin, Italy

ABSTRACT

In recent years, *multi-factor strategies* have gained increasing popularity in the financial industry, as they allow investors to have a better understanding of the risk drivers underlying their portfolios. Moreover, such strategies promise to promote diversification and thus limit losses in times of financial turmoil. However, recent studies have reported a significant level of redundancy between these factors, which might enhance risk contagion among multi-factor portfolios during financial crises. Therefore, it is of fundamental importance to better understand the relationships among factors.

Empowered by recent advances in *causal structure learning* methods, this paper presents a study of the causal structure of financial risk factors and its evolution over time. In particular, the data we analyze covers 11 risk factors concerning the US equity market, spanning a period of 29 years at daily frequency.

Our results show a statistically significant sparsifying trend of the underlying causal structure. However, this trend breaks down during periods of financial stress, in which we can observe a densification of the causal network driven by a growth of the out-degree of the market factor node. Finally, we present a comparison with the analysis of factors cross-correlations, which further confirms the importance of causal analysis for gaining deeper insights in the dynamics of the factor system, particularly during economic downturns.

Our findings are especially significant from a risk-management perspective. They link the evolution of the causal structure of equity risk factors with market volatility and a worsening macroeconomic environment, and show that, in times of financial crisis, exposure to different factors boils down to exposure to the market risk factor.

CCS CONCEPTS

• **Mathematics of computing** → *Causal networks*; Time series analysis; • **Applied computing** → *Economics*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICAF’21, November 3–5, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9148-1/21/11...\$15.00

<https://doi.org/10.1145/3490354.3494370>

KEYWORDS

Causal Discovery, Structure Learning, Networks Dynamics, Risk Premia

ACM Reference Format:

Gabriele D’Acunto, Paolo Bajardi, Francesco Bonchi, and Gianmarco De Francisci Morales. 2021. The Evolving Causal Structure of Equity Risk Factors. In *2nd ACM International Conference on AI in Finance (ICAF’21)*, November 3–5, 2021, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3490354.3494370>

1 INTRODUCTION

Multi-factor investing strategies have gained wide adoption during the last decade, as they allow investors to have a better understanding of the risk drivers underlying a portfolio. Such strategies promise to promote diversification and thus limit drawdown during financial turmoils [32, 34]. However, out of hundreds existing factors, only a small number is truly significant in explaining the cross-section of stock returns [19, 21]. Therefore, an important open question across both financial research and industry concerns factors redundancy. In particular, the adoption by researchers of different testing approaches (e.g., panel vs. cross-sectional regression), and the presence of model selection biases as those entailed by omitted variables, may contribute to the publication of new redundant factors. Motivated by these observations, we aim at gaining insights into the underlying dynamics of risk factor interactions by leveraging recent advances in *causal structure learning* [48]. More precisely, starting from the known results about correlations among factors [19, 21], in the spirit of Reichenbach’s *principle of common cause* [41], we investigate whether there is an underlying causal structure within the universe of considered financial factors, and how this structure evolves over time.

Specifically, our computational task is as follows: we are given as input a dataset $\mathcal{Y} \in \mathbb{R}^{N \times T}$ composed by N time series of length T of factors returns, and we want to learn a graph (as the one in Figure 1) representing the causal relations among the factors values along time. More in details we are interested in understanding whether \mathcal{Y} admits a *functional representation* in which each y_t^i (i.e., the value of the factor y^i at timestamp t) is determined by at most L past values of a set of other factors named *parents*:

$$y_t^i = f^i \left((\mathcal{P}_L^i)_{t-L}, \dots, (\mathcal{P}_0^i)_t, \epsilon_t^i \right), \quad i \in \{1, \dots, N\}. \quad (1)$$

Here ϵ_t^i represents the noise term, while $(\mathcal{P}_l^i)_{t-l}$ indicates the set of risk factors that cause y_t^i with lag $l \in \mathbb{N}_0$. When $l > 0$, $(\mathcal{P}_l^i)_{t-l}$

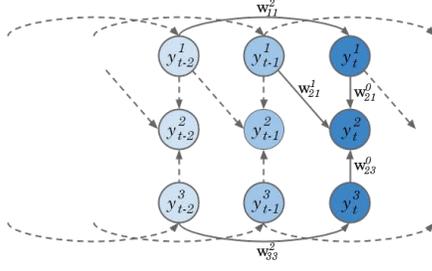


Figure 1: An example of the causal graph among equity risk factors along time with maximum lag $L = 2$.

may contain y_{t-l}^i as well. Conversely, as we cannot observe causal effects from present to past, $(\mathcal{P}_0^i)_t$ can neither contain y_t^i nor y_t^j if $y_t^i \in (\mathcal{P}_0^j)_t$, otherwise it would be impossible to define the direction of the causal relation. More in general, we require the set of Equations (1) to be *acyclic*, which means that feedback loops among variables are forbidden. This assumption plays a key role in causal inference since it allows to set causes apart from effects.

In case the aforementioned functional forms are assumed to be linear, by using matrix notation, Equations (1) can be compactly written as a *structural vector autoregressive model* (SVAR)

$$\mathbf{y}_t = \sum_{l=0}^L \mathbf{W}^l \mathbf{y}_{t-l} + \boldsymbol{\epsilon}_t, \quad (2)$$

where $\mathbf{y}_t \in \mathbb{R}^N$ is a column vector constituted by the observations at time t , $\mathbf{W}^l \in \mathbb{R}^{N \times N}$ are the matrices of lagged causal effects with lag l up to maximum lag L , such that $w_{ij}^l \neq 0$ iff $y_{t-l}^j \in (\mathcal{P}_l^i)_{t-l}$. In addition, the matrix of instantaneous causal effects \mathbf{W}^0 respects the acyclicity requirement mentioned above. Finally, $\boldsymbol{\epsilon}_t \in \mathbb{R}^N$ is the column vector of random disturbances at time t .

We remark that Equation (2) should not be read as a usual system of equations, but rather as a set of functions describing how certain factors determine others. Indeed, the model is said to be *structural* since it allows to compute variables (effects) by means of linear functions of other endogenous variables (causes), by taking into account both instantaneous and lagged relations (also referred to as inter-layer connections). In particular, the previous model can be thought of as a combination of a *structural equation model* (SEM [38]) and a *vector autoregressive model* (VAR [46]).

Finally, note that Equation (2) entails a *directed acyclic graph* (DAG) in which there is a weighted edge from $y_{t-l}^j \in (\mathcal{P}_l^i)_{t-l}$ (with $l \geq 0$) to y_t^i : this is the causal graph of factors along time as depicted in Figure 1. In order to study the evolution of such non-stationary system, we estimate Equation (2) by adopting a sliding window approach and performing a regression analysis on the inferred causal networks. Our data covers 11 risk factors concerning the US equity market, over a period of 29 years at daily frequency. Our main results can be summarized as follows:

- Causal interactions between factors exhibit a statistically significant sparsification trend along time, with anomalies in periods of financial turmoil.

- We expose a relationship between the density of the causal networks and investor sentiment. In particular, we show that periods of worsening sentiment and business cycle phase (see Section 3.2) are associated with a densification of the causal network. This phenomenon is driven by a growth in the out-degree of the market risk factor.
- Finally, we conduct a comparative study between causation and correlation among factors. Our findings highlight how causal analysis better describes the importance of the market factor among the considered risk factors. Besides, while according to correlation analysis, the business cycle indicator is not related to the evolution of the factorial system, the study of causal structures reveals a statistically significant relationship with a 95% confidence level.

The rest of the paper is organized as follows. Section 2 provides the background about factor investing and causal discovery. Section 3 describes the data used to conduct the analysis and presents our methodology. Section 4 reports the results of our analysis. Finally, Section 5, provides a discussion on our results together with some direction for future investigation.

2 RELATED WORK

An equity risk factor is a variable able to explain the cross-section of expected stock returns, i.e., how the expected return varies among stocks. Its significance is usually assessed via the usage of linear regression models [12]. Factor investing is well rooted in finance, in particular it dates back to the first asset pricing model, the *Capital Asset Pricing Model* (CAPM), introduced by Sharpe [43]. CAPM looks at stock returns through the exposure to one factor, the *market beta*, and introduces a precise definition of risk and how it drives expected stock returns. Successively, Fama and French [17] found out that the exposure to small cap stocks (SMB) and cheaper, under-priced stocks (HML) provide two additional sources of risk not captured by CAPM. Based on these evidences, an enormous amount of research has been produced on factor investing, and hundreds of potential factors have been proposed [13, 22, 26, 35]. This abundance of proposed factors has opened an important question across both financial research and industry concerning their redundancy. As already mentioned, Feng et al. [19] and Harvey and Liu [21] recently pointed out that only a small number of the existing factors is truly significant in explaining stock returns cross-section. Our study of risk factor interactions fits into this stream of research, and in particular tackles the problem by exploiting recent advances in causal learning [37].

The possibility of examining a system under interventions (i.e., counterfactual analysis) makes causal analysis a powerful tool for studying complex systems [39]. However, in many cases, the underlying causal structure is unknown, and it is not possible to carry out randomized experiments in order to study the system at hand under distribution changes. Therefore, the interest in inferring causal structures from observational data, also known as *causal structure learning*, has been significantly growing during recent years.

Causal structure learning algorithms can be classified into three main families: (i) *constraint-based approaches*, which make use of conditional independence tests to establish the presence of an edge between two nodes [29, 47]; (ii) *score-based methods*, which use

several search procedures in order to optimize a certain score function [10, 23, 28]; (iii) *structural causal models*, which express a variable at a certain node as a function of its parents [8, 27, 40, 44, 45]. Additionally, as highlighted in Section 1, whenever we are dealing with a process that evolves over time, temporal ordering drives the causal inference procedure. Nevertheless, the main issue concerns the estimation of instantaneous effects, which must satisfy the acyclicity requirement. With regards to Equation (2), this means that W^0 must entail the structure of a DAG.

However, it is common to deal with non-Gaussian data. For instance, this is the case of equity time series which show *heteroscedasticity* (i.e., the variance of the stock returns varies over time) and *volatility clustering* (i.e., large (small) swings in stock prices tend to group together) [6].

This observation implies that there is additional information, not described by the covariance matrix, that can be exploited to retrieve W^0 . Consequently, by leveraging a non-gaussianity assumption of ϵ_t , a series of linear non-Gaussian methods to estimate the model in Equation (2) have been proposed in the past years [31, 36]. In particular, in this study we employ VAR-LiNGAM [31], as provided by the python package `lingam`¹ made available by authors.

Econophysics, is an interdisciplinary effort devoted at analyzing the risk contagion among financial institutions by representing the financial system as a network [4]. Our analysis differs from *Econophysics* as we do not focus on financial institutions, and instead aim to understand the dynamics of the factorial network, which captures sources of risk largely accepted and widely used by investors. In addition, we look for causal relationships between the observed variables by means of a machine learning causal model that, differently from existing work [5], allows us to evaluate the presence of instantaneous causal effects as well.

3 DATA AND METHODOLOGY

In this section, we first introduce the financial risk factors and provide some information about the factor dataset \mathcal{Y} . Then, we focus on the additional variables selected as indicators of changes in both investor expectation and business cycle evolution. Finally, we present the adopted causal inference approach and introduce the *generalized linear model* (GLM) regression we employ to analyze the model results.

3.1 Financial factors

We consider 11 risk factors at daily frequency concerning the US equity market. The choice to use equity factors from the latter market is driven by the greater availability of data and the higher presence of results in the financial literature that can be compared to those we obtain. More precisely, we deal with 7306 daily observations spanning 29 years, from 2 January 1991 to 31 December 2019. We include in our analysis the following published risk factors, gathered directly from the websites of authors: *Excess Market Return* (Mkt-RF) [33], *Small Minus Big* (SMB) and *High Minus Low* (HML) [17], *Momentum* (UMD) [9], *HML Devil* (HML Dev) [2], *Robust Minus Weak* (RMW) and *Conservative Minus Aggressive* (CMA) [18], *HXZ Investment* (R-IA) and *HXZ Profitability*

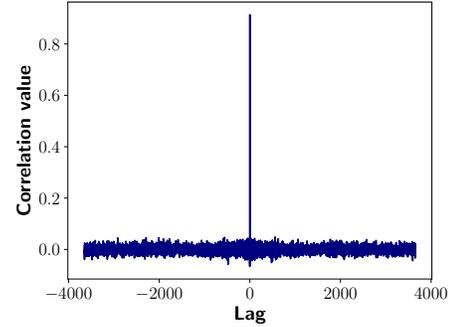


Figure 2: Cross-correlation function between CMA (Conservative Minus Aggressive) and R-IA (HXZ Investment).

(R-ROE) [25], *Betting Against Beta* (BAB) and *Quality Minus Junk* (QMJ) [3].

Summary statistics are provided in Table 1. All factors display positive annualised average compounded returns over the considered period. Moreover, the volatility value significantly varies among factors. In particular, Mkt-RF shows the highest value whereas the investment attitude relating factors (CMA and R-IA) display the lowest ones. Overall, according to risk-adjusted returns, BAB and R-ROE factors turn out to be the highest-performing risk premia. The previous observation is also supported by the Sortino ratio.² We also report the release date of each factor, and note that the majority of the factors were published after 2010.

Finally, by inspecting the cross-correlation function (CCF), we observe significant values, especially within factors which aim to capture the same anomaly in stock returns. Moreover, cross-correlation tends to be higher at lag $l = 0$ and then to significantly drop in almost all cases. As an example, Figure 2 depicts the CCF for the pair of previously-mentioned factors, CMA and R-IA. In particular, the CCF shows a peak at lag $l = 0$ of about 0.90 and then plummets for higher order lags. Thus, previous observations concerning the CCF behaviour suggest that there is an important ongoing associative dynamic among factors.

3.2 Fear and business cycle indicators

In order to relate the evolution of risk factor interactions to both stock market volatility and macroeconomic downturn, in addition to factors data, we consider two indicators defined on the VIX Index and the yield spread, respectively. The former index measures 30-day-ahead investors expectation of US equity market volatility. In particular, the VIX Index is widely referred as the *fear index*, since it is an indicator of market stress and financial turmoil. The latter spread is a macroeconomic indicator which is largely used to predict recessions [16]. In particular a growth in the difference between 3-month and 10-year US Treasury yield is linked to a worsening of the macroeconomic environment [15]. We obtain data concerning these indexes from the CBOE and FRED repositories, respectively.

¹<https://github.com/cdt15/lingam>

²A metric to evaluate the risk-adjusted performance of a portfolio discounting for its downside standard deviation.

Table 1: Summary statistics for the analyzed factors at daily frequency. Average compounded return, volatility, risk adjusted return, and Sortino Ratio are annualised.

	Mkt-RF	SMB	HML	RMW	CMA	R-IA	R-ROE	BAB	HML-dev	UMD	QMJ
Avg Comp. Ret. (%)	7.68	0.67	2.10	3.89	2.14	2.23	5.35	9.43	0.65	5.17	4.46
Volatility (%)	17.46	9.17	9.61	7.29	6.53	6.60	7.43	11.01	10.46	13.39	7.97
Risk Adj. Ret. (%)	0.44	0.07	0.22	0.53	0.33	0.34	0.72	0.86	0.06	0.39	0.56
Sortino Ratio (%)	0.62	0.10	0.32	0.79	0.47	0.48	1.04	1.21	0.09	0.53	0.83
Skew	-0.15	-0.22	0.43	0.26	-0.44	-0.72	-0.11	-0.34	0.52	-0.23	0.20
Kurtosis	8.34	3.91	9.05	7.71	11.44	15.92	5.84	11.72	11.47	24.57	8.70
1st %-ile (%)	-2.98	-1.46	-1.67	-1.28	-1.10	-1.08	-1.41	-2.20	-1.73	-2.52	-1.30
5th %-ile (%)	-1.72	-0.91	-0.83	-0.66	-0.57	-0.57	-0.68	-0.97	-0.88	-1.20	-0.71
Min	-8.95	-4.71	-4.39	-3.02	-5.94	-6.88	-3.96	-6.29	-7.00	-9.46	-3.74
Max	11.35	3.78	4.83	4.49	2.53	2.75	3.26	7.94	6.35	14.53	5.04
Publication year	1972	1993	1993	2015	2015	2014	2014	2014	2013	1997	2013

Starting from these two indexes, we build *fear* and *business cycle* z-scores as follows. We first define the ΔVIX *historical expected short-fall* for the k -th inference sample period S_k : $\Delta\text{VIX-ES}_k = \mathbb{E}[\mathbf{V}_k^{95}]$, defined over $\mathbf{V}_k^{95} = \{v_i \in S_k | v_i > v_k^{95}\}$, where v_k^{95} is the 95-percentile value of the percent change between VIX Index closing and opening daily values. By computing $\Delta\text{VIX-ES}_k$ we quantify the extreme values of the daily volatility swing over the inference period S_k . Next, we measure the extent to which such a value is unusual with respect to past observations. Therefore, we define:

$$\text{f-zscore} = \frac{\Delta\text{VIX-ES}_k - \mu_k^{10Y}}{\sigma_k^{10Y}}$$

with μ_k^{10Y} and σ_k^{10Y} being the 10-year rolling average and standard deviation of $\Delta\text{VIX-ES}_k$ respectively.

For the *business cycle* z-score, we first evaluate the extreme values of the 3M10Y yield spread for sample S_k , i.e., $3\text{M10Y-ES}_k = \mathbb{E}[\mathbf{B}_k^{95}]$, defined over $\mathbf{B}_k^{95} = \{b_i \in S_k | b_i > b_k^{95}\}$, where, b_k^{95} is the 95-percentile value of the difference between the 3-month and 10-year US Treasury daily rates. Thus, we define the z-score as:

$$\text{bc-zscore} = \frac{\Delta 3\text{M10Y-ES}_k - \mu_k^{10Y}}{\sigma_k^{10Y}}$$

where μ_k^{10Y} and σ_k^{10Y} represent the 10-year rolling average and standard deviation of $\Delta 3\text{M10Y-ES}_k$ respectively.

3.3 Causal inference procedure and regression model

As mentioned earlier, to study the dynamic of the causal structure along time, we adopt a sliding window approach. More precisely, we divide the overall analysis period into windows S_k of length 18 months each, with a sliding step of 3 months, obtaining 111 inference periods S_k .

As described in Section 2, we apply the VAR-LiNGAM algorithm to infer the causal model. In particular, the algorithm first fits a VAR model on the data, and then estimates on the regression residuals a linear non-Gaussian causal inference method, the DirectLiNGAM [45]. Other existing linear non-Gaussian approaches leverage *independent component analysis* (ICA [30]) to estimate

the matrix of instantaneous causal effect \mathbf{W}^0 . The DirectLiNGAM model was proposed to solve the potential convergence issues of ICA-based methods [24]. After the fit of VAR model on data, suppose to regress the residuals associated with factor j on those of factor i , $\forall j \in \{1, \dots, N\} | j \neq i$. Then, the residual z_i is exogenous to the system if it is independent of the regression residual $r_j^i = z_j - (\text{cov}(z_j, z_i) / \text{var}(z_i))z_i$. The algorithm starts with an empty causal ordering set \mathcal{O} and, iteratively, appends the variable which is the most independent of its residual. The procedure stops when $N - 1$ insertions have been made.

In each sample period, we apply the model by selecting the number of lags according to the BIC criterion [42]: the resulting maximum lag L is equal to 1 for every sample. Subsequently, we validate the estimated causal coefficients by running a permutation test, i.e., resampling with replacement, with a significance level equal to 95%. In addition, the total number of permuted samples per period is 5000, the length of the generated samples equals that of the inference period (18 months) and we do not apply any thresholding to the resulting significant coefficients. Since we are interested in comparing the information provided by causal inference with that coming from correlation analysis, the same methodology is used to estimate correlation networks, by replacing the estimation of Equation (2) with the Pearson correlation coefficient.

Once we retrieve both causal and correlation network structures, we analyse their temporal evolution by means of a regression analysis. We set as dependent variable d the number of network edges and as covariates the following three variables: time (measured in days), f-zscore, and bc-zscore. Moreover, since $d \in \mathbb{N}_0$, we employ a *Poisson log-linear model* specified by the following GLM [1]:

$$\log(d) = \beta_0 + \sum_{i=1}^3 \beta_i \cdot f_i, \quad (3)$$

from which we have $d = e^{\beta_0} \cdot \prod_{i=1}^3 e^{\beta_i \cdot f_i}$, where f_i are the regressors mentioned before. Therefore, according to Equation (3), a unit increase in the independent variable f_i is associated with a multiplicative effect e^{β_i} on d . As a consequence, if $\beta_i = 0$, the growth of f_i does not affect that of d . Furthermore, if $\beta_i > 0$ then d increases as f_i grows, and conversely if $\beta_i < 0$ it decreases.

4 RESULTS

In this section we provide the results of the analysis. The inferred networks for causal and correlation structures are shown in Figure 3. For readability, only three different inference samples are shown: (i) before the publication of the majority of the factor models; (ii) during the Global Financial Crisis (GFC), located between 2007 and 2009; (iii) after the publication of all factor models.

The networks are constituted by 22 nodes, split in accordance to time ordering. The factors are sorted vertically according to their publication date, from the oldest to the newest. In addition, edges associated with a positive weight are shown in grey, while those with a negative one are given in red. Thicker lines indicate a higher weight of the edge, and thus a stronger association of the factors. Finally, in order to make figures easier to read, throughout the paper results relating to correlation are shown in purple and those relating to causation in blue.

We notice a significant variability in both correlation and causal network structures across the periods. As time goes by, the number of edges decreases and the networks become sparser. Such phenomenon is more pronounced for inter-layer relationships. Figure 3 highlights the key role of market factor, which impacts 9 factors out of 11 during the GFC.

As far as the stability of the inferred networks is concerned, Figure 4 shows the evolution over time of Jaccard score between two networks estimated on adjacent periods. Such a metric is used in order to quantify the matching ratio between the edge sets \mathcal{E}_{k-1} and \mathcal{E}_k of two consecutive structures. In particular, it is defined as:

$$Jacc(\mathcal{E}_{k-1}, \mathcal{E}_k) = \frac{|\mathcal{E}_{k-1} \cap \mathcal{E}_k|}{|\mathcal{E}_{k-1} \cup \mathcal{E}_k|}$$

and it holds $Jacc(\mathcal{E}_{k-1}, \mathcal{E}_k) \in [0, 1]$. In order to smooth the score, we apply a 1-year rolling average filter: as we use a step size of 3 months, the average of the last four values is shown. The Jaccard score is already normalized and is not affected by the size of the sets, however when the numbers are very small it could be misleading. To better interpret the plot, the cardinality of the edge set $|\mathcal{E}_k|$ is given on the right y-axis of each chart. When the correlation networks are considered, the aforementioned score is more stable over time and is much higher than the one returned by causal structures. In both cases, we observe the presence of a sparsifying trend over time. Interestingly, for causal networks, this trend breaks down during the GFC, and the Jaccard score sharply increases.

We further analyze the temporal trend of network density by means of a regression analysis. In particular, we relate the number of edges to time, fear, and business cycle indicators through the estimation of Equation (3). Results related to correlation structures are reported in Table 2.

Considering the overall relations, both time and f-zscore are statistically significant at 99% level and are related to the sparsification of the network. It is worth noticing that time is measured in days, and thus the magnitude of the estimated coefficient is expected to be very small. Relating the value of the coefficient to the length of the considered time window, we obtain that every 18 months the number of edges in the correlation structure is reduced by approximately 4%. With regard to investors future expectation, a growth of one standard deviation in f-zscore is associated with a reduction of almost 4% in the number of edges as well. On the contrary, the

Table 2: Results of the GLM regression of the number of edges in correlation networks. Bold coefficients are statistically significant at 99% level.

Relation type	variable	coef.	std. err.	p-value
Overall	intercept	4.7419	0.024	0.000
	time	-7.054e-05	4.3e-06	0.000
	f-zscore	-0.0402	0.009	0.000
	bc-zscore	-0.0039	0.010	0.687
Instantaneous	intercept	3.7993	0.035	0.000
	time	4.07e-06	5.72e-06	0.477
	f-zscore	-0.0340	0.012	0.004
	bc-zscore	-0.0038	0.012	0.761
Lagged	intercept	4.3698	0.034	0.000
	time	-0.0002	6.76e-06	0.000
	f-zscore	-0.0718	0.014	0.000
	bc-zscore	-0.0229	0.015	0.131

Table 3: Results of the GLM regression of the number of edges in causal networks. Bold coefficients are statistically significant at 99% level.

Relation type	variable	coef.	std. err.	p-value
Overall	intercept	2.8670	0.067	0.000
	time	-0.0002	1.29e-05	0.000
	f-zscore	0.1732	0.027	0.000
	bc-zscore	0.0782	0.030	0.010
Instantaneous	intercept	-10.5020	4.188	0.012
	time	0.0008	0.000	0.102
	f-zscore	1.0139	0.357	0.005
	bc-zscore	-0.4617	0.415	0.266
Lagged	intercept	2.8790	0.067	0.000
	time	-0.0002	1.29e-05	0.000
	f-zscore	0.1640	0.027	0.000
	bc-zscore	0.0727	0.030	0.017

bc-zscore is never statistically significant. Thus, the results suggest that the correlation structure does not show a link with changes in macroeconomic conditions.

In addition to overall relations, we analyse instantaneous and lagged interactions separately. In particular, while the f-zscore remains significant in both cases, time is only relevant for lagged relations. Regression results and Figure 5a (in which we provide the fit of the observational data) show that, even though we observe a slight decrease during stress periods, the level of instantaneous association is pretty stable over time. Thus, the overall statistical significance of time is due to the disappearance of inter-layer links.

As far as causal networks are concerned, Table 3 provides the results of the regression analysis. By focusing on the overall relations, we see that both time and fear index are statistically significant at 99% level. Here, the network sparsification is faster, i.e., every 18 months the total number of connections decreases by about 11%. On the contrary, an increase of one standard deviation in the fear index is associated with a growth of almost 19% in the number of edges. This finding explains the break in the sparsifying trend during periods of financial stress shown in Figure 4. Regarding the business cycle indicator, it is significant at 95% level, with the sign of the regression coefficient in accordance with that shown by the fear index. Therefore, the analysis of causal structures shows that a worsening in the macroeconomic environment is also linked to an increase of causal relations among factors. However, the strength of the effect is much smaller than that of the fear index.

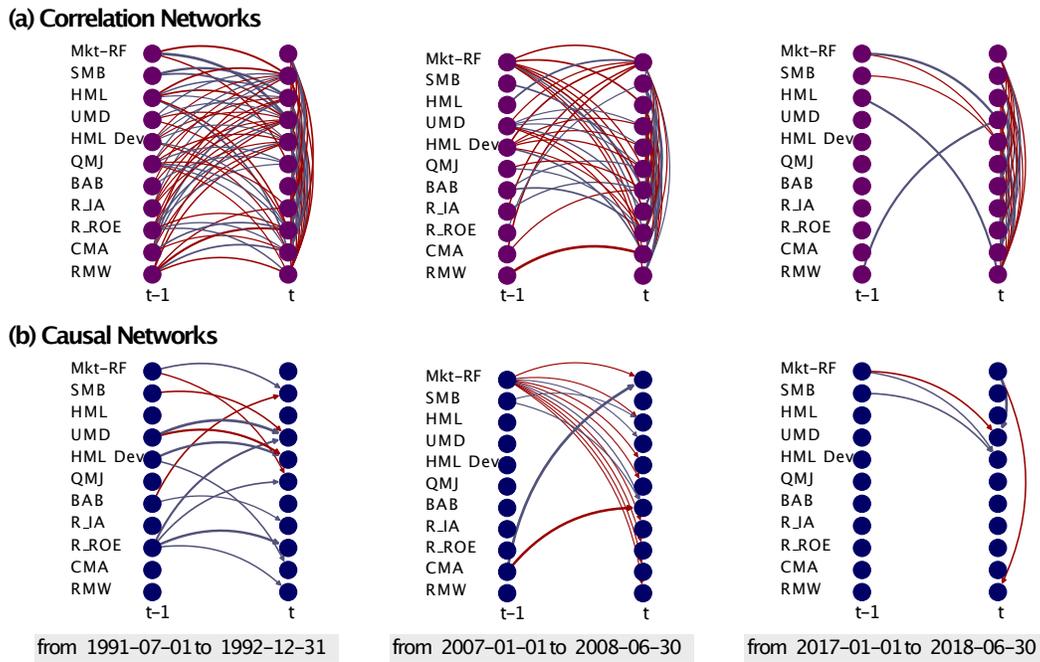


Figure 3: Correlation (a) and causal (b) structures inferred over three different periods. Factors are sorted in ascending order from the oldest to the newest. Grey edges represent positive weights, red edges negative ones.

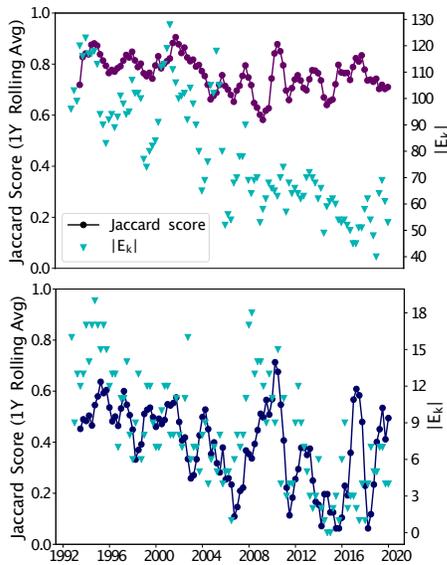


Figure 4: Behaviour over time of 1-year rolling average Jaccard score between consecutive correlation networks (top) and causal networks (bottom). The right axis shows the number of network edges.

By analysing instantaneous relations apart from lagged ones, we notice that the statistical significance of the fear index is preserved, whereas time is only relevant for lagged connections. The estimated

Table 4: Results of the GLM regression of the out-degree of the market risk factor in causal networks.

variable	coef.	std. err.	p-value
intercept	1.0963	0.144	0.000
time	-0.0001	2.57e-05	0.000
f-zscore	0.4382	0.058	0.000
bc-zscore	-0.0790	0.063	0.210

coefficient indicates that the network of lagged interactions thins out with the same intensity as that of all the connections. Such a phenomenon is illustrated in Figure 5b: the trend of the causal structure mainly consists of lagged causal interactions while instantaneous effects only appear in recent years. Interestingly, the temporal trend of network sparsification is similar to the one shown above for lagged connections in correlation networks.

Finally, we focus on the role of the market risk factor node within the network of causal relations among risk factors. As shown in Figure 6, the out-degree of the corresponding node remains in the range of 2 most of the time, then dramatically increases during periods of crisis, as observed during the GFC and more recently during 2018. In the latter example the stock market suffered heavy losses: first a volatility shock occurred in early February; subsequently the market plummeted in the last quarter due to both US-China trade war and the slowdown in economic growth.

In order to investigate the evolution over time of market node out-degree, Table 4 reports the results of a regression analysis. We set as a dependent variable the aforementioned out-degree and as covariates time, f-zscore, and bc-zscore. We use again the Poisson

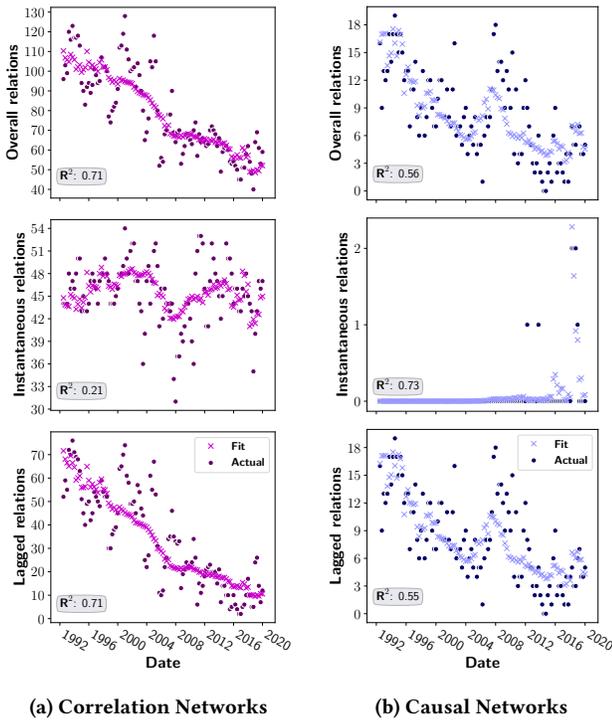


Figure 5: Evolution of the estimated number of significant edges in both (a) correlation and (b) causal networks.

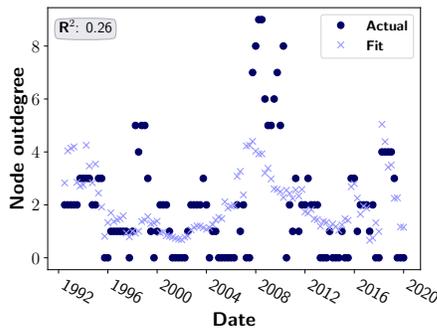


Figure 6: Evolution over time (and GLM fit) of the out-degree of the market risk factor in causal networks.

log-linear model described by Equation (3). The association with the fear indicator is statistically significant at 99% level: an increase of one standard deviation in the latter is linked to a strong growth of almost 55% in the node out-degree. Furthermore, time turns out to be a significant feature, which is consistent with the already underlined presence of the sparsifying trend in causal relations.

5 DISCUSSION

The results shown in the present work highlight the continuously-changing nature of risk factors interactions. Taking such behaviours

into account is of paramount importance for implementing an effective risk management process in multi-factor investing: given the high non-stationarity of the system at hand, it is fundamental to develop robust causal structure learning models for dealing with small samples, and to analyze the causal interactions of risk premia at a finer grain.

Sparsification. As far as the changing in factors relationships over time is concerned, our results support, from a causal perspective, the evidence of factor redundancy that have been provided by recent findings [19, 21]. The analysis of the inferred causal networks suggests that, during past years the factorial system was driven by one-lag causal interactions whereas, more recently, instantaneous relationships among those risk factors have appeared. Therefore, a proper factor causal model needs to take into account instantaneous effects as well. In addition, both causation and correlation analyses support the presence of a sparsifying trend in inter-layer connections. In particular, we hypothesize that the loss of memory of the system may be due to an increase in the sophistication of the market participants, who are able to react faster to external stimuli [7, 11].

Factor unveiling. To better characterize the evolution over time of the system under consideration, we checked whether the phenomenon of unveiling a factorial model would affect the causal structure. The rationale behind such a test is that, after a factor becomes known, investors start betting on it, and then factor relations might be altered. However, we did not find any statistically significant evidence for the link between unveiling a factor and a change in the causal structure around it. In order to identify the drivers of the aforementioned sparsifying trend, it might be useful to analyze the system at hand by using higher frequency data and, in addition, to inspect possible links with the recent commodification of factorial strategies within the US market.

Financial and economic stress. With regard to the behavior of the factorial system during stress periods, inspecting both causation and correlation results provides a richer view of the underlying dynamics. Our findings show that, during financial turmoil, the level of instantaneous association among factors slightly decreases, and the factorial system becomes driven by the market factor. As a consequence, the causal structure becomes denser and more stable due to the increase of the out-degree of the market node, i.e., the influence that the market has on all the other factors. This result is of paramount importance for investors, since it shows that, during stress periods, the exposure to different factors reverts to a simple exposure to market risk. Our finding echoes recent ones which provide evidence for the market factor being by far the dominant one [20, 21]. Furthermore, the reported relationship between the change in the VIX Index and both correlation and causal structures reinforces existing evidence in financial literature concerning the link between the VIX Index and factors returns [14].

Finally, we have analyzed the relation between the change in macroeconomic conditions, as measured by the spread of the yield curve, and the evolution of the factorial system. Also in this case we can appreciate the benefit of looking at causation. According to correlation analysis alone, the business cycle indicator is not associated with the dynamics of the analyzed system. Conversely, the regression analysis which relates the business cycle indicator

to the causal structure of the factorial system displays a statistically significant relationship with a 95% confidence level. Indeed, similarly to the results of the volatility analysis, the worsening of macroeconomic environment is associated with a growth in the number of network arcs. Therefore, looking at the causation allows to better inspect the evolution of the system during negative economic phases.

Future work. The results in this paper contribute to make some progress in understanding the relationships among risk factors. However, several questions remain open. As an example, it would be interesting to study the interactions among risk factors belonging to different equity markets. Moreover, our analysis concerns only the equity asset class. Thus, enlarging the considered set of factors, by including risk premia concerning other asset classes as well, could help in taking into consideration also inter-asset class dynamics. Finally, by construction, causal networks enable studying the response of the system at hand under interventions. Therefore, it would be interesting to exploit the attained results to setup a suitable stress testing procedure for multi-factor portfolios.

ACKNOWLEDGMENTS

The authors acknowledge the support from Intesa Sanpaolo Innovation Center. The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

REFERENCES

- [1] Alan Agresti. 2018. *An introduction to categorical data analysis*. John Wiley & Sons. 74–90 pages.
- [2] Clifford Asness and Andrea Frazzini. 2013. The devil in HML's details. *The Journal of Portfolio Management* 39, 4 (2013), 49–68.
- [3] Clifford S Asness, Andrea Frazzini, and Lasse Heje Pedersen. 2019. Quality minus junk. *Review of Accounting Studies* 24, 1 (2019), 34–112.
- [4] Marco Bardoscia, Paolo Barucca, Stefano Battiston, Fabio Caccioli, Giulio Cimini, Diego Garlaschelli, Fabio Saracco, Tiziano Squartini, and Guido Caldarelli. 2021. The Physics of Financial Networks. *arXiv preprint arXiv:2103.05623* (2021).
- [5] Monica Billio, Mila Getmansky, Andrew W Lo, and Loriana Pelizzon. 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of financial economics* 104, 3 (2012), 535–559.
- [6] Tim Bollerslev. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics* 31, 3 (1986), 307–327.
- [7] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. 2014. High-frequency trading and price discovery. *The Review of Financial Studies* 27, 8 (2014), 2267–2306.
- [8] Peter Bühlmann, Jonas Peters, Jan Ernest, et al. 2014. CAM: Causal additive models, high-dimensional order search and penalized regression. *Annals of statistics* 42, 6 (2014), 2526–2556.
- [9] Mark M Carhart. 1997. On persistence in mutual fund performance. *The Journal of finance* 52, 1 (1997), 57–82.
- [10] David Maxwell Chickering. 2002. Optimal structure identification with greedy search. *Journal of machine learning research* 3, Nov (2002), 507–554.
- [11] Tarun Chordia, T Clifton Green, and Badrinath Kottimukkalur. 2018. Rent seeking by low-latency traders: Evidence from trading on macroeconomic announcements. *The Review of Financial Studies* 31, 12 (2018), 4650–4687.
- [12] John H Cochrane. 2009. The Cross-section: CAPM and Multifactor Models. In *Asset pricing (Revised edition)*. Princeton university press, Chapter 20, 435–449.
- [13] John H Cochrane. 2011. Presidential address: Discount rates. *The Journal of finance* 66, 4 (2011), 1047–1108.
- [14] Robert B Durand, Dominic Lim, and J Kenton Zumwalt. 2011. Fear and the Fama-French factors. *Financial Management* 40, 2 (2011), 409–426.
- [15] Arturo Estrella and Gikas A Hardouvelis. 1991. The term structure as a predictor of real economic activity. *The Journal of Finance* 46, 2 (1991), 555–576.
- [16] Arturo Estrella and Frederic S Mishkin. 1996. The yield curve as a predictor of US recessions. *Current issues in economics and finance* 2, 7 (1996).
- [17] Eugene F Fama and Kenneth R French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of financial economics* 33, 1 (1993), 3–56.
- [18] Eugene F Fama and Kenneth R French. 2015. A five-factor asset pricing model. *Journal of financial economics* 116, 1 (2015), 1–22.
- [19] Guan hao Feng, Stefano Giglio, and Dacheng Xiu. 2020. Taming the factor zoo: A test of new factors. *The Journal of Finance* 75, 3 (2020), 1327–1370.
- [20] Stefano Giglio and Dacheng Xiu. 2017. *Inference on risk premia in the presence of omitted factors*. Technical Report. National Bureau of Economic Research.
- [21] Campbell R. Harvey and Yan Liu. 2021. Lucky factors. *Journal of Financial Economics* (2021). <https://doi.org/10.1016/j.jfineco.2021.04.014>
- [22] Campbell R Harvey, Yan Liu, and Heqing Zhu. 2015. ... and the cross-section of expected returns. *The Review of Financial Studies* 29, 1 (2015), 5–68.
- [23] David Heckerman, Dan Geiger, and David M Chickering. 1995. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine learning* 20, 3 (1995), 197–243.
- [24] Johan Himberg, Aapo Hyvärinen, and Fabrizio Esposito. 2004. Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage* 22, 3 (2004), 1214–1222.
- [25] Kewei Hou, Chen Xue, and Lu Zhang. 2015. Digesting anomalies: An investment approach. *The Review of Financial Studies* 28, 3 (2015), 650–705.
- [26] Kewei Hou, Chen Xue, and Lu Zhang. 2017. *Replicating Anomalies*. Technical Report. National Bureau of Economic Research.
- [27] Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. 2008. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems* 21 (2008), 689–696.
- [28] Biwei Huang, Kun Zhang, Yizhu Lin, Bernhard Schölkopf, and Clark Glymour. 2018. Generalized score functions for causal discovery. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1551–1560.
- [29] Biwei Huang, Kun Zhang, Jiji Zhang, Joseph Ramsey, Ruben Sanchez-Romero, Clark Glymour, and Bernhard Schölkopf. 2020. Causal discovery from heterogeneous/nonstationary data. *Journal of Machine Learning Research* 21, 89 (2020), 1–53.
- [30] Aapo Hyvärinen. 1999. Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks* 10, 3 (1999), 626–634.
- [31] Aapo Hyvärinen, Kun Zhang, Shohei Shimizu, and Patrik O Hoyer. 2010. Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research* 11, 5 (2010).
- [32] Antti Ilmanen and Jared Kizer. 2012. The Death of Diversification Has Been Greatly Exaggerated. *The Journal of Portfolio Management* 38, 3 (2012), 15–27.
- [33] Michael C Jensen, Fischer Black, and Myron S Scholes. 1972. The capital asset pricing model: Some empirical tests. In *Studies in the Theory of Capital Markets*. Praeger Publishers Inc.
- [34] Philipp J Kremer, Andreea Talmaciu, and Sandra Paterlini. 2018. Risk minimization in multi-factor portfolios: What is the best strategy? *Annals of Operations Research* 266, 1 (2018), 255–291.
- [35] R David McLean and Jeffrey Pontiff. 2016. Does academic research destroy stock return predictability? *The Journal of Finance* 71, 1 (2016), 5–32.
- [36] Alessio Moneta, Doris Entner, Patrik O Hoyer, and Alex Coad. 2013. Causal inference by independent component analysis: Theory and applications. *Oxford Bulletin of Economics and Statistics* 75, 5 (2013), 705–730.
- [37] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [38] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press. 33–41 pages.
- [39] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. 2017. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press. 1–14 pages.
- [40] Jonas Peters, Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. 2014. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research* (2014).
- [41] Hans Reichenbach. 1956. *The Direction of Time*. University of California Press.
- [42] Gideon Schwarz et al. 1978. Estimating the dimension of a model. *Annals of statistics* 6, 2 (1978), 461–464.
- [43] William F Sharpe. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance* 19, 3 (1964), 425–442.
- [44] Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. 2006. A Linear Non-Gaussian Acyclic Model for Causal Discovery. *J. Mach. Learn. Res.* 7 (Dec. 2006), 2003–2030.
- [45] Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvärinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, and Kenneth Bollen. 2011. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *The Journal of Machine Learning Research* 12 (2011), 1225–1248.
- [46] Christopher A Sims. 1980. Macroeconomics and reality. *Econometrica: journal of the Econometric Society* (1980), 1–48.
- [47] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. 2000. *Causation, prediction, and search*. MIT press.
- [48] Matthew J Vowels, Necati Cihan Camgoz, and Richard Bowden. 2021. D'ya like DAGs? A Survey on Structure Learning and Causal Discovery. *arXiv preprint arXiv:2103.02582* (2021).