
Electronic Thesis and Dissertation Repository

April 2013

News, Copulas and Independence

Ivan Medovikov

The University of Western Ontario

Supervisor

Dr. John Knight

The University of Western Ontario

Graduate Program in Economics

A thesis submitted in partial fulfillment of the requirements for the degree in Doctor of Philosophy

© Ivan Medovikov 2013

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>

 Part of the [Econometrics Commons](#)

Recommended Citation

Medovikov, Ivan, "News, Copulas and Independence" (2013). *Electronic Thesis and Dissertation Repository*. 1176.
<https://ir.lib.uwo.ca/etd/1176>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact tadam@uwo.ca.

NEWS, COPULAS AND INDEPENDENCE
(Thesis format: Monograph)

by

Ivan Medovikov

Graduate Program in Economics

A thesis submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

The School of Graduate and Postdoctoral Studies
The University of Western Ontario
London, Ontario, Canada

© Ivan Medovikov 2013

Acknowledgements

I am very grateful to a number of people who helped me get to this point in my career. I would like to thank the staff and the faculty of the Department of Economics for their support and input, especially Yvonne Adams, Igor Livshits, Jim MacGee, Maria Ponomareva, Gregory Pavlov. I am also grateful to Ricardas Zitikis (Department of Statistics), Hao Yu (Department of Statistics), Youngki Shin, Chris Bennett and Tim Conley for their insightful comments, to Tony Wirjanto (Department of Statistics, University of Waterloo) for his excellent suggestions, and to Martijn van Hasselt for his enduring support which was always available when it was most needed. I am especially thankful to my advisor, John Knight, who has my greatest appreciation and admiration, who continues to inspire me, and whom I see not only as an advisor, but as a mentor, and a friend.

Personal acknowledgement: I would also like to acknowledge my spouse, Ksenia Bushmeneva, and thank her for her love, care, and support throughout these years.

To my parents.

Abstract

The essence of almost any multivariate econometric analysis is the analysis of dependence between sets of random quantities of interest. A series of new, powerful and omnibus non-parametric tools for the analysis of dependence emerged recently in the econometric literature, which are based on the statistical theory of copulas. A copula function completely, and in the case of random variables with continuous marginal distribution functions, uniquely characterizes their interdependence. The term "copula" emphasizes the manner in which it "couples" the marginal distributions into joint by specifying the dependence structure. This dissertation consists of four chapters and aims to contribute to both the theory and the applications of copulas to problems in economics, econometrics and finance.

The second chapter is empirical in nature. It adopts the copula approach to the analysis of the relationship between macroeconomic news and financial markets. A large body of literature which studies the impact of news on equity markets exists. A common finding reported in such studies is that news relating to dividends, profits, and the overall state of the economy explains only a small part of the variability in the aggregate stock market returns. A common limitation of such studies is the use of the narrow measures of news, such as, for example, scheduled releases of economic data. The second chapter proposes a broad measure of macroeconomic news which is termed the "Macroeconomic News Index". The index is based on a manual review and classification of thousands of news releases. Using the copula approach, new findings relating to the relationship between macroeconomic news measured using the Index and the stock markets are revealed. In particular, it is found that macroeconomic news has a much larger impact on the equity markets than reported in earlier studies, and that the relationship is highly asymmetric.

The third chapter is theoretical, and aims to improve the existing non-parametric copula-based tools which help formally and rigorously establish the presence of dependence in the data. It provides an extension to the independence test statistic proposed in Genest and Remillard (2004) and recently Kojadinovic and Holmes (2009), which is obtained through the introduction of a weighted functional norm. The addition of the weights creates a channel through which the power properties of the test can be manipulated. The choice of the weights which favors observations closer to the median of the distribution is shown have the ability to give the statistic a significant power advantage. Quessy (2010) recently showed how a test for independence can be used to test for the goodness of fit of parametric copulas. The addition of the weights may be particularly beneficial in this context, since it permits the imposition of asymmetric losses which arise as a consequence of modeling error. The third chapter provides additional results which enable the application of the test statistic to the problems which involve estimated quantities such as regression model residuals. In an illustrative application of these results, the test is used to probe the presence of conditional heteroscedasticity in a linear regression model, and is shown to have a power advantage over the test of White (1980) in several settings. This appears to be one of the first applications of the copula theory to the problem of residual-based testing which deserves closer attention and future work.

The test statistic proposed in the third chapter cannot be directly applied to the analysis of time series. The fourth chapter provides a serial extension to the statistic, which is closely related to the serial test of Kojadinovic and Yan (2009). The ability to adjust the power of the test through the different choices of the weights is retained in the serial case. The fourth chapter

further extends the results of Quessy (2010) to the serial setting, which permits the application of the statistic to the testing for the goodness of fit of serial copulas. One limitation of the statistics of Kojadinovic and Holmes (2009), Kojadinovic and Yan (2009), and the weighted statistic proposed in the third chapter is the lack of standardization, meaning that they cannot be used as measures of dependence. An upper bound for the statistic is derived in Chapter 4, which is subject to the empirical marginal copulas. A standardized version of the statistic is proposed, which can serve as an omnibus measure of vectorial dependence. A computational formula for the new copula-based dependence measure is provided.

Keywords: Copulas, independence, testing, serial independence, measures of dependence, macroeconomic news, market efficiency, heteroskedasticity.

Contents

Acknowledgements	ii
Abstract	iv
List of Figures	ix
List of Tables	x
List of Appendices	xi
1 Introduction	1
2 The Marco-economic News Index	5
2.1 Introduction	5
2.2 Relevant Literature	8
2.3 The Macroeconomic News Index	9
2.3.1 Indexing Methodology	10
2.3.2 Employment News Index And The Labor Market News	12
2.3.3 Housing News Index And The Housing Market News	15
2.3.4 Industrial News Index And Industry News	17
2.3.5 Energy News Index, Crude Oil Market And OPEC-Related News	19
2.3.6 The Macroeconomic News Index	21
2.4 The Impact Of News	23
2.4.1 The Linear Model	23
2.4.2 The Copula Approach	24
2.5 Conclusion	26
3 An Omnibus Copula Test of Independence and Goodness of Fit	28
3.1 Introduction	28
3.1.1 Background on Copulas in Econometrics	31
3.1.2 Outline of the Chapter	32
3.2 The Non-Linear Dependence Problem	32
3.3 The Copula Model Selection Problem	34
3.4 Non-Parametric Tests for Independence	34
3.4.1 Empirical distribution function tests	35
3.4.2 Empirical copula tests	35
3.4.3 Empirical copula process	36

3.4.4	Copula test statistics	38
3.5	The Weighted Test Statistic	40
3.5.1	Weighted empirical process	41
3.5.2	The weighting function	44
3.5.3	Asymptotic distribution and the critical values	45
3.5.4	Closed-form expression	47
3.5.5	Weighted tests as a unifying class	48
3.6	Power Study and Practical Implementation	49
3.7	Applications and Extensions	51
3.7.1	Testing Copula Goodness of Fit	53
3.7.2	Testing independence of regressors and model errors	54
3.7.3	Copula-based test for conditional heteroscedasticity	58
3.7.4	Residual tests using parametric copula models	60
3.8	Summary and Concluding Remarks	60
4	A Generalized Measure of Serial Dependence	63
4.1	Introduction	63
4.2	The Test Statistic	66
4.2.1	Serial Copula	67
4.2.2	Weak convergence of $C_n(\mathbf{u})$ under the H_0 of independence	68
4.2.3	The Weighted Test Statistic	70
4.2.4	The Weighting Function	70
4.2.5	The Dependence Structure of Some Common Time-Series	71
4.2.6	Normalized Statistic	73
	Computational Formula for the $\iota_{n,s}$	75
	Measures of Dependence and the Shuffles of M	76
4.2.7	Closed-Form Expression	77
4.2.8	The Critical Values	78
4.2.9	Power Properties	78
4.3	Applications	81
4.3.1	Testing for Vectorial Serial Independence in Financial Data	81
4.3.2	A Goodness of Fit Test for Archimedian Serial Copulas	82
4.3.3	Interval Estimation of Serial Copula Parameters	85
4.3.4	A Note About the Interpretation of the Weighting Function	86
4.3.5	Applied Example: Portfolio Value-at-Risk Computation	87
4.4	Conclusion	88
	Bibliography	89
A	Proofs of Theorems	94
A.0.1	Proof of Lemma 3.5.1	94
A.0.2	Proof of Theorem 3.5.2	94
A.0.3	Proof of Theorem 3.5.4	95
A.0.4	Proof of Theorem 3.7.4	97
A.0.5	Proof of Theorem 3.7.5	97

A.0.6	Proof of Lemma 3.7.2	98
A.0.7	Proof of Theorem 4.2.3	98
A.0.8	Proof of Theorem 4.2.4	98
A.0.9	Proof of Theorem 4.2.5	99

Curriculum Vitae		102
-------------------------	--	------------

List of Figures

2.1	Employment News Sub-Index, January 1999 - December 2008	13
2.2	Number of Employment-Related Releases, January 1999 - December 2008 . . .	14
2.3	Housing News Sub-Index, January 1999 - December 2008	16
2.4	Industry News Sub-Index, January 1999 - December 2008	17
2.5	Number of Industry-Related Releases, January 1999 - December 2008	18
2.6	Energy News Sub-Index, January 1999 - December 2008	20
2.7	Number of Energy-Related Releases, January 1999 - December 2008	21
2.8	Macroeconomic News Index, January 1999 - December 2008	22
2.9	Dependence map of the MNI and real market returns.	25
3.1	Independence and symmetric tail dependence.	33
3.2	Convergence of finite sample distributions of I_n^h to asymptotic.	46
3.3	Local power of I_n^h and I_n^l	50
3.4	Local power of \tilde{I}_n^h and of the White test.	59
4.1	Empirical Dependence Maps.	73

List of Tables

2.1	News counts by category, January 1999 to August 2009.	11
2.2	Effect of Employment News on Real Market Returns.	15
2.3	Effect of housing news on real market returns.	17
2.4	Effect of industry news on real market returns.	19
2.5	Effect of energy news on real market returns.	22
2.6	Effect of Macroeconomic News Releases on Aggregate Market Returns.	24
3.1	Bi-variate asymptotic critical values for I_n^w	46
3.2	Empirical Type I error rates.	51
3.3	Power of the copula goodness of fit test based on I_n^h	54
4.1	Power of the $I_n(g^1(\mathbf{u}))$ and the $I_n(w(\mathbf{u}))$ statistics. $n = 25, s=100$ samples (Gumbel Copula).	80
4.2	The power of the $\bar{I}_n(g^1(\mathbf{u}))$ and the $\bar{I}_n(t(\mathbf{u}))$ statistics. $n = 25, s = 100$ samples (Gumbel copula). Embedding dimension $p = 2$, series dimension $q = 2$	80
4.3	Power of the $I_n(g^1(\mathbf{u}))$ and the $I_n(t(\mathbf{u}))$ statistics. $n = 25, s = 100$ samples (t copula). Embedding dimension $p = 2$, series dimension $q = 1$	80
4.4	Variable names for the leading stock market indexes.	81
4.5	Stock market index groupings.	81
4.6	p -values from the test of H_0 of serial independence of some of the series in Table 4.5 for embedding dimensions from $p = 2$ to $p = 6$	82

List of Appendices

Appendix A Proofs of Theorems	94
---	----

Chapter 1

Introduction

The essence of almost any multivariate econometric analysis is the analysis of dependence between sets of random quantities of interest. Econometric models therefore serve as cameras through which dependence can be identified and analysed, and as with any camera, the picture which the viewer sees is in part a product of its design. The classical linear regression model, for example, can accurately resolve the presence of linear correlation, but has a limited field of view when exposed to some non-linear forms of dependence. Among all econometric techniques, non-parametric methods have the ability to detect the greatest variety of forms of dependence with minimal assumptions, giving the viewer perhaps not the sharpest picture, but the broadest possible field of view.

A series of new and powerful non-parametric tools emerged recently in the econometric literature, which are based on the statistical theory of copulas. A copula function completely, and in the case of random variables with continuous marginal distribution functions, uniquely characterizes interdependence. The term "copula" emphasizes the manner in which it "couples" the marginal distributions into joint by specifying the dependence structure. The copula can easily be estimated empirically, allowing the accurate and complete mapping of the dependence in the data with minimal assumptions. Originally emerging in the statistical literature, copulas are now widely used in derivatives pricing, risk management, portfolio design, equity

and property research, and are slowly making their way into the econometric mainstream.

This dissertation consists of four chapters and aims to contribute to both the theory and the applications of copulas to problems in economics, econometrics and finance. The second chapter is empirical in nature. It adopts the copula approach to the analysis of the relationship between macroeconomic news and financial markets. Lasting debate about market efficiency exists in the literature. Following the efficient market hypothesis, fluctuations in real aggregate equity prices should be attributed to the arrival of new information about the underlying state of the economy and the expected dividends. Contrary to the hypothesis, empirical studies typically find that the so-called "fundamental" factors such as macroeconomic news and dividends account for only a small share of the variation of the U.S. aggregate equity returns. The presence of significant unexplained variation is often taken as evidence of irrational or speculative behaviour dominating the markets. The key limitation of existing studies is the use of narrow measures of economic news, such as, for example, scheduled statistical releases of important economic data. While significant, economic data releases omit important information such as political developments, natural disasters and other economic shocks. The second chapter introduces an alternative broad and comprehensive measure of news which is termed the "Macroeconomic News Index", or MNI for short. The Index is based on a review and quantification of a very large number of economic news releases carried by a major international news agency, the Reuters Newswire. Using an empirical estimate of the copula of the MNI and the U.S. aggregate stock market returns, new findings about the impact of economic news on the financial market are revealed.

The third chapter is theoretical, and aims to improve the existing non-parametric copula-based tools which help formally and rigorously establish the presence of dependence in the data. Genest and Remillard (2004) and recently Kojadinovic and Holmes (2009) developed powerful statistics which can be used to test for stochastic independence between sets of random variables directly from the sample. Their tests are pivotal, meaning that no distributional knowledge or assumption is required to carry them out. The ability to detect dependence in

a distribution and model-free way is helpful since it enables identification of the relationship of interest in the data prior to the development of a full structural model. The test has additional applications outside of non-parametric data analysis, since many other economic and econometric hypotheses can be formulated as hypotheses of independence. The third chapter provides an extension to the test of Kojadinovic and Holmes (2009) which is obtained through the introduction of a weighting function into the definition of the copula-based test statistic. The addition of the weights is shown to give the statistic a significant power advantage in several settings. In general, the weighting function allows for the adjustment of the power properties of the test through the alternative choices of the weights, enabling the researcher to "tune" the test based on the problem at hand. Convergence results and a computational formula for the test statistic are provided for a class of continuous, non-negative, integrable and additively-separable weighting functions.

While the test statistic provided in chapter three works well with independently-drawn observations, it cannot be applied directly to the analysis of time-series. A surprisingly small number of tools is available for the analysis of temporal dependence in vector time series, and even among the existing tests, few are able to distinguish between the contemporaneous dependence within the vector series, the temporal dependence within the univariate series, and the vectorial serial dependence. The fourth chapter provides an extension to the weighted test statistic, enabling its application to the analysis of vector time series of any dimension. The key advantage of the copula approach in this context is the ability to test for serial independence between vectors, without contaminating the results by the presence of dependence within a vector or serial dependence within univariate series. The statistic is closely related to that of Kojadinovic and Yan (2009), and is capable of identifying vectorial dependence at any order of lag. As in chapter three, the computational formula and convergence results for the statistic are provided.

The weighted statistic proposed in chapter three, as well as in Kojadinovic and Holmes (2009) and Kojadinovic and Yan (2009) share a common limitation in that they cannot serve as

measures of dependence. This limitation is addressed in the second half of chapter four which derives an upper bound for the serial version of the weighted test statistic, which is subject to the empirical marginal copulas of the vectors involved in the test. The existence of the upper bound enables the standardization of the test statistic, which can now be constrained to the $[0, 1]$ interval, turning it into a meaningful and a truly general distribution and model-free measure of dependence. An exact computational formula for this new measure is provided.

Chapter 2

The Marco-economic News Index

2.1 Introduction

Scheduled releases of various economic data are commonly used to assess the impact of macroeconomic news on financial markets. While many scheduled statistical releases do move the markets, they account only for a fraction of news about the underlying economic variables which hits the market wire. Comments made by senior government officials, statisticians, prominent economists and business figures, political events, natural disasters, and even long-term weather forecasts by Farmer's Almanac or the thickness of Federal Reserve Chairman's briefcase before the Open Market Committee's meetings¹, are all signals (of varying precision) about the underlying economic conditions to which the markets may react, and which often appear to be at least as important to the markets as economic data themselves.

Perhaps due to the large volume of economy-related releases which the major news agencies carry, and to the many issues which arise when quantifying largely qualitative news such

¹Farmer's Almanac long-term weather forecasts are often carried by Reuters News Wire and are seen by many market commentators as signals of, for example, future performance of the construction industry, or future energy prices. The thickness of Alan Greenspan's briefcase before the FOMC meetings formed the basis for the so-called Briefcase Indicator popularized by CNBC: a thick briefcase was viewed as an indicator of much paperwork and hence of some concern about the economy, and indicated a heightened likelihood of an interest rate change. In reality, the thickness of the briefcase, as was revealed by Greenspan after his retirement from the Federal Reserve, was determined by the presence (or absence) of lunch.

as expert opinions or geo-political events, there has not been a study in the literature which assesses the impact of a broader macro-economic news on the financial markets.

In this paper I construct a comprehensive measure of U.S. macroeconomic news which I call the Macroeconomic News Index, or MNI for short, and use it to assess the impact of economic news on the U.S. real aggregate equity market returns. The index is based on a large number of macro economy-related news stories carried by major international news agency, the Reuters Newswire, and quantifies both the volume and the information content of a wide range of news releases covering key U.S. macroeconomic indicators. I use the index to estimate the amount of variation in the U.S. real aggregate stock market returns which can be attributed to the inflow of economic news.

Understanding the relationship between market returns and the inflow of information which is often viewed as "fundamental" by the market participants can contribute to a lasting debate over the efficiency of the financial markets. Recent turmoil in the financial markets and the apparent failure of the markets to price correctly mortgage-backed securities brings up new questions about market efficiency in a broader sense. On one hand multiple price anomalies discovered over the years such as predictability of stock returns based on price to earnings ratios (Basu (1977)), price to book value (Fama and French (1992)), IPO characteristics (Ritter (1991)), market over-reaction to significant price movements (Bondt and Thaler (1985)), the tendency of a price to revert to its mean (Poterba and Summers (1988)), or to follow the daily sentiment in the social media (Bollen, Mao, and Zeng (2011)), and an apparent volatility of stock prices in excess of what is justified by dividends (Shiller (1981), Leroy and porter (1981) and Grossman and Shiller (1981)) suggest that significant market inefficiencies may be present. On the other hand many of these effects disappeared shortly after they were discovered, suggesting that any abnormal profit opportunities were quickly competed away, lending support to the efficient markets view.

A common approach to the problem of testing for market efficiency involves searching for factors which tend to move but not necessarily to predict stock prices. A common view

is that on the aggregate, equity prices in an efficient market should be moved by "fundamentals" - macroeconomic variables which have a direct effect over future dividends, and by other similarly-relevant news signals, but not by "behavioural" factors such as the hours of sunshine in a given day, or the short-term social media sentiment. A consistent finding in the literature is that macroeconomic fundamentals, at least when measured using historical economic data or releases of new economic data, do not explain a significant amount of the variation of the aggregate stock market returns. This fact is often used as evidence against market efficiency.

This study contributes to the earlier literature in several ways. Firstly, Macroeconomic News Index presented in Section 2.3 is arguably the broadest and most comprehensive quantitative measure of macro economy-related news available today. The monthly index is based on a large number of daily news releases carried by major news agencies, and is a quantification of the majority of the U.S. employment, housing, construction, and industry-related news as well as of a large share of energy-related news releases which hit the market wire over the past ten years. The news releases are processed in their original form, and the resulting series can be viewed as real-time data, meaning that unlike the macroeconomic data which are subject to numerous subsequent revisions, the resulting index is based on information exactly as it was available to the market participants on the day of its release. Using the index, we reveal new information about the degree and the nature of interdependence between macroeconomic news and the financial markets.

Secondly, since the index provides an observable, quantitative measure of news polarity and volume, it can be used in many other econometric applications beyond market efficiency studies.

The rest of this chapter is organized as follows. Section 2.2 reviews the existing literature on the effects of macroeconomic news on financial markets. Section 2.3 describes the historical news data sources as well as the collection, classification and aggregation methods used in the construction of the index. It also reviews general properties of the raw data and of the resulting index. The relationship between macroeconomic news and aggregate stock market returns is

investigated in Section 2.4. Section 2.5 concludes.

2.2 Relevant Literature

A rich body of literature on the impact of various economic news on financial markets exists. Virtually all studies use scheduled economic data releases as the measure of news. Pearce and Roley (1985) are among the first to study the effects of scheduled money supply, inflation, and real economy-related statistical releases on the U.S. stock market returns, and find that while the markets react significantly to the unexpected components of monetary announcements, there was no evidence that news of real activity have an effect. Jain (1988) finds similar results using industrial production, unemployment and producer price announcements as measures of real economy-related news. The lack of early evidence to suggest that fundamental macroeconomic news influences the stock prices can be attributed to the non-linearities in the true relationship between news and market returns. Cutler, Poterba, and Summers (1989) find that stock returns respond to industrial production, but the relationship holds for only a part of their sample. McQueen and Roley (1993) find that a relationship between macroeconomic news and prices emerges once the different stages of the business cycle are taken into account. In particular, McQueen and Roley (1993) show that the markets tend to respond negatively to favorable economic news during period of high economic activity. Exploring the non-linearities further, Flannery and Protopapadakis (2002) find that out of the 17 macro series announcements they consider, balance of trade, employment report and housing starts have a significant impact on the conditional volatility of market returns. Similar results are reported by Graham, Nikkinen, and Sahlstrom (2003), who find that the employment report, the National Association of Purchasing Managers manufacturing data, Producer Price Index, Import / Export Price Index and the Employment Cost Index appear to influence the markets. More recently, Boyd, Hu, and Jagannathan (2005) find a significant market reaction to the Bureau of Labor Statistics (BLS) monthly announcements of the unemployment rate, particularly on a smaller time interval.

In other financial markets, the effects of macroeconomic news on the high-frequency behavior of the exchange rates is documented by Almeida, Goodheart, and Payne (1998) and by Ederington and Lee (1993) and on interest rates by Becker, Finnery, and Kopecky (1995), among others.

2.3 The Macroeconomic News Index

The Macroeconomic News Index proposed in this section aims to capture and quantify both the polarity and the volume of news, and is based on a broad range of news releases covering a variety of U.S. macroeconomic indicators. The series are constructed at monthly frequency, beginning in January 1999 and ending in December 2008, and are a result of manual aggregation and processing of 13,813 daily news releases carried by major news agencies which feed daily information to the U.S. financial markets. In particular, stories covering the U.S. labor market, the U.S. housing and construction, and the U.S. industrial production are used to measure the information which the markets receive about the underlying state of the national economy. Scheduled data releases by the U.S. Department of Commerce covering new home sales and new housing starts, Bureau of Labor Statistics employment reports and the FRB's industrial production and capacity utilization figures have all been shown to trigger a market reaction (for example, see Boyd, Hu, and Jagannathan (2005) and Flannery and Protopapadakis (2002)), which motivates the inclusion of employment, industry and construction-related news groups into the index.

Swings in energy costs are often viewed as another important source of stock market movements. Stories covering global energy markets, with particular focus on the crude oil market news and news relating to the Organization of Petroleum Exporting Countries, are also incorporated into the news index. I use Dow Jones Factiva as the source for historical news releases, and the Factiva's Dow Jones Intelligent Indexing service to identify news stories as belonging to one of the three news categories (U.S. employment, industrial production or housing

market-related news).

2.3.1 Indexing Methodology

To minimize duplication, and to reasonably limit the total amount of news stories to be indexed, I restrict attention to employment, housing and construction, and industrial production-related releases carried by the Reuters Newswire. As a news agency, Reuters acts as a primary news source for other news outlets, and will typically be among the first to transmit newsworthy information, albeit often in condensed form.

I use the Dow Jones Energy Service archive as the source of historical energy news, and restrict my attention to the crude oil market and OPEC-related news.

For each of the four news groups (U.S. employment, U.S. industrial production, U.S. housing market and global crude oil market and OPEC news) the news releases are manually reviewed, and each release is first classified either as a "major" or a "minor", based on the perceived significance of the news, and then subsequently as a "favorable", "unfavorable" or a "neutral". For example, a Reuters release announcing better-than-expected national employment numbers will be classified as a "major favorable". Similarly, an unexpected announcement by the U.S. Secretary of Energy about a significant release of the U.S. strategic oil reserves will be recorded as "major" energy news. It is important to note that each news story is classified based on its perceived polarity and significance at the time – for example, an announcement of favorable jobless claims numbers may still be recorded as negative, if the claims reduction was significantly below what was expected at the time.

While any such classification is bound to be subjective, sections 2.3.2, 2.3.3, 2.3.4 and 2.3.5 outline the classification guidelines for each of the news groups in more detail. For each of the four news groups, Table 2.1 shows the total number of treated news releases in each of the six categories and four news groups.

The news index is based on the four sub-indexes which aggregate the relevant daily information for each of the four news groups into a monthly index. News sub-index for news

	Category						Total
	Major			Minor			
News group	Favorable	Unfavorable	Neutral	Favorable	Unfavorable	Neutral	Total
Employment	448	568	50	1751	1614	245	4676
Housing	479	427	55	248	209	48	1466
Ind. prod.	753	422	46	368	338	22	1949
Energy	1682	938	449	955	1105	593	5722
Total	3362	2355	600	3322	3266	908	13813

Table 2.1: News counts by category, January 1999 to August 2009.

group $i \in \{\text{labor, energy, construction, industry}\}$ for a given month t is based on the ratio p_t of favorable stories to the total number of non-”neutral” news releases in a month, and on the total amount of news releases n_t , and is constructed as:

$$N_t^i = (p_t - 0.5) \sqrt{n_t} \quad (2.1)$$

The ratio p_t captures the polarity of news in month t and shows whether, on average, good or bad news dominated the month. The total number of releases n_t measures the news volume.

Minor news releases receive half of the weight of the major news in the index. Linear model parameter estimates reported in Section 2.4 suggest that ”minor” news releases appear to have, on average, half as big of an impact on market returns as news stories classified as ”major”, which motivates such weighting. Following such definition, index values of zero indicate a mixed news-month, regardless of news volume. Larger positive index values are recorded when largely favorable news stories dominated the headlines, and the overall volume of news was high. Similarly, small positive values of the index represent monthly news coverage which was, on average, favorable, but had a smaller news volume. Negative values of the index are recorded when unfavorable stories dominate the news.

In general, a news release is viewed as ”major” if it contains important information about the national economy which is likely to have an immediate and significant effect on the market. A release is viewed as ”favorable” if it indicates either improving economic conditions

(for employment, housing and industry news groups), or an increase in price (in energy markets), or raises the outlook for future economic performance. For example, news of increasing OPEC quotas or decreasing compliance among members of the cartel, or of the decision of the U.S. Energy Secretary to release large portions of the U.S. Strategic Petroleum Reserve to the markets are added to the "major unfavorable" category of the energy news group, since they both indicate a raising supply and a likely lower oil price. Similarly, news of increasing non-farm payrolls or dropping weekly jobless claims, or comments made by senior BLS officials indicating "strong" or "improving" labor market are recorded in the "major favorable" category of employment news group. Guidelines which were used to classify news releases in each of the four news groups, and the nature of news stories which are included into the index, are reviewed in more detail in Subsections 2.3.2, 2.3.5, 2.3.3 and 2.3.4.

2.3.2 Employment News Index And The Labor Market News

Figure 2.1 shows the monthly values of employment news sub-index between January 1999 and August 2008. Note that the negative news trend clearly identifies the two "bad news" regimes during the 2001-2002 and the more recent 2008- recession.

Figure 2.2 shows the number of employment-related news stories carried by the Reuters News wire every month between January 1999 and December 2008.

Roughly two thirds of employment-related news releases cover scheduled U.S. Labor Department announcements reporting changes in monthly non-farm payrolls, national employment and unemployment rates, and weekly jobless claims, as well as reports by the regional Bureaus of Unemployment Services showing regional unemployment rates. Remaining news is driven by items such as commentary and predictions by senior BLS, White house or Federal Reserve officials, forecasts based on Reuters opinion polls, forecast and comments by notable analysts, and revisions to earlier figures issued by the BLS.

Stories covering nation-wide data releases, such as non-farm payrolls and unemployment, polls of forecasters predicting key monthly figures, and senior official comments, are treated

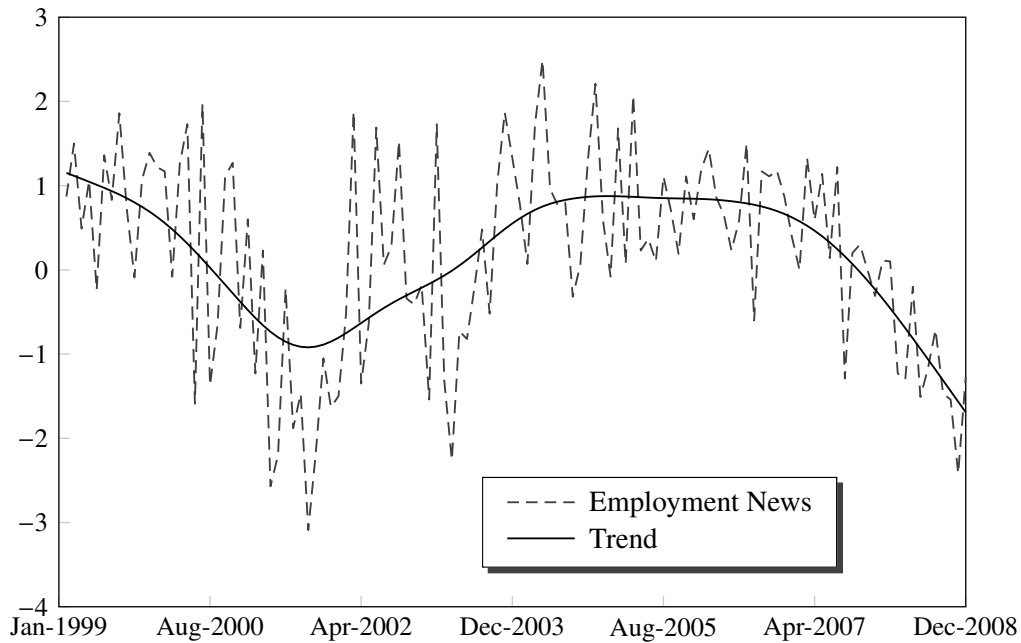


Figure 2.1: Employment News Sub-Index, January 1999 - December 2008

as "major" releases. News on regional statistics or high frequency data, such as weekly jobless claims, along with individual analyst comments, are recorded as "minor" releases.

A particular difficulty of quantifying news over longer periods of time is the changing composition and nature of daily news. For example, in the later part of the sample, regional news are much less frequently carried by the wire, while expert opinions become more prevalent.

To test whether there is a broad relationship between employment news and market returns, I regress aggregate monthly market returns on raw news counts in each of the six employment news categories. Dividend-inclusive return on a value-weighted NYSE market portfolio, deflated by monthly change in the consumer price index is used as a measure of monthly real aggregate return throughout this paper. Value-weighted market return data are taken from the Standard and Poor's COMPUSTAT database.

Table 2.2 shows parameter estimates for the corresponding linear model. Employment news counts in "major neutral", "minor unfavorable" and "minor neutral" categories appear insignificant, and are excluded from the model. News counts in "major favorable", "minor favorable" and "major unfavorable" categories on the other hand are significant at below 5%.

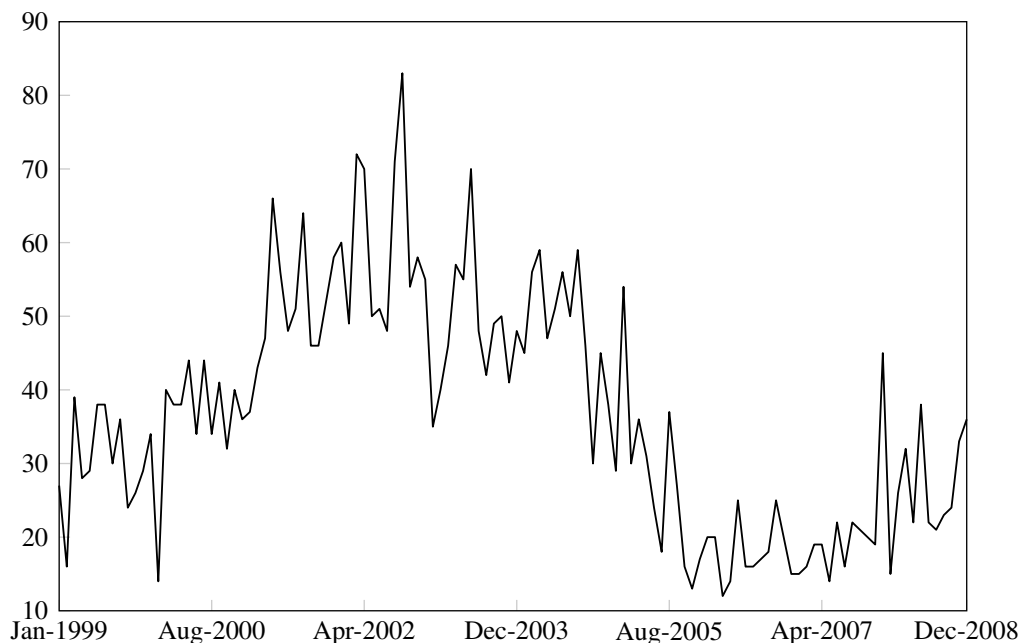


Figure 2.2: Number of Employment-Related Releases, January 1999 - December 2008

Raw employment news counts appear to be able to explain between four and nine percent of variation in monthly real aggregate U.S. stock returns ², which is comparable to the explanatory power of contemporaneous monthly payroll figures over market returns reported in other studies.

Table 2.2 parameter estimates suggest that on average, a major favorable employment-related news release tends to decrease monthly market returns. This could be due to the market's view that a tightening labor market signals inflationary pressure and increases the chance of a real interest rate hike. The estimate of the coefficient on major unfavorable employment news count is also negative, indicating that the news of, for example, falling non-farm payrolls are also seen as negative, perhaps now on fundamental rather than inflationary basis. To that extent, there appears to be no good employment news. The somewhat-puzzling coefficients could be explained by non-linearities in the true relationship between employment news and market returns. Boyd, Hu, and Jagannathan (2005) find that the markets tend to react differently to

²An R-squared of 0.09 is achieved when regressing monthly returns on employment news counts, but excluding months in which there was no relevant news stories (news count of zero), which is equivalent to assuming that the absence of news is itself not news.

Dependent variable:	Real value-weighted return to market portfolio			
Method:	Least squares			
Date range:	January 1999 - December 2008			
Included observations:	120			
R^2	0.0532			
Adjusted R^2	0.0370			
Explanatory variables:	Coefficient	Std. Error	T-Stat	P-Value
Major favorable	-0.002318	0.001144	-2.026	0.045
Major unfavorable	-0.001658	0.000784	-2.115	0.037
Minor favorable	0.000894	0.000443	2.016	0.046

Table 2.2: Effect of Employment News on Real Market Returns.

BLS statistical releases during economic contractions and expansions. In particular, they find that an announcement of the raising unemployment is good news for the markets during expansions, but bad news during contractions. The negative coefficients on both favorable and unfavorable labor market news may therefore be a product of the linear model's restrictions.

2.3.3 Housing News Index And The Housing Market News

The monthly values of the Housing News sub-Index are shown in Figure 2.3. While the index fluctuates a lot, favorable stories dominate the news throughout the early parts of the sample. A significant and persistent decline of the index begins in October 2007 and coincides with the start of sub-prime mortgage crisis in the U.S. The volume of housing and construction-related news releases remains roughly constant over time, with an average of twelve stories per month.

The housing news sub-index is based on a broad range of housing and construction-related news. In particular, news on the new and existing home sales, housing completions, new building permits and housing starts reported by the U.S. Department of Commerce, National Association of Homebuilders housing index, long-term weather forecasts, and Reuters opinion polls which publicize market expectations, together with comments by senior White House or Department of Commerce officials are added to the "major" news counts of the construction news group. Revisions to key housing and construction data are also treated as "major" news.

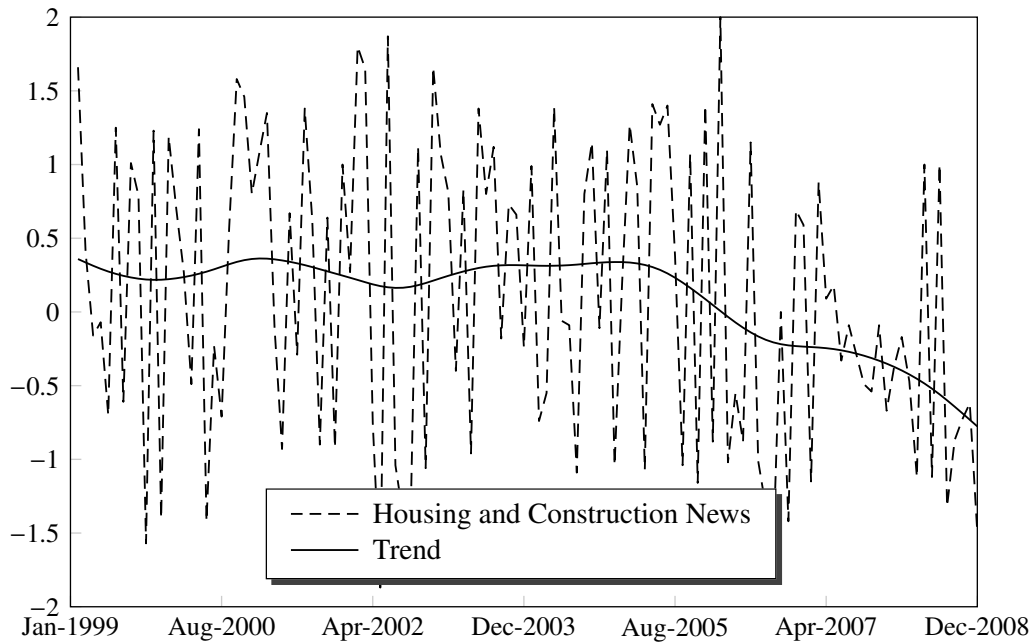


Figure 2.3: Housing News Sub-Index, January 1999 - December 2008

For example, a news release which announces new data showing a rise in housing starts, or comments by a senior White House official pledging to increase home ownership is recorded as "major favorable" news.

On the other hand, other indirect news signals such as commodity market news which is interpreted by commentators to be representative of the situation in the housing market, such as, for example, news of significant changes in timber prices on the Chicago Mercantile Exchange, or changes to mortgage refinance and other mortgage indexes, changes in state-level construction spending, National Association of Realtors home affordability index, rents index reported by the Department of Housing and Urban Development, Canadian house starts and even changes to the American Institute of Architecture's architectural billings index, are all recorded as minor housing and construction news.

Table 2.3 shows parameter estimates of a linear model where monthly real aggregate return is regressed on news counts in the "major favorable", "minor favorable" and "minor neutral" categories of the housing news group.

"Major unfavorable", "major neutral" and "minor unfavorable" news appear insignificant

Dependent variable:	Real value-weighted return to market portfolio			
Method:	Least squares			
Included observations:	120			
Date range:	January 1999 - December 2008			
R^2	0.0769			
Adjusted R^2	0.0612			
Explanatory variables:	Coefficient	Std. Error	T-Stat	P-Value
Major favorable	0.002256	0.001091	2.069	0.041
Minor favorable	-0.003518	0.001735	-2.028	0.045
Minor neutral	-0.007580	0.003030	-2.502	0.014

Table 2.3: Effect of housing news on real market returns.

and were eliminated from the regression. Leaving only the news counts which appear significant at or below 5% level of significance, construction and housing market news appear to be able to explain approximately 7% of the variation in aggregate market returns.

2.3.4 Industrial News Index And Industry News

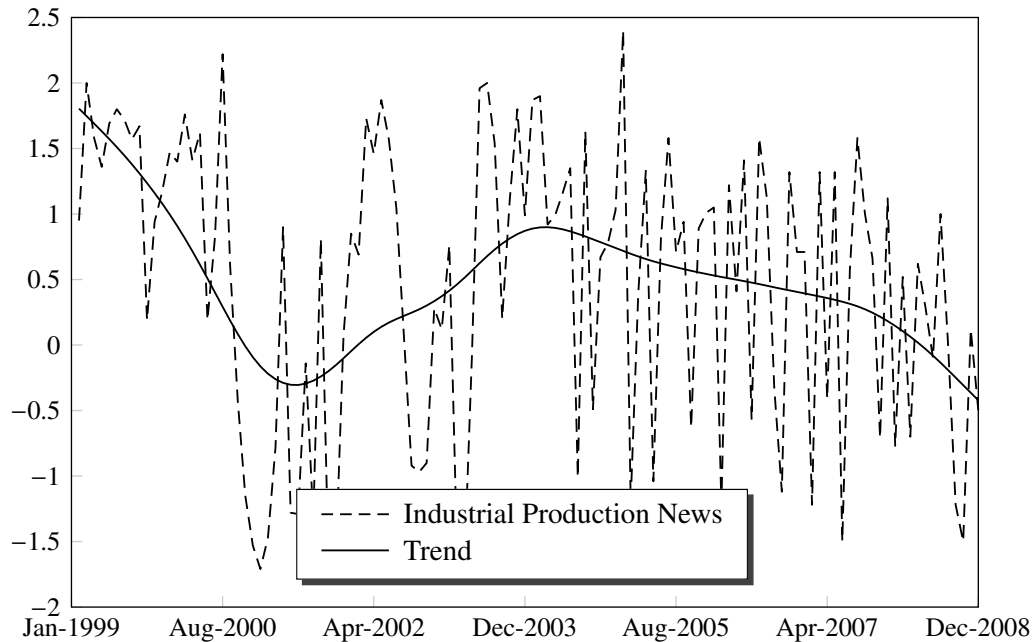


Figure 2.4: Industry News Sub-Index, January 1999 - December 2008

Just like housing and employment, the Industrial News Sub-index is based on a range of

industry-related news releases carried by the Reuters Newswire between January 1999 and August 2009. Figure 2.4 shows the monthly values of the index, and Figure 2.5 shows monthly news volume.

The industry news sub-index is based on large number of news releases covering the statistical releases made by the National Association of Purchasing Managers, comments on the state of the economy made by the NAMP officials, scheduled releases of manufacturing production figures and durable goods orders, Federal Reserve industrial output data releases, Chicago FED National Activity Index, and factory activity indexes, non-farm productivity measures, proprietary industrial indexes, Reuters survey results containing analyst expectations, and revisions to published data, which are all treated as "major" news.

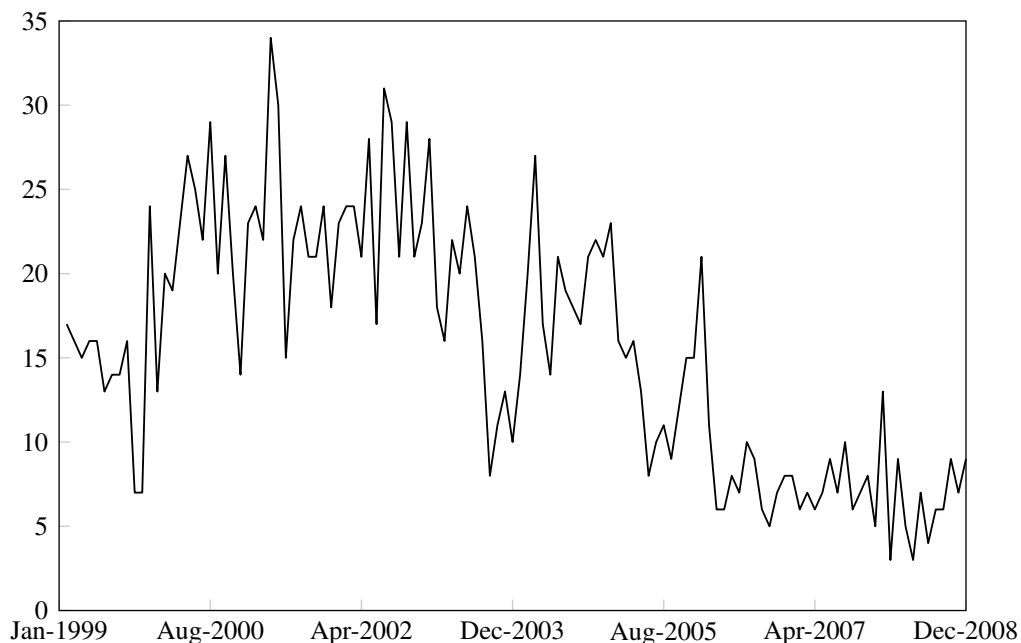


Figure 2.5: Number of Industry-Related Releases, January 1999 - December 2008

Publicized comments made by individual analysts, regional FRB reports announcing changes in manufacturing activity and durable good orders, such as, for example, the Chicago Midwest Manufacturing Index, the Mid-Atlantic Manufacturing Report (Business Conditions Report) published by Philadelphia FRB, the Richmond FRB's monthly manufacturing and services index, the Morgan Stanley Global Purchasing Managers Index, and the Manufacturing Alliance

Industry reports, among others, are also included in the index, and are recorded as "minor" news releases.

Table 2.4 shows linear model estimates, where monthly return is regressed on "major favorable" and "minor favorable" industry-related news counts. News in the remaining four industry-related news categories appear to be insignificant. Major favorable and minor favorable news releases are only borderline significant. Overall, industry-related news appear to explain only 1% of variation in aggregate stock market returns.

Dependent variable:	Real value-weighted return to market portfolio			
Method:	Least squares			
Date range:	January 1999 - December 2008			
R-squared	0.0187			
Adjusted R-squared	0.0104			
Included observations:	120			
Explanatory variables:	Coefficient	Std. Error	T-Stat	P-Value
Major favorable	0.001324	0.000894	1.482	0.141
Major neutral	-0.002430	0.001594	-1.524	0.130

Table 2.4: Effect of industry news on real market returns.

2.3.5 Energy News Index, Crude Oil Market And OPEC-Related News

Figure 2.6 shows the monthly values of the Energy News Sub-Index. Energy news are classified from the standpoint of energy producers, meaning that news indicating market tightening is recorded as favorable. An alternative index centered around energy consumers can be easily derived as the negative of the current series. Crude oil prices which persistently rose throughout the decade partly explain largely favorable composition of energy-related news since 1999. Figure 2.7 shows the amount of energy-related news stories which were treated to construct the monthly energy sub-index. The news volume appears to be steadily raising in the later part of the sample, which may be due to the increasing media attention toward the crude oil market as it rallied to its historical high during this time.

Other than the periodic quota announcements which tend to follow the OPEC meetings,

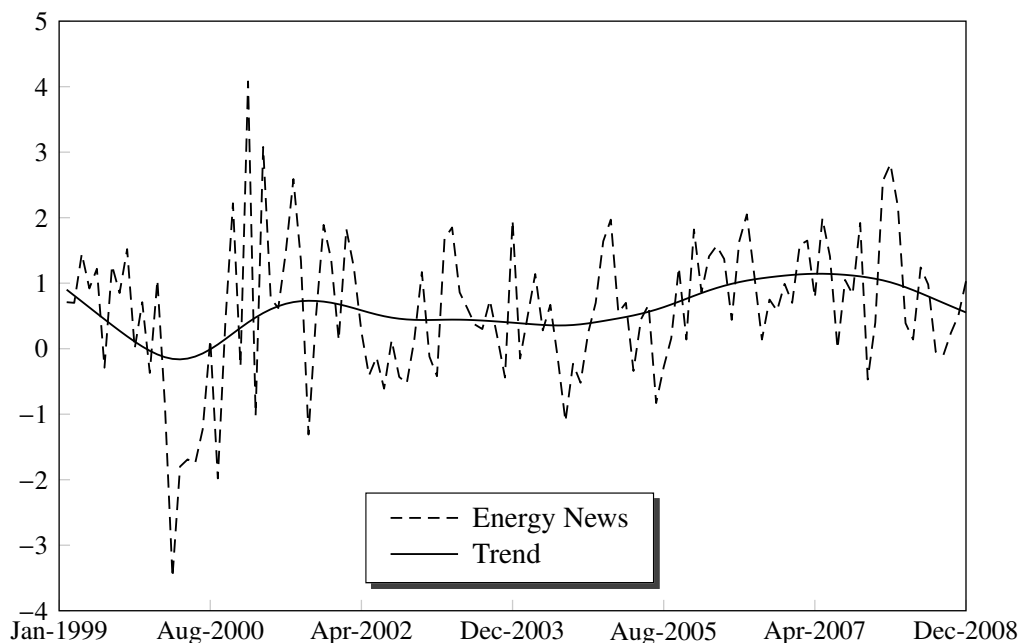


Figure 2.6: Energy News Sub-Index, January 1999 - December 2008

there are few significant crude oil data which are released on a schedule. To that extent, the Energy News Sub-Index is based on perhaps the broadest mix of news stories. At the core of the index are the Dow Jones Energy News Service releases which cover crude oil production and OPEC-related issues. Among these, announcements of changes to OPEC output quotas, geo-political events and natural hazards, such as wars, strikes, terrorist acts and hurricanes, which are likely to affect production of crude, decisions and commentary by senior White House officials on the fate of the U.S. Strategic Petroleum Reserves, estimates, comments and announcements made by the senior members of the International Energy Agency, OPEC quota compliance rates, comments made by oil ministers of the OPEC countries, and even the relevant long-term weather forecasts, particularly the hurricane season forecasts published by the Farmers Almanac, are all treated as "major" energy-related news releases.

Similarly, announcements of output hikes or cuts by individual OPEC member states, local geo-political events which signal supply disruptions, and decisions, announcements or comments made by either members of IEA, U.S. Administration or Legislature indicating pressure on OPEC to change output quotas, such as, for example, legislative initiatives to impose sanc-

tions on member states, are treated as "minor" releases.

For example, a news release announcing a joint OPEC decision to significantly cut production quotas, or a formation of a major hurricane in the Gulf of Mexico, is classified as a "major favorable" to the price of crude. News of a possible future shock to the production of crude, such as, for example, verbal comments made by the U.S. Energy Secretary urging OPEC members to increase quotas, is classified as "minor unfavorable".

Table 2.5 shows model parameter estimates where market return is regressed on "major favorable", "major neutral" and "minor favorable" energy news counts. Raw energy-related news story counts appear to be able to explain approximately 7% of the variation in monthly returns.

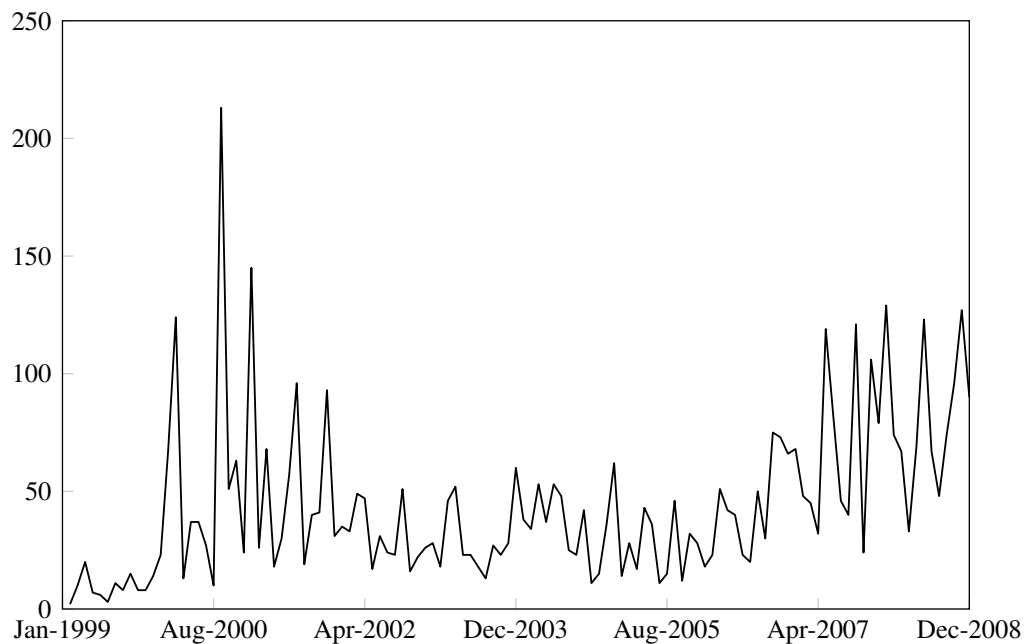


Figure 2.7: Number of Energy-Related Releases, January 1999 - December 2008

2.3.6 The Macroeconomic News Index

The final Macroeconomic News Index (MNI) is defined as the sum of employment, housing, industrial production and energy news sub-indexes presented in Sections 2.3.2 through

Dependent variable:	Real value-weighted return to market portfolio			
Method:	Least squares			
Date range:	January 1999 - December 2008			
R-squared	0.0781			
Adjusted R-squared	0.0623			
Included observations:	120			
Explanatory variables:	Coefficient	Std. Error	T-Stat	P-Value
Major favorable	0.000737	0.000403	1.830	0.070
Major neutral	-0.002003	0.000918	-2.183	0.031
Minor favorable	-0.000997	0.000569	-1.753	0.082

Table 2.5: Effect of energy news on real market returns.

2.3.5. Figure 2.8 shows the monthly values of the Macroeconomic News Index between January 1999 and December 2008. While the definition of the index is kept simple, the MNI is able to trace the key macroeconomic events of the past decade, such as the early 2000s recession and the more recent slowdown following the sub-prime crisis of 2007 fairly accurately. Unlike the historical macroeconomic data, however, the index provides a real-time and a highly-comprehensive measure of macroeconomic news.

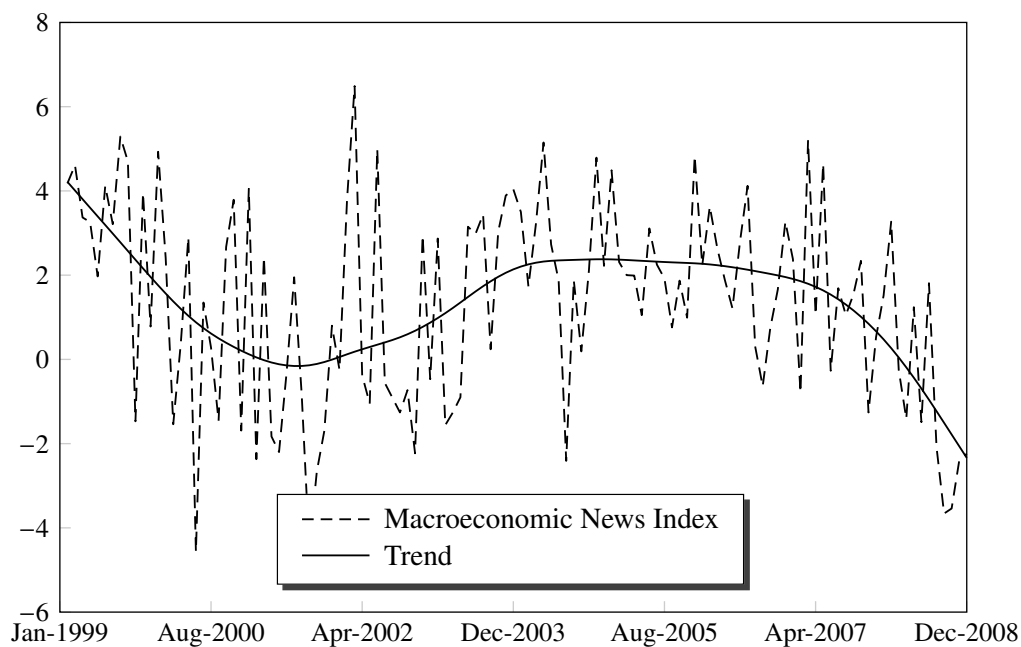


Figure 2.8: Macroeconomic News Index, January 1999 - December 2008

2.4 The Impact Of News

The following section assesses the relationship between the Macroeconomic News Index and the aggregate stock market returns. Firstly, Subsection 2.4.1 aims to establish the degree to which the macroeconomic news and the financial markets are connected. Secondly, Subsection 2.4.2 adopts the copula approach in order to precisely map the nature of this dependence.

2.4.1 The Linear Model

To assess the explanatory potential of macroeconomic news on stock market returns, I estimate a linear model where monthly real aggregate return is regressed on a combination of news counts from all four news groups. Beginning with all 24 news counts (six individual news categories for each of the four news groups), I re-estimate the model, eliminating variables which do not appear to be significantly different from zero, and while such elimination improves the adjusted R-squared. From the initial 24, only 11 news categories survive the elimination. Ten are significant at below 10% level of significance, and one is borderline significant, but still improves the adjusted R^2 .

Table 2.6 shows the model estimates. Macroeconomic news, measured by the classified news counts in the employment / unemployment, industrial production, housing / construction and energy costs categories explain about a fourth to a third of the variation in monthly aggregate real stock returns. This appears to be the highest explanatory power of macro news over market returns reported in the literature to date, which suggests that broader measures of news such as the Macroeconomic News Index contain information which is significant to the markets, and which is not captured by scheduled macroeconomic data releases.

While news counts in individual categories can explain some market returns, aggregating news into employment, housing, industrial production and energy news sub-index using (2.1) does not yield interesting results. While individual news are significant, the news sub-indexes appear insignificant, at least in the linear model setting.

Dependent variable:	Real value-weighted return to market portfolio			
Method:	Least squares			
R-squared	0.3046			
Adjusted R-squared	0.2408			
Included observations:	120			
Explanatory variables:	Coefficient	Std. Error	T-Stat	P-Value
Employment major favorable	-0.309521	0.105281	-2.940	0.004
Employment major unfavorable	-0.175322	0.071706	-2.445	0.016
Employment minor favorable	-0.807097	0.372161	-2.169	0.032
Energy major favorable	0.072231	0.043402	1.664	0.099
Energy major neutral	-0.223742	0.085758	-2.609	0.010
Energy minor favorable	-0.130286	0.053455	-2.437	0.016
Housing major favorable	0.368552	0.115326	3.196	0.002
Housing major unfavorable	0.235956	0.126412	1.867	0.065
Housing minor unfavorable	0.676294	0.220999	3.060	0.003
Housing minor neutral	-1.038506	0.295768	-3.511	0.001
Industry major neutral	-0.458655	0.314821	-1.457	0.148

Table 2.6: Effect of Macroeconomic News Releases on Aggregate Market Returns.

2.4.2 The Copula Approach

While the MNI can explain a significant percentage of the variability in the U.S. aggregate stock returns, it is important to understand the nature of this relationship in finer detail. This section adopts the copula approach in order to map the dependence between the MNI and the market returns. It is now well understood that the dependence structure of any pair of continuous random variables is completely and uniquely characterized by the copula function of the corresponding joint probability distribution.

Independence of any two random variables is characterized by the so-called independence copula. In order to study the nature of the dependence between the MNI and the market returns, I construct what I refer to as the "dependence map" of these data. In particular, I first obtain an estimate of the empirical copula of the MNI and of the market return series, and then calculate the differences between the empirical copula and the independence copula across the support of the joint distribution. The presence of differences at a particular point in the distribution indicates interdependence between the Macroeconomic News Index and the market.

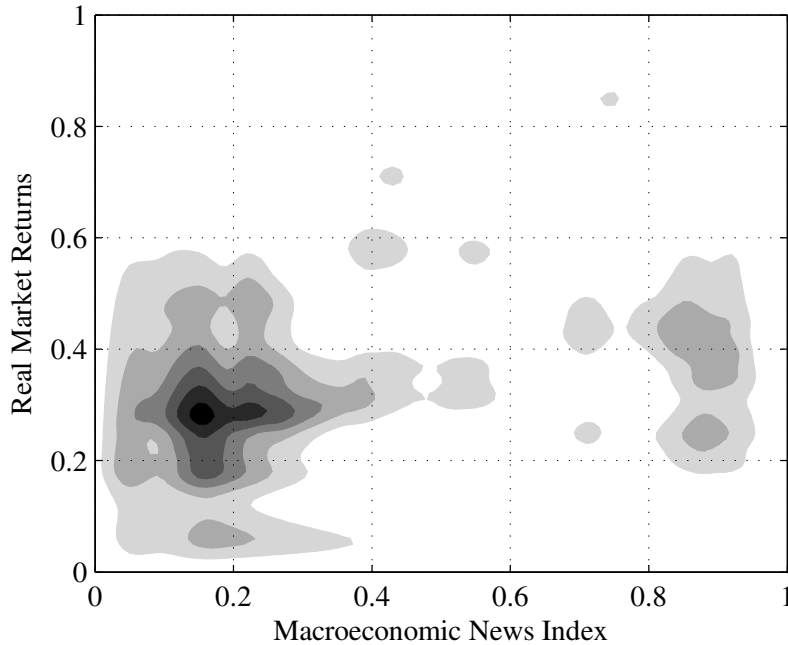


Figure 2.9: Dependence map of the MNI and real market returns.

Figure 2.9 shows such differences for the empirical copula of the MNI and the market returns. The units on the horizontal and vertical axes are the percentiles from the corresponding joint probability distribution. For example, a point $[0.5, 0.5]$ on this plot indicates the median of the joint distribution, while point $[0, 0]$ indicates the leftmost tail, and $[1, 1]$ the rightmost.

Gray shading indicates differences between the empirical copula of the data, and the independence copula. Darker shades indicate larger differences, meaning that stronger interdependence is present. Several interesting conclusions can be drawn from Figure 2.9. The dependence between MNI and the market is highly asymmetric around the median. In particular, the clustering of the data is strongest in the lower-left tail of the joint distribution, suggesting the existence of a particularly high interdependence between negative macroeconomic news and market returns.

2.5 Conclusion

The impact of macroeconomic news on the financial markets has been widely studied in the literature. Traditionally, releases of important economic data are used as a measure news, yet scheduled releases account for only a fraction of significant macro economy-related news stories carried by the daily news wires. In this paper I assess the impact of a broader measure of economic news on the U.S. stock markets by constructing what I call the Macroeconomic News Index, or MNI for short.

The Index is based on the manual review and classification of 13,813 daily economy-related news releases carried by the Reuters News Wire and the Dow Jones Energy Service between January 1999 and August 2009. In particular, MNI captures and quantifies both the polarity and the volume of U.S. labor market, housing and construction, and industrial product news, as well as the global energy market news. Using a simple linear model I assess the explanatory power of news on the U.S. stock market returns and find that at a monthly frequency, macroeconomic news can explain up to 30% of variation in real aggregate market returns. Such explanatory power is significantly higher than that reported in earlier studies, suggesting that broader measures of news do carry significant information well beyond what is captured by scheduled releases of macroeconomic data.

News in the narrowly-classified categories which form the basis of the aggregated news indexes significantly outperform indexes themselves in all models estimated here, suggesting that the method of aggregation I choose may not be the optimal. This could be one venue for future research.

The precise nature of the relationship between macroeconomic news and the market returns is further analysed using the copula approach. It is now well understood that the so-called copula function of any pair of continuous random variables completely and uniquely characterized the nature of their interdependence. Using an estimate of the corresponding empirical copula, I construct what I call a "dependence map" of the data, revealing that firstly, the interdependence between macroeconomic news and the financial markets appears highly asymmetric, and sec-

only, that dependence is particularly strong in the lower-left tail of the joint distribution, that is, between the negative macroeconomic news and market returns.

Chapter 3

An Omnibus Copula Test of Independence and Goodness of Fit

3.1 Introduction

A growing body of empirical evidence confirms the presence of highly non-linear and asymmetric dependence in many important economic and financial data. For example, Beine, Cosma, and Vermeulen (2010) document non-linear co-movements among extreme values in the tails of joint distributions of some of the common macroeconomic and financial series. Patton (2006) uncovers asymmetric dependence in the exchange rate data that is inconsistent with the commonly used normal or Student's t distributions. Cappiello, Engle, and Sheppard (2006) find asymmetric dynamics in the correlations of global equity and bond returns. Knight, Lizieri, and Satchell (2005) show the existence of lower-tail dependence between the real estate and equity markets. Many more examples of non-linear dependence in financial data can be found in Ang and Chen (2002), Ang and Bekaert (2002), Bae, Karolyi, and Stulz (2003) and Campbell, Koedijk, and Kofman (2002).

Econometric models with a richer dependence structure that accommodate non-linear and asymmetric relationships are still emerging and tend to have better fit and forecasting power

(for example, see Xu, Knight, and Wirjanto (2011)). In many cases, the development of such models stems from discoveries of previously unknown non-linearities or asymmetries in the data. Despite the recent empirical work, many interesting relationships may still be hidden, since common tests for independence, for example those based on Pearson correlation, Kendall's τ and Spearman's ρ -statistics as well as normal scores are generally not omnibus tests. While effective at detecting specific types of alternatives to independence, they may have no power with lesser known forms of non-linear dependence¹. Section 3.2 gives an example involving symmetric tail dependence which standard tests including regression-based tests cannot detect.

This chapter proposes a new test for independence which can identify co-movements of any form and at any point in the distribution without relying on distributional knowledge or assumptions. The test statistic is based on a weighted empirical copula process with attractive finite sample properties and is amenable to computation. While many economic hypotheses can be reduced to testing for independence, the proposed test has important applications beyond non-parametric data analysis. For example, Quessy (2010) proposes a new goodness of fit procedure for Archimedian copulas as an application for a test of independence. We show that by applying the new test in a manner similar to Quessy (2010), it can be a versatile and powerful tool for copula model selection.

The contribution of this chapter is two-fold. Firstly, this chapter proposes a class of non-parametric copula tests for independence that are omnibus, pivotal², and in certain settings are shown to have power advantage over comparable unweighted test statistic proposed by Kojadinovic and Holmes (2009). A key innovation is the use of a weighted functional norm when constructing the test statistic. In addition to the ability to deliver gains in power, the new class unifies classic independence test statistics based on the distribution function with the more recent copula testing literature. In particular, we show that the distribution function-based

¹For a review of some lesser-known concepts of dependence such as quadrant dependence, tail dependence and likelihood ratio dependence see Lehmann (1966).

²The test is pivotal for a chosen weighting function, and if none of the variables is of dimension greater than one.

test statistic of Blum, Kiefer, and Rosenblatt (1961) is a special case of a weighted copula test proposed here defined by the appropriate choice of a weighting function³.

Secondly, new theoretical results are provided that establish convergence of the copula estimator and of the test statistic in situations when a directly observed sample is unavailable and is replaced by estimated quantities. This permits testing for independence between unobserved random variables such as regression model errors and stochastic regressors using estimates and corresponding empirical pseudo-copula⁴. Wu (1973) notes that many economic and econometric hypotheses reduce to the null of independence between regressors and model disturbances. For example, hypotheses of the absence of measurement errors, the presence of recursions in simultaneous equations models and auto-correlations in a cross-section dynamic model can be restated in this form⁵. In general, the usual t and F tests are only applicable if stochastic regressors are independent from model errors⁶. This chapter extends the copula approach to the problem of testing for independence of regressors and model disturbances and provides a comprehensive way of carrying out these tests. To illustrate one application, we restate the hypothesis of conditional homoscedasticity in a linear regression model as the null of independence and show that the new copula-based test can significantly outperform the White (1980) test, especially in small samples. This result is surprising given the non-parametric nature of the test and suggests that the copula approach to residual-based testing has potential. To my knowledge, this chapter represents the first application of copulas to the problem of testing for heteroscedasticity.

³I thank Ricardas Zitikis (Department of Statistics and Actuarial Sciences, The University of Western Ontario) for this insight

⁴Here 'pseudo-copula' means empirical copula based on estimated quantities, or 'pseudo-observations'. This is different from the definition of 'pseudo-copula' by Fermanian and Wegkamp (2012).

⁵See Wu (1973) and Zellner, Kmenta, and Dreze (1966) for details.

⁶e.g. see Goldberger (1964) p. 268

3.1.1 Background on Copulas in Econometrics

Copulas are becoming increasingly popular in econometrics and finance. The copula function combines marginal distributions into joint by specifying the way in which two or more variables are "coupled" together. From an independence testing standpoint, the copula gives a complete description of the dependence structure of the data and uniquely characterizes stochastic independence. This makes it natural to base tests for independence on the empirical copula process.

A key advantage of the copula approach is that it allows the building of complex multivariate distributions by modeling the dependence separately from the marginal behavior of the variables. In many situations economic theory leads to marginal distributions that are difficult to combine due to the lack of an appropriate family of joint distribution functions. This situation frequently arises in microeconometrics, especially in models of discrete or limited choice⁷. For example, Munkin and Trivedi (1999) use copulas to construct joint distributions of discrete event counts that are otherwise restrictive and difficult to estimate.

Copulas have multiple applications in finance, particularly in portfolio risk management, pricing of derivatives and modeling financial contagion. Rosenberg and Schuermann (2006) apply copulas to a risk management problem for a mixed asset class portfolio with high risk of joint extreme events due to tail-dependent asset returns. Cherubini, Luciano, and Vecchiato (2004) consider, in great detail, the application of copulas to derivative pricing. Rodriguez (2007) adopts a copula approach to testing for financial contagion. In other areas of econometrics, Smith (2003) uses copulas to construct models for data that may suffer from selection bias, Chen and Fan (2006) study the estimation of copula-based stationary Markov models which allow separating temporal dependence from marginal behavior of the series. Sancetta and Satchell (2004) develop a theory of approximations for multivariate distributions using a special family of copulas. Patton (2009) and Fan (2010) give an excellent review of the use of copulas in econometrics and finance.

⁷For a discussion of the copula approach to modeling joint parametric distributions and its advantages over simulation-based methods see Trivedi and Zimmer (2007).

3.1.2 Outline of the Chapter

The rest of this chapter is organized as follows. Section 3.2 illustrates the problem of testing for non-linear dependence and gives an example of tail dependence that is difficult to detect using the usual methods. Section 3.3 discusses the copula selection problem and shows how independence tests can be used to test the goodness of fit of a copula model. Section 3.4 gives an overview of currently-available non-parametric tests of independence and introduces basic copula function theory. Section 3.5 presents the new test statistic and establishes its convergence. Section 3.6 studies the local power of the test. Section 3.7 provides several examples of the applications of the test, including testing for the copula goodness of fit. It also establishes convergence of the copula estimator and of the test statistic in situations when direct sampling is replaced by an estimation procedure and shows further applications of the test to the problem of testing for independence of stochastic regressors and model errors. In an illustrative application of these results, the test is used to detect conditional heteroscedasticity in a linear regression model. Section 3.8 concludes. Mathematical proofs are grouped in the Appendix A.

3.2 The Non-Linear Dependence Problem

In many situations common linear and rank correlation statistics and regression-based tests may have no power at detecting non-linear dependence. Consider the following example involving symmetric tail dependence between two uncorrelated variables. I use copulas to construct two bi-variate distributions $F(X_1, Y_1)$ and $G(X_2, Y_2)$ such that the variables are marginally normally distributed with zero mean and unit variance and are uncorrelated. Figure 3.1 shows scatter plots of two random samples from F and G , a thousand observations each.

In the first case, X_1 and Y_1 are also jointly normal and hence independent. In the second case, I use a bi-variate t copula to combine the marginals of X_2 and Y_2 . The resulting joint distribution is non-normal, and despite the lack of correlation and marginal normality, X_2 and Y_2 are

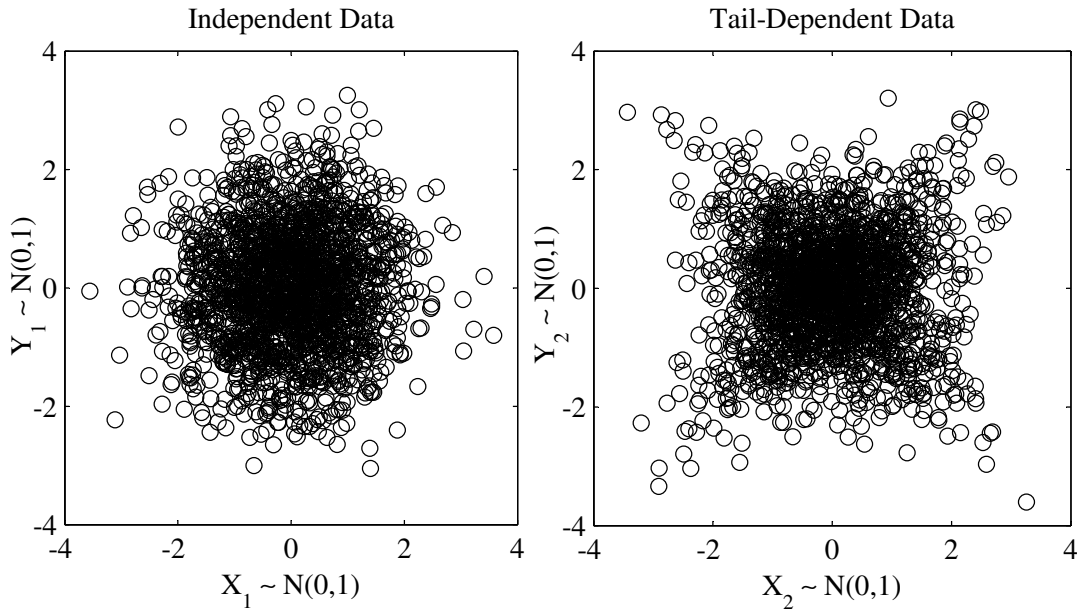


Figure 3.1: Independence and symmetric tail dependence.

not independent but connect through all four tails of their joint distribution. Rank correlations such as Spearman's ρ and Kendall's τ are zero both in F and in G . Rank correlation measures concordance and discordance in the data, where pairs of observations (x_i, y_i) and (x_j, y_j) are said to be concordant if $x_i < x_j$ and $y_i < y_j$ or if $x_i > x_j$ and $y_i > y_j$. Similarly, (x_i, y_i) and (x_j, y_j) are discordant if $x_i < x_j$ and $y_i > y_j$ or if $x_i > x_j$ and $y_i < y_j$. Kendall's τ and Spearman's ρ measure the probability of concordance minus the probability of discordance, and capture the tendency of the data to ordering. Sample measures of linear and rank correlations in both cases are not statistically different from zero, and it is therefore not possible to distinguish between a tail-dependent sample and independent data using these statistics.

The usual regression approach has little power in this setting. At 5% size both t and F tests produce no evidence of significant relationship between the X_2 and Y_2 even when fitting a third-degree polynomial in X_2 through Y_2 . Applying likelihood and score tests is also problematic since the distribution G does not belong to any common multi-variate family and is difficult to characterize without adopting the copula approach.

Two classes of non-parametric tests provide an omnibus solution to this problem. The first uses the sample distribution function. The second and a more recent approach uses sample estimate of the underlying copula. Both types of tests are "blanket" and are able to detect any dependence of any form, including tail dependence as in the example above.

3.3 The Copula Model Selection Problem

Another useful application of a test for independence is to the problem of selecting an appropriate parametric copula to model the data. In most applications, the choice is between several parametric families of copulas which belong to the so-called Archimedean class. This class encompasses over twenty families of the best known and most widely used multivariate copulas which have the capacity to accommodate the widest range of dependence structures. A unique feature of Archimedean copulas is that each copula family is identified by a corresponding *generator function*. Quessy (2010) shows that it is possible to test the hypothesis of the goodness of fit for families of Archimedean copulas using a test for independence by transforming the data in a non-obvious yet simple way with the corresponding generator function. Under the null that the true copula belongs to a hypothesized family, the transformed series are independent. As such, non-parametric tests for independence provide a versatile and powerful approach to the copula model selection problem.

3.4 Non-Parametric Tests for Independence

This section reviews the currently-available non-parametric tests for independence, including the distribution function-based statistic of Blum, Kiefer, and Rosenblatt (1961) and the copula-based test of Genest and Remillard (2004) and Kojadinovic and Holmes (2009) and introduces the basic copula function theory.

3.4.1 Empirical distribution function tests

Let $(X_1, \dots, X_p) = \mathbf{X}$ be continuous real-valued random variables of dimensions d_1, \dots, d_p respectively with joint distribution function F and possibly multi-variate marginals $F_j(X_j)$, $j \in S = \{1, \dots, p\}$. Variables (X_1, \dots, X_p) are independent if it is possible to express the joint distribution as product of the marginals. Let

$$T_n(x) = \hat{F}_n(x) - \prod_{j=1}^p \hat{F}_{j,n}(x_j), \quad (3.1)$$

where $\hat{F}_n(x)$ and $\hat{F}_{j,n}(x_j)$ are the empirical distribution functions of \mathbf{X} and X_j 's. Under H_0 of independence, $E[T_n(x)] = 0$, $\forall x$. Blum, Kiefer, and Rosenblatt (1961) and Hoeffding (1948) characterize, in the bi-variate case, the limiting distribution of the Kolmogorov-Smirnov and Cramer-von Misés functionals of $T_n(x)$, given respectively by:

$$A_n = \sup |T_n(x)| \quad (3.2)$$

$$B_n = \int [T_n(x)]^2 d\hat{F}_n(x) \quad (3.3)$$

Both statistics are consistent against any alternative to independence, and their limiting properties do not depend on the marginal distributions if none of the X 's is of dimension greater than one.

3.4.2 Empirical copula tests

A new class of powerful tests for independence emerged recently from the work of Deheuvels (1979), Genest and Remillard (2004), Kojadinovic and Holmes (2009) and Quessy (2010). The tests use the empirical copula function and much like the distribution function-based statistics, are pivotal (in the scalar case and for a set p) and capable of detecting departures from independence of any form. Sklar (1959) showed that any joint distribution function H of x and y can be written in terms of the marginal distributions $F(x)$ and $G(y)$, and of the so-called cop-

ula $C(u)$ as $H(x, y) = C(F(x), G(y))$. Copula C uniquely describes the way in which marginal distributions are 'coupled' with the joint. As such, the copula provides a complete description of the dependence between x and y , but carries no information about their marginal behavior. Since H takes values in $[0, 1]$ and $u = F(x)$ and $v = G(y)$ are uniform on $[0, 1]$, copula $C : [0, 1] \times [0, 1] \rightarrow [0, 1]$ is itself a joint distribution function on the unit square. This result extends to the more general case as $F(\mathbf{X}) = C(F_1(X_1), \dots, F_d(X_p))$, where $C : [0, 1]^d \rightarrow [0, 1]$ is a distribution function of $d = d_1 + d_2 + \dots + d_p$ uniform random variables, where d_j is the dimension of X_j . Note that for any X_j such that $d_j > 1$, we can use copula theory to further decompose the multivariate marginal distribution function F_j of vector X_j in terms of the distributions of the components of X_j and its *marginal copula* $C(u^{(j)})$. When working with high-dimensional distributions, copulas allow the separate modeling of the marginal behavior of the variables, the dependence *within* groups of variables (random vectors), and the dependence *between* groups of variables, giving an unprecedented richness of distributional structure and complete tractability. In higher-dimensional cases, independence of random vectors is uniquely characterized by the *product copula* $C_{\perp}(u) = \prod_{k=1}^p C(u^{(k)})$. Letting the

$$D(C)(u) = C(u) - \prod_{k=1}^p C(u^{(k)}), \quad (3.4)$$

we have that under the hull of independence between (X_1, \dots, X_p) , the $D(C)(u) = 0, \forall u$, and it is possible to define test statistics using the Cramer-von Misés and Kolmogorov-Smirnov functionals of the underlying empirical process. We review the current copula tests for independence in the next two sub-sections.

3.4.3 Empirical copula process

In what follows, we adopt the notation of Kojadinovic and Holmes (2009). Given a random sample $\{(X_{i1}, \dots, X_{id})\}_{i=1}^n$ from F , where $d = d_1 + d_2 + \dots + d_p$, it is possible to estimate copula $C(u)$ using the *empirical copula estimator* $C_n(u)$. First studied by Deheuvels (1979), the estimator

is commonly defined as:

$$C_n(u) = \frac{1}{n} \sum_{i=1}^n \prod_{j=1}^d \mathbf{1}(\hat{U}_{ij} \leq u_j), \quad (3.5)$$

where $\hat{U}_{ij} = \frac{1}{n} \sum_{l=1}^n \mathbf{1}(X_{lj} \leq X_{ij})$, for $j \in \{1, \dots, d\}$, and $u \in [0, 1]^d$. This amounts to working with the ranks $R_{ij} = n\hat{U}_{ij}$. The scaled and weighted version of (3.5) gives rise to the associated *empirical copula process* \mathbb{C}_n , which establishes the limiting properties of $C_n(u)$ using the functional delta method based on the notion of Hadamard differentiability:

$$\mathbb{C}_n = \sqrt{n} [C_n(u) - C(u)], \quad (3.6)$$

for $u \in [0, 1]^d$. van der Vaart and Wellner (1996) establish weak convergence of (3.6) in $\ell^\infty([a, b]^2)$ for $0 < a < b < 1$, and Fermanian, Radulovic, and Wegkamp (2004) show weak convergence in $\ell^\infty([0, 1]^2)$. Theorem 1 of Kojadinovic and Holmes (2009), which I restate here without proof, summarizes the results found in van der Vaart and Wellner (1996), Tsukahara (2005), Gaenssler and Stute (1987), Ruschendorf (1976) and Fermanian, Radulovic, and Wegkamp (2004) and shows weak convergence of (3.6) in $\ell^\infty([0, 1]^p)$. It is a version of Donsker-Skorohod-Kolmogorov theorem for $\ell^\infty([0, 1]^p)$, where $\ell^\infty([0, 1]^p)$ is understood to be the space of all bounded real-valued functions on $[0, 1]^d$ equipped with uniform metric.

Theorem 3.4.1 *Suppose that $C(u)$ is such that for every $i = 1, \dots, d$, $\delta_i C(u)$ exist and is continuous on the set $\{u \in [0, 1]^2 : 0 < u_i < 1\}$. Then, the empirical copula process*

$$\mathbb{C}_n = \sqrt{n} [C_n(u) - C(u)]$$

converges weakly in $\ell^\infty([0, 1]^p)$ to a tight centered Gaussian process

$$g(u) = \mathcal{B}(u) - \sum_{i=1}^d \partial_i C(u) \mathcal{B}(1, \dots, 1, u_i, 1, \dots, 1), \quad u \in [0, 1]^d$$

where $\mathcal{B}(u)$ is a tight centered Gaussian process on $[0, 1]^d$ with covariance function $E[\mathcal{B}(u)\mathcal{B}(u')] =$

$C(u \wedge u') - C(u)C(u')$, where \wedge denotes component-wise minimum.

Note that the first term of the limit $g(u)$ is independent of the marginals, and is not affected by the replacement of the true marginals with the empirical counterparts. The terms $-\delta_i C(u)\mathcal{B}(u)$ arise due to the replacement of the true quantiles $F_i^{-1}(u_i)$ with the empirical quantiles $F_{n,i}^{-1}(u_i)$. The negative sign is due to the Bahadur-Kiefer asymptotic equivalence of $\sqrt{n}(G_{n,i}^{-1}(u_i) - u_i)$ and $-\sqrt{n}(G_{n,i}^{-1}(u_i) - u_i)$. The partials $\delta_i C(u)$ capture the sensitivity due to deviations away from the margin F_i ⁸.

3.4.4 Copula test statistics

The goal is to test for mutual independence of p continuous random vectors X_1, \dots, X_p of dimensions d_1, \dots, d_p . As before, we let $S = \{1, \dots, p\}$, and $d = d_1 + \dots + d_p$. Let integers $b_j = \sum_{k=1}^j d_k$, $\forall j \in S$, with the convention that $b_0 = 0$. Note that $b_j = b_{j-1} + d_j$. Given a vector $u \in [0, 1]^d$ and a set $B \subset S$, define vector $u^{\{B\}} \in [0, 1]^d$ for any $i \in \{1, \dots, d\}$ as:

$$u_i^{\{B\}} = \begin{cases} u_i & \text{if } i \in \cup_{j \in B} \{b_{j-1} + 1, \dots, b_j\}, \\ 1 & \text{if otherwise.} \end{cases} \quad (3.7)$$

This formalizes the definition of the marginal copula of X_k as $C(u^{\{k\}})$, $u \in [0, 1]^d$ for $k \in S$. Recall the distance functional $D(C)(u)$ that shows the difference between the copula of \mathbf{X} and independence copula $C_{\perp}(u) = \prod_{k=1}^p C(u^{\{k\}})$:

$$D(C)(u) = C(u) - \prod_{k=1}^p C(u^{\{k\}}), \quad \forall u \in [0, 1]^d. \quad (3.8)$$

Mutual independence among X_1, \dots, X_p occurs when $D(C)(u) = 0$, $\forall u \in [0, 1]^d$. The Cramer-von Misés and Kolmogorov-Smirnov-type test statistics for independence that use $D(C)(u)$ are,

⁸I thank Dr. Tony Wirjanto, Department of Statistics, University of Waterloo, for pointing these features of $g(u)$ out.

respectively,

$$I = \int_{[0,1]^d} D(C)(u)^2 du \quad (3.9)$$

and

$$J = \sup |D(C)(u)|. \quad (3.10)$$

These statistics are clearly zero when (X_1, \dots, X_p) are independent. While both statistics lead to a consistent test, the focus has been on the Cramer-von Misés functionals of $D(C)(u)$ as they tend to have power advantage over supremum statistics. To define the empirical counterpart of I , first let

$$\mathbb{C}_{n,\perp} = \sqrt{n} \left[C_n(u) - \prod_{j=1}^p C_n(u^{(k)}) \right], \quad (3.11)$$

which is the *independence empirical copula process*. Integrating $\mathbb{C}_{n,\perp}$ over the domain of the copula gives the sample version of I :

$$I_n = \int_{[0,1]^d} D(C_n)(u)^2 du. \quad (3.12)$$

The statistic I_n is considered briefly in Kojadinovic and Holmes (2009) and is included in the power study by Quessy (2010), but it is not the focus of the copula testing literature. Deheuvels (1979) shows that under the null of mutual independence of X_1, \dots, X_p , the empirical process (3.6) can be decomposed using the Möbius inversion formula into $2^d - d - 1$ mutually-independent sub-processes $\sqrt{n}M_A(C_n)$, $A \subseteq \{1, \dots, d\}$, and $|A| > 1$ which converge jointly to tight centred and mutually-independent Gaussian processes. Under mutual independence of X_1, \dots, X_p , the functional $M_A(C)(u) = 0$, $\forall u \in [0, 1]^d$ and $A \subseteq \{1, \dots, d\}$ such that $|A| > 1$. Genest and Remillard (2004) propose testing for independence using the $2^d - d - 1$ test statistics of the form

$$M_{A,n} = \int_{[0,1]^d} \sqrt{n}M_A(C_n)(u)^2 du \quad (3.13)$$

where the map $M_A(f)(x) : l^\infty([0, 1]^d) \rightarrow l^\infty([0, 1]^d)$ is given by

$$M_A(f)(x) = \sum_{B \subseteq A} (-1)^{|A|-|B|} f(x^B) \prod_{k \in A \setminus B} f(x^{(k)}), \quad x \in [0, 1]^d, \text{ and } P_s = \{B \subseteq S : |B| > 1\}.$$

Since $|S| = d$, the set P_s contains $2^d - d - 1$ elements. Approximate critical values can be obtained numerically using the distribution of the test statistics $M_{A,n}$, for $A \in P_s$ computed from randomized samples, which, under the null of independence, is approximately uniform, and the corresponding p -values are independent, meaning that they can be aggregated so that to establish an overall significance using the Fisher's p -value combination method. Rejection occurs when at least one of the $M_{A,n}$'s exceeds a critical value, and it is possible to determine the overall size of the test by combining the $2^d - d - 1$ individual p -values. Note that if any of the X_k is of size $d_k \geq 2$, the statistic $M_{A,n}$ is a functional of $C_n(u^{(k)})$ even if X_1, \dots, X_p are mutually independent, and is not distribution-free.

Kojadinovic and Holmes (2009) extend this approach to testing for independence of random vectors, Kojadinovic and Yan (2009) provide a serial extension to these statistics. Quessy (2010) studies the asymptotic power and proposes a new goodness-of-fit procedure for Archimedian family of copulas using statistics $M_{A,n}$.

3.5 The Weighted Test Statistic

This chapter proposes a class of statistics that are similar to I_n but use another version of the independence empirical copula process. In particular, we use a weighting function $w(u) : [0, 1]^d \rightarrow \mathfrak{R}^+$ to alter the scaling of the process $C_{n,\perp}(u)$. Throughout the chapter I refer to this new class of tests for independence as the *weighted copula tests*, or simply *W-tests*. The weighting function serves as an index of the test statistics that belong to this class. Convergence results of this section hold for an arbitrary continuous and integrable $w(u)$, meaning that any such weighting function defines a consistent *W-test* with possibly different local power properties. This suggests several interesting questions about the existence (or non-existence)

of the optimal W -test, circumstantial choice of weights, and the possibly of a semi-parametric estimation of the weighting function. I select a particular weighting function and consequently a W -test for the power study in Section 3.6 and show that it has local power advantage over the unweighted tests. Section 3.5.2 discusses the choice of a weighting function in more detail. Section 3.5.5 shows that the statistic I_n of Kojadinovic and Holmes (2009) and the classical statistic of Blum, Kiefer, and Rosenblatt (1961) B_n are, in fact, also W -tests defined by a different choice of the weights.

3.5.1 Weighted empirical process

To develop the test statistic, first consider the altered empirical independence copula process

$$\mathbb{W}_n = \sqrt{n} \left[\left(C_n(u) - \prod_{j=1}^p C_n(u^{(j)}) \right) \sqrt{w(u)} \right] = \sqrt{n} \left[\nu(C_n)(u) - \nu\left(\prod_{j=1}^p C_n\right)(u) \right], \quad (3.14)$$

where $\nu(f)(u)$ is a map from $l^\infty([0, 1]^d)$ to $l^\infty([0, 1]^d)$ given by

$$\nu(f)(x) = f(x) \sqrt{w(x)}, \text{ where } w(x) \in l^\infty([0, 1]^d). \quad (3.15)$$

The weighted Cramer-von Misés test of independence this chapter proposes takes the following form:

$$I^w = \int_{[0,1]^d} \left[C(u) - \prod_{j=1}^p C(u^{(j)}) \right]^2 w(u) du, \quad (3.16)$$

and the sample analog I_n^w based on \mathbb{W}_n is

$$I_n^w = n \int_{[0,1]^d} \left[C_n(u) - \prod_{k=1}^p C_n(u^{(k)}) \right]^2 w(u) du \quad (3.17)$$

To characterize the asymptotic distribution of I_n^w , first define the following difference functional $\phi : l^\infty([0, 1]^d) \rightarrow l^\infty([0, 1]^d)$ given by:

$$\phi(f)(x) = \left(f(x) - \prod_{j=1}^p f(x^{(k_j)}) \right) \sqrt{w(x)}, \text{ for } x \in [0, 1]^d. \quad (3.18)$$

The map $\phi(f)(x)$ takes some bounded continuous function $f(x)$ as input and returns another bounded continuous function given by (3.18). Note that when this map is applied to copula C , the result is a weighted copula distance function which shows weighted differences between copula C and independence copula C_\perp :

$$\phi(C)(u) = \left(C(u) - \prod_{j=1}^p C(u^{(k_j)}) \right) \sqrt{w(u)} \quad (3.19)$$

Applying this to the empirical copula C_n we have that

$$\phi(C_n)(u) = \left(C_n(u) - \prod_{j=1}^p C_n(u^{(k_j)}) \right) \sqrt{w(u)}, \quad (3.20)$$

and the associated empirical process is given by

$$\mathbb{W}_n = \sqrt{n} \phi(C_n)(u) = \sqrt{n} \left[C_n(u) - \prod_{j=1}^p C_n(u^{(k_j)}) \right] \sqrt{w(u)}, \quad (3.21)$$

meaning that we can use the functional Delta-method to obtain the weak limit of the weighted empirical copula process \mathbb{W}_n from the basic limiting properties of the copula estimator C_n . The weak convergence of the weighted empirical processes has been considered in the univariate case by Shorack and Wellner (2009) and Koul (2002) and in the multivariate case by Vanderzande (1981) and van der Vaart and Wellner (1996). Lemma 3.5.1 and subsequently Theorem 3.5.2 formally establish weak convergence of \mathbb{W}_n .

Lemma 3.5.1 *The map ϕ is Hadamard-differentiable tangentially to $l^\infty([0, 1]^d)$ and its deriva-*

tive at $f \in l^\infty([0, 1]^d)$ is:

$$\phi'_f(a)(x) = \left(a(x) - \sum_{i=1}^p a(x^{(i)}) \prod_{j=1, j \neq i}^p f(x^{(j)}) \right) \sqrt{w(x)} \quad (3.22)$$

Proof See Appendix A.

Theorem 3.5.2 *Assume that C has continuous partial derivatives. Then, if X_1, \dots, X_p are mutually independent, $\phi(C)(u) = 0$, $u \in [0, 1]^d$, and the empirical process $\mathbb{W}_n = \sqrt{n}\phi(C_n)(u)$, $u \in [0, 1]^d$ converges weakly in $l^\infty([0, 1]^d)$ to the tight centered Gaussian process*

$$\phi'_C(g)(u) = \left(g(u) - \sum_{j=1}^p g(u^{(j)}) \prod_{i=1, i \neq j}^p C(u^{(i)}) \right) \sqrt{w(u)},$$

where $g(u)$ is defined as in Theorem 3.4.1, with covariance function

$$\begin{aligned} E[\phi'_C(u)\phi'_C(v)] &= \left(C(u \wedge v) - C(u)C(v) \right. \\ &\quad - 2 \sum_{i=1}^p [C(u^{(i)} \wedge v^{(i)}) - C(u^{(i)})C(v^{(i)})] \prod_{j=1, j \neq i}^p C(u^{(j)}) \\ &\quad \left. + \sum_{k=1}^p \sum_{l=1}^p (C(u^{(k)} \wedge v^{(l)}) - C(u^{(k)})C(v^{(l)})) \prod_{i=1, i \neq k}^p C(u^{(i)}) \prod_{j=1, j \neq l}^p C(v^{(j)}) \right) \sqrt{w(u)w(v)}. \end{aligned}$$

Proof See Appendix A.

Note that since the copula function is bounded, the limiting covariance of \mathbb{W}_n is also bounded under H_0 for any $w(u) \in l^\infty([0, 1]^d)$, which ensures that the asymptotic distribution of the test statistic I_n^w is non-degenerate. Corollary 3.5.3 characterizes the limiting distribution of I_n^w .

Corollary 3.5.3 *Suppose that C has continuous partial derivatives. Then, if X_1, \dots, X_p are mutually independent, I_n^w converges in distribution to Z given by*

$$Z = \int_{[0,1]^d} [\phi'(g)(u)]^2 du$$

Proof Similar to Corollary 9 of Kojadinovic and Holmes (2009), convergence is established as a direct consequence of Theorem 3.5.2 and the continuous mapping theorem.

3.5.2 The weighting function

While the convergence results of this section hold for an arbitrary weighting function $w(u) \in l^\infty([0, 1]^d)$, the focus in this chapter is on the W -test defined by the weighting function

$$h(u) = \sum_{k=1}^d (u_k - u_k^2). \quad (3.23)$$

The corresponding W -statistic this chapter proposes is

$$I_n^h = \int_{[0,1]^d} \left(C(u) - \prod_{j=1}^p C(u^{(j)}) \right) \sum_{k=1}^d (u_k - u_k^2) du. \quad (3.24)$$

Weighted metrics are widely used in estimation and testing, but this chapter is the first to address the issue of weighting in the context of non-parametric copula-based tests for independence. The question of the existence (or non-existence) of the optimal weighting function and of the contextual choice of weights is outside of the scope of this chapter. The objective here is to document the ability to manipulate the power properties of the independence test through the alternative choices of the weights. Section 3.6 shows that in some circumstances, the weighted test statistic I_n^h has significant power advantage over the unweighted counterpart.

In practice, the preferred choice of the weights is circumstantial, and depends on the underlying nature of the dependence in the data. For example, copulas belonging to certain families such as Gumbel and Clayton copulas have elliptical contours and will deviate from the independence copula the most at around the median of the joint distribution, which is the point $u = (0.5, \dots, 0.5)$. Copula families belonging to the so-called Archimedian class will tend to share this property. The weighting function $h(u)$ places larger weight on observations closer to the median, and in the case in which the data originate from a distribution with an Archimedian

copula, will likely lead to a power advantage.

In other cases, particularly in the case of extreme tail dependence in the absence of linear correlation, weights $h(u)$ which favour the median may in fact lead to a power loss.

The issue surrounding the choice of the weights is much less clear in the absence of any prior knowledge or assumptions about the nature of the dependence in the data. One approach could be to consider the weights which are inversely-proportional to the limiting variance function of $\phi'_C(g)(u)$, which would lead to a limiting process with a unit variance throughout $[0, 1]^d$. The resulting test statistic will be in the spirit of Anderson and Darling (1952).

Another relevant issue to consider are the restrictions which the copula function theory imposes on the distance functional $D(C)(u)$. Any copula C must lie between the so-called upper and lower Fréchet-Hoeffding bounds $C^U(u) = \min(u_1, u_2, \dots, u_d)$ and $C_L(u) = \max(u_1 + u_2 + \dots + u_d - d + 1, 0)$, which too are copulas. That is, for any C , $C_L(u) \leq C(u) \leq C^U(u)$, $\forall u \in [0, 1]^d$ and therefore $(C_L(u) - C_\perp(u))^2 \leq D(C)(u)^2 \leq (C^U(u) - C_\perp(u))^2$. Since all copulas are grounded (see Nelsen (2006)), if any element of u' is zero, $C^U(u') = C_L(u') = C_\perp(u') = 0$, and $D(C)(u') = 0$. This means the most extreme observations in the sample carry no information about dependence. On the other hand, it is possible to show that $(C^U(u) - C_\perp(u))^2$ symmetrically increases away from the tails and is maximal at the median of the distribution, where the largest deviations from the independence copula are possible.

3.5.3 Asymptotic distribution and the critical values

Note that the copula C of \mathbf{X} and the marginal copulas $C(u^{(k)})$ enter into the expression for the asymptotic covariance function in Theorem 3.5.2. If all of the (X_1, \dots, X_p) are scalars, that is, if $d_j < 2$, $\forall j \in S$, under the null of independence the copula $C(u)$ reduces to the product copula $C(u)_\perp = \prod_{j=1}^d u^{(j)}$. The marginal copulas become $C(u^{(j)}) = u^{(j)}$, which makes the test statistic distribution-free. The only parameters that affect the asymptotic distribution of I_n^w in the scalar case are the weighting function $w(u)$ and the dimension d of the vector \mathbf{X} , and it is possible to tabulate the critical values for a given choice of weights and the number of the variables. Table

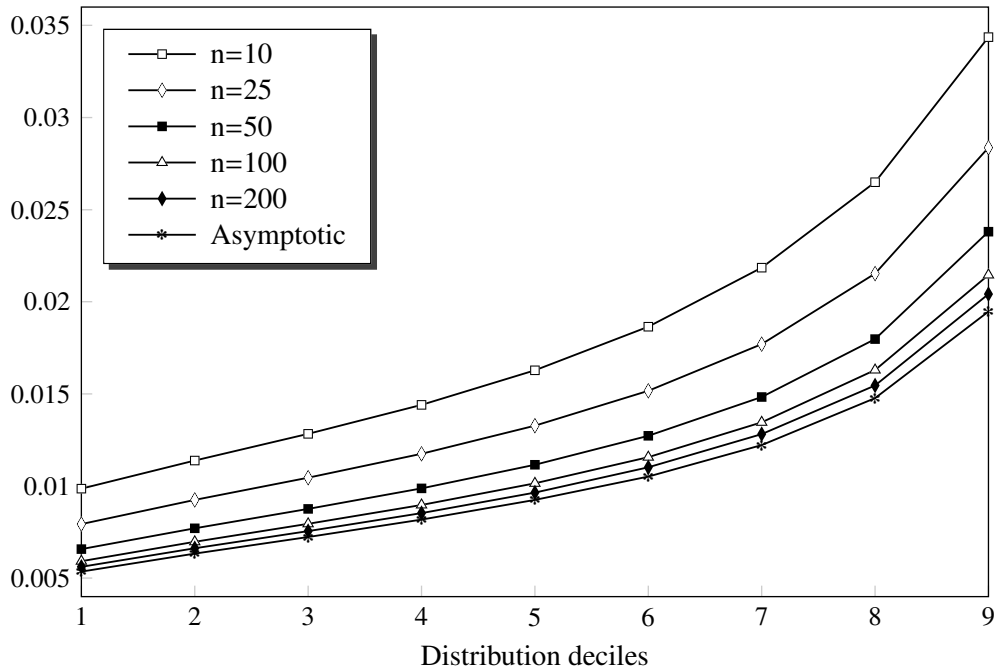


Figure 3.2: Convergence of finite sample distributions of I_n^h to asymptotic.

3.1 shows the asymptotic critical values for I_n^h in the bi-variate case.

$Prob\{x \geq x^*\} = \alpha$							
Significance Level (α)							
	0.3	0.25	0.2	0.15	0.1	0.05	0.01
x^*	0.0122	0.0134	0.0148	0.0167	0.0195	0.0246	0.0375

Table 3.1: Bi-variate asymptotic critical values for I_n^w

To characterize convergence of the finite sample distributions to the asymptotic under the null of independence, I obtain quantiles of the exact sampling distributions for sample sizes $n=15, 25, 50, 100$ and $n=200$ through simulation. Figure 3.2 shows that sampling distributions converge quickly, suggesting that the asymptotic critical values may be used even in small samples of approximately 50 to 100 observations.

If at least one of the X_j 's is a random vector, the asymptotic distribution of the test statistic depends on the marginal copula $C(u^{(j)})$, and the test is no longer pivotal. In the vector case, or when working with samples smaller than 50 observations, it is possible to obtain the critical values numerically by randomly permutating the observations in the sample. This approxi-

mates the finite sample distribution of the statistic under H_0 , conditional on the observed data. Section 3.6 describes the permutation procedure that gives the approximate p -values in detail, and Kojadinovic and Holmes (2009) establish its consistency. While it may seem that obtaining critical values by permutation is computationally costly, the procedure is in fact highly amenable to computation since exact computational formula for the integral I_n^h exists. The next section gives a closed-form expression for the test statistic.

3.5.4 Closed-form expression

Even though the proposed class of statistics is of Cramer-von Misés type, no numerical integration is required to carry out the test. The empirical copula $C_n(u)$ is a sum of products of indicator functions, and since no arguments are shared across the product terms, it is possible to express and evaluate the entire integral precisely using summation. Theorem 3.5.4 gives an exact computational formula for the weighted copula test statistic and for an arbitrary integrable and additively separable weighting function $W(u)$ and shows that the entire integral is only as difficult to evaluate as the weighting function itself.

Theorem 3.5.4 *Let $W(u) : [0, 1]^d \rightarrow \mathfrak{R}^+$ be an integrable weighting function that can be separated additively as $W(u) = \sum_{j=1}^d w(u_j)$. Then,*

$$\begin{aligned}
I_n^W &= \frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \sum_{m=1}^d \left[\prod_{j=1, j \neq m}^d (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \tilde{w}(\hat{u}_{im} \vee \hat{u}_{lm}, 1) \\
&- \frac{2}{n^p} \sum_{i=1}^n \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\
&\times \left[\sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right] \\
&+ \frac{1}{n^{2p-1}} \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\
&\times \left[\sum_{i=1}^n \sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right].
\end{aligned}$$

where $\tilde{w}_k(a, b)$ is the definite integral of $w_j(u_j)$ from a to b , and \vee is the $\max()$ operator.

Proof See Appendix A.

This makes it possible to obtain computational formulas for W -tests defined by additively-separable weighting functions without the need to re-derive the entire expression. For example, to get the closed-form for I_n^h , we only need to find the integral of the summands $(u_j - u_j^2)$ of $h(u)$ which is

$$\tilde{h}(a, b) = \int_a^b (u_j - u_j^2) du_j = \frac{1}{2}(b^2 - a^2) + \frac{1}{3}(a^3 - b^3). \quad (3.25)$$

The expression for I_n^h is easily obtained by replacing $\tilde{w}(\hat{u}_{ik} \vee \hat{u}_{lk}, 1)$ in Theorem 3.5.4 by $\tilde{h}(\hat{u}_{ik} \vee \hat{u}_{lk}, 1)$.

3.5.5 Weighted tests as a unifying class

Interestingly, the W -tests appear to act as a unifying class between the classic non-parametric tests that use the distribution function and the more recent copula-based statistics. The simplest choice of the weights is $w^1(u) = 1, \forall u \in [0, 1]^d$, which leads to the unweighted test statistic considered in Kojadinovic and Holmes (2009) and Genest and Remillard (2004). Another

interesting choice of the weights is $w'(u) = [f_1(F_1^{-1}(u_1)) \times \dots \times f_p(F_p^{-1}(u_p))]^{-1}$, where f_j denotes the probability density of X_j , and the corresponding W -test statistic is

$$I^{w'} = \int_{[0,1]^d} \left(C(u) - \prod_{j=1}^p C(u^{(j)}) \right) w'(u) du. \quad (3.26)$$

Let $u_j = F_j(X_j)$ and $X_j = F_j^{-1}(u_j)$, then $dX_j = [f_j(F_j^{-1}(u_j))]^{-1} du_j$. Then, by a change of variables,

$$I^{w'} = \int_{[0,1]^d} \left[C(u) - \prod_{j=1}^p u_j \right]^2 w'(u) du \quad (3.27)$$

$$= \int_{[0,1]^d} \left[C(u_1, \dots, u_p) - \prod_{j=1}^p u_j \right]^2 \prod_{j=1}^p [f_j(F_j^{-1}(u_j))]^{-1} du_1 \dots du_p \quad (3.28)$$

$$= \int_{\Omega_1 \times \dots \times \Omega_p} \left[F(F_1^{-1}(u^{(1)}), \dots, F_p^{-1}(u^{(p)})) - F_1(X_1)F_2(X_2) \dots F_p(X_p) \right]^2 d\mathbf{X} \quad (3.29)$$

$$= B, \quad (3.30)$$

which is the multivariate Blum, Kiefer, and Rosenblatt (1961) statistic

$$B = \int_{\Omega_1 \times \dots \times \Omega_p} \left[F(X_1, \dots, X_p) - F_1(X_1)F_2(X_2) \dots F_p(X_p) \right]^2 d\mathbf{X}, \quad (3.31)$$

and Ω_j is the support of $F_j(X_j)$. It is therefore possible to view the Blum, Kiefer, and Rosenblatt (1961) and the newer copula-based tests as members of the same class of weighted copula tests, where the difference between the two statistics is in the respective choice of a weighting function.

3.6 Power Study and Practical Implementation

To assess the performance of the proposed W -test, we construct local power curves for the scalar version of the weighted statistic I_n^h and its unweighted counterpart I_n^1 defined by the uniform weights $w^1(u) = 1, \forall u \in [0, 1]^2$, for a simulated bivariate distribution with stan-

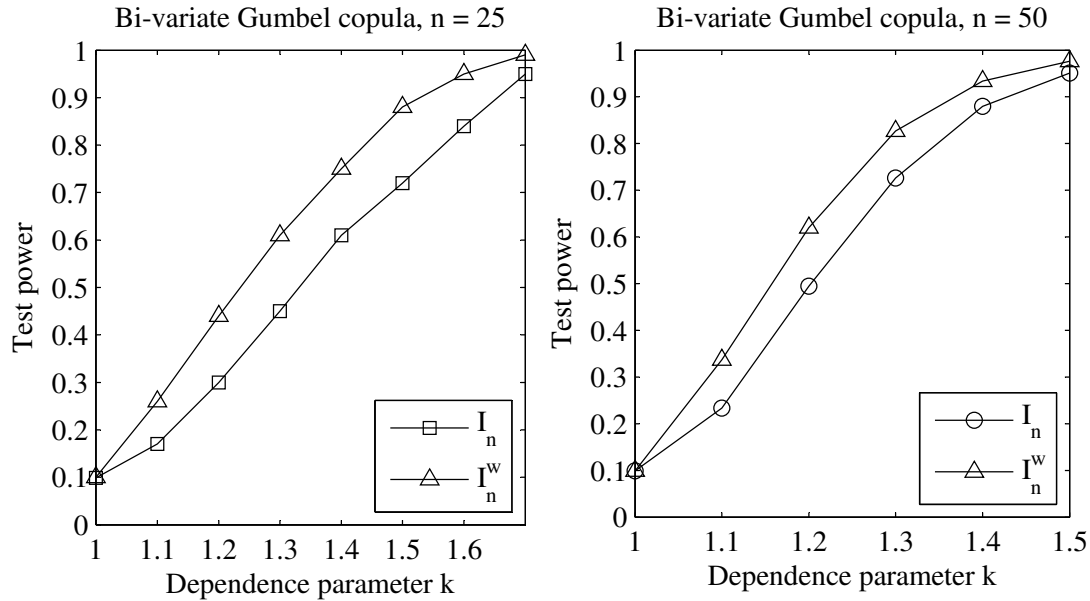


Figure 3.3: Local power of I_n^h and I_n^l .

dard normal marginals and one-parameter Gumbel copula. Note that since the test operates on uniformly distributed percentile ranks \hat{u}_i , and since in the bi-variate case the statistic is distribution-free, the power does not depend on the choice of the margins. Gumbel copula is widely used for the modelling of both linear and non-linear dependence, including tail dependence. It belongs to a class of Archimedean copulas, which Cherubini, Luciano, and Vecchiato (2004) refer to as comprehensive copulas due to their ability to accommodate the maximum range of dependence structures through different choices of parameter values. The power of the test in this simulated distribution shows its performance under a range of alternatives to independence.

Figure 3.3 shows the power of the weighted and unweighted scalar test statistics (I_n^h and I_n^l respectively) in bi-variate samples $\{(X_i, Y_i)\}_{i=1}^n$ of size $n = 25$ and $n = 50$ with a Gumbel copula.

The dependence parameter k determines the strength of the relationship between X_i and Y_i and is such that the variables are independent when $k = 1$. Test power is the percent of correct

rejections of H_0 of independence in 1,000 identically distributed samples. In each case, the same sample is used to calculate both test statistics, and the only difference between the two tests is the presence of the weighting function.

The tests are carried out at $\alpha = 10\%$ size, and we adopt the same numerical procedure as Kojadinovic and Holmes (2009) to find the approximate p -values. Given a random sample $\{(X_{i1}, \dots, X_{ip})\}_{i=1}^n$ from \mathbf{X} , approximate p -values for I_n^h can be obtained as follows:

1. Let $I_{n,0}^h$ denote the value of the test statistic found using the original sample.
2. Generate N random permutations of the indexes $i \in \{1, \dots, n\}$ of X_i and for each permutation calculate the value of the statistic $I_{n,i}^h$ using the new pseudo-independent sample.
3. Find the approximate p -value for I_n^h as

$$p = \frac{1}{N+1} \left\{ \frac{1}{2} + \sum_{i=1}^N \mathbf{1}(I_{n,0}^h \leq I_{n,i}^h) \right\}. \quad (3.32)$$

Kojadinovic and Holmes (2009) establish consistency of this resampling procedure. We verify that we obtain a good control of empirical Type I error in all simulations. Table 3.2 shows empirical Type I error rates in samples of size $n = 50$ from a distribution with Gumbel copula.

	Type I Error
Unweighted test	0.0992
Weighted test	0.0990

Table 3.2: Empirical Type I error rates.

3.7 Applications and Extensions

The test has many important applications and warrants several interesting extensions. Firstly, it can serve as a powerful tool for non-parametric data analysis since it does not require distributional assumptions and can detect any form of dependence. Given a sufficient sample, the test

will always identify a relationship between the data series if it exists. In situations where the parameters of a structural model are found to be insignificant, the test may indicate whether the lack of explanatory power is due to the absence of a statistical relationship between the regressors and the dependent variable, or due to mis-specification of the parametric model. For example, if some X appears insignificant, but the test rejects independence between X and regressand Y , the lack of significance is due to model restrictions, but not to the absence of an underlying economic relationship between Y and X . As such, when a researcher sets to empirically test some economic hypothesis, the test can be used to first determine whether or not the variables of interest are indeed related, before developing the full structural model to answer how they are related.

Secondly, many important hypotheses in econometrics reduce to the null of independence. Wu (1973) notes that tests for measurement errors, recursions in simultaneous equation models and the presence of auto-correlation in cross-section dynamic models can all be reduced to tests for independence between regressors and model disturbances. In general, the usual t and F tests are consistent only when regressors and errors are independent. Another particularly useful application of a test for independence is to the problem of copula model selection and goodness of fit. Quessy (2010) proposes a new and versatile goodness of fit procedure for Archimedean copulas using a non-parametric test for independence which is based on a particular transformation of the data. Using a result of Barbe, Genest, Ghoudi, and Remillard (1996), Quessy (2010) shows that under the null that the true copula belongs to a particular Archimedean family, data transformed using the generator function of that copula family are independent. The next two sections explore the application of the independence test to other problems such as the testing for conditional heteroscedasticity, and for the goodness of fit of Archimedean copulas.

3.7.1 Testing Copula Goodness of Fit

Suppose that the objective is to develop a parametric model of the joint distribution of a random vector \mathbf{X} . The copula approach allows the building of an extremely rich yet tractable dependence structure in high dimensions, but poses a problem of choosing the appropriate parametric family of copulas. Suppose we hypothesize that the copula of \mathbf{X} is of the Archimedean type. The Archimedean class encompasses over twenty families of best known and most widely used multivariate parametric copulas which have the capacity to accommodate the widest range of dependence structures. A unique feature of Archimedean copulas is that each family is identified by its corresponding *generator function* $\psi(u)$. Let $C(u, v)$ be some bi-variate copula of Archimedean class. It is then possible to write C in terms of the corresponding generator $\psi(u)$ as

$$C(u, v) = \psi^{[-1]}(\psi(u) + \psi(v)), \quad (3.33)$$

where $\psi^{[-1]}$ denotes the pseudo-inverse of ψ given by

$$\psi^{[-1]}(t) = \begin{cases} \psi^{-1}(t), & 0 \leq t \leq \psi(0), \\ 0, & \psi(0) \leq t \leq \infty. \end{cases} \quad (3.34)$$

For example, the generator $\phi_c(t) = \frac{1}{\theta}(t^{-\theta} - 1)$ defines the entire family of bi-variate Clayton copulas, and generator $(-\ln(t))^\theta$ defines bi-variate Gumbel copulas. The Joe, Frank, Ali-Mikhail-Haq, Gumbel-Barnett, Gumbel-Hougaard families of copulas all belong to Archimedean class and are uniquely identified by their corresponding generator functions.

Suppose that we wish to test the H_0 that a copula C of a random vector \mathbf{X} belongs to some Archimedean family of copulas $\mathcal{H} = \{C(\theta, \phi)\}$ defined by generator ϕ and indexed by parameter θ . Then, the generator $\phi(t)$ and θ uniquely identify C . Suppose $\{X_{1,i}, \dots, X_{p,i}\}_{i=1}^n$ is a random sample of \mathbf{X} , and $\{\hat{u}_{1,i}, \dots, \hat{u}_{p,i}\}_{i=1}^n$ are the corresponding percentile ranks. Further, let

$$V_{k,i} = \frac{\phi(\hat{u}_{k,i})}{\sum_{k=1}^p \phi(\hat{u}_{k,i})}, \quad (3.35)$$

for $k = 1, \dots, p, i = 1, \dots, n$ and

$$Z_i = \sum_{k=1}^p \phi(\hat{u}_{k,i}). \quad (3.36)$$

Quessy (2010) and Barbe, Genest, Ghoudi, and Remillard (1996) show that under $H_0 : C \in \mathcal{H}$, variables Z_i and $V_{k,i}$ are independent and it is possible to use I_n^h to test the fit of C to the data. To assess the power of I_n^h as a test for copula goodness of fit, we simulate 1,000 bi-variate samples of sizes $n = 25, 50, 75$ and 100 (a total of 4,000 samples) with standard normal margins and Gumbel copula with dependence parameter $\theta = 1.5$. We transform the data using generator functions for Clayton and Ali-Mikhail-Haq families of copulas and use I_n^h to test the H_0 that the true copula belongs to one of these families. Approximate p -values are obtained through permutation, and the power of the test is the percent of correct rejections of Clayton or Ali-Mikhail-Haq families in 1,000 samples. For n large enough, the probability of correct rejection approaches 1 verifying that I_n^h is consistent, and we numerically confirm the lack of bias in each case. Table 3.3 shows that I_n^h has good power properties when it is used as a test of goodness of fit for Archimedean copulas.

Copula family under H_0	Sample size			
	n=25	n=50	n=75	n=100
Clayton	20%	39%	60%	77%
Ali-Mikhail-Haq	16%	17%	35%	38%

Table 3.3: Power of the copula goodness of fit test based on I_n^h .

3.7.2 Testing independence of regressors and model errors

This section establishes the convergence of both the copula estimator and of the weighted test statistic in situations when the sample is no longer observed directly but replaced by estimated quantities, or *pseudo-observations*. The results allow the use of the counter-part to I_n^h based on estimates for the testing for independence between stochastic regressors and model errors. Proofs are given for the bi-variate case for simplicity, but it is relatively straightforward to extend these results to vectorial independence.

Consider a linear regression model $Y = X\beta + \epsilon$, $X, Y, \epsilon \in \mathbb{R}$, and let $Z = (Y, X)$. Let $\epsilon = H(Z) = Y - X\beta$. Model residuals $\hat{\epsilon} = H_n(Z) = Y - X\hat{\beta}_n$ are *pseudo-observations* of ϵ , where $\hat{\beta}_n$ is some estimate of β . We are interested to test the H_0 that X and ϵ are independent using the observed quantities $(X_i, \hat{\epsilon}_i)$, $i = 1..n$. Let K_n be the empirical distribution function of $\hat{\epsilon}_i$'s. The next lemma establishes conditions for weak convergence of the empirical process $\mathbb{K}_n(t) = \sqrt{n} \{K_n(t) - K(t)\}$, where K is the distribution function of ϵ .

Lemma 3.7.1 *Suppose that the law of ϵ admits density $k(t)$ which is continuous on the support $T \subset \mathbb{R}$, and suppose that $E[|X|] < \infty$. Then, if $\hat{\beta}_n \rightarrow \beta$, where \rightarrow denotes convergence in probability, then $\mathbb{K}_n(t)$ converges weakly in $D(T)$ to continuous stochastic process $\mathbb{K}(t)$ given by:*

$$\mathbb{K}(t) = B(K(t)) - \{E[X]\beta\} k(t) \quad (3.37)$$

where $B(t)$ is a Brownian bridge.

Proof A complete proof is given in Theorem 2.1 and Hypothesis I and II of Ghoudi and Remillard (1998).

The next lemma establishes convergence of the empirical distribution function of X and ϵ based on pseudo-observations.

Lemma 3.7.2 *Suppose that the conditions of Lemma 3.7.1 are satisfied, and let $\tilde{K}_n(x, e) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_i \leq x, \hat{\epsilon}_i \leq e)$ be the empirical joint distribution function of (X, ϵ) . Then, the empirical process $\tilde{\mathbb{K}}_n = \sqrt{n} \{ \tilde{K}_n(x, e) - \tilde{K}(x, e) \}$, where $\tilde{K}(x, e)$ is the joint distribution function of (X, ϵ) , converges weakly in $D(\bar{\mathbb{R}})$ to a continuous zero-mean Gaussian process $d(z, e)$ with bounded covariance function $\mathbb{D}(s, t)$.*

Proof See Appendix A.

Let the map $\psi(H)(u, v)$ from bivariate distribution functions H on \mathbb{R}^2 to bivariate distribution functions on $[0, 1]^2$, be defined as follows:

$$\psi(H)(u, v) = H(H^{-1}(u, \infty), H^{-1}(\infty, v)).$$

The following lemma (van der Vaart and Wellner (1996), Lemma 3.9.28), which we restate here without proof, establishes Hadamard-differentiability of $\psi(H)(u, v)$ at H tangentially to $C(\mathbb{R}^2)$, where $C(T)$ is the space of all continuous functions from T to \mathbb{R} .

Lemma 3.7.3 *Suppose that for any p, q , s.t. $0 < p < q < 1$, H is a distribution function on \mathbb{R}^2 with marginal distribution functions F and G that are continuously differentiable on the intervals $[F^{-1}(p) - \delta, F^{-1}(q) + \delta]$ and $[G^{-1}(p) - \delta, G^{-1}(q) + \delta]$ with positive derivatives f and g , for some $\delta > 0$. Furthermore, assume that $\partial H/\partial x$ and $\partial H/\partial y$ exist and are continuous on the product of these intervals. Then, the map $\psi(H)(u, v)$ is Hadamard-differentiable tangentially to $C(\mathbb{R}^2)$ and its derivative is given by:*

$$\begin{aligned} \psi'_H(a)(u, v) &= a(F^{-1}(u), G^{-1}(v)) \\ &- \frac{\partial H}{\partial x}(F^{-1}(u), G^{-1}(v)) \frac{a(F^{-1}(u), \infty)}{f(F^{-1}(u))} \\ &- \frac{\partial H}{\partial y}(F^{-1}(u), G^{-1}(v)) \frac{a(\infty, G^{-1}(v))}{g(G^{-1}(v))}. \end{aligned}$$

Let $\tilde{C}(u, v)$, $u \in [0, 1]^d$, $v \in \mathbb{R}$ denote the copula of (X, ϵ) , and its empirical counterpart based on pseudo-observations $\hat{\epsilon}$ be given by:

$$\tilde{C}_n(u, v) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(F_n^{-1}(X_i) \leq u) \mathbf{1}(G_n^{-1}(\hat{\epsilon}_i) \leq v). \quad (3.38)$$

The next theorem gives the weak limit of the empirical process based on (3.38) and establishes convergence of the empirical copula estimator $\tilde{C}_n(u, v)$ based on pseudo-observations.

Theorem 3.7.4 *Suppose that X_1, \dots, X_n are i.i.d, and $\hat{\epsilon} = Y - X\hat{\beta}$. Further, suppose that*

the hypotheses of Lemmas 3.7.1, 3.7.2 and 3.7.3 are satisfied. Then, the empirical process $\tilde{C}_n = \sqrt{n} \{ \tilde{C}_n(u, v) - \tilde{C}(u, v) \}$ converges weakly in $l^\infty([0, 1]^2)$ to a continuous stochastic process $\zeta(u, v) = \psi'_H(d)(u, v)$, where ψ'_H is given by Lemma 3.7.3, and the process $d(x, e)$ is given in Lemma 3.7.2.

Proof See Appendix A.

When the dimension of X is greater than one, independence under H_0 is characterized by $\tilde{C}(u, v) = \tilde{C}_x(u)v$, where $\tilde{C}_x(u)$ is the copula of X . The next theorem establishes the limiting properties of the empirical process \tilde{C}_n^I given by

$$\tilde{C}_n^I = \sqrt{n} \left\{ \left(\tilde{C}_n(u, v) - \tilde{C}_{n,x}(u)v \right) \sqrt{w(u, v)} \right\} \quad (3.39)$$

under H_0 , where $\tilde{C}_{n,x}(u)$ is the empirical copula of X .

Theorem 3.7.5 *Suppose that X_1, \dots, X_n are i.i.d, and $\hat{\epsilon} = Y - X\hat{\beta}$. Further, suppose that the hypotheses of Lemmas 3.7.1, 3.7.2 and 3.7.3 are satisfied, and that X and ϵ are independent. Then, empirical process \tilde{C}_n^I converges weakly in $l^\infty([0, 1]^{d+1})$ to a continuous stochastic process $\zeta(u, v)$ with bounded covariance function.*

Proof See Appendix A.

Corollary 3.7.6 *If \tilde{C} has continuous partial derivatives, the statistic given by*

$$\tilde{I}_n^w = n \int_{[0,1]^{d+1}} \left(\tilde{C}_n(u, v) - \tilde{C}_{n,x}(u)v \right)^2 w(u, v) dudv$$

converges in distribution to

$$\int_{[0,1]^{d+1}} \left(\phi'_{\tilde{C}_n}(\zeta)(t) \right)^2 dt$$

Proof The result follows by Theorem 3.7.5 and the continuous mapping theorem.

As with I_n^w , the asymptotic distribution of \tilde{I}_n^w does not depend on the distribution of X when it is scalar, and it is possible to tabulate the asymptotic critical values. In vectorial or small sample case, we use the same permutation procedure as in Section 3.6 to obtain approximate p -values by conditioning on the observed sample. While we do not formally establish consistency of the resampling procedure in this case, we numerically verify consistency and the lack of bias for \tilde{I}_n^w in all subsequent power studies when using the approximate p -values.

3.7.3 Copula-based test for conditional heteroscedasticity

This section illustrates an application of \tilde{I}_n^w to the testing for conditional heteroscedasticity by re-writing the H_0 of homoscedasticity as the null of independence between the model disturbances and regressors. Consider a simple linear regression model $Y_i = \alpha_0 + \beta_0 X_i + \epsilon_i$. Conditional heteroscedasticity is present when the variance of the error terms in the model depends on the values on X_i : $E[\epsilon^2|X_i] = \sigma_\epsilon^2(X_i)$. The results of Theorem 3.7.5 and of Corollary 3.7.6 allow testing for the presence of heteroscedasticity in model errors by operating on the empirical copula of model regressors X and the fitted model residuals $\hat{\epsilon}$. Under the null of homoscedasticity, we have that $\beta_0 = 0$, in which case X and ϵ^2 are independent and we can use \tilde{I}_n^w to test the $H_0 : X \perp \epsilon^2$.

We run a series of simulations to assess the local power of the copula-based test and compare the performance to that of the classical test of White (1980). To obtain local power in small samples, we generate two sets of bivariate samples (X, Y) of size $n = 25$ such that $X_i \sim N(\mu_x, \sigma_x^2)$, $\epsilon_i \sim N(0, \sigma_\epsilon^2(X_i))$ and $Y_i = \alpha + \beta X_i + \epsilon_i$. In the first set, linear heteroscedasticity arises from conditioning the error variance on X_i according to $\sigma_\epsilon^2(X_i) = \alpha_0 + \beta_0 X_i$. In the second set, we simulate non-linear heteroscedasticity of the form $\sigma_\epsilon^2(X_i) = \alpha_0 + \beta_0 X_i^2$. Parameter β_0 controls the degree of heteroscedasticity in the data such that setting $\beta_0 = 0$ leads to homoscedastic errors, and $\beta_0 > 0$ to increasing degree of dependence. Copula-based test-statistic \tilde{I}_n^w for $H_0 : \beta_0 = 0$ is calculated using the squared fitted OLS model residuals $\hat{\epsilon}^2$ and the X 's and the critical values are obtained by random permutation, conditional on the observed sample, as

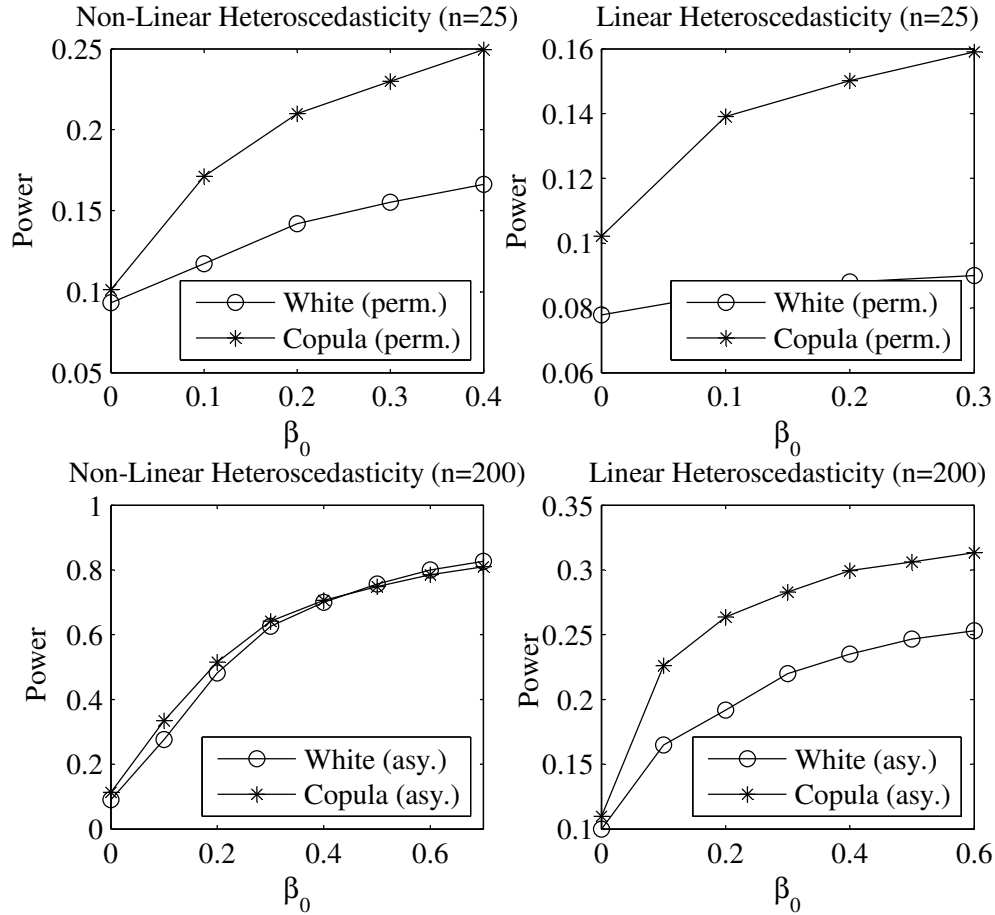


Figure 3.4: Local power of \tilde{I}_n^h and of the White test.

in the i.i.d. case. The White (1980) test serves as the benchmark for the copula-based test statistic. Testing for $H_0 : \beta_0 = 0$ is done in the usual way by first obtaining fitted residuals $\hat{\epsilon}_i$ from the regression of Y on X , and then by estimating the following equation: $\hat{\epsilon}_i^2 = \gamma_0 + \gamma_1 X_i + \gamma_2 X_i^2 + u_i$. The test statistic is then nR^2 , which is $\chi^2(2)$ asymptotically. Since the sample size is small, we use permutation to obtain approximate p -values for \tilde{I}_n^w , and to ensure comparability of the results across the two tests, we carry out an alternative version of the White test using the permutation approach to obtain approximate p -values from the finite-sample distribution of nR^2 instead of the asymptotic critical values.

Since test performance is likely to change in larger samples, we conduct a separate simula-

tion study to assess test power in linear and non-linear heteroscedasticity settings for samples of size $n = 200$, now using asymptotic critical values to conduct both tests. We adopt the following parametrization in all simulations in this section: $\{\mu_x = 5, \sigma_x^2 = 1, \alpha = 0.5, \beta = 3, \alpha_0 = 1\}$.

Figure 3.4 shows the power of the tests at detecting both linear and non-linear heteroscedasticity for values of β_0 ranging from 0 (homoscedastic samples) to 0.6. The copula-based test has a significant power advantage over both the asymptotic and the finite-sample versions of the White (1980) test, except in the case of non-linear heteroscedasticity in samples of $n = 200$, where the power of the copula test matches that of the White (1980). Note that in all simulations all tests maintain their nominal size, and hence the power advantage does not arise from bias due to over-rejection. Tests are also consistent, and achieve full power for all tried values of $\beta_0 > 0$ in large enough samples.

3.7.4 Residual tests using parametric copula models

The gain in power is non-trivial and surprising, given the non-parametric nature of the proposed test. It seems that operating on the copula in principal can lead to power advantages over current residual testing methods. Since, to my knowledge, copula residual-based testing of regression assumptions was not addressed in the literature prior to this work, this calls for further research into residual testing which utilizes copula theory. For example, it is straightforward to construct a parametric copula-based test for independence of regressors and disturbances by selecting a known copula family, in which case certain parameter values characterize independence uniquely. Depending on the estimation procedure, this may lead to a consistent and pivotal test, both asymptotically and in finite samples.

3.8 Summary and Concluding Remarks

This chapter proposes a new class of non-parametric tests for stochastic independence which are based on the weighted empirical copula process and have a power advantage over compara-

ble non-parametric methods. Unlike the usual correlation or regression-based approaches, the proposed tests can detect dependence of any form, including non-linear or asymmetric dependence, at any point in the distribution. The tests have attractive finite sample properties, and are amenable to computation. In particular, this chapter shows that the asymptotic critical values may be used in samples as small as 50 observations, and when using asymptotics is impossible or undesirable, approximate p -values can be obtained numerically using a straightforward permutation procedure. Exact computational formula for the test statistic is also provided.

The contribution of this chapter is two-fold. Firstly, it is the first to consider a weighted empirical process in the context of non-parametric copula tests for independence. The weighting function serves as an index for a class of weighted test statistics proposed in this chapter and aside from the power advantage, has an added benefit of unifying the classical non-parametric tests for independence based on distribution function with the newer copula testing literature. In particular, the distribution function-based statistic of Blum, Kiefer, and Rosenblatt (1961) and the copula-based statistic of Kojadinovic and Holmes (2009) and Genest and Remillard (2004) are shown to be members of the same class of the weighted copula tests proposed here, where each statistic can be obtained as a special case by choosing the appropriate weighting function.

Secondly, this chapter establishes convergence of the copula estimator and of the weighted test statistics in situations when estimated quantities are used instead of direct observations. This permits testing for independence between unobserved random variables, such as regression model errors. Many important tests in econometrics can be reduced to tests for independence, particularly between regression model errors and the regressors. In an illustrative application of these results, we use the proposed statistic to test for conditional heteroscedasticity in a linear regression model using the empirical copula based on model residuals and the regressors. We find that the copula test has power advantage over the test of White (1980), especially in small samples. This is the first application of copulas to testing for heteroscedasticity, and the results suggest that more powerful parametric tests may be possible that use the

copula of the residuals. While in this chapter the regression residuals are used as examples of pseudo-observations, results extend to other estimated quantities, such as realized volatility or model forecasts and are therefore applicable to other problems.

Another useful application of the tests is to the problem of model selection and goodness of fit for Archimedian copulas. Under the null that the copula function belongs to a particular family, transforming the data using the corresponding generator function yields independent series. This chapter shows that the proposed statistic has good power properties when used as a test of copula goodness of fit.

This chapter warrants several interesting extensions. First, it would be useful to extend the proposed test to the case of conditional independence. Patton (2006) introduces the concept of conditional copula, and it is fairly straightforward to modify the proposed statistic so that it can be used to test for independence of random vectors conditional on some probabilistic statement.

The issue of existence (or non-existence) of the optimal weighting function, and of contextual choice of weights is also left unexplored. While the results provided here hold for an arbitrary integral and bounded weighting function, it would be useful to study the problem of choosing the weights in more detail.

Another interesting extension is to testing for serial dependence. The auto-correlation function is widely used for non-parametric analysis of time series and model selection, but it is only capable to detecting linear serial dependence. The weighted test statistic can be extended to the serial case in a manner similar to Kojadinovic and Yan (2009) and provide a "blanket" test of serial dependence. This direction is pursued in Chapter 4.

Chapter 4

A Generalized Measure of Serial Dependence

4.1 Introduction

Non-parametric testing for independence received considerable attention in the statistical and econometric literature. Among the most well-known tests are the statistics based on the empirical distribution function of Hoeffding (1948) and Blum, Kiefer, and Rosenblatt (1961), with the time-series extensions provided in Skaug and Tjostheim (1993) and Delgado (1996). These tests are powerful and omnibus in the sense that they are capable of detecting departures from independence of any form with minimal assumptions, but are generally not applicable to the analysis of vector time series.

To highlight the difficulty associated with the testing for vectorial serial independence, consider a stationary ergodic sequence of continuous random vectors $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots$. For every i , let $X_i^{(1)}, \dots, X_i^{(q)}$ denote the q components of $\mathbf{X}_i \in \mathbb{R}^q$. Given the *embedding dimension* $p > 0$, the aim is to test the H_0 that vectors $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$ are independent. Since the pairwise independence does not generally imply joint independence, we cannot test this H_0 by performing a series of bivariate tests for independence of $X_i^{(j)}$ from $X_i^{(k)}$ and of $X_j^{(i)}$ from $X_k^{(i)}$, $j \neq k$. In other

words, the serial independence of the univariate series $X_1^{(j)}, X_2^{(j)}, X_3^{(j)}, \dots$, for $j = 1..q$, together with independence of scalars $X_i^{(1)}, \dots, X_i^{(q)}$ within \mathbf{X}_i , generally does not imply the lack of serial dependence of the vector series $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots$

The non-parametric tools available for the testing of such multivariate serial dependence are much less developed. Earlier exceptions to this are the rank-based multivariate Wald-Wolfowitz test of Hallin and Puri (1995), and the portmanteau test of Hosking (1980). Recently, a new class of powerful tests for independence, including a test for vectorial serial independence has emerged, which stem from the theory of copula functions. It is now widely recognized that the so-called copula of the distribution provides a complete, and in the case of continuous random variables, a unique description of the dependence structure of the data. Following the seminal work of Deheuvels (1981) and Genest and Remillard (2004), Kojadinovic and Holmes (2009) propose a test for independence of random vectors which is based on the Möbius decomposition of the underlying empirical copula process. The test is further extended to the multivariate serial case by Kojadinovic and Yan (2009).

This chapter proposes a class of copula-based tests for serial independence of vector time series which are closely related to the Cramér-von Mises statistic of Kojadinovic and Yan (2009), but are specified using an alternative weighted functional norm. The Kojadinovic and Yan (2009) statistic emerges as a special member of this new class, defined by the uniform weighting of the data. Selecting the alternative weighting produces other statistics within this class with possibly different local power properties. We show the convergence of these statistics under general assumptions, and provide a computational formula. While, except in the simplest case, the statistic is not distribution-free, the approximate p -values can be found using a simple and computationally efficient permutation procedure.

The addition of the weighting function permits the unequal weighting of the deviations of the empirical copula from independence, and is beneficial in several ways. Firstly, we show that the alternative weighting can give the statistic a power advantage over the uniform weights. This power gain depends on the specific choice of the weights and on the underlying structure

of the dependence, suggesting that the issues surrounding the choice of the weights, such as the question of existence or non-existence of the optimal weighting function, or of the contextual choice of the weights, should be investigated further. Secondly, it may be possible to learn additional information about the nature of the dependence through the different choices of the weights. In this context, the weighting function can be thought of as a tuning parameter through which the local power properties of the statistic could be varied, and made such that the test is most powerful when particular type of dependence is present, such as, for example, dependence through the tails. Lastly, the ability to adjust the power through the weights can be useful in other applications of the statistic. Following the work of Barbe, Genest, Ghouli, and Remillard (1996), Quessy (2010) showed how a test for independence can be reformulated as a test for the goodness of fit of Archimedian copulas. Section 4.3.2 generalizes the results of Quessy (2010) to the serial case. The addition of the weights provides a channel for controlling the power of this goodness of fit statistic, which is useful in practical applications, such as, for example, portfolio risk management, where the weights permit the modelling of asymmetric losses. Copula theory is widely used for the modelling of the aggregate risks of a mixture of interdependent risky assets. From the practitioner's standpoint, the estimate of such aggregate risk is of prime importance, but is subject to several sources of uncertainty, one of which is the risk arising from the possible misspecification of the copula model which describes the interdependence between the assets in a portfolio and therefore aggregates the risks. The practitioner's preferences over the sources of such modelling error are highly asymmetric. When the model misspecification is due to the failure of the model to correctly capture the clustering of the gains, the consequences of such error are not as severe as when the clustering of losses is modelled incorrectly. To penalize the wrongful modelling of the losses more than of the gains, the portfolio design literature often resorts to the utility theory by explicitly specifying the practitioners preferences. The addition of a weighting function creates another channel for the imposition of asymmetric penalties for the misfitting of the certain parts of the distribution, which is highly tractable.

In the third chapter of this dissertation, as in the other copula-based independence-testing literature, the values of the independence test statistics found using different samples are generally not comparable. This is due to the absence of a known upper bound for these statistics, meaning that they cannot be used as meaningful measures of vectorial dependence. This chapter calculates such upper bound, which is subject to the empirical marginal copulas of the vectors involved in the test. The resulting standardized version of the test statistic can serve as a meaningful measure of dependence *between* random vectors. The proposed measure is closely related to the copula-based dependence measure of Gaißer, Ruppert, and Schmid (2010), which captures dependence *within* a random vector. We derive a computational formula for the new standardized statistic. Such derivation was previously advocated in Kojadinovic and Yan (2009).

The rest of this chapter is organized as follows. Section 4.2 introduces the weighted test statistic and investigates its convergence properties. The upper bound on the statistic subject to empirical marginal copulas is derived in Section 4.2.6. Section 4.3 reviews the various applications of the statistic, including the generalization of the Quesy (2010) copula goodness of fit procedure to the serial copula case. Section 4.4 concludes. Proofs are presented in the Appendix A.

4.2 The Test Statistic

In this section we adopt the notation of Kojadinovic and Yan (2009). Consider a stationary ergodic sequence of q -dimensional continuous random vectors $\mathbf{X}_1, \mathbf{X}_2, \dots$. Let $F(\mathbf{x})$ be the joint c.d.f. of \mathbf{X}_i , and $F^{(q)}(X^{(q)})$ be the marginals, where $X^{(q)}$ is the q -th component of vector \mathbf{X}_i . For any integer $k > 0$, denote the set $\{1, \dots, k\}$ by $[k]$, with the convention that $[0] = \emptyset$, and let $p > 1$ be the *embedding dimension*. Letting $n' = n + p - 1$, the empirical distribution function of \mathbf{X}_i is given by

$$F_n(\mathbf{x}) = \frac{1}{n'} \sum_{i=1}^{n'} \prod_{j=1}^q \mathbf{1}(X_i^{(q)} \leq x^{(q)}), \mathbf{x} \in \mathbb{R}^q, \quad (4.1)$$

with the corresponding empirical marginal c.d.f.'s denoted by $F_n((x^{(1)}), \dots, F_n(x^{(q)}))$, and $\mathbf{1}(\cdot)$ is the indicator function. Given $X_i^{(j)}$, $i \in [n']$, $j \in [q]$, the empirical ranks are $R_i^{(j)} = n' F_n^{(j)}(X_i^{(j)})$.

4.2.1 Serial Copula

Letting $\mathbf{Y}_i = (\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$, denote its distribution function by $H(\mathbf{Y})$. It is now well understood that the dependence structure of \mathbf{Y} is completely and uniquely characterized by its copula C . Following a result by Sklar (1959), this joint distribution function can be written in terms of the marginals, and the serial copula C^s as $H(\mathbf{Y}) = C^s(F(\mathbf{X}_1), \dots, F(\mathbf{X}_p))$. Vectors $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$ are independent if, and only if, the copula C coincides with the independence copula C_\perp^s .

To further characterise the H_0 of independence in terms of the underlying copula, additional notation is needed. Given some $B \subseteq [p]$ and $\mathbf{u} \in [0, 1]^{pq}$, let $\mathbf{u}_{\{B\}} \in [0, 1]^{pq}$ be given by

$$\mathbf{u}_{\{B\}}^{(i)} = \begin{cases} u^{(i)} & \text{if } i \in \cup_{j \in B} \{(j-1)q + 1, \dots, jq\}, \\ 1 & \text{if otherwise.} \end{cases} \quad (4.2)$$

Furthermore, given $\mathbf{u} \in [0, 1]^{pq}$ and $i \in [p]$, define a sub-vector $\mathbf{u}_{\langle j \rangle} \in [0, 1]^{pq}$ of u as $u_{\langle j \rangle}^{(i)} = u^{i+(j-1)q}$, $j \in [q]$. We then have that under the H_0 of independence,

$$\left(C^s(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{\langle j \rangle}) \right) = 0, \forall \mathbf{u} \in [0, 1]^{pq}, \quad (4.3)$$

where $C^s(\mathbf{u}_{\langle j \rangle})$, for $j \in [p]$, are the marginal copulas of pq -dimensional vectors $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$, for $i \in [n]$, and $\prod_{j=1}^p C^s(\mathbf{u}_{\langle j \rangle})$ is the independence copula C_\perp . Given a sample of pq -dimensional observations Y_i , $i = 1, \dots, n$, we can estimate the serial copula $C^s(\mathbf{u})$ by

$$C_n^s(\mathbf{u}) = \frac{1}{n} \sum_{i=1}^n \prod_{j=1}^p \prod_{k=1}^q \mathbf{1}(F_n^{(k)}(X_{i+j-1}^k) \leq u_{\langle j \rangle}^{(k)}) = \frac{1}{n} \sum_{i=1}^n \prod_{j=1}^p \prod_{k=1}^q \mathbf{1}(R_{i+j-1}^{(k)} \leq n' u_{\langle j \rangle}^{(k)}), \quad (4.4)$$

for any $\mathbf{u} \in [0, 1]^{pq}$. This gives rise to the multivariate serial empirical copula process

$$C_n(\mathbf{u}) = \sqrt{n} \left(C_n^s(\mathbf{u}) - \prod_{j=1}^p C_n^s(\mathbf{u}_{(j)}) \right), \mathbf{u} \in [0, 1]^{pq}, \quad (4.5)$$

which serves as the basis for the test statistic considered in Section 2.4 of Kojadinovic and Yan (2009).

4.2.2 Weak convergence of $C_n(\mathbf{u})$ under the H_0 of independence

This section establishes the convergence of the process (4.5) when $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$ are independent. Let $l^\infty([0, 1]^{pq})$ be the space of all bounded real-valued functions on $[0, 1]^{pq}$. Theorem 4.2.1 from Kojadinovic and Yan (2009), which we restate here without proof, establishes the convergence of the serial copula estimator $C_n^s(\mathbf{u})$ under the H_0 .

Theorem 4.2.1 *Suppose that C has continuous partial derivatives. Then, under serial independence of $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$, the process $\sqrt{n} \left(C_n^s(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{\langle j \rangle}) \right), \mathbf{u} \in [0, 1]^{pq}$, converges weakly in $l^\infty([0, 1]^{pq})$ to the tight centred Gaussian process*

$$\mathbb{C}(\mathbf{u}) = \mathbb{H}(\mathbf{u}) - \sum_{i=1}^p \prod_{k=1, k \neq i}^p C^s(\mathbf{u}_{\langle k \rangle}) \sum_{j=1}^q \partial_j C^s(\mathbf{u}_{\langle i \rangle}) \mathbb{H}(1, \dots, 1, u_{\langle i \rangle}^{(j)}, 1, \dots, 1), \mathbf{u} \in [0, 1]^{pq}, \quad (4.6)$$

where $\partial_j C^s$ is the j 'th partial derivative of C^s , and $\mathbb{H}(\mathbf{u}), \mathbf{u} \in [0, 1]^{pq}$, is a tight centered Gaussian process with covariance function

$$E[\mathbb{H}(\mathbf{u})\mathbb{H}(\mathbf{u}')] = \gamma_0(\mathbf{u}, \mathbf{v}) + \sum_{\delta=1}^{p-2} [\gamma_\delta(\mathbf{u}, \mathbf{v}) + \gamma_\delta(\mathbf{v}, \mathbf{u})], \mathbf{u}, \mathbf{v} \in [0, 1]^{pq}, \quad (4.7)$$

where

$$\gamma_\delta(\mathbf{u}, \mathbf{v}) = \prod_{j \in [p] \setminus [p-\delta]} C^S(\mathbf{u}_{<j>}) \prod_{j \in [\delta]} C^S(\mathbf{v}_{<j>}) \left[\prod_{j \in [p-\delta]} C^S(\mathbf{u}_{<j>} \wedge \mathbf{v}_{<j+\delta>}) \right. \quad (4.8)$$

$$- \sum_{i \in [p-\delta]} C^S(\mathbf{u}_{<i>} \wedge \mathbf{v}_{<i+\delta>}) \prod_{j \in [p-\delta] \setminus \{i\}} C^S(\mathbf{u}_{<j>}) C^S(\mathbf{v}_{<j+\delta>}) \quad (4.9)$$

$$\left. + (p - \delta - 1) \prod_{j \in [p-\delta]} C^S(\mathbf{u}_{<j>}) C^S(\mathbf{v}_{<j+\delta>}) \right]. \quad (4.10)$$

Theorem 4.2.1 establishes the consistency of the estimator $C_n^s(\mathbf{u})$ of the underlying copula C^s of \mathbf{Y}_i under the null. By applying the functional Delta-method of van der Vaart and Wellner (1996) (Theorem 3.9.4), Theorem 2 of Kojadinovic and Yan (2009), stated here without proof, establishes convergence of the serial independence copula process $C_n(\mathbf{u})$:

Theorem 4.2.2 *Suppose that the copula C is continuously-differentiable. Then, under the H_0 of serial independence, the empirical process $\sqrt{n} \left(C_n^s(\mathbf{u}) - \prod_{j=1}^p C_n^s(\mathbf{u}_{(j)}) \right)$, $\mathbf{u} \in [0, 1]^{pq}$, converges weakly in $l^\infty([0, 1]^{pq})$ to a tight centered Gaussian process*

$$\mathbb{C}^I(\mathbf{u}) = \phi'(\mathbb{C})(\mathbf{u}) = \mathbb{H}(\mathbf{u}) - \sum_{k=1}^p \mathbb{H}(\mathbf{u}_{\{k\}}) \prod_{j=1, j \neq k}^p C^S(\mathbf{u}_{<j>}), \mathbf{u} \in [0, 1]^{pq}, \quad (4.11)$$

with the covariance function given by

$$E[\mathbb{C}^I(\mathbf{u})\mathbb{C}^I(\mathbf{v})] = E[\mathbb{H}(\mathbf{u})\mathbb{H}(\mathbf{v})] = \gamma_0(\mathbf{u}, \mathbf{v}) + \sum_{\delta=1}^{p-2} [\gamma_\delta(\mathbf{u}, \mathbf{v}) + \gamma_\delta(\mathbf{v}, \mathbf{u})], \mathbf{u}, \mathbf{v} \in [0, 1]^{pq}, \quad (4.12)$$

where γ_δ is defined in the same way as in Theorem 4.2.1.

While a Cramér-von Mises statistic based on (4.5) is briefly considered in Kojadinovic and Yan (2009), their focus is on a set of statistics derived from the further decomposition of the empirical process $C_n(\mathbf{u})$ into independent sub-processes.

4.2.3 The Weighted Test Statistic

This paper proposes an extension to the Cramér-von Mises statistic of Kojadinovic and Yan (2009) based on $C_n(\mathbf{u})$, obtained through the use of a weighted functional norm. Let $\mathcal{G}([0, 1]^{pq}) = \{g \in L([0, 1]^{pq})\}$ be a set of bounded real-valued functions on $[0, 1]^{pq}$ that are also positive, continuous and integrable. The focus in this paper is on a class of Cramér-von Mises statistics indexed by a weighting function $g(\mathbf{u}) \in \mathcal{G}([0, 1]^{pq})$:

$$I_n(g) = n \int_{[0,1]^{pq}} \left(C_n^s(\mathbf{u}) - \prod_{k=1}^p C_n^s(\mathbf{u}_{\{k\}}) \right)^2 g(\mathbf{u}) d\mathbf{u}, \mathbf{u} \in [0, 1]^{pq}. \quad (4.13)$$

Let $\mathcal{H}(\mathcal{G})$ be the set of all test-statistics given by (4.13). The statistic of Kojadinovic and Yan (2009) can be viewed as a special element of this set, defined by uniform weights $g^1(\mathbf{u}) = 1, \forall \mathbf{u} \in [0, 1]^{pq}$. The next theorem establishes the convergence of $I_n(g)$ under the null, for an arbitrary continuous and integrable weighting function.

Theorem 4.2.3 *Suppose that the copula C is continuously-differentiable. Then, for any $g(\mathbf{u}) \in \mathcal{G}$, the test statistic $I_n(g)(\mathbf{u})$ converges in distribution to $I(g)(\mathbf{u})$ given by*

$$I(g) = \int_{[0,1]^{pq}} \left(\mathbb{H}(\mathbf{u}) - \sum_{k=1}^p \mathbb{H}(\mathbf{u}_{\{k\}}) \prod_{j=1, j \neq k}^p C^s(\mathbf{u}_{\langle j \rangle}) \right)^2 g(\mathbf{u}) d\mathbf{u}, \quad (4.14)$$

where $\mathbb{H}(\mathbf{u})$ is a tight centered Gaussian process as in Theorem 4.2.1.

Proof: Appendix A.

Note that the restrictions imposed on g ensure that $E[I(g)(\mathbf{u})I(g)(\mathbf{v})] < \infty, \forall \mathbf{u}, \mathbf{v} \in [0, 1]^{pq}$. Also note that since the limiting distribution depends on the marginal copulas $C^s(\mathbf{u}_{\langle j \rangle}), j \in [p]$, the statistic is not distribution-free.

4.2.4 The Weighting Function

The addition of a weighting function permits the unequal weighting of the deviations of the empirical copula from the independence, which is useful in several ways. Firstly, in Section

4.2.9 we show that the local power properties of the test can be influenced by the choice of the weights. In particular, we document the ability of the non-uniform weights to give the statistic a power advantage. The weighting function acts as a tuning parameter in this context, and could possibly be used to gain additional insights into the specific nature of the dependence. For example, we may choose $g(\mathbf{u})$ such that the resulting statistic has power properties which favour a particular form of dependence such as dependence through the tails, or through a certain quadrant of the distribution. Lastly, Section 4.3 shows how the statistic can be used to conduct goodness of fit tests for Archimedean serial copulas. As a result, the power of the goodness of fit test can also be influenced through the different choices of the weights. This can be useful for some applications, particularly risk management, where it may be important to fit certain parts of the distribution especially well. Section 4.3.3 further shows how the inverted statistic can be used as an interval estimator of the serial copula parameters. The weighting function $g(\mathbf{u})$ could be used in a similar way to obtain better interval estimates, or to make estimation more sensitive to the deviations from the H_0 at certain parts of the distribution.

4.2.5 The Dependence Structure of Some Common Time-Series

This section investigates the nature of the dependence in some of the common time series. Note that when the embedding dimension is $p = 2$, we can estimate and visualize the distribution of the differences between the serial copula of the data and the independence copula. This distance $d_n(\mathbf{u})$ is given by

$$d_n(\mathbf{u}) = \left(C_n^s(\mathbf{u}) - \prod_{j=1}^p C_n^s(\mathbf{u}_{(j)}) \right)^2, \mathbf{u} \in [0, 1]^{pq}, \quad (4.15)$$

which can be thought of as the *serial dependence map* of the data. For any $\tilde{\mathbf{u}} \in [0, 1]^{pq}$, larger values of $g_n(\tilde{\mathbf{u}})$ indicate a stronger tendency of the data to cluster together at that particular point in the distribution. For example, for univariate time-series, observing that $g(0.5, 0.5) > g(0.1, 0.1)$ suggests that the dependence between observations around the median is stronger

than in the lower left tail.

We obtain empirical dependence maps $g_n(\mathbf{u})$ for several popular time-series from economics and nature, as well as for a simulated $AR(1)$ process. The time-series used are the (i) annual trappings of the Canadian Lynx between 1821-1934, which have been analysed extensively, (ii) a simulated univariate $AR(1)$ process with a positive autocorrelation, (iii) daily changes in the spot price for Brent crude, and (iv) quarterly changes in the Canadian real gross domestic product between 1961 and 2011.

This chapter seems to be the first to address the issue of weighting in the context of copula-based tests for serial independence. No attempt is made here to search for a weighting function that is optimal. The choice of weights, however, is not random, but is motivated by the findings and the discussion in this section.

Figure 4.1 shows the corresponding maps of the dependence between realizations of the series and first lags. Darker shading indicates larger values of $g(\mathbf{u})$, and hence a stronger tendency of the data to cluster. Contours are smoothed for plotting using a Gaussian filter. Firstly, note that some series appear to show evidence of asymmetric dependence. For example, unlike the median-centric map for the $AR(1)$ process, the dependence map for the Canadian Lynx series is shifted away from the median towards the upper-left tail, suggesting that population changes can be predicted with higher degrees of accuracy in the years when population growth was below the median. Secondly, for all of the series, the deviations from the independence copula $g(\mathbf{u})$ are greatest around the median, and fade towards the tails.

The focus in the rest of this chapter is on the test statistic $I_n(w) \in \mathcal{G}$, which we define by choosing the following weights:

$$w(\mathbf{u}) = \sum_{i=1}^{pq} (u^{(i)} - (u^{(i)})^2), \mathbf{u} \in [0, 1]^{pq}. \quad (4.16)$$

The weighting function $w(\mathbf{u})$ is an inverted elliptic paraboloid, where the maximum is attained at a point $\bar{\mathbf{u}} = (0.5, \dots, 0.5)$, which is the median of the distribution of Y_i . The function $w(\mathbf{u})$

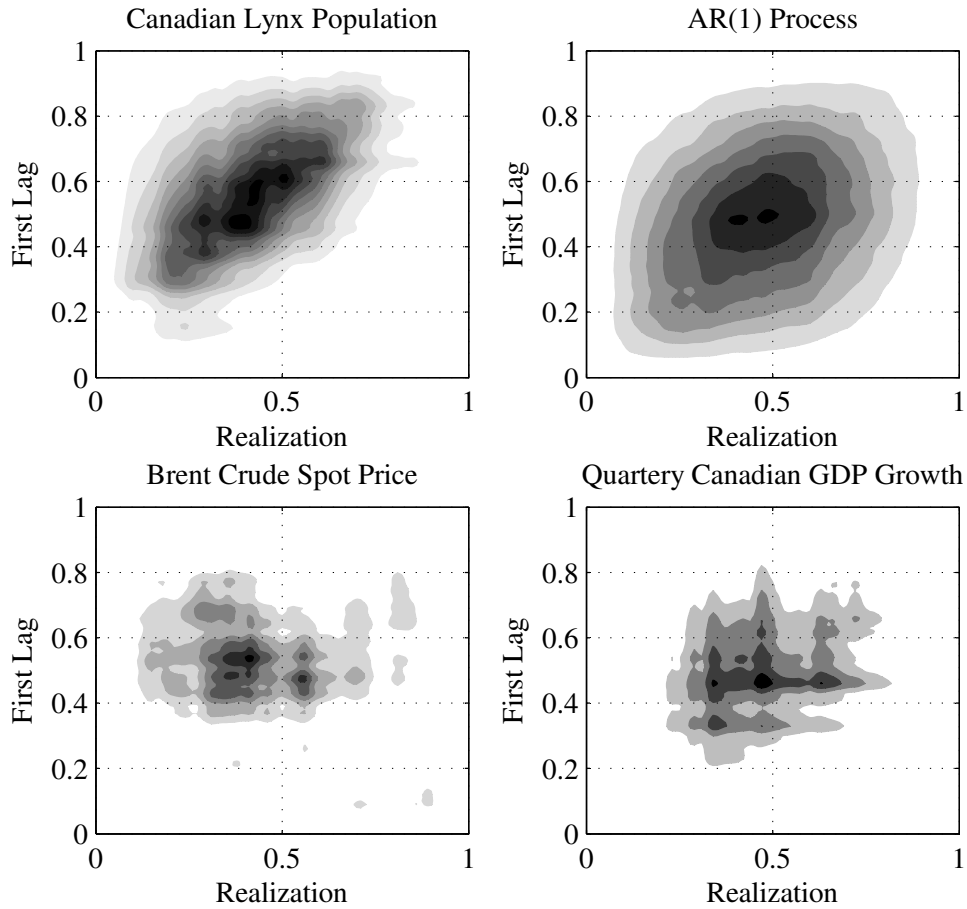


Figure 4.1: Empirical Dependence Maps.

places progressively larger weights on the observations which are closer to the median, emphasizing the parts of the distribution where the clustering of the data is expected to be strongest.

4.2.6 Normalized Statistic

Regardless of the choice of the weights, the proposed statistic cannot be used as a measure of dependence, since the values of $I_n(g)$ obtained during separate tests are generally not comparable. In this section we aim to establish an upper bound on the variable $I(g)$ and develop a normalized version of the statistic. The bound on $I(g)$ is subject to marginal copu-

las $C^s(\mathbf{u}_{<j>})$, and therefore depends on the distribution of \mathbf{Y}_i . This issue is closely related to the so-called *Fréchet Problem* studied extensively in the statistical literature. In our case, we use the fact that any copula must lie within the so-called Fréchet-Hoeffding bounds $W(\mathbf{u})$ and $M(\mathbf{u})$, which themselves are copulas. In particular, for any pq -copula $C^s(\mathbf{u})$, it must be that $W(\mathbf{u}) \leq C^s(\mathbf{u}) \leq M(\mathbf{u})$, $\forall \mathbf{u} \in [0, 1]^{pq}$ (see Section 2.2 in Nelsen (2006) for additional details on the Fréchet-Hoeffding copula bounds). The inequality must also hold for the independence copula $C^\perp(\mathbf{u}) = \prod_{j=1}^p C^s(\mathbf{u}_{<j>})$. Combining these two observations, we then have that

$$(C^s(\mathbf{u}) - C^\perp(\mathbf{u}))^2 \leq \max\left((W(\mathbf{u}) - C^\perp(\mathbf{u}))^2, (M(\mathbf{u}) - C^\perp(\mathbf{u}))^2\right), \mathbf{u} \in [0, 1]^{pq}. \quad (4.17)$$

Then, for any $g(\mathbf{u})$, it must be that

$$\int_{[0,1]^{pq}} \left(C^s(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{<j>}) \right)^2 g(\mathbf{u}) d\mathbf{u} \leq \quad (4.18)$$

$$\int_{\Omega_1} \left(W(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{<j>}) \right)^2 g(\mathbf{u}) d\mathbf{u} + \int_{\Omega_2} \left(M(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{<j>}) \right)^2 g(\mathbf{u}) d\mathbf{u}, \quad (4.19)$$

where $\Omega_1 = \{\mathbf{u} \in [0, 1]^{pq} : (W(\mathbf{u}) - C^\perp(\mathbf{u}))^2 \geq (M(\mathbf{u}) - C^\perp(\mathbf{u}))^2\}$, and $\Omega_2 = [0, 1]^{pq} \setminus \Omega_1$. If the weighting function $g(\mathbf{u})$ is symmetric around the median, the right-hand side of the inequality (4.18) is less than

$$D(C^s) = \int_{[0,1]^d} \left(M(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{<j>}) \right)^2 g(\mathbf{u}) d\mathbf{u}. \quad (4.20)$$

Then, the variable

$$\iota_p = \left[\int_{[0,1]^{pq}} \left(C^s(\mathbf{u}) - \prod_{j=1}^p C^s(\mathbf{u}_{<j>}) \right)^2 g(\mathbf{u}) d\mathbf{u} \right] \times D^{-1}(C^s) \quad (4.21)$$

is constrained to the interval $[0, 1]$, and is such that $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$ are independent when $\iota_p = 0$, and are increasingly dependent when $\iota_p > 0$. We call the set $\{\iota_p, p = 2, 3, \dots\}$ the *serial dependence function (SDF)* of the series $\{\mathbf{X}_t\}_{t=0}^n$, which shows the amount of serial dependence

for different embedding dimensions (lags) of the series. The *SDF* is conceptually similar to the *ACF*, but unlike the latter, it may be used with the time series of arbitrary dimension, and is also capable of detecting non-linear forms of serial dependence, which may be difficult or impossible to detect using the *ACF*.

The measure $\iota_{n,p}$ is closely related to the multivariate version of Hoeffding's Phi-Square proposed by Gaißer, Ruppert, and Schmid (2010), defined as

$$\Phi^2(C) = h_2(d) \int_{[0,1]^d} (C(\mathbf{u}) - \Pi(\mathbf{u}))^2 d\mathbf{u}, \quad (4.22)$$

with the normalization factor given by

$$h_d(d) = \left(\int_{[0,1]^d} (M(\mathbf{u}) - \Pi(\mathbf{u}))^2 d\mathbf{u} \right)^{-1}, \quad (4.23)$$

where $\Pi(\mathbf{u}) = \prod_{j=1}^d u_j$ is the product copula and $M(\mathbf{u})$ is the Fréchet-Hoeffding upper bound given by $M(\mathbf{u}) = \min(u_1, \dots, u_d)$. The key difference between the $\Phi^2(C)$ and $\iota_{n,s}$, aside from the applicability of $\iota_{i,s}$ to the serial data, is the specification of the independence copula, which in the case of $\iota_{n,s}$ is a functional of the marginal copulas $C^s(u^{(j)})$. This enables $\iota_{n,s}$ to serve as a meaningful measure not only of the dependence *within* a vector, but also of the dependence *between* vectors.

Computational Formula for the $\iota_{n,s}$

In the case of the $\Phi^2(C)$, the denominator $h_d(d)$ depends on d only, and the exact computation formula for $h_d(d)$ and $\Phi^2(C)$ is derived in Gaißer, Ruppert, and Schmid (2010). In our case, the denominator D is the product of the marginal copulas $C^s(\mathbf{u}_{<j>})$. In this sub-section we derive a computational formula for the $D(C^s)$, subject to the empirical marginal copulas $C_n^s(\mathbf{u}_{(j)})$, and for the case of the uniform weights. Such derivation of the bounds which are subject to the empirical marginals was advocated in Section 2.4 of Kojadinovic and Yan (2009). In conjunction with the results from Section 4.2.7, this derivation gives a complete computational

formula for $\iota_{n,s}$.

Theorem 4.2.4 *Given a sequence of Y_i , $i = 1..n$ and letting $\hat{u}_{i,j} = R_i^{(j)}/n' = F_n(X_i^{(j)})$ and $d = pq$, the upper bound*

$$D(C_n^s) = \int_{[0,1]^d} \left(M(\mathbf{u}) - \prod_{j=1}^p C_n^s(\mathbf{u}_{\{j\}}) \right)^2 d\mathbf{u}, \quad (4.24)$$

where $M(\mathbf{u}) = \min\{u_1, \dots, u_d\}$ is the upper Fréchet-Hoeffding copula bound, and $C_n^s(\mathbf{u}_{\{j\}})$ are the empirical marginal copulas, can be found as

$$D(C_n^s) = \frac{2}{(d+1)(d+2)} \quad (4.25)$$

$$- \frac{2}{n^p} \sum_{s=1}^d \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_p=1}^n \frac{1}{2} (1 - \hat{u}_{i_{k_s}, s}^2) \prod_{j=1, j \neq s}^d (1 - \hat{u}_{i_{k_j}, j}) \quad (4.26)$$

$$+ \frac{1}{n^{2p}} \prod_{k=1}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{i,j} \wedge \hat{u}_{l,j}), \quad (4.27)$$

where \wedge denotes the max operator, and $k_j \in [p] : j \in \{b_{k-1} + 1, \dots, b_k\}$.

Proof: Appendix A.

Measures of Dependence and the Shuffles of M

A note about additional caveat surrounding the interpretation of the SDF should be added. The assertion that larger values of ι_p represent stronger dependence is only valid when the copula C^s cannot be represented as a shuffle of M (for a formal definition and discussion of shuffles of M see Section 3.2.3 in Nelsen (2006)). If indeed C^s is such shuffle, the relationship between $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$ is entirely deterministic, and the copula C^s need not deviate from C^\perp by much to have such a case. In fact, C^s may be a shuffle of M , yet uniformly approximate the independence copula arbitrarily closely. In practice, this means $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$ may be a chaotic process, yet be indistinguishable from a random noise. We may then have ι_p be arbitrarily close to zero, yet the data be generated by a deterministic law. Note that this caveat is not unique to

ι_p , but equally applies to any other measures of dependence, and tests for independence. See Nelsen (2006) and Mikusiński, Sherwood, and Taylor (1991) for additional discussions.

4.2.7 Closed-Form Expression

Theorem 4.2.5 *Let $h(\mathbf{u}) \in \mathcal{G}([0, 1]^{pq})$ be an arbitrary weighting function that can also be written as a sum of pq terms as*

$$h(\mathbf{u}) = \sum_{j=1}^{pq} \tilde{w}(u^{(j)}), \quad (4.28)$$

where \tilde{w} is an arbitrary function in $\mathcal{G}([0, 1])$. Then, letting $d = pq$, the statistic $I_n(h)$ can be calculated from the normalized ranks $\hat{u}_{i,j} = R_i^{(j)}/n' = F_n(X_i^{(j)})$ as

$$\begin{aligned} I_n(h) &= \frac{1}{n} \sum_{i=1}^n \sum_{l=1}^n \sum_{m=1}^d \left[\prod_{j=1, j \neq m}^d (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \tilde{w}(\hat{u}_{im} \vee \hat{u}_{lm}, 1) \\ &\quad - \frac{2}{n^p} \sum_{i=1}^n \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\ &\quad \times \left[\sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right] \\ &\quad + \frac{1}{n^{2p-1}} \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\ &\quad \times \left[\sum_{i=1}^n \sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right]. \end{aligned}$$

where $\bar{w}(a, b)$ is the definite integral of $\tilde{w}(u^{(j)})$ from a to b , and \vee is the $\max()$ operator.

Proof: Appendix A.

This allows for an easy substitution of alternative weighting functions, and a closed-form expression for any statistic in $\mathcal{H}(\mathcal{G})$ is easy to obtain, as long as the weighting function can be expressed as a sum of univariate weighting functions $\tilde{w}(u^{(j)}) \in \mathcal{G}([0, 1])$.

4.2.8 The Critical Values

Recall from Theorem 4.2.3 that the marginal copulas $C^s(u_{<j>})$, $j \in [p]$, affect the limiting distribution of the test statistic. If the embedding dimension is $p = 2$, the marginal copulas reduce to $C^s(u_{<j>}) = u_{<j>}$, and the statistic is distribution-free. In such case it is possible to tabulate the exact asymptotic critical values for a given choice of $g(\mathbf{u}) \in \mathcal{G}$. When the embedding dimension is $p > 2$, the statistic is no longer distribution-free. It is possible to obtain approximate critical values from the finite-sample distribution of the statistic, conditional on the observed sample, through a random permutation of sample indexes. We use the following algorithm to obtain approximate p -values for the statistic, and numerically verify that the test maintains its nominal size for various values of p and n . The algorithm is not costly in terms of computation time due to the availability of the computational formula for the test statistic.

1. Let $I_{n,0}^h$ denote the value of the test statistic found using the original sample.
2. Generate N random permutations of the indexes $i \in \{1, \dots, n\}$ of Y_i and for each permutation calculate the value of the statistic $I_{n,i}^h$ using the new pseudo-independent sample.
3. Find the approximate p -value for I_n^h as

$$p = \frac{1}{N+1} \left\{ \frac{1}{2} + \sum_{i=1}^N \mathbf{1}(I_{n,0}^h \leq I_{n,i}^h) \right\}. \quad (4.29)$$

4.2.9 Power Properties

In this section we study the power of the weighted statistic $I_n(w(\mathbf{u}))$, specified by the median-centric weights $w(\mathbf{u}) = \sum_{i=1}^{pq} (u^{(i)} - (u^{(i)})^2)$, and of the $I_n(w^-(\mathbf{u}))$, where the weights $w^-(\mathbf{u}) = 1 - (u^{(i)} - \sum_{i=1}^{pq} (u^{(i)})^2)$ represent a horizontal reflection of $w(\mathbf{u})$, and emphasize the tails of the distribution instead. We benchmark the test power against that of the uniformly-weighted statistic $I_n(g^1(\mathbf{u}))$ of Kojadinovic and Yan (2009). We also track the performance of a scalar

version of the test statistic specified as

$$\tilde{I}_n(w) = n \int_{[0,1]^d} \left(C_n^s(\mathbf{u}) - \prod_{j=1}^{pq} u^{(j)} \right)^2 w(\mathbf{u}) d\mathbf{u}. \quad (4.30)$$

Unlike the $I_n(w)$, which tests for independence between vectors $(\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$, the $\tilde{I}_n(w)$ tests for independence within the pq -matrix $Y_i = (\mathbf{X}_i, \dots, \mathbf{X}_{i+p-1})$. The difference is due to the specification of the independence copula, which, in the latter case, is specified as the product of the components of \mathbf{u} . The statistic $\tilde{I}_n(w)$ is therefore not capable of distinguishing between the temporal and the serial dependence of the vector series. To assess the impact of the weights on the test power, we therefore track the performance of the following statistics: the median-weighted $I_n(w(\mathbf{u}))$ and $\tilde{I}_n(w(\mathbf{u}))$, the uniformly-weighted $I_n(g^1(\mathbf{u}))$ and $\tilde{I}_n(g^1(\mathbf{u}))$, and the tail-weighted $I_n(w^-(\mathbf{u}))$.

We simulate realizations of both univariate and multivariate series $\mathbf{X}_1, \mathbf{X}_2, \dots$, where the serial dependence is characterized through multidimensional parametric Gumbel and Student t copulas. The resulting series have a rich dependence structure. For example, data generated through a Gumbel copula will have upper-tail dependence, while the series with a t copula will be dependent through all of the tails, in addition to being serially correlated. We set the embedding dimension to $p = 2$ for simplicity. The copula parameter θ (or ρ for the Student t copula) controls the amount of both serial and temporal dependence in the series, where larger values give stronger relationship between \mathbf{X}_t and \mathbf{X}_{t+1} . As before, the test power is the percent of correct rejections in $s = 100$ realizations of the simulated process $\{\mathbf{X}_t\}_{t=0}^{\infty}$.

Table 4.1 shows the local power of $I_n(w)$ and $I_n(g^1)$ statistics in samples of size $n = 25$, with embedding dimension $p = 2$ and series size $q = 1$, with data generated with a bi-variate Gumbel copula.

Table 4.2 shows the local power of $\tilde{I}_n(w)$ and $\tilde{I}_n(g^1)$ in samples of size $n = 25$, with embedding dimension $p = 2$ and series size $q = 2$, with data generated with a Gumbel 4-copula.

Table 4.3 shows the local power of $I_n(w^-)$ and $I_n(g^1)$ statistics in samples of size $n = 25$,

θ	Unweighted	Median Weights	Power Gain
1	0.1	0.09	-
1.1	0.55	0.54	-
1.2	0.88	0.86	-

Table 4.1: Power of the $I_n(g^1(\mathbf{u}))$ and the $I_n(w(\mathbf{u}))$ statistics. $n = 25$, $s=100$ samples (Gumbel Copula).

θ	Unweighted	Median weights	Power Gain
1	0.09	0.09	-
1.2	0.34	0.38	0.04
1.4	0.72	0.73	0.01
1.6	0.96	0.96	-

Table 4.2: The power of the $\bar{I}_n(g^1(\mathbf{u}))$ and the $\bar{I}_n(t(\mathbf{u}))$ statistics. $n = 25$, $s = 100$ samples (Gumbel copula). Embedding dimension $p = 2$, series dimension $q = 2$.

with embedding dimension $p = 2$ and series size $q = 1$, with data generated with a bi-variate Student t copula.

ρ	Unweighted	Tail Weights	Power Gain
0	0.1	0.09	-
0.2	0.14	0.17	0.03
0.3	0.22	0.24	0.02
0.4	0.41	0.44	0.03
0.5	0.5	0.52	0.02
0.6	0.67	0.69	0.02
0.8	0.91	0.93	0.02
1	0.99	0.99	-

Table 4.3: Power of the $I_n(g^1(\mathbf{u}))$ and the $I_n(t(\mathbf{u}))$ statistics. $n = 25$, $s = 100$ samples (t copula). Embedding dimension $p = 2$, series dimension $q = 1$.

4.3 Applications

4.3.1 Testing for Vectorial Serial Independence in Financial Data

In this section we use the statistic $I_n(w)$ to study the dependence structure of some real-world financial series. In particular, we test for serial independence of the daily return series of the leading equities indexes in the Americas, Europe and Asia, selecting three indexes from each region. We first study each index individually, and then combine the indexes into multiple vector series. To make the referencing of these groupings easier, we attach a variable name to each of the indexes in our sample. Table 4.4 shows the nine stock market indexes which we selected, and the corresponding labels for the univariate series. Table 4.5 shows the grouping of the indexes into vector series by their geographic origin.

Equities Index	Variable Name
S&P 500	$\mathbf{X}_t^{(1)}$
TSX Composite	$\mathbf{X}_t^{(2)}$
Bovespa	$\mathbf{X}_t^{(3)}$
FTSE 100	$\mathbf{X}_t^{(4)}$
CAC 40	$\mathbf{X}_t^{(5)}$
DAX Index	$\mathbf{X}_t^{(6)}$
Nikkei 225	$\mathbf{X}_t^{(7)}$
All Ordinaries	$\mathbf{X}_t^{(8)}$
Shanghai Composite	$\mathbf{X}_t^{(9)}$

Table 4.4: Variable names for the leading stock market indexes.

Index Group	Variable Vector
Americas	$\mathbf{X}_t^{[A]} = (\mathbf{X}_t^{(1)}, \mathbf{X}_t^{(2)}, \mathbf{X}_t^{(3)})$
Europe	$\mathbf{X}_t^{[E]} = (\mathbf{X}_t^{(4)}, \mathbf{X}_t^{(5)}, \mathbf{X}_t^{(6)})$
Asia & Australia	$\mathbf{X}_t^{[U]} = (\mathbf{X}_t^{(7)}, \mathbf{X}_t^{(8)}, \mathbf{X}_t^{(9)})$
World	$\mathbf{X}_t = (\mathbf{X}_t^{(1)}, \mathbf{X}_t^{(2)}, \dots, \mathbf{X}_t^{(9)})$

Table 4.5: Stock market index groupings.

Daily closing values of the indexes from April 23, 2011 to April 23, 2012 are used to calculate the series of returns vectors \mathbf{X}_t . For each of the series in Tables 4.4 and 4.5 we calculate

the statistic $I_n(g)$ and the associated approximate p -value for embedding dimensions ranging from $p = 1$ to $p = 6$. Note that increasing the embedding dimension p greatly increases the computational complexity of the problem. For example, for $p = 6$, testing for serial independence of the vector of world stock market returns $\mathbf{X}_t \in \mathbb{R}^9$ amounts to working with vectors in \mathbb{R}^{54} .

Table 4.6 shows the p -values for the test of H_0 of serial independence of some of the series from Table 4.5 for various embedding dimensions. The p -values which are lower than 0.2 are highlighted in bold. Interestingly, there appears to be some evidence of serial dependence present in the S&P 500 and the TSX Composite return series at daily lag of order $p = 4$.

p	$\mathbf{X}_t^{(1)}$	$\mathbf{X}_t^{(2)}$	$\mathbf{X}_t^{(3)}$	$\mathbf{X}_t^{(4)}$	$\mathbf{X}_t^{[A]}$
2	0.87	0.64	0.86	0.76	0.69
3	0.43	0.34	0.77	0.3	0.44
4	0.14	0.17	0.74	0.36	0.15
5	0.41	0.52	0.62	0.31	0.86
6	0.05	0.5	0.47	0.43	0.75

Table 4.6: p -values from the test of H_0 of serial independence of some of the series in Table 4.5 for embedding dimensions from $p = 2$ to $p = 6$.

4.3.2 A Goodness of Fit Test for Archimedean Serial Copulas

Copula models are becoming widely used in economics and finance, and the problem of selecting an appropriate copula family to model the data received considerable attention recently. Quessy (2010) proposed a new goodness of fit procedure for Archimedean copulas that uses a test for independence to reject certain parametric copula families. This procedure is highly tractable and has good power properties, but it cannot be directly applied to time-series data. This section extends the results of Quessy (2010) to the serial setting, enabling the application of the weighted serial test statistic proposed in this chapter to the testing for copula goodness of fit.

Let $C_\theta^s(\mathbf{u})$, $\mathbf{u} \in [0, 1]^{pq}$, be a copula of Archimedean type with a parameter vector θ . We can

then write $C_\theta^s(\mathbf{u})$ in terms of its generator $\psi_\theta(u)$ as $C_\theta^s(\mathbf{u}) = \psi_\theta^{-1}(\psi_\theta(u^{(1)}) + \dots + \psi_\theta(u^{(pq)}))$, such that $\psi_\theta(1) = 0$, $\psi_\theta(0) = \infty$, and $(-1)^k \partial^k \psi_\theta^{-1}(u) / \partial u^k \geq 0$, $\forall k \in \mathbb{N}$. The generator ψ_θ uniquely defines families of Archimedean copulas, up to a value of the parameter θ . For example, we can obtain a one-parameter pq -dimensional Clayton copula by setting $\psi_\theta(t) = t^{-\theta} - 1$, for $\theta > 0$, as $C_\theta^s = ((u^{(1)})^{-\theta} + (u^{(2)})^{-\theta} + \dots + (u^{(pq)})^{-\theta} - pq + 1)^{-1/\theta}$. Other popular Archimedean copulas such as, for example, Gumbel, Frank, Joe, and Ali-Mikhail-Haq families are defined in the same way through a different choice of ψ_θ .

Barbe, Genest, Ghoudi, and Remillard (1996) showed that the variables $\mathbf{V} = (\psi(u^{(1)})/Z, \dots, \psi(u^{(pq)})/Z) \in \mathbb{R}^{pq}$ and $Z = \sum_{j=1}^{pq} \psi(u^{(j)})$ are independent. This permits the application of a test for independence to the problem of testing for copula goodness of fit. Let $\mathcal{A}(\psi_\theta)$ be a family of pq -copulas of Archimedean type with generator ψ_θ and parameter vector θ . We can then use the weighted statistic $I_n(g)$ given by (4.13) to test the $H_0 : C^s \in \mathcal{A}(\psi_\theta)$ by testing for independence of \mathbf{V} and Z .

Note that the \mathbf{u} 's are rarely observable in practice, since this requires the marginal distributions to be known. As in the previous sections, the sample versions of \mathbf{u} 's, and hence of \mathbf{V} and Z can be obtained from sample ranks. Given the sequence \mathbf{Y}_i , $i \in [n]$, define $\hat{\mathbf{u}}_i = (R_i, \dots, R_{i+pq-1})/n'$, where R_i is an empirical ranks vector associated with X_i as before. Further, let $l_{\theta,j}(\mathbf{u}) = \psi_\theta(\mathbf{u}^{(j)}) / \sum_{k=1}^{pq} \psi_\theta(\mathbf{u}^{(k)})$, $j \in [pq]$, and $l_{\theta,pq+1}(\mathbf{u}) = \sum_{k=1}^{pq} \psi_\theta(\mathbf{u}^{(k)})$. Then, calculate sample versions of \mathbf{V} and Z as $\mathbf{V}_i = (l_{\hat{\theta},1}(\hat{\mathbf{u}}_i), l_{\hat{\theta},2}(\hat{\mathbf{u}}_i), \dots, l_{\hat{\theta},pq}(\hat{\mathbf{u}}_i))$ and $Z_i = l_{\hat{\theta},pq+1}(\hat{\mathbf{u}}_i)$, for $i \in [n]$, where $\hat{\theta}$ is an estimate of the copula parameter vector and ψ_θ is the generator chosen under H_0 .

To establish the convergence of the underlying empirical process, and ultimately of the test statistic $I_n(g)$ when it is used in this context, first define the empirical distribution function of \mathbf{V}_i and Z as

$$M_n(\mathbf{v}, z) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(\mathbf{V}_i \leq \mathbf{v}, Z_i \leq z), \mathbf{v} \in \mathbb{R}^{pq}, z \in \mathbb{R}, \quad (4.31)$$

where $\mathbf{1}(\cdot)$ is an indicator function, and inequality $\mathbf{V}_i \leq \mathbf{v}$ is component-wise. The empirical

copula of \mathbf{V} and Z is then given by

$$C_n(\mathbf{u}_1, u_2) = M_n(M_n^{-1}(\mathbf{u}_{1,1}), M_n^{-1}(\mathbf{u}_{1,2}), \dots, M_n^{-1}(\mathbf{u}_{1,pq}), M_n^{-1}(u_2)), \mathbf{u}_1 \in \mathbb{R}^{pq}, u_2 \in \mathbb{R}, \quad (4.32)$$

where for $i \in [pq]$, $\mathbf{u}_{1,i} \in \mathbb{R}^{pq} = (1, 1, \dots, \mathbf{u}_1^{(i)}, \dots, 1, 1)$. Since under the $H_0 : C^s \in \mathcal{A}(\psi_\theta)$, we have that $C(\mathbf{u}_1, u_2) = C(\mathbf{u}_1, 1)u_2$, the associated empirical process is given by

$$\mathbb{E}_n(\mathbf{u}_1, u_2) = \sqrt{n} (C_n(\mathbf{u}_1, u_2) - C_n(\mathbf{u}_1, 1)C_n(\mathbf{1}, u_2)), \quad (4.33)$$

and the weighted goodness of fit statistic by

$$\bar{I}_n(g) = n \int_{[0,1]^{pq+1}} (C_n(\mathbf{u}_1, u_2) - C_n(\mathbf{u}_1, 1)C_n(\mathbf{1}, u_2))^2 g(\mathbf{u}_1, u_2) d\mathbf{u}_1 du_2. \quad (4.34)$$

As before, we can use the closed-form expression given in Section (4.2.7) to calculate the value of $\bar{I}_n(g)$ using our sample of $Y_i, i \in [n]$. The next two theorems establish the weak convergence of $\mathbb{E}_n(\mathbf{u}_1, u_2)$ and the limiting distribution of $\bar{I}_n(g)$, and hence ensure the validity of this method.

Theorem 4.3.1 *Under $H_0 : C^s \in \mathcal{A}(\psi_\theta)$, assuming that the Hypothesis I and II of Ghoudi and Rémillard (2004) hold, the empirical process $\mathbb{E}_n(\mathbf{u}_1, u_2)$ converges weakly in $\mathcal{D}([0, 1]^{pq+1})$ to a Gaussian process given by*

$$\mathbb{E}(\mathbf{u}_1, u_2) = \mathbb{B}(\mathbf{u}_1, u_2) - \sum_{j=1}^{pq} k_j(u_1^{(j)}) E\{\mathbb{L}_j(\mathbf{U}) \mathbf{I}(\mathbf{U} \leq \mathbf{u}_1, U \leq u_2) | \mathbf{U}^{(j)} = u_1^{(j)}\} \quad (4.35)$$

$$+ k_{pq+1}(u_2) E\{\mathbb{L}_{pq+1}(\mathbf{U}) \mathbf{I}(\mathbf{U} \leq \mathbf{u}_1, U \leq u_2 | U = u_2)\}, \quad (4.36)$$

where $\mathbb{B}(\mathbf{u}_1, u_2)$ is the weak limit of $\sqrt{n}(M_n - M)$ and \mathbb{L}_j is defined as in Lemma 3 of Quessy (2010).

Proof: Proposition 5 of Quessy (2010) and Hypotheses I, II and Theorem 2.4 of Ghoudi and Rémillard (2004).

Theorem 4.3.2 *Under $H_0 : C^s \in \mathcal{A}(\psi_\theta)$ and the assumptions of Theorem 4.3.1, if copula C^s has continuous partial derivatives, the goodness of fit statistic $\bar{I}_n(g)$ converges in distribution to the distribution of*

$$\int_{[0,1]^{p+q+1}} \mathbb{E}(\mathbf{u}_1, u_2)^2 g(\mathbf{u}) d\mathbf{u}. \quad (4.37)$$

Proof: The proof follows from Theorem 4.3.1 and the continuous mapping theorem (Theorem 1.3.6 in van der Vaart and Wellner (1996)).

A complete goodness of fit testing procedure for serial copulas involves the following steps:

1. Obtaining a maximum likelihood estimate $\hat{\theta}$ of the copula parameter vector under H_0
2. Transforming the data and calculating \mathbf{V}_i and Z_i using the generator ψ chosen under H_0
3. Using the statistic $\bar{I}_n(g)$ to test for independence of \mathbf{V}_i and Z_i

The use of the weighting function g makes it possible to impose heavier penalties for misfitting certain parts of the copula, such as, for example, the tails. This may be particularly useful in applications such as risk management, where it is more important to have an accurate model of the clustering of extreme shocks. In this respect, the weighting function represents a potentially-useful tuning parameter.

4.3.3 Interval Estimation of Serial Copula Parameters

The goodness of fit testing procedure can be reversed to obtain estimates of the α -confidence interval of the parameter θ under H_0 . Assuming that the copula C^s indeed belongs to a known Archimedean family $\mathcal{A}(\psi_\theta)$, the parameter value θ_1 is in the α -confidence interval if we cannot reject independence between \mathbf{V}_i and Z_i following a transformation with ψ_{θ_1} at a level of significance α . The estimate of the confidence interval can therefore be constructed by drawing a sufficiently large number of the values of θ either at random, or from a sufficiently fine grid,

and subsequently testing for independence of the transformed data. Collecting the values of θ for which the test fails to reject gives the estimate of the interval.

Note that as with the goodness of fit testing, the weighting function allows the conditioning of the estimation on the specifics of the problem at hand. Selecting the weights which, for example, make the statistic $I_n(g)$ more sensitive to tail dependence will make the interval estimate respond stronger to observations indicating tail dependence. The local power properties of $I_n(g)$ embedded in $g(\mathbf{u})$ will be inherited by the resulting interval estimator, and the interval width may be directly affected by the choice of g .

4.3.4 A Note About the Interpretation of the Weighting Function

In the context of interval estimation and testing for copula goodness of fit, the weighting function $g(\mathbf{u})$ can be thought to represent the preferences of the practitioner over the possible sources of the modelling error, which are likely asymmetric. Following an earlier risk-management example, the consequences of wrongfully modelling the clustering of losses in a portfolio of risky assets are potentially a lot more severe than the consequences of mis-modelling the clustering of gains. Such asymmetric preferences are often accounted for by explicitly specifying the practitioner's utility function and making it a part of the portfolio design problem. The same preferences could be accounted for by selecting an appropriate $g(\mathbf{u})$ which will penalize mis-fitting certain parts of the distribution more than the others. This approach can help avoid many of the complications which arise from the inclusion of utility theory into the portfolio optimization problem, and needs to be investigated further. In particular, it would be useful to understand how the deviations from the H_0 of the copula goodness of fit translate into the deviation from the H_0 of independence of the transformed variables Z and \mathbf{V} , which could be a useful guide for the selection of the weights in this context.

4.3.5 Applied Example: Portfolio Value-at-Risk Computation

One particularly fruitful application of the theory of copulas is to the problem of calculating the Value-at-Risk (VaR) of a portfolio of risky assets. Given a confidence level α , the corresponding Value-at-Risk is the amount VaR_α such that the portfolio losses L exceed VaR_α with probability α . More formally, VaR_α is the solution to the following equation:

$$Pr(L \leq VaR_\alpha) = \alpha. \quad (4.38)$$

For simplicity, suppose that the portfolio contains of only two assets. Let X and Y represent their continuous returns, and the weight on X be $\beta \in (0, 1)$. Portfolio return is $L = \beta X + (1 - \beta)Y$. Assuming that $X, Y \in \mathbb{R}$, the distribution function of L is given by

$$K(l) = Pr(L \leq l) = Pr(\beta X + (1 - \beta)Y \leq l) \quad (4.39)$$

$$= \int_{-\infty}^{+\infty} Pr\left(X \leq \frac{1}{\beta}l - \frac{(1 - \beta)}{\beta}y, Y = y\right) f_y(y) dy \quad (4.40)$$

$$= \int_{-\infty}^{+\infty} C_y\left(F_x\left(\frac{1}{\beta}l - \frac{(1 - \beta)}{\beta}y\right), F_y(y)\right) f_y(y) dy, \quad (4.41)$$

where F_x, F_y are the marginal distribution functions, f_y is the density of Y , and C_y is the partial derivative of the copula C of X and Y with respect to the second argument. This result follows from the Sklar's theorem (see Remark 2.2 and Section 2.3.4 in Cherubini, Luciano, and Vecchiato (2004)).

Since in practice the true copula C of X and Y is unknown, to solve the problem we need to select the appropriate parametric family of copulas and estimate the value of the parameter, which introduces additional errors. The statistic $I_n(g)$ can be used both for copula selection and for the estimation of the copula parameter bounds. By selecting the weighting function $g(\mathbf{u})$ which places higher weights on the values of X and Y in the lower tail of the joint distribution we impose additional penalties for mis-fitting the model in the region where Type II errors are most costly.

4.4 Conclusion

This chapter provides an extension to the vectorial serial independence test statistic of Kojadinovic and Yan (2009) which is obtained through the introduction of a weighted function norm. The non-uniform weighting of the data enables the adjustment of the power properties of the test through the alternative choices of the weights. The chapter also provides a serial extension to the copula goodness of fit procedure proposed in Quessy (2010), which enables the application of the weighted independence test statistic to the problem of serial copula selection. The addition of the weights is particularly useful in this context, since it permits the imposition of asymmetric losses which arise due to modeling error.

No attempt is made to study the power properties of the serial copula goodness of fit procedure using the weighted test statistic, or to investigate the practical implications of the alternative choices of the weights. As in the third chapter, the question of existence of the optimal weighting function, now in the serial context, is also left unanswered. In that sense, significant scope for future theoretical and applied work relating to the choice of the weights exists.

While the copula-based vectorial independence test statistics proposed in the literature and in this dissertation are omnibus, in a sense that that can detect dependence of any form, they lack standardization and therefore cannot serve as meaningful measures of vectorial dependence. This chapter derives an upper bound on the weighted test statistic subject to the empirical marginal copulas, effectively turning it into a proper measure of dependence. While the computational formula for this measure is provided, no attempt is made to take it to the data, once again leaving significant scope for future work which would focus on the applications of the measure to a very broad range of empirical problems in economics, econometrics, and finance.

Bibliography

- ALMEIDA, A., C. GOODHEART, AND R. PAYNE (1998): “The Effects of Macroeconomic News on High Frequency Exchange Rate Behavior,” *Journal of Financial and Quantitative Analysis*, 33(3), 383–408.
- ANDERSON, T., AND D. DARLING (1952): “Asymptotic theory of certain” goodness of fit” criteria based on stochastic processes,” *The Annals of Mathematical Statistics*, 23(2).
- ANG, A., AND G. BEKAERT (2002): “International asset allocation with regime shifts,” *Review of Financial Studies*, 15, 1137–1187.
- ANG, A., AND J. CHEN (2002): “Asymmetric correlations of equity portfolios,” *Journal of Financial Economics*, 63, 443–494.
- BAE, K.-H., A. KAROLYI, AND R. STULZ (2003): “A new approach to measuring financial contagion,” *Review of Financial Studies*, 16, 717–763.
- BARBE, P., C. GENEST, K. GHOUDI, AND B. REMILLARD (1996): “On Kendall’s Process,” *Journal of Multivariate Analysis*, 58, 197–229.
- BASU, S. (1977): “The investment performance of common stocks relative to their price-earnings ratios: a test of the efficient markets,” *Journal of Finance*, 32(3), 663–682.
- BECKER, K. G., J. E. FINNERY, AND K. J. KOPECKY (1995): “Domestic macroeconomic news and foreign interest rates,” *Journal of International Money and Finance*, 14(6), 763–783.
- BEINE, M., A. COSMA, AND R. VERMEULEN (2010): “The dark side of global integration: Increasing tail dependence,” *Journal of Banking and Finance*, 34(1), 184–192.
- BLUM, J., J. KIEFER, AND M. ROSENBLATT (1961): “Distribution free tests of independence based on the sample distribution function,” *Annals of Mathematical Statistics*, 32, 485–498.
- BOLLEN, J., H. MAO, AND X. ZENG (2011): “Twitter mood predicts the stock market,” *Journal of Computational Science*, 2(1), 1–8.
- BONDT, W. D., AND R. H. THALER (1985): “Does the stock market overreact?,” *Journal of Finance*, 40(3), 793–805.
- BOYD, J. H., J. HU, AND R. JAGANNATHAN (2005): “The Stock Market’s Reaction to Unemployment News: Why Bad News Is Usually Good For Stocks,” *The Journal of Finance*, 60(2), 649–672.

- CAMPBELL, R., K. KOEDIJK, AND P. KOFMAN (2002): "Increasing correlation in bear markets: a downside risk perspective," *Financial Analysis Journal*, 58(87-94).
- CAPPIELLO, L., R. ENGLE, AND K. SHEPPARD (2006): "Asymmetric dynamics in the correlations of global equity and bond returns," *Journal of Financial Econometrics*, 4, 537–572.
- CHEN, X., AND Y. FAN (2006): "Estimation of copula-based semiparametric time series models," *Journal of Econometrics*, 130(2), 307–335.
- CHERUBINI, U., E. LUCIANO, AND W. VECCHIATO (2004): *Copula Methods in Finance*. John Wiley & Sons, England.
- CUTLER, D. M., J. M. POTERBA, AND L. H. SUMMERS (1989): "What moves stock prices?," *Journal of Portfolio Management*, 15(3), 4–12.
- DEHEUVELS, P. (1979): "La fonction de dependance empirique et se propriers: Un test non-parametrique dindependance," *Academie Royale de Belgique. Bulletin de la Classe des Sciences*, 65, 274–292.
- DEHEUVELS, P. (1981): "An asymptotic decomposition for multivariate distribution-free tests of independence," *Journal of Multivariate Analysis*, 11(1), 102–113.
- DELGADO, M. (1996): "Testing Serial Independence Using the Sample Distribution Function," *Journal of Time Series Analysis*, 17(3), 271–285.
- EDERINGTON, L. H., AND L. H. LEE (1993): "How Markets Process Information: News Releases and Volatility," *The Journal of Finance*, 48(4), 1161–1191.
- FAMA, E., AND K. FRENCH (1992): "The cross section of expected stock returns," *Journal of Finance*, 47, 427–466.
- FAN, Y. (2010): *Encyclopedia of Quantitative Finance* chap. Copulas in Econometrics. John Wiley & Sons, Chichester, UK.
- FERMANIAN, J.-D., D. RADULOVIC, AND M. WEGKAMP (2004): "Weak convergence of empirical copula processes," *Bernoulli*, 10(5), 847–860.
- FERMANIAN, J.-D., AND M. WEGKAMP (2012): "Time-dependent copulas," *Journal of Multivariate Analysis*, 110, 19–29.
- FLANNERY, M. J., AND A. A. PROTOPAPADAKIS (2002): "Macroeconomic Factors Do Influence Aggregate Stock Returns," *Review of Financial Studies*, 15(3), 751–182.
- GAENSSLER, P., AND W. STUTE (1987): "Seminar on Empirical Processes," DMV Sem. 9. Basel: Birkhauser.
- GAIBER, S., M. RUPPERT, AND F. SCHMID (2010): "A multivariate version of Hoeffdings Phi-Square," *Journal of Multivariate Analysis*, 101(10), 2571–2586.

- GENEST, C., AND B. REMILLARD (2004): "Tests of independence and randomness based on the empirical copula process," *Test (Sociedad de Estadística e Investigación Operativa)*, 13(2), 335–369.
- GHOUDI, K., AND B. REMILLARD (1998): "Empirical processes based on pseudo-observations," in *Proceedings of the International Conference on Asymptotic Methods in Statistics*, Amsterdam. Elsevier.
- GHOUDI, K., AND B. RÉMILLARD (2004): "Empirical processes based on pseudo-observations. II. The multivariate case," *Asymptotic methods in stochastics*, 44, 381–406.
- GOLDBERGER, A. S. (1964): *Econometric Theory*. John Wiley & Sons, New York.
- GRAHAM, M., J. NIKKINEN, AND P. SAHLSTROM (2003): "Relative importance of scheduled macroeconomic news for stock market investors," *Journal of Economics and Finance*, 27(2), 153–165.
- GROSSMAN, S. J., AND R. J. SHILLER (1981): "The determinants of the variability of stock market prices," *American Economic Review*, 71, 222–227.
- HALLIN, M., AND M. PURI (1995): "A Multivariate Wald-Wolfowitz Rank Test against Serial Dependence," *Canadian journal of statistics*, 23(1), 55–65.
- HOEFFDING, W. (1948): "A nonparametric test of independence," *Annals of Mathematical Statistics*, 19, 546–557.
- HOSKING, J. (1980): "The multivariate portmanteau statistic," *Journal of the American Statistical Association*, 75(371), 602–608.
- JAIN, P. C. (1988): "Responses of Hourly Stock Prices and Trading Volume to Economic News," *The Journal of Business*, 61(2), 219–231.
- KNIGHT, J., C. LIZIERI, AND S. SATCHELL (2005): "Diversification When It Hurts? The Joint Distribution of Real Estate and Equity Markets," *Journal of Property Research*, 22(4), 309–323.
- KOJADINOVIC, I., AND M. HOLMES (2009): "Tests of independence among continuous random vectors based on Cramer-von Mises functionals of the empirical copula process," *Journal of Multivariate Analysis*, 100, 1137–1154.
- KOJADINOVIC, I., AND J. YAN (2009): "Tests of serial independence for continuous multivariate time series based on a Mobius decomposition of the independence empirical copula process," *Annals of the Institute of Statistical Mathematics*, 10, 347–373.
- KOUL, H. (2002): *Weighted empirical processes in dynamic nonlinear models*. Springer.
- LEHMANN, E. (1966): "Some concepts of dependence," *The Annals of Mathematical Statistics*, 37(5), 1137–1153.

- LEROY, S., AND R. PORTER (1981): "Stock price volatility: a test based on implied variance bounds," *Econometrics*, 49, 97–113.
- MCQUEEN, G., AND V. V. ROLEY (1993): "Stock Prices, News, and Business Conditions," *The Review of Financial Studies*, 6(3), 683–707.
- MIKUSIŃSKI, P., H. SHERWOOD, AND M. TAYLOR (1991): *Probabilistic interpretations of copulas and their convex sums*. In: Dall'Aglio G, Kotz S, Salinetti G (eds) *Advances in Probability Distributions with Given Marginals*. Kluwer Dodrecht.
- MUNKIN, M., AND P. TRIVEDI (1999): "Simulated maximum likelihood estimation of multivariate mixed-poisson regression models with application," *Econometric Journal*, 1, 1–21.
- NELSEN, R. B. (2006): *An Introduction To Copulas*. Springer Science, 233 Springer Street, New York, USA.
- PATTON, A. (2009): *Handbook of Financial Time Series* vol. 5, chap. Copula-Based Models for Financial Time Series. Springer.
- PATTON, A. J. (2006): "Modelling Asymmetric Exchange Rate Dependence," *International Economic Review*, 47(2), 527–556.
- PEARCE, D. K., AND V. V. ROLEY (1985): "Stock Prices and Economic News," *The Journal of Business*, 58(1), 49–67.
- POTERBA, J., AND L. SUMMERS (1988): "Mean reversion in stock prices: evidence and implications," *Journal of Financial Economics*, 22, 26–59.
- QUESSY, J.-F. (2010): "Applications and asymptotic power of marginal-free tests of stochastic vectorial independence," *Journal of Statistical Planning and Inference*, 140, 3058–3075.
- RITTER, J. R. (1991): "The long-run performance of initial public offerings," *Journal of Finance*, 46(1), 3–27.
- RODRIGUEZ, J. C. (2007): "Measuring Financial Contagion: A Copula Approach," *Journal of Empirical Finance*, 14(3), 401–423.
- ROSENBERG, J. V., AND T. SCHUERMAN (2006): "A general approach to integrated risk management with skewed, fat-tailed risks," *Journal of Financial Economics*, 79, 569–614.
- RUSCHENDORF, L. (1976): "Asymptotic distributions of multivariate rank order statistics," *The Annals of Statistics*, pp. 912–923.
- SANCETTA, A., AND S. SATCHELL (2004): "The Bernstein copula and its applications to modeling and approximations of multivariate distributions," *Econometric Theory*, 20, 535–562.
- SHILLER, R. J. (1981): "Do stock prices move too much to be justified by subsequent movements in dividends?," *American Economic Review*, 73(3), 421–436.

- SHORACK, G., AND J. WELLNER (2009): *Empirical Processes With Applications To Statistics*. Society for Industrial and Applied Mathematics.
- SKAUG, H. J., AND D. TJOSTHEIM (1993): "A Nonparametric Test of Serial Independence Based on the Empirical Distribution Function," *Biometrika*, 80(3), 593–602.
- SKLAR, A. (1959): "Fonctions de repartition a n dimensions et leurs marges," *Publications de l'Institut de Statistique de l'Universite de Paris*, 8, 229–231.
- SMITH, M. D. (2003): "Modelling sample selection using Archimedean copulas," *Econometrics Journal*, 6, 99–123.
- TRIVEDI, P., AND D. ZIMMER (2007): *Copula Modeling: An Introduction for Practitioners*. World Scientific Publishing.
- TSUKAHARA, H. (2005): "Semiparametric estimation in copula models," *The Canadian Journal of Statistics*, 33(357-375).
- VAN DER VAART, A., AND J. WELLNER (1996): *Weak Convergence and Empirical Processes*. Springer-Verlag, New York.
- VANDERZANDEI, A. J. (1981): "Some results for the weighted empirical process concerning the law of the iterated logarithm and weak convergence," *Dissertation Abstracts International Part B: Science and Engineering*, 4(7).
- WHITE, H. (1980): "A Heteroscedasticity-Consistent Covariance Matrix and a Direct Test for Heteroscedasticity," *Econometrica*, 48, 817–838.
- WU, D.-M. (1973): "Alternative tests of independence between stochastic regressors and disturbances," *Econometrica*, 41(4), 733–750.
- XU, D., J. KNIGHT, AND T. WIRJANTO (2011): "Asymmetric Stochastic Conditional Duration Model - A Mixture of Normals Approach," *Journal of Financial Econometrics*, 9(3), 469–488.
- ZELLNER, A., J. KMENTA, AND J. DREZE (1966): "Specification and estimation of Cobb-Douglas production function models," *Econometrica*, 34, 784–795.

Appendix A

Proofs of Theorems

A.0.1 Proof of Lemma 3.5.1

Proof Let $t_n, n = 1, 2, \dots$ be a sequence of reals converging to zero, and let $a_n, n = 1, 2, \dots$ be a sequence of maps in $l^\infty([0, 1]^d)$ converging to $a \in l^\infty([0, 1]^d)$ such that $f + t_n a_n \in l^\infty([0, 1]^d)$. Then, from the definition of

$$\begin{aligned}
 & \frac{\phi(f + t_n a_n)(x) - \phi(f)(x)}{t_n} = \\
 &= \lim_{n \rightarrow \infty} \frac{\left\{ f(x) + t_n a_n(x) - \prod_{j=1}^p (f(x) + t_n a_n(x^{[j]})) - f(x) + \prod_{j=1}^p f(x^{[j]}) \right\} \sqrt{w(x)}}{t_n} \\
 &= a(x) \sqrt{w(x)} \\
 &- \lim_{n \rightarrow \infty} \left\{ \left(\prod_{k=1}^p f(x^{[k]}) + t_n \sum_{k=1}^p a_n(x^{[k]}) \prod_{j=1, j \neq k}^p f(x^{[j]}) \right. \right. \\
 &+ \left. \left. \sum_{k=2}^{p-1} t_n^k \left(\sum_{l=1}^p f(x^{[l]}) \prod_{i=1, i \neq l}^p a_n(x^{[i]}) \right) + t_n^p \prod_{k=1}^p a_n(x^{[k]}) - \prod_{k=1}^p f(x^{[k]}) \right) / t_n \right\} \sqrt{w(x)} \\
 &= \left(a(x) - \sum_{k=1}^p a_n(x^{[k]}) \prod_{j=1, j \neq k}^p f(x^{[j]}) \right) \sqrt{w(x)}.
 \end{aligned}$$

A.0.2 Proof of Theorem 3.5.2

Proof By applying the functional delta method as in van der Vaart and Wellner (1996) (Theorem 3.9.4) with Hadamard-differentiable map ϕ and its Hadamard-derivative $\phi'_C(u)$ given by Lemma 3.5.1 to the process (3.6) with weak limit $g(u)$ given in Theorem 3.4.1 we have that $\sqrt{n}[\phi(C_n)(u) - \phi(C)(u)] \rightsquigarrow \phi'_C(g)(u)$. The limiting covariance function is:

$$\begin{aligned}
E[\phi'_C(u)\phi'_C(v)] &= \\
& E \left[\left(g(u) - \sum_{j=1}^p g(u^{(j)}) \prod_{i=1, i \neq j}^p C(u^{(i)}) \right) \left(g(v) - \sum_{j=1}^p g(v^{(j)}) \prod_{i=1, i \neq j}^p C(v^{(i)}) \right) \sqrt{w(u)w(v)} \right] \\
&= \left(E[g(u)g(v)] - \sum_{i=1}^p E[g(u)g(v^{(i)})] \prod_{j=1, j \neq i}^p C(v^{(j)}) - \sum_{i=1}^p E[g(v)g(u^{(i)})] \prod_{j=1, j \neq i}^p C(u^{(j)}) \right. \\
&\quad \left. + \sum_{k=1}^p \sum_{l=1}^p E[g(u^{(k)})g(v^{(l)})] \prod_{i=1, i \neq k}^p C(u^{(i)}) \prod_{j=1, j \neq l}^p C(v^{(j)}) \right) \sqrt{w(u)w(v)}.
\end{aligned}$$

Using the result from Theorem 3.4.1 that $E[\mathcal{B}(u)\mathcal{B}(u')] = C(u \vee u') - C(u)C(u')$ and our assumption of independence we get the desired expression for the covariance function.

A.0.3 Proof of Theorem 3.5.4

Proof

$$I_n^w = n \int_{[0,1]^d} \left(C_n(u)^2 W(u) - 2C_n(u) \prod_{k=1}^p C_n(u^{(k)}) W(u) + \left[\prod_{k=1}^p C_n(u^{(k)}) \right]^2 W(u) \right) du$$

. Integrating the first term we have

$$\begin{aligned}
& \int_{[0,1]^d} C_n(u)^2 W(u) du = \int_{[0,1]^d} \left(\frac{1}{n} \sum_{i=1}^n \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \frac{1}{n} \sum_{m=1}^n \prod_{l=1}^d \mathbf{1}(\hat{u}_{lm} \leq u_m) \right) W(u) du \\
&= \int_{[0,1]^d} \left(\frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) W(u) du \\
&= \frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \sum_{m=1}^d \int_{[0,1]^d} \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) W(u) du \\
&= \frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \sum_{m=1}^d \left[\prod_{j=1, j \neq m}^d (1 - \hat{u}_{ik} \vee \hat{u}_{lk}) \right] \tilde{w}(\hat{u}_{im} \vee \hat{u}_{lm}, 1).
\end{aligned}$$

Further, for $s \in \{1, \dots, d\}$, let $k_s \in S = \{1, \dots, p\}$ be such that $s \in \{b_{k-1} + 1, \dots, b_k\}$. Then, for the second term we get

$$\begin{aligned}
& \int_{[0,1]^d} C_n(u) \prod_{k=1}^p C_n(u^{[k]}) W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] \prod_{k=1}^p C_n(u^{[k]}) W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{k=1}^p \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] \prod_{k=1}^p C_n(u^{[k]}) W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{k=1}^p C_n(u^{[k]}) \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{k=1}^p \left(\frac{1}{n} \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] W(u) du \\
&= \frac{1}{n^{p+1}} \int_{[0,1]^d} \left[\sum_{i=1}^n \prod_{k=1}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] \left(\sum_{s=1}^d w(u_s) \right) du \\
&= \frac{1}{n^{p+1}} \int_{[0,1]^d} \sum_{i=1}^n \sum_{s=1}^d w(u_s) \left[\prod_{k=1}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] du \\
&= \frac{1}{n^{p+1}} \sum_{i=1}^n \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \int_{[0,1]^{d^k}} \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) du^{[k]} \right] \\
&\quad \times \left[\int_{[0,1]^{d^{k_s}}} \sum_{l=1}^n \prod_{j=b_{k_s-1}+1}^{b_{k_s}} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) w(u_s) du^{[k_s]} \right] \\
&= \frac{1}{n^{p+1}} \sum_{i=1}^n \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\
&\quad \times \left[\sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right].
\end{aligned}$$

Lastly, integrating the third term we have

$$\begin{aligned}
& \int_{[0,1]^d} \left(\prod_{k=1}^p C_n(u^{(k)}) \right)^2 W(u) du \\
&= \int_{[0,1]^d} \prod_{k=1}^p \left(\frac{1}{n} \sum_{i=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right) \prod_{k=1}^p \left(\frac{1}{n} \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) W(u) du \\
&= \int_{[0,1]^d} \prod_{k=1}^p \left(\frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) \sum_{s=1}^d w(u_s) du \\
&= \frac{1}{n^{2p}} \int_{[0,1]^d} \sum_{s=1}^d w(u_s) \left[\prod_{k=1}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] du \\
&= \frac{1}{n^{2p}} \int_{[0,1]^d} \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] \\
&\quad \times \left[w(u_s) \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k_s-1}+1}^{b_{k_s}} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] du \\
&= \frac{1}{n^{2p}} \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\
&\quad \times \left[\sum_{i=1}^n \sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right].
\end{aligned}$$

A.0.4 Proof of Theorem 3.7.4

Proof Since by Lemma 3.7.2, $\sqrt{n} \{ \tilde{K}_n(x, e) - \tilde{K}(x, e) \} \rightsquigarrow d(z, e)$, by applying the functional Delta-method (van der Vaart and Wellner (1996), Theorem 3.9.4) combined with the result from Lemma 3.7.3 we have that

$$\begin{aligned}
& \sqrt{n} \{ \psi'_h(\tilde{K}_n)(x, e) - \psi'_H(\tilde{K})(x, e) \} \\
&= \sqrt{n} \{ \tilde{K}_n(F_n^{-1}(u), G_n^{-1}(v)) - \tilde{K}(F^{-1}(u), G^{-1}(v)) \} \\
&= \sqrt{n} \{ \tilde{C}_n(u, v) - \tilde{C}(u, v) \} \rightsquigarrow \psi'_H(d)(u, v).
\end{aligned}$$

A.0.5 Proof of Theorem 3.7.5

Proof Apply $\phi(a)(x)$ to \tilde{C}_n to get

$$\begin{aligned}
& \sqrt{n} \{ \phi(\tilde{C}_n)(u, v) - \phi(\tilde{C})(u, v) \} = \\
&= \sqrt{n} \{ (\tilde{C}_n(u, v) - \tilde{C}_{x,n}(u)v) \sqrt{w(u, v)} - (\tilde{C}(u, v) - \tilde{C}_x(u)v) \sqrt{w(u, v)} \} \\
&= \sqrt{n} \{ (\tilde{C}_n(u, v) - \tilde{C}_{x,n}(u)v) \sqrt{w(u, v)} \}
\end{aligned}$$

under H_0 . Applying the Delta-method once more (van der Vaart and Wellner (1996), Theorem 3.9.4) we get that $\tilde{C}_n^I \rightsquigarrow \phi'_{\tilde{C}_n}(\zeta)(u, v) = \zeta(u, v)$, with covariance function given by $E[\zeta(u, v)\zeta(u', v')] = E[\phi'_{\tilde{C}_n}(\zeta)(u, v)\phi'_{\tilde{C}_n}(\zeta)(u', v')]$...

A.0.6 Proof of Lemma 3.7.2

Proof The proof follows from Lemma 3.7.1 and from Donsker's theorem.

A.0.7 Proof of Theorem 4.2.3

The proof follows directly from Theorem 4.2.2, and the Continuous Mapping Theorem.

A.0.8 Proof of Theorem 4.2.4

$$D(C_n^S) = \int_{[0,1]^d} \left(M(\mathbf{u}) - \prod_{k=1}^p C_n^S(\mathbf{u}_{\{j\}}) \right)^2 d\mathbf{u} \quad (\text{A.1})$$

$$= \int_{[0,1]^d} M^2(\mathbf{u}) d\mathbf{u} - 2 \int_{[0,1]^d} M(\mathbf{u}) \prod_{k=1}^p C_n^S(\mathbf{u}_{\{j\}}) d\mathbf{u} + \int_{[0,1]^d} \left(\prod_{j=1}^p C_n^S(\mathbf{u}_{\{j\}}) \right)^2 d\mathbf{u}. \quad (\text{A.2})$$

For the first term, from Gaißer, Ruppert, and Schmid (2010) (Result A.1, Appendix A) we have that

$$\int_{[0,1]^d} M^2(\mathbf{u}) d\mathbf{u} = \int_{[0,1]^d} (\min\{u_1, \dots, u_d\})^2 d\mathbf{u} = \frac{2}{(d+1)(d+2)}. \quad (\text{A.3})$$

The last integral can be obtained either as a corollary to Theorem 4.2.5, by substituting unit weights $w(\mathbf{u}) = 1$, $\forall \mathbf{u} \in [0, 1]^d$, into the computational formula, or using Proposition 10 of

Kojadinovic and Holmes (2009). Lastly, for the second term, we have:

$$\int_{[0,1]^d} M(\mathbf{u}) \prod_{j=1}^p C_n^s(\mathbf{u}_{(j)}) d\mathbf{u} \quad (\text{A.4})$$

$$= \int_{[0,1]^d} \min\{u_1, \dots, u_d\} \prod_{k=1}^d \left(\frac{1}{n} \sum_{i=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{i,j} \leq u_j) \right) d\mathbf{u} \quad (\text{A.5})$$

$$= \frac{1}{n^d} \sum_{s=1}^d \int_0^1 \int_{u_s}^1 \dots \int_{u_s}^1 u_s \prod_{k=1}^d \left(\sum_{i=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{i,j} \leq u_j) \right) du_1, \dots, du_s \quad (\text{A.6})$$

$$= \frac{1}{n^d} \sum_{s=1}^d \int_0^1 \int_{u_s}^1 \dots \int_{u_s}^1 \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_p=1}^n u_s \prod_{k=1}^d \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{i_k,s,j} \leq u_j) du_1, \dots, du_s \quad (\text{A.7})$$

$$= \frac{1}{n^d} \sum_{s=1}^d \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_p=1}^n \int_{\hat{u}_{i_1,1}}^1 \int_{\hat{u}_{i_2,2}}^1 \dots \int_{\hat{u}_{i_p,d}}^1 u_s du_1, \dots, du_d \quad (\text{A.8})$$

$$= \frac{1}{n^d} \sum_{s=1}^d \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_p=1}^n \frac{1}{2} (1 - \hat{u}_{i_k,s,s}^2) \prod_{j=1, j \neq s}^d (1 - \hat{u}_{i_k,j,j}). \square \quad (\text{A.9})$$

A.0.9 Proof of Theorem 4.2.5

$$I_n^w = n \int_{[0,1]^d} \left(C_n(u)^2 W(u) - 2C_n(u) \prod_{k=1}^p C_n(u^{(k)}) W(u) + \left[\prod_{k=1}^p C_n(u^{(k)}) \right]^2 W(u) \right) du$$

. Integrating the first term we have

$$\begin{aligned} \int_{[0,1]^d} C_n(u)^2 W(u) du &= \int_{[0,1]^d} \left(\frac{1}{n} \sum_{i=1}^n \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \frac{1}{n} \sum_{l=1}^n \prod_{m=1}^d \mathbf{1}(\hat{u}_{lm} \leq u_m) \right) W(u) du \\ &= \int_{[0,1]^d} \left(\frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) W(u) du \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \sum_{m=1}^d \int_{[0,1]^d} \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) W(u) du \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \sum_{m=1}^d \left[\prod_{j=1, j \neq m}^d (1 - \hat{u}_{ik} \vee \hat{u}_{lk}) \right] \tilde{w}(\hat{u}_{im} \vee \hat{u}_{lm}, 1). \end{aligned}$$

Further, for $s \in \{1, \dots, d\}$, let $k_s \in S = \{1, \dots, p\}$ be such that $s \in \{b_{k-1} + 1, \dots, b_k\}$. Then, for the second term we get

$$\begin{aligned}
& \int_{[0,1]^d} C_n(u) \prod_{k=1}^p C_n(u^{(k)}) W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{j=1}^d \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] \prod_{k=1}^p C_n(u^{(k)}) W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{k=1}^p \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] \prod_{k=1}^p C_n(u^{(k)}) W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{k=1}^p C_n(u^{(k)}) \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] W(u) du \\
&= \int_{[0,1]^d} \left[\frac{1}{n} \sum_{i=1}^n \prod_{k=1}^p \left(\frac{1}{n} \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right] W(u) du \\
&= \frac{1}{n^{p+1}} \int_{[0,1]^d} \left[\sum_{i=1}^n \prod_{k=1}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] \left(\sum_{s=1}^d w(u_s) \right) du \\
&= \frac{1}{n^{p+1}} \int_{[0,1]^d} \sum_{i=1}^n \sum_{s=1}^d w(u_s) \left[\prod_{k=1}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] du \\
&= \frac{1}{n^{p+1}} \sum_{i=1}^n \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \int_{[0,1]^{d^k}} \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) du^{(k)} \right] \\
&\quad \times \left[\int_{[0,1]^{d^{k_s}}} \sum_{l=1}^n \prod_{j=b_{k_s-1}+1}^{b_{k_s}} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) w(u_s) du^{(k_s)} \right] \\
&= \frac{1}{n^{p+1}} \sum_{i=1}^n \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\
&\quad \times \left[\sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right].
\end{aligned}$$

Lastly, integrating the third term we have

$$\begin{aligned}
& \int_{[0,1]^d} \left(\prod_{k=1}^p C_n(u^{(k)}) \right)^2 W(u) du \\
&= \int_{[0,1]^d} \prod_{k=1}^p \left(\frac{1}{n} \sum_{i=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \right) \prod_{k=1}^p \left(\frac{1}{n} \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) W(u) du \\
&= \int_{[0,1]^d} \prod_{k=1}^p \left(\frac{1}{n^2} \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right) \sum_{s=1}^d w(u_s) du \\
&= \frac{1}{n^{2p}} \int_{[0,1]^d} \sum_{s=1}^d w(u_s) \left[\prod_{k=1}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] du \\
&= \frac{1}{n^{2p}} \int_{[0,1]^d} \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] \\
&\quad \times \left[w(u_s) \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k_s-1}+1}^{b_{k_s}} \mathbf{1}(\hat{u}_{ij} \leq u_j) \mathbf{1}(\hat{u}_{lj} \leq u_j) \right] du \\
&= \frac{1}{n^{2p}} \sum_{s=1}^d \left[\prod_{k=1, k \neq k_s}^p \sum_{i=1}^n \sum_{l=1}^n \prod_{j=b_{k-1}+1}^{b_k} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right] \\
&\quad \times \left[\sum_{i=1}^n \sum_{l=1}^n \left(\prod_{j=b_{k_s-1}+1, j \neq s}^{b_{k_s}} (1 - \hat{u}_{ij} \vee \hat{u}_{lj}) \right) \tilde{w}(\hat{u}_{is} \vee \hat{u}_{ls}, 1) \right]. \square
\end{aligned}$$

Curriculum Vitae

Name: Ivan Medovikov **Born:** 1984 in Kirov, Russia

Post-Secondary Education and Degrees: The University of Western Ontario, London, Ontario, Canada
2004 - 2007 Hons. B.A. (Economics)

The University of Western Ontario, London, Ontario, Canada
2008 - 2009 M.A. (Economics)

The University of Western Ontario, London, Ontario, Canada
2009 - 2013 Ph.D. (Economics)

Honours and Awards: Ontario Graduate Scholarship (May 2011)
Western Social Science Alumni Association Award (April 2011)
Western Social Science Research Fellowship (May 2010)
Western Graduate Research Fellowship (May 2007)

Related Work Experience: Assistant Professor of Economics
Brock University, St. Catharines, Ontario, Canada
2012 -

Teaching Assistant
The University of Western Ontario, London, Ontario, Canada
2008 - 2012

Research Assistant
Bank of Canada, Ottawa, Ontario, Canada
2007

Publications:

Keshishbanoosy, R., St-Amant, P., Ball, D., Medovikov, I., A Money and credit real-time database for Canada, Bank of Canada Review, Summer 2008, pp. 55-64.