Essays on Modelling state-dependent dynamics: Applications to financial time series

by

Konstantin Kuck

A thesis submitted to the
Faculty of Business, Economics and Social Sciences,
in partial fulfillment of the requirements for the Degree
Doctor oeconomiae (Dr. oec.)
at the
University of Hohenheim

Date of submission: September 5, 2019

This thesis has been accepted in 2019 as a dissertation to attain the degree of *Doctor oeconomiae (Dr. oec.)* by the Faculty of Business, Economics and Social Sciences at the University of Hohenheim.

Date of defense: November 26, 2019

Professor Dr Karsten Hadwich, Dean, University of Hohenheim
Professor Dr Tereza Tykvová, Chair of the doctoral committee, University of Hohenheim
Professor Dr Robert C. Jung, Supervisor, University of Hohenheim
Professor Dr Dirk G. Baur, Second Supervisor, University of Western Australia
PD Dr Thomas Dimpfl, University of Tübingen (on behalf of Professor Dr Dirk G. Baur)

Acknowledgements

While working on this thesis, I received essential support from many people to whom I would like to express my gratitude: First of all, I want to thank my supervisor Professor Dr Robert Jung for his guidance, advice and encouragement and for providing a pleasant research environment. And it were his interesting lectures which initiated my interest for time series and financial econometrics. Further, I want to thank my second supervisor Professor Dr Dirk G. Baur: I am grateful for many interesting discussions, his advice and his commitment to our joint research projects. Likewise, I want to thank him for inviting me to Perth for a stimulating and motivating research stay. Moreover, I thank Professor Dr Tereza Tykvová for joining the committee and serving as chair for the defense. Many thanks also go to PD Dr Thomas Dimpfl for joining the committee on short notice. Further, I owe gratitude to Dr Robert Maderitsch for his ongoing commitment to our joint research projects, even after he had left academia. Together with him I co-authored various research papers, two of which are part of this thesis. Special thanks go to Dr Karsten Schweikert for thorough discussions and his advice that helped to sharpen the exposition and improve the quality of this thesis. Special thanks also go to Dr Vincent Dekker: I benefited a lot from his substantial contributions and our discussions, which did not necessarily end with working hours. I also owe thanks to my former and present colleagues at the Department of Econometrics and Statistics who contributed to this thesis in manifold ways: Kerstin Bubeck, Dr Stephanie Glaser, Markus Mößler, Dr Alexander Schmidt and Susanne Schultz-Roth. Last, and not least, I am deeply grateful to my parents. Their loving support has been essential for the success of my work.

Essays on Modelling state-dependent dynamics: Applications to financial time series

Konstantin Kuck

Abstract. This thesis explores state-dependence in the context of financial market dynamics and cross-market linkages. Time-varying behaviour of financial markets is widely observed and implies that their price dynamics are characterized by *state-dependence* with regard to changing economic conditions. From a statistical perspective, this means that the (inter-)dependencies of financial variables are *non-linear* and cannot be adequately described in the context of linear models. Using non-linear econometric models like quantile (auto)regression and Markov-switching models, this thesis focuses on the following issues:

1. Are the dynamics among crude oil prices stable or time-varying? Are the crude oil markets generally integrated or 'regionalized'? Is there a leading benchmark price?

2. How are the volatility dynamics of crude oil and precious metals affected by the level of volatility? Are there differences between crude oil and precious metals?

3. How fast do investors react to negative shocks in the equity market? Do negative shocks in the equity market affect the volatility of gold and what are the implications for the role of gold as a safe haven?

4. What can be learned from intra-day data about temporal dependencies and information processing in the foreign exchange (FX) market?

All markets and assets considered are economically and financially important and, therefore, insights regarding the issues outlined above are relevant for international investors and producers, but might be of interest for policy makers as well.

Crude oil is an important economic input factor, and therefore, understanding interdependencies of crude oil benchmark prices is of direct relevance with regard to production planning, asset allocation and risk management. In particular, it is of interest if there is a benchmark price which 'globally' leads the crude oil price formation process in the sense that it reflects all available information first and, therefore, can be used for an assessment of the *global* price for crude oil. In this context, the question whether the crude oil market is integrated or 'regionalized' is of relevance as well since it implies if changing conditions in one region affect other geographical regions through the presence of arbitrage opportunities. Applying a Markov-switching vector error correction model to five crude oil benchmark prices (WTI, Brent, Bonny Light, Dubai, Tapis), we find the world crude oil market to be integrated, although the *degree* of integration is varying over time. More specifically, we identify three distinct regimes characterized by different price dynamics. First, we find that the world crude oil market was regionalized throughout the 1980s but became integrated in the subsequent years. Second, whilst the crude oil market is generally integrated after the 1980s, the degree of integration varies with the level of global economic uncertainty. Further, our empirical results suggest that there is no benchmark which universally leads the other crude oil prices at all times. Therefore, rather the system of benchmark prices should be considered for a precise assessment of the global crude oil price.

Apart from crude oil, commodity markets in general attracted the attention of institutional investors in recent years, known in the literature as 'financialization' of commodity markets. In contrast to commodity producers and processors who aim to hedge price risks, institutional investors perceive commodity futures itself as an asset which they use in portfolio diversification. Since the volatility behaviour might vary across commodities due to their different characteristics (e.g., crude oil is a non-durable 'consumption' good whereas gold is durable and provides a 'store of value'), a profound understanding of their individual volatility dynamics may help to limit losses and to reduce the overall risk of a portfolio under

changing market conditions. Focusing on gold, silver and crude oil futures, we find the dynamics of commodity volatility to be state-dependent with regard to the level of current volatility: Our findings from a quantile heterogeneous autoregressive model of realized volatility (Q-HAR-RV) suggest that information generated over the medium-term becomes more important when uncertainty is high. From an economic viewpoint, this is consistent with the idea of shifts in investor attention and investment horizons depending on uncertainty and market risk. We also document differences in the (average) volatility dynamics between commodities, which seem to be associated with their individual characteristics: For crude oil, past monthly volatility is of particular importance which might be related to its role as an economic input factor and the rather long-term time horizon of investors. By contrast, in the case of gold and silver, past daily and monthly volatility appear to be equally important for current volatility which may reflect that gold and silver have industrial use but are important 'investment' assets as well.

Moreover, the relationship between gold and equity is of interest for investors in the context of risk management since empirical evidence suggests that gold holds its value in presence of a negative shock in the stock market. This phenomenon, referred to as the *safe haven effect* of gold, implies that losses due to falling equity prices could be limited by diversifying the portfolio with gold. However, in using gold as a safe haven, it is necessary to understand how strong and how fast the gold price reacts to a negative shock in stock prices. In addition, the response in the gold volatility, is of relevance as well since it affects the risk of the portfolio.

Using high frequency intra-day gold and S&P500 data, we find that on average, gold prices do not move in tandem with stocks when the equity market is in distress. Further, we show that gold holds its value also in the overnight period following a negative shock in the equity market. Although ex-post observed negative shocks in equity prices accumulate over the trading day, extremely strong negative 5 minute returns lead to a positive reaction of the gold price, implying a fast reaction of the gold price to extreme equity price declines. Further, based on a quantile regression model for gold returns measured over the stock trading hours, we document that the absolute size of returns does not increase due to a negative shock in the equity market. In contrast, an analysis of realized gold variances based on the heterogeneous autoregressive model of realized volatility with exogenous regressors (HAR-RV-X) reveals contemporaneously higher realized variances of gold when the equity market is in distress. However, the implications for the safe haven effect of gold appear to be limited for investors with daily and longer trading horizon since the variation in daily returns appears not to be affected. Since realized measures also capture the influence of intra-daily returns, we interpret the higher realized variance as increased intra-daily uncertainty in the gold market.

The foreign exchange (FX) market 'connects' all financial markets around the globe since transactions in stock, bond or commodity markets might require currency trades as a by-product and changes in FX rates may affect the value of international investment portfolios, at least in the short-run. Therefore, a profound understanding of the FX rate short-term dynamics is of relevance not only for market makers and currency traders in order to provide liquidity but for high frequency traders exposed to currency risk as well. Using quantile autoregression techniques, we show that the intra-day FX market dynamics are characterized by state-dependence. More precisely, non-linear temporal dependence is present in FX returns observed at various intra-day sampling frequencies from 10 minutes to 3 hours. Moderate changes in an FX rate exhibit statistically significant negative serial correlation, whereas strong intra-day appreciations and depreciations over a short time interval tend to be characterized by the absence of autocorrelation. This finding may be attributed to *diverging opinions* of traders regarding the impact of news on the direction of the FX rate, but also suggests a tendency for an immediate adjustment when unambiguous news – reflected by extreme FX rate returns – arrive. Further, we show that the dependence structure is not affected by diurnal volatility variation due to the level of trading activity and is stable with regard to increased daily FX return volatility.

The central message behind all studies is identical: The dynamics of financial markets are char-

acterized by state-dependence which must be taken into account in risk management and the design of forecasting models. From a statistical perspective, our findings suggest that results based on linear models may be highly inaccurate.

Essays on Modelling state-dependent dynamics: Applications to financial time series

Konstantin Kuck

Kurzzusammenfassung. Ziel dieser Arbeit ist die empirische Untersuchung von Zustandsabhängigkeiten in Finanzmarktdynamiken und Beziehungen zwischen Finanzmärkten. Zeitabhängiges Verhalten von Finanzmärkten ist allgemein zu beobachten und weist darauf hin, dass deren Dynamiken durch Zustandsabhängigkeit in Hinsicht auf sich verändernde wirtschaftliche Bedingungen gekennzeichnet sind. Aus statistischer Perspektive bedeutet dies, dass die (Inter-)Dependenzen zwischen Finanzmarktvariablen Nicht-Linearitäten aufweisen, welche in linearen Modellen nicht angemessen beschrieben werden können. Diese Arbeit verwendet daher nicht-lineare ökonometrische Ansätze, wie Quantils(auto)regression und sogenannte Markov-switching Modelle, um folgenden Fragen zu untersuchen:

1. Sind die Dynamiken von Rohölpreisen zeitstabil oder zeitvariabel? Ist der Rohölmarkt im Allgemeinen integriert oder fragmentiert? Ist einer der Referenzpreise Preissetzer?

2. Wie wirkt sich das Niveau der Volatilität von Rohöl und Edelmetallen auf deren Preisdynamiken aus? Gibt es Unterschiede zwischen Rohöl und Edelmetallen?

3. Wie schnell reagiert der Goldpreis auf negative Schocks im Aktienmarkt? Beeinflussen negative Schocks im Aktienmarkt die Volatilität von Gold und welche Implikationen ergeben sich daraus für die „Safe Haven"-Eigenschaft von Gold?

4. Welche Erkenntnisse hinsichtlich Zeitabhängigkeit in Wechselkursrenditen und Informationsverarbeitung im Wechselkursmarkt lassen sich aus intra-täglichen Daten gewinnen?

Alle betrachteten Märkte sind aus ökonomischer Sicht von Bedeutung und die in dieser Arbeit untersuchten Fragen für Investoren, Produzenten aber auch für politische Entscheidungsträger von Interesse.

Rohöl ist ein wesentlicher wirtschaftlicher Inputfaktor. Ein umfassendes Verständnis der Beziehungen zwischen einzelnen Rohölpreisen ist daher für die Produktionsplanung, die Allokation von Ressourcen und das Risikomanagement von großer Bedeutung. Von besonderem Interesse ist dabei, ob es einen Rohölpreis gibt, der Marktinformationen zu jeder Zeit zuerst widerspiegelt und damit für die Bewertung des *globalen* Rohölpreises herangezogen werden kann. Darüberhinaus ist die Frage, ob der Rohölmarkt *integriert* oder *fragmentiert* ist von Bedeutung: In einem integrierten Markt beeinflussen veränderte Bedingungen in einer Region über Arbitragemöglichkeiten auch andere Regionen, wohingegen deren Einfluss in einem fragmentierten Markt auf eine Region beschränkt ist. Die Ergebnisse aus einem Markov-switching Vektor Fehlerkorrekturmodells für fünf Rohöl-Referenzpreise (WTI, Brent, Bonny Light, Dubai, Tapis) zeigen, dass der Rohölmarkt generell integriert ist, das Maß der Integration aber über die Zeit hin variiert. Wir identifizieren drei Regime, die unterschiedliche Preisdynamiken aufweisen und Veränderungen im Rohölmarkt abbilden. Erstens, der Rohölmarkt war in 1980er Jahren ein fragmentierter Markt der sich zu Beginn der 1990er Jahre zu einem integrierten Markt entwickelt hat. Zweitens variiert das Maß der Integration über die Zeit und hängt von der wirtschaftlichen Unsicherheit ab. Desweiteren deuten unsere Ergebnisse darauf hin, dass keiner der untersuchten Referenzpreise als Preissetzer fungiert. Für eine akkurate Einschätzung des globalen Rohölpreises sollten daher verschiedene Rohölpreise betrachtet werden.

Neben dem Rohölmarkt haben Rohstoffmärkte allgemein die Aufmerksamkeit institutioneller Investoren auf sich gezogen, was in der Literatur häufig als „financialization" der Rohstoffmärkte bezeichnet wird. Im Gegensatz zu Rohstoffproduzenten und Rohstoffverarbeitern, deren Handelsmotiv in der

Absicherung von Rohstoffpreisrisiken besteht, betrachten institutionelle Investoren Rohstoff-Futures als Finanzanlage, die sie zur Diversifikation ihrer Portfolios einsetzen. Da sich die Dynamik der Volatilität aufgrund ihrer unterschiedlichen Eigenschaften zwischen den einzelnen Rohstoffen potenziell unterscheidet (z.B. ist Rohöl ein „Verbrauchsgut" während Gold eine Wertanlage darstellt), ist ein fundiertes Verständnis der individuellen Volatilitätsdynamiken hilfreich, um Verluste und das Risiko des gesamten Portfolios unter wechselnden Marktbedingungen zu mindern. Unsere Untersuchung der Gold-, Silber- und Rohöl-Volatilität zeigt deren Zustandsabhängigkeit hinsichtlich der gegenwärtigen Volatilität. Mittels eines sogenannten Quantile Heterogeneous Autoregressive Model of Realized Volatility (Q-HAR-RV) zeigen wir, dass in Zeiten hoher Unsicherheit (hoher Volatilität) Informationen, die über einen wöchentlichen Zeitraum generiert wurden, substantiell an Bedeutung gewinnen. Aus ökonomischer Sicht steht dies in Einklang mit der Idee, dass die Aufmerksamkeit von Investoren sowie deren Anlagehorizont im Zeitverlauf schwanken und sich dem Marktrisiko anpassen. Wir dokumentieren zudem Unterschiede in den „durchschnittlichen" Volatilitätsdynamiken, die in Beziehung zu deren individuellen Charakteriska zu stehen scheinen: Im Fall von Rohöl ist für die gegenwärtige Volatilität vor allem die Volatilität des vergangenen Monats von Relevanz, was mit der wirtschaftlichen Relevanz von Rohöl zusammenhängen könnte und auf einen eher langen Anlagehorizont der im Rohölmarkt aktiven Investoren hinweist. Im Gegensatz dazu ist für die Volatilität von Gold und Silber die Volatilität des vergangenen Tages und Monats gleichermaßen relevant. Dies könnte die Tatsache widerspiegeln, dass beide Edelmetalle sowohl in der Industrie eingesetzt werden, also auch wichtige Finanzanlagen darstellen.

Darüberhinaus ist die Beziehung zwischen dem Gold- und Aktienmarkt im Kontext von Risikomanagement von besonderer Relevanz, da Gold dazu tendiert, seinen Wert zu behalten, wenn es am Aktienmarkt zu starken Verlusten kommt. Dieses empirische Phänomen, was in der Literatur als "Safe-Haven-Effekt" bezeichnet wird, impliziert, dass Verluste aufgrund fallender Aktienpreise durch Diversifikation des Portfolios durch Gold gemindert werden könnten. Dafür ist es jedoch erforderlich, zu verstehen, wie stark und wie schnell der Goldpreis auf negative Schocks im Aktienmarkt reagiert. In diesem Zusammenhang erscheint auch die Reaktion der *Goldvolatilität* von Bedeutung, da diese das Risiko des Portfolios beeinflusst. Unter Verwendung intra-täglicher Daten von Gold und S&P500 zeigen wir, dass Gold durchschnittlich seinen Wert behält wenn die Aktienpreise massiv fallen. Wir zeigen darüberhinaus, dass Gold in dieser Situation seinen Wert auch nach Handelsschluss am Aktienmarkt behält. Obwohl die ex-post beobachteten extremen Verluste im Aktienmarkt über einen längeren Zeitraum während des Handelstages entstehen, führen (relativ) extrem negative fünf-minütliche Aktienrenditen zu einer positiven Reaktion des Goldpreises, was auf eine schnelle Reaktion des Goldpreises hindeutet. Mittels Quantilsregression zeigen wir darüberhinaus, dass die absolute Größe der über den Handelszeitraum des Aktienmarktes gemessenen Goldpreisveränderungen nicht durch Schocks im Aktienmarkt beeinflusst wird. Die Analyse der realisierten Varianzen von Gold in einem Heterogeneous Autoregressive Model of Realized Volatility mit exogenen Regressoren (HAR-RV-X) zeigt jedoch eine deutlich höhere kontemporäre Volatilität von Gold, an Tagen, an denen die Aktienpreise massiv fallen. Allerdings dürfte der Einfluss auf den Safe Haven Effekt von Gold, insbesondere für Investoren mit einem Anlagehorizont von mehr als einem Tag begrenzt sein, da die Gesamtpreisveränderung von Gold von Schocks im Aktienmarkt unbeeinflusst zu bleiben scheint. Da Realisierte Varianzen den Einfluss intra-täglicher Renditen auf die Volatilität berücksichtigen, deuten unsere Ergebnisse darauf hin, dass Schocks im Aktienmarkt mit höherer (intra-täglicher) Unsicherheit am Goldmarkt einhergehen.

Der Wechselkursmarkt (FX Markt) „verknüpft" die übrigen Finanzmärkte, da Transaktionen in Aktien-, Anleihe- oder Rohstoffmärkten auch Devisengeschäfte erfordern können. Zudem können Wechselkursschwankungen den Wert international investierter Portfolios, zumindest kurzfristig, beeinflussen. Fundierte Kenntnisse über intra-tägliche Dynamiken am FX Markt sind daher nicht nur für Marketmaker und Devisenhändler, sondern auch für Tageshändler die Währungsrisiken ausgesetzt sind, von Bedeutung. Basierend auf Quantilsauto-regressions-Modellen zeigen wir, dass die intra-täglichen Wechselkursrenditen

nicht-lineare Zeitabhängigkeiten aufweisen: Moderate Wechselkursrenditen sind signifikant negativ autokorreliert während starke Auf- oder Abwertungen keine Autokorrelation aufweisen. Dieser Befund kann auf *divergierende Einschätzungen* der Marktteilnehmer hinsichtlich der Auswirkung neuer Informationen auf den Devisenkurs zurückgeführt werden. Er zeigt jedoch auch, dass es bei Eintreten eindeutiger Nachrichten – reflektiert durch extreme Wechselkursrenditen – zu einer sofortigen Anpassung des Wechselkurses kommt. Diese Abhängigkeitsstruktur scheint weder durch Schwankungen in der Handelsaktivität noch durch erhöhte tägliche Volatilität der Devisenrendite beeinflusst zu werden.

Zentraler Befund der Arbeit ist: Die Dynamiken der Finanzmärkte sind geprägt von Zustandsabhängigkeit, die beim Risikomanagement und in Prognosemodellen berücksichtigt werden muss. Aus statistischer Sicht deuten die Ergebnisse darauf hin, dass Resultate die auf linearen Modellen basieren sehr ungenau sein können.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis investigates the time-varying dynamics of financial markets and cross-market linkages. The relevance of a profound understanding of financial market behaviour under different market conditions, and particularly in times of financial distress, was highlighted by the events of the Global Financial Crisis, which hit financial markets and subsequently affected the real sectors of economies around the globe.

From an econometric perspective, financial assets are characterized by various variables, e.g. prices, returns and volatilities, which can be measured at different sampling frequencies. Therefore, the temporal and the cross-market dependencies are multifaceted: For instance, (temporal) dependencies might be present in some variables (e.g. prices or volatilities) but absent in other variables (e.g. returns). In addition, although returns can typically be directly observed and analyzed at various sampling frequencies, ranging from low (e.g. monthly, weekly, daily) to high and ultra-high frequencies (e.g. intra-daily and tick-data), the return volatility is latent and must be estimated from return data or can be derived from high-low price ranges. Further, since financial markets might operate at small time intervals, dependencies in financial variables might be confined to particular sampling frequencies. In other words, (temporal) depedencies in financial variables might only be observable in high frequency intra-day data (transaction data, price quotes) but not in data sampled at lower frequencies due to the aggregation of information. Financial data sampled at (ultra-)high frequency, however, have many unattractive characteristics, including measurement errors, microstructure noise or missing data, and need to be cleaned before any econometric analysis. Moreover, financial variables exhibit different statistical properties which must be taken into account in the statistical analysis and thus add an additional level of complexity: Prices are non-stationary, financial returns are usually uncorrelated and characterized by conditional heteroscedasticity, whereas their volatilities are stationary but exhibit long memory, i.e. they are highly persistent but do not have a unit root. Apart from these aspects, the observed time-varying behavior of financial markets suggests that their dynamics are characterized by *state-dependence* with regard to changing economic conditions. From a statistical viewpoint, this implies that the (inter-)dependencies of financial variables are *non-linear*, and therefore, cannot be adequately captured by linear models or measures of linear dependence. More precisely, if the dynamics and relationships exhibit non-linearities, linear models may provide only an incomplete description of the dependencies over both time and across assets.

Therefore, this thesis employs non-linear econometric models such as quantile (auto-)regression in the spirit of Koenker and Bassett (1978) and Koenker and Xiao (2006) and Markov-Switching models

according to Hamilton (1989) to study in detail the temporal dependencies in financial time series as well as financial market interdependencies. The choice of quantile regression techniques is motivated by Baur (2013) who demonstrates that quantile regression provides a useful modeling and estimation method for both the *structure* and the *degree* of dependence inherent in financial time series. In contrast to ordinary least squares (OLS) regression which is focused on the conditional mean of the dependent variable, the quantile regression framework allows to model specific quantiles of the dependent variable as a linear function of the explanatory variable(s). From an economic perspective, different quantiles of the dependent variable may represent different 'states' or market conditions, which makes quantile (auto)regression a useful tool in the context of this thesis. Since the model can be estimated individually for any (conditional) percentile of interest, quantile regression allows to uncover non-linear dependencies with respect to the level of the dependent variable, and to precisely estimate the dependence structure from the data without the need to assume in advance the number of 'states'. Another, different approach to capture and describe non-linearities is provided by the Markov-switching model introduced by Hamilton (1989). Like in the threshold autoregressive (TAR) model by Tong (1983), the basic idea of the Markov-switching model is that a process can be adequately described by a 'piece-wise' linearization. Non-linearities are incorporated by allowing for linear approximations in a (pre-specified) number of different states (or regimes) where the variable that governs the transition between the different states is unobserved. This is in contrast to the TAR model where the state is governed by an observable ('state') variable and depends on pre-specified thresholds thereof. Compared to quantile regression, Markov-switching and TAR models are hence more flexible in the sense that also non-linearities which are not directly governed by the dependent variable can be described.

Using the techniques outlined above, this thesis presents a differentiated analysis of temporal dependencies as well as cross-market linkages of different financial and commodity markets. Specifically, the following issues and aspects of financial market dynamics and interdependencies are studied:

1. **Are the dynamics among crude oil prices stable or time-varying? Are the crude oil markets generally integrated or 'regionalized'? Is there a leading benchmark price?**

2. **How are the volatility dynamics of crude oil and precious metals affected by the level of volatility? Are there differences between crude oil and precious metals?**

3. **How fast do investors react to negative shocks in the equity market? Do negative shocks in the equity market affect the volatility of gold and what are the implications for the role of gold as a safe haven?**

4. **What can be learned from intra-day data about temporal dependencies and information processing in the foreign exchange (FX) market?**

All markets and assets considered are economically and financially important and thus, insights regarding the issues outlined above are relevant for international investors and producers, but might be of interest for policy makers as well.

Crude oil, for instance, is used for transportation and heating fuel and, hence, is an important economic input factor. Moreover, it is employed in electricity generation and the making of chemicals, plastics, and synthetic materials. A solid understanding of the interplay among crude oil benchmark prices from different regions is of relevance with regard to crude oil pricing. In particular, it is of interest

if there is a benchmark price that 'globally' leads the crude oil price formation process in the sense that it reflects all available information first and, therefore, can be used for a precise assessment of the *global* price for crude oil. The question whether the crude oil market is integrated or 'regionalized' has been examined extensively in the literature (see e.g. Adelman, 1984; Weiner, 1991; Ji and Fan, 2015). Specifically, in an integrated market, changes in (market) conditions in one region are assumed to affect other geographical regions through arbitrage opportunities. Answers to these issues are of direct relevance for both producers and investors exposed to crude oil prices with regard to production planning, asset allocation and risk management. There are important policy implications as well, since the members of the International Energy Agency are required to stockpile certain amounts of crude oil in order to provide emergency crude oil in times of disruptive shocks. In a fragmented market, where the effects of supply shocks are restricted to one region, higher reserves would have to be stockpiled than in an integrated market where arbitrage opportunities lead to the supply of, then relatively cheaper, oil from other production sites.

In addition to the relevance of crude oil as an economic input factor and the resulting importance of a precise assessment of the world crude oil price, commodity markets in general gained financial importance in recent years. Traditionally, mainly commodity producers, commodity processors and specialized investors were active in commodity futures markets (Adams and Glück, 2015). While the trading incentive of commodity producers and processors is to hedge their exposure to price fluctuations, the motivation of specialized investors is to earn compensation for providing insurance against price risk. Further, commodity producers and processors have a 'real' and rather long-term interest behind their transactions. Within the last decades, however, institutional investors became interested in commodities as a financial asset (Domanski and Heath, 2007), which is referred to as the 'financialization' of commodity markets in the literature (see e.g. Tang and Xiong, 2012; Cheng and Xiong, 2014). In contrast to the participants traditionally active in commodity markets, institutional investors typically have no interest in the delivery of the underlying commodity, but consider the futures contract itself as a financial asset which they use for portfolio diversification. Hence, a precise description of state-dependence in the volatility dynamics of different commodities is not only important to commodity producers and processors but of direct relevance for financial investors as well. Moreover, the volatility dynamics may differ across commodities due to their different characteristics (for example, crude oil as non-durable 'consumption' good and gold as a 'store of value'). Therefore, a profound understanding of commodity volatility dynamics may help to limit losses and to reduce the overall risk of a portfolio under changing market conditions.

The relation between gold and other financial markets – especially the equity market – has attracted substantial attention in both the academic literature and financial media (see e.g. Baur and McDermott, 2010; Baur and Lucey, 2010). This interest seems mainly motivated by the empirical phenomenon that gold holds its value or exhibits positive returns in presence of a negative shock in equity prices, which is typically referred to as the *safe haven effect* of gold. Naturally, this observed state-dependent relationship is of great interest for investors in the context of risk management as it implies that extreme losses in the equity market could be limited by holding gold. However, in using gold to limit losses when the equity market is in distress, it is necessary to understand how strong and how fast the gold price reacts to a negative shock in stock prices. While the relationship between gold and stock returns is an important aspect of the safe haven effect, their volatility-relationship, and in particular the volatility of

gold in response to negative shocks in the equity market, is of relevance as well. In the context of an equity portfolio diversified with gold, the benefit of a state-dependent negative correlation (absence of correlation) between the two assets may be affected by the volatility of gold in response to a negative shock in the equity market (Baur, 2012). More specifically, the risk of a portfolio composed of the two assets might become higher when the stock market is in distress which may be particularly unattractive in times of uncertainty. Therefore, a comprehensive analysis of the gold-stock return and volatility relationship, will contribute to the understanding of the safe haven effect.

The foreign exchange (FX) market 'connects' all financial markets around the globe since transactions in stock, bond or commodity markets might require currency trades as a by-product. Also, changes in the FX rate may affect the value and profitability of international investment portfolios, at least in the short-run. A profound understanding of the FX rate short-term dynamics, therefore, is of relevance not only for market makers and currency traders in order to provide liquidity but for high frequency traders exposed to currency risk as well. The corresponding statistical feature is (non-)linear temporal dependence. The question whether financial returns exhibit temporal dependence is a key topic in financial economics and has been investigated extensively in the literature (for a review of studies for the FX market see Charles, Darné, and Kim, 2012). Fama (1965) and Samuelson (1965) assume financial markets to be weak-form efficient, which means that returns are characterized by the absence of temporal dependence, implying that future prices are purely unpredictable based on past information. Although there are episodes where FX returns are found to exhibit temporal dependence, the FX market generally seems to be very well-described by weak-form efficiency (Charles et al., 2012). The efficient market hypothesis (EMH), however, is typically focused on the (linear) temporal association of returns which is only one aspect of the dynamics and, in particular, does not consider the pricing behaviour and information processing when more extreme news arrive. Hence, a more differentiated perspective on this issue seems useful, especially for an assessment of the magnitude of future FX appreciations and depreciations at different intra-day time intervals.

This thesis is divided into seven chapters. Chapter 1 motivates the research questions and presents the structure of the thesis. Chapters 2 to 6 consist of individual research papers which can be read independently. In Chapter 2, **A Markov regime-switching model of crude oil market integration**[1], the *globalization-regionalization hypothesis* for the world crude oil market is revisited. We examine long-run equilibrium relationships between major crude oil prices – WTI, Brent, Bonny Light, Dubai and Tapis – and focus on the adjustment behaviour following disequilibrium states. We account for a changing adjustment behaviour over time by using a Markov-switching vector error correction model. Our overall findings suggest that the crude oil market is globalized. Dubai turned out to be the only weakly exogenous price in all regimes, indicating its important role as a benchmark price. Furthermore, an interesting finding of our study is that the degree of market integration seems to be connected to global economic uncertainty.

Whilst Chapter 2 is focused on the linkages among different regional crude oil prices, Chapter 3, **The Quantile-Heterogeneous Autoregressive Model of Realized Volatility: New Evidence from Commodity Markets**[2] widens the perspective and provides a comprehensive view on volatility dynamics in

[1]This chapter is joint work with Karsten Schweikert and has been originally published as KUCK, K. AND K. SCHWEIKERT (2017). "A Markov regime-switching model of crude oil market integration," *Journal of Commodity Markets*, **6**: 16–31. DOI 10.1016/j.jcomm.2017.03.001

[2]This chapter is joint work with Robert Maderitsch and has been originally published as KUCK, K. AND R. MADERITSCH (2019). "The Quantile-Heterogeneous Autoregressive Model of Realized Volatility: New Evidence from

precious metals and crude oil markets and contrasts them to the S&P 500. Using high frequency futures data, we construct realized volatilities and estimate (Quantile) Heterogeneous Autoregressive models for the daily volatility of gold, silver and crude oil futures. We model realized volatility as a linear function of lagged realized volatility, measured over different time resolutions to explicitly account for the potentially heterogeneous impact of market participants with different trading motives and investment horizons. Using quantile regression allows us to identify potential non-linearities and asymmetries in the dependence on short-, mid- and long-term volatilities with respect to different levels of current volatility. We document considerable changes in the relative importance of short-, mid-, and long-term volatility components under varying market conditions. The identified patterns are remarkably similar across the three assets. Specifically, past daily and monthly volatility have a strong positive impact on today's volatility, when current volatility is low (lower quantiles of the volatility distribution). The effect of past weekly volatility, however, increases distinctly from intermediate to higher quantiles of the conditional volatility distribution. The results might indicate considerable investor attention shifts and changes in the proportions of traders with different time horizons.

Chapters 4 and 5 explore the relationship between the equity and gold market. In Chapter 4, **The Timing of the Flight to Gold: An Intra-day Analysis of Gold and the S&P500**[3], we 'zoom in' and use high frequency intra-day gold and S&P500 data covering the period from 2007 to 2018 to investigate when and how fast gold prices react to extreme negative shocks in the equity market. Our empirical analysis reveals three new features of gold: First, extreme negative 5-min S&P500 returns lead to a positive reaction of the gold price. Second, on days with extreme price declines in the stock market, gold continues to increase post US stock trading hours. Third, daily extreme negative equity returns accrue comparatively slowly over several hours. The findings show that there is a fast reaction of gold prices to extreme negative stock returns consistent with a flight to gold.

Chapter 5, **Gold volatility and the safe haven effect**, investigates the contemporaneous response of gold returns and volatility to equity returns of different magnitudes. Whilst studies investigating the role of gold as safe haven asset typically focus on returns, we also investigate the relationship between gold volatility and stock returns. In particular, we are interested in the behaviour of gold volatility on days with negative shocks in the equity market since it is of direct relevance for the effectiveness of gold as a safe haven asset from a portfolio perspective. We find that gold acts as weak safe haven for equity in the sense that it does not move in tandem with stocks when the equity market is in distress. However, we also show that the volatility of gold spot and futures is influenced by the state of the stock market. More specifically, comparing the findings for contemporaneous returns and realized variances of gold, we reveal that the higher volatility reflects increased uncertainty and is not caused by the magnitude of aggregate gold price change over the trading day in presence of a negative shock in the equity market.

Chapter 6 focuses on the FX market which is of particular relevance since it 'connects' all international financial and commodity markets around the globe. Specifically, Chapter 6, **Intra-day dynamics of exchange rates: New evidence from quantile regression**[4], again 'zooms in' and provides a compre-

Commodity Markets," in *Financial Mathematics, Volatility and Covariance Modelling*, Volume 1, ed. by Julien Chevallier, Stéphane Goutte, David Guerreiro, Sophie Saglio, Bilel Sanhaji, New York, NY: Routledge. DOI 10.4324/9781315162737

[3]This chapter is joint work with Dirk G. Baur and has been originally published as BAUR, D. G. AND K. KUCK (2019). "The Timing of the Flight to Gold: An Intra-day Analysis of Gold and the S&P500," *Finance Research Letters*. DOI 10.1016/j.frl.2019.05.005

[4]This chapter is joint work with Robert Maderitsch and has been originally published as KUCK, K. AND R. MADERITSCH (2019). "Intra-day dynamics of exchange rates: New evidence from quantile regression," *The Quarterly Review of Economics and Finance*, **71**, 247–257. DOI 10.1016/j.qref.2018.09.001

hensive description of the intra-day dynamics of major US-Dollar spot exchange rates. We use quantile autoregression to investigate the presence of (non-)linear temporal dependence in foreign exchange returns at various intra-daily time-horizons, ranging from ten minutes up to three hours. Specifically, we investigate an 11-year long sample of non-intermittent high frequency returns for the Euro (EUR), the British Pound (GBP) and the Japanese Yen (JPY) against the US-Dollar (USD). This allows us to take a long-term perspective and to consider potential changes in the dynamics across different market conditions. In contrast to previous studies, we find the temporal dependence of intra-daily foreign exchange returns to be non-linear and symmetrically U-shaped. Specifically, we observe pronounced negative autocorrelation for moderate USD appreciations and depreciations (central quantiles). For extreme positive and negative USD movements, we detect positive autocorrelation. This symmetric non-linear form of temporal dependence is remarkably stable across different exchange rates and states of the market. It appears to be a unique feature of foreign exchange returns and might be related to the fundamental 'two-sidedness' of foreign exchange markets.

Finally, Chapter 7 reviews the key findings and provides a critical assessment of the studies in this thesis.

Chapter 2

A Markov regime-switching model of crude oil market integration†

2.1 Introduction

The discussion on whether world crude oil markets are globalized or regionalized has received a great deal of attention in recent years. Adelman (1984) described the world crude oil market as 'one great pool'. Changes in market conditions in one region are then expected to affect other geographical regions immediately. An existing price differential in local oil markets that exceeds the transportation costs of third party exporters gives rise to arbitrage opportunities. The subsequent supply pressure is expected to close the difference in prices. The idea of 'one great pool' was challenged by Weiner (1991) who finds empirical support for a high degree of regionalization. His findings imply that the world crude oil market is fragmented and the effects of price shocks to regional crude oil prices are restricted to this specific regional market.

This initial discussion has triggered numerous empirical studies, among them Guelen (1999), Fattouh (2010), Reboredo (2011) and Ji and Fan (2015), that tackle the 'globalization-regionalization' hypothesis from different angles. The majority of recent studies finds evidence for a globalized crude oil market. However, the structure of the market does not seem to be stable over time.

Our paper contributes to the literature by proposing a regime-switching model for the long-run relationships among benchmark crude oil prices. This allows us to relax the assumption of constant dynamics over the sample period which has to hold for linear cointegration models. More specifically, we apply a Markov-switching vector error correction model (MSVECM) to capture changing roles of crudes in the world crude oil market and a changing degree of market integration. This enables us to identify regime-shifts from the data without the need to pre-specify structural breaks. We aim to account for increasingly volatile crude oil prices and changing economic and geopolitical conditions over a sample reaching from 1987 to 2015. Our data-set consists of five major crude oil benchmark prices – WTI, Brent, Bonny Light, Dubai and Tapis – representative of five crude oil producing regions.

The question whether the crude oil market is globalized or regionalized has important policy implications. Developed countries hold strategic petroleum reserves to provide emergency crude oil in times

†This chapter is joint work with Karsten Schweikert and has been originally published as KUCK, K. AND K. SCHWEIKERT (2017). "A Markov regime-switching model of crude oil market integration," *Journal of Commodity Markets*, **6**: 16–31. DOI 10.1016/j.jcomm.2017.03.001

of disruptive supply shocks. Members of the International Energy Agency are required to stockpile crude oil equal to 90 days of prior year's net oil imports[1]. If effects of supply shocks were restricted to one region, higher reserves would have to be stockpiled than in a globalized market where arbitrage opportunities lead to supply of cheaper oil from other production sites.

Furthermore, a precise assessment of the market behaviour is needed to anticipate the scope of new energy policies. Energy markets are currently experiencing fundamental changes since production of giant oil fields declines (Höök, Hirsch, and Aleklett, 2009) whereas new technologies, like hydraulic fracturing, are used to revitalize existing oil fields. Also, the interest in renewable energy has recently increased as might be reflected by the renewable energy directive of the European Union (European Commission, 2016). The decision to invest in the energy sector requires an accurate prediction of future crude oil prices. Focussing on the classical benchmarks (WTI and Brent) or only on local benchmark prices might prevent assessing the correct market behaviour if they do not reflect global supply and demand.

Moreover, a precise assessment of crude oil prices is needed for hedging purposes and the pricing of other derivatives related to crude oil prices. It is therefore of interest which benchmark price reflects crude oil market developments first and leads the pricing process. This may become even more important since activity in commodity exchange contracts has risen in recent years which is discussed under the term 'financialization' of commodity markets in the literature (see, for example, Buyuksahin and Harris, 2011 and Tang and Xiong, 2012). Although activity in crude oil exchange trading has increased accordingly, trading physical oil is still carried out in large quantities and is non-transparent to the public. In practice, price reporting agencies, like Platts, provide assessments of benchmark crude oil prices. The prices in the physical oil market are collected by a window or market-on-close process in which bids, offers and the trade volume are assessed and prices are published as an end-of-day value. This leads to price-discovery which rests on voluntary and selective disclosure by market participants as well as subjective judgement of the price reporting agency. Although WTI, Brent and Dubai are considered to be the most important crude oil benchmarks, there is no universally recognized *global* crude oil spot price. Market agents exposed to crude oil price risks, therefore, are particularly interested in how different crude oil benchmarks interact and which of them responds fastest to changing conditions on the crude oil market.

The remainder of the paper is organized as follows. Section 2.2 describes the structure of the world crude oil market and the role of benchmark prices. In Section 2.3, we review the literature on crude oil market integration, Section 2.4 outlines the econometric framework used in the empirical part of the paper, Section 2.5 reports the results of the empirical application, Section 2.6 relates our findings to previous studies and Section 2.7 concludes.

2.2 Market structure and the role of benchmark prices

Internationally traded crude oil comes in different qualities and characteristics. Lighter crude oils yield a higher percentage of gasoline and diesel fuel than heavier crudes (usually measured in American Petroleum Institute (API) gravity). Since sulphur is an undesirable component, 'sour' crudes with a higher sulphur level are less sought after than 'sweet' crudes. Generally, light and sweet crudes are

[1]The International Energy Agency (IEA) was founded in the wake of the first oil crisis. Historically, the majority of member states were net oil importers. Net exporters are exempt from this requirement. Although the role of the US as a net importer has to be reconsidered, following the resurgence of shale oil fields, the largest crude oil stockpiles are concentrated in the US.

Figure 2.1
Time series plots for regional crude oil price series (WTI, Brent, Dubai, Bonny Light, Tapis).

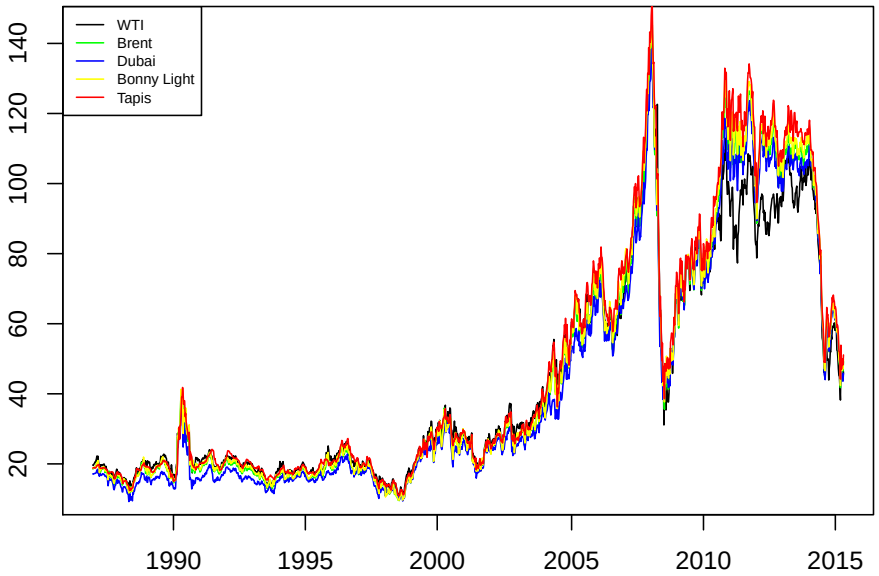

Source: DATASTREAM (see Section 2.5.1).

priced at a premium relative to heavy and sour crudes. Buyers and sellers of crude oil rely on the use of benchmark crude oils (price markers) to price the different types of crude oil. These benchmarks typically exhibit the following properties: First, the volume of production must be sufficiently large to ensure physical liquidity. Second, the oilfield has to be located in a geopolitically and financially stable region to encourage market interactions. Third, delivery points have to be provided at locations suitable for trade with other market hubs to enable arbitrage. Finally, a diverse ownership of production should be present to prevent market interference and price manipulation. In practice, however, major crude oil benchmarks do not fulfil all the requirements equally. Non-benchmark crudes are priced relative to the benchmark crude at a premium or discount depending on their quality. This is known as formula pricing.

Brent is the reference for about 65% of crude oil traded around the globe according to the Intercontinental Exchange, whereas WTI is the dominant benchmark in the US (Intercontinental Exchange, 2016). Dubai is the main reference for Persian Gulf oil delivered to the Asian market. Bonny Light is a benchmark for West African oil fields and Tapis serves as a benchmark crude for the Asian Pacific region. Figure 2.1 shows the trajectories of the five benchmark prices from 1987 to 2015. The amount of oil production over time is depicted in Figure 2.2.

Originally, crude oil extracted from the Brent oilfield, which was discovered in 1971, formed the Brent benchmark (API gravity of 38.3° and 0.37% sulphur). Production from the Brent oilfield started to decline in the mid-1980s which led to volatile prices. Commingling Brent with oil produced in the Ninian oil field, also located in the North Sea, alleviated this problem temporarily. A further decline in production led to the inclusion of oil from the Forties, Oseberg and Ekofisk fields (Fattouh, 2006). Today, the production is still declining (see Figure 2.2) and a substantial share of Europe's crude oil supply comes from Russia, which raises the question whether Brent has retained its role as a benchmark price.

The North American crude oil West Texas Intermediate (WTI), which has an API of 39.6° and contains 0.24% sulphur, making it a light and sweet crude, is transported from the extraction sites via

Figure 2.2
Temporal evolution of crude oil production in five production sites denoted in thousand barrels per day.

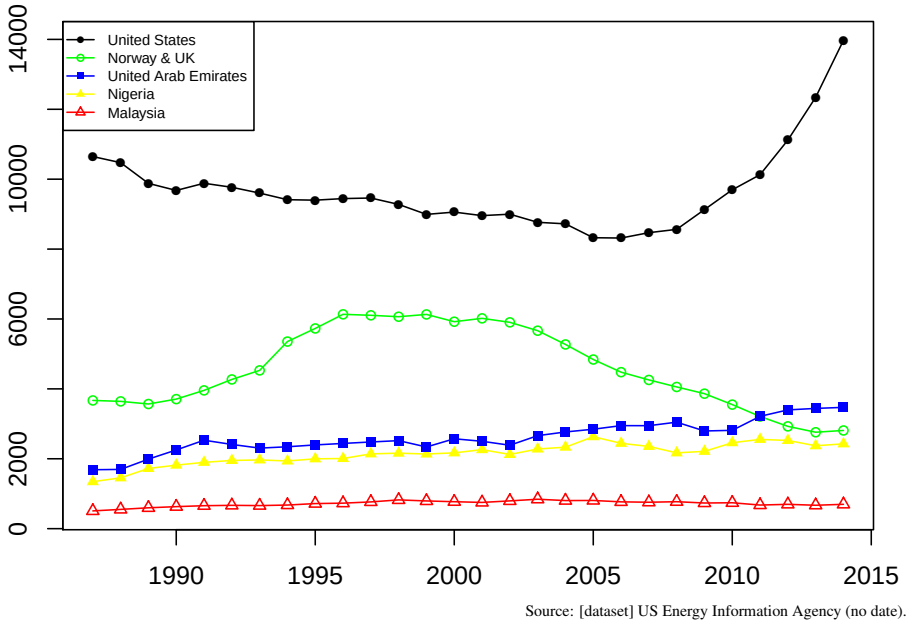

pipelines to Cushing, Oklahoma. In 1983, NYMEX chose Cushing as the official delivery point for its light sweet crude futures contract which in turn connects the oil fields to refineries and ports. Following the explosive growth in production from shale oil fields, the Cushing pipeline nexus has turned out to be a bottleneck. Oil is transported to Cushing in large quantities but the ill-equipped infrastructure delayed the distribution of oil. Consequently, the build-up in inventory caused WTI to trade at a discount compared to other benchmark crude oils and to decouple from the world crude oil market. This phenomenon is known in the literature as the 'broken benchmark' (Fattouh, 2007, 2010; Ji and Fan, 2015). If WTI was considered the global price setter, a decoupling effect would severely impair effective hedging against risks related to energy prices and would lead to incorrect pricing of other derivatives based on crude oil.

WTI and Brent held a constant price differential until around 2010. Historically, WTI traded at a premium compared to Brent, attributed to the fact that WTI is the lighter and sweeter crude oil. Beginning in 2010, the spread has been reversed. The hydraulic fracturing boom in the US helped to increase the US crude oil production by 75% from 2008 to 2014 according to the US Energy Information Agency (US Energy Information Agency, 2016) and subsequently ensured full inventories. Hydraulic fracturing is not utilized with the same intensity in the oil fields of the North Sea. A significant widening of the price differential can be observed after the shale oil boom in the US picked up speed. Moreover, the US ban on crude oil exports during our observational period may have prohibited the reduction of overcapacities through international trade[2].

Dubai is of slightly lower grade than WTI or Brent. An API gravity of 31° and 2% sulphur makes Dubai a medium heavy and sour crude. It comprises of crudes from different oil fields in Dubai, Oman and Abu Dhabi. Despite the existence of other regional crudes with a larger physical base, Dubai serves as a benchmark price for oil extracted in the Gulf region.

Bonny Light is a sweet but medium heavy crude oil (API 33.4°, 0.16% sulphur). The Bonny Light

[2]The US have lifted the crude oil export ban in January 2016.

production is concentrated in the onshore and offshore areas of the Niger Delta of Nigeria. West African crude oil is mostly refined outside the region, in Asia, Europe and the US. Violent conflicts in the Niger region led to temporary disruption of the oil production in September 2004.

Tapis is produced offshore in the South China Sea (the Seligi, Guntong, Tapis, Semangkok, Irong Barat, Tebu, and Palas fields). It is of the highest quality with an API gravity of 45.2° and low sulphur content (0.03%).

Historically, none of the five benchmark prices in our study has emerged as a universally recognized global price setter. A price setter is defined as a price that influences other prices in the same category directly or indirectly without being influenced itself. In terms of our empirical application which focuses on a cointegrated system, a price setter can be identified as a variable which does not adjust to deviations from the long-run equilibrium which is instead maintained by the remaining variables. The price setter takes the role of a *lead* variable whereas the remaining variables act as *lag* variables.

We believe that focussing on benchmark prices reduces the problem encountered by studies involving both benchmark and non-benchmark prices (Wlazlowski, Hagströmer, and Giulietti, 2011; Candelon, Joëts, and Tokpavi, 2013): Non-benchmark prices are priced in relation to the regional benchmark with price adjustments made depending on quality and transportation costs (formula pricing). While we expect the benchmark/non-benchmark relation to be strong, we are primarily interested in the relationship between geographically separated markets. Only if we find long-run co-movement and short-run adjustments among prices without a formula pricing relationship, we can argue in favour of a globalized crude oil market.

2.3 Literature

After Adelman (1984) and Weiner (1991) initiated the discussion on the integration of international crude oil markets, a substantial body of literature on the subject has emerged. Empirical studies mostly employ cointegration models to assess the relations among crude oil prices. For instance, Rodriguez and Williams (1993) aim to test the 'one great pool' hypothesis using a cointegration analysis for monthly data from 1982 to 1992. They claim to find evidence for integrated crude oil markets by rejecting the hypothesis of no cointegration which implies the presence of a long-run stable relationship among regional crude oil prices. However, Weiner (1993) emphasizes that, although prices follow a common trend, the short-run dynamics are important to characterize the relationship among regional prices. More precisely, Weiner (1993) argues that only price reactions to changes in other crude oil prices in the short-run should lead to a rejection of the 'regionalization' hypothesis. He criticizes the use of linear cointegration models which are not able to capture the true dynamics of a changing world crude oil market.

Guelen (1999) tries to account for structural change by applying cointegration models to subsamples of falling and rising crude oil prices. He finds evidence for stronger co-movement in periods of increasing prices, implying that linear cointegration models indeed are not well-suited for the analysis of price dynamics in global crude oil markets. Further, he finds that WTI and Brent take the role of global benchmark prices. Bentzen (2007) specifies a vector error correction model for daily crude oil prices from the Middle East, North America and the North Sea. Using data from January 1988 to December

2004, evidence is found for a globalized market with an increasing role of OPEC prices, thereby reducing the strength of WTI and Brent as global benchmarks.

Hammoudeh, Ewing, and Thompson (2008) and Fattouh (2010) use threshold cointegration models to capture a potentially non-linear relationship among crude oil prices. More specifically, Hammoudeh et al. (2008) examine the relationship among four benchmark prices (WTI, Brent, Dubai, Maya) based on daily data from 1990 to 2006. They use momentum threshold autoregressive (MTAR) models which allow for different adjustment depending on whether the spread between crudes is widening or narrowing. While all price pairs are cointegrated, Brent and WTI are found to be leading the pricing process in the long-run. Instead, Fattouh (2010) analyzes crude oil price differentials at a weekly frequency from 1997 to 2008 using threshold autoregressive (TAR) models. Prices of crude oils with a similar quality show a strong comovement over the sample whereas divergence of prices for crudes of different qualities can be observed.

Liu, Chen, and Wan (2013) investigate the role of China in the world crude oil market. Since China is one of the major oil importers with increasing demand in recent years, China's energy policy has an important influence on regional crude oil prices. If price changes of the regional benchmark, Daqing, were transmitted to world crude oil prices, indications of market integration would be found. However, the results of a threshold VECM reveal only a one-directional effect from world crude oil markets to the regional Daqing benchmark. Wilmot (2013) focusses on the Canadian-US market integration. He argues that the 'globalization' hypothesis also requires that a long-run relationship among secondary 'non-benchmark' crudes exists. Evidence from a cointegration analysis of Edmonton Par, a light crude, and Western Canadian Select, a heavy crude, and its US (Mexican) analogues, confirm a long-run relationship. However, the analysis reveals a structural break in the cointegrating vector and the breakpoint is determined to coincide with the Financial Crisis.

More recently, Ji and Fan (2015) investigate the long-run equilibrium relationships among the five major regional crude oil benchmarks (WTI, Brent, Dubai, Bonny Light, Tapis) by using a VECM combined with a directed acyclic graph technique. Based on tests for the presence of structural breaks, they split their sample at the break point in October 2010. They find that WTI was a price setter before 2010 while Brent is in a leading role since 2011. Tapis has always been a price taker whereas Dubai and Bonny Light have taken both roles at times. Mann and Sephton (2016) use band-TAR threshold cointegration models to examine the long-run relationships between WTI and Brent and WTI and Oman. They find these crude oil price pairs to be tied together in the long-run. Since each price adjusts to the long-run equilibrium at some point, they conclude that a unique global benchmark prices does not exist.

Additionally, there are further studies that focus on the changing conditions on the crude oil market. Reboredo (2011) models the dependence structure between crude oil benchmark prices using a copula approach. Upper and lower tail dependence is found, suggesting that benchmark crude oils boom and crash simultaneously. This is considered evidence for a globalized world crude oil market. Candelon et al. (2013) examine causal linkages at regional oil markets when prices are on average extremely high or low. The study reveals benchmark prices besides WTI and Brent. Moreover, market integration is found to be weaker during extreme times. Instead of Candelon et al. (2013)'s set of 32 different crudes, Lu, Hong, Wang, Lai, and Liu (2014) restrict their analysis to four benchmark prices (WTI, Brent, Dubai, Tapis) and find a stronger market integration after disruptive events take place. Zhang and Zhang (2015)

employ a Markov-switching autoregressive model to investigate the short-run dynamics between Brent and WTI. They find three price regimes which are characterized by different dynamics.

In all, evidence is mounting that crude oil markets are 'globalized'. Crude oil prices seem to hold long-run equilibrium relationships. However, the degree of market integration does not seem to be stable over time.

2.4 Econometric methodology

The long-run and short-run dynamics of the crude oil prices, collected in a vector y_t, are modelled using a vector error correction model (VECM). The model assumes that the prices are linked by stable long-run relationships. However, the variables deviate from these equilibrium relationships in the short-run due to random shocks. Maintaining the long-run relationships requires that deviations are corrected by the variables in the short-run. Put differently, the variables are said to adjust to equilibrium errors. Following Johansen (1988)'s notation, the linear VECM is given as

$$\Delta y_t = \mu + \Pi y_{t-1} + \sum_{i=1}^{p-1} \Gamma^i \Delta y_{t-i} + u_t, \tag{2.1}$$

where y_t is a $N \times 1$ vector of $I(1)$ variables, μ is a vector of drift parameters and u_t is a vector of white noise error terms. The $k \times k$ parameter matrix $\Pi = \alpha\beta'$ captures both the long-run equilibrium relations and the adjustment behaviour. The matrix β contains r cointegrating vectors and α carries the loadings in each of the r vectors.

A particular feature of the linear VECM is that it assumes constancy of all parameters in its data generating process. Certainly, this assumption appears to be restrictive in the context of a volatile crude oil market. Previous studies described relevant disruptive events concerning the energy market (see Lu et al. (2014) for a list of events from 2002 to 2011), and specific issues on the crude oil market, for example WTI, as a 'broken benchmark'. These events are likely to induce structural changes in the relations among crude oil prices. Although we expect the crude oil prices to maintain constant long-run equilibria since crude oils are close substitutes[3], the roles of crude oils in the market, for example, switching from price takers to prices setters and vice versa, might change over time. Particularly, a decoupling of WTI from the world crude oil market might have led to exogeneity of WTI for this period. We therefore study the evolution of the adjustment coefficients while the long-run equilibrium relationships are assumed to stay constant over time.

To account for potential time-varying adjustment, we apply a Markov-switching VECM (MSVECM) to the data. Markov-switching models in a time series econometrics framework were introduced by Hamilton (1989) and the MSVECM used in this paper was proposed by Krolzig (1997). We consider a q-regime VECM which allows the parameters to be state-dependent. The MS(q)-VECM takes the form of

$$\Delta y_t = \mu_{s_t} + \Pi_{s_t} y_{t-1} + \sum_{i=1}^{p-1} \Gamma^i_{s_t} \Delta y_{t-i} + u_t, \qquad u_t | s_t \sim N(0, \Sigma_{s_t}), \tag{2.2}$$

where μ_{s_t} are state-dependent drift terms, Π_{s_t} is the state-dependent long-run impact matrix, $\Gamma^i_{s_t}$ are state-dependent short-run dynamics and the error terms have a normal distribution conditional on the state s_t.

[3]Differences in quality (density and sulphur content) are reflected in discount or premium prices.

A Cholesky decomposition of the error term variance-covariance matrix gives $\Sigma = LS^2L'$ where L is a normalized lower triangular matrix and S is diagonal. The error term variance can either be restricted to stay fixed over all states, $\Sigma_{s_t} = \Sigma$ for all $s_t = 1,2,\ldots,q$, or change over states. We distinguish between a switching scale, $\Sigma_{s_t} = LS^2_{s_t}L'$, and a fully switching variance, where each element of Σ_{s_t} is switching according to s_t, $\Sigma_{s_t} = L_{s_t}S^2_{s_t}L'_{s_t}$. A fully switching variance-covariance matrix comes at the cost of an increasing number of parameters that have to be estimated.

The state of the data-generating process is governed by a latent integer state variable s_t. The probability that s_t attains some particular value $j \in \{1,2,\ldots,q\}$ depends only on the most recent value s_{t-1}:

$$P(s_t = j|s_{t-1} = i, s_{t-2} = k, \ldots) = P(s_t = j|s_{t-1} = i) = p_{ij} \quad \forall i,j = 1,2,\ldots,q. \tag{2.3}$$

Such a process is described as an q-state Markov chain with constant transition probabilities $p_{ij} > 0$, $\sum_{j=1}^{q} p_{ij} = 1$ (Hamilton, 1994). We assume the Markov chain to be irreducible and ergodic, which means that each regime can be reached from any previous regime (absence of absorbing states) and no regime has a periodic occurrence.

The state-dependent long-run impact matrix Π_{s_t} is decomposed in the constant cointegrating vectors and the state-dependent weighting matrix α_{s_t},

$$\Pi_{s_t} = \alpha_{s_t}\beta', \tag{2.4}$$

where α_{s_t} contains the state-dependent adjustment coefficients which measure the reaction to deviations from the long-run equilibria for each regime. In our application, we are particularly concerned with the evolution of the adjustment coefficients over time and regimes. The adjustment coefficients can be interpreted in the context of a lead-lag relationship among the crude oil prices. If one of our crudes was a global price setter, it would not adjust to deviations from the long-run equilibrium induced by random shocks. The price setting crude thus takes the role of a *lead* variable. Analyzing the long-run relationships among crude oil prices via a MSVECM provides further insights in the structure of the world crude oil market since it enables us to identify exogenous benchmark prices under particular regimes of the process.

The dynamic properties are further investigated by observing the behaviour of the system after shocks to variables of the system using regime-specific orthogonalized impulse response functions. For this matter, we need to transform the VECM representation given in (2.2) to a vector moving average (VMA) representation,

$$y_t = u_t + \Psi^1_{s_t}u_{t-1} + \Psi^2_{s_t}u_{t-2} + \Psi^3_{s_t}u_{t-3} + \ldots \tag{2.5}$$

Since the error terms u_t are correlated with each other, we use the Cholesky decomposition of the regime-specific error term variance-covariance matrix again and construct orthogonalized impulse response functions,

$$IRF^1_{s_t}(\hat{\theta}) = \widehat{L}_{s_t}, \quad IRF^2_{s_t}(\hat{\theta}) = \widehat{\Psi}^1_{s_t}\widehat{L}_{s_t}, \quad \ldots, \quad IRF^h_{s_t}(\hat{\theta}) = \widehat{\Psi}^{h-1}_{s_t}\widehat{L}_{s_t}, \tag{2.6}$$

where $\hat{\theta}$ denotes the entirety of all estimated parameters.

Naturally, the number of parameters to estimate increases with the number of states which are specified in the MSVECM, so that a parsimonious model specification leads to a maximum of two or three

states. However, the exact number of states is usually not known a priori and has to be jointly selected with additional variables, that is, further lags to capture short-run dynamics. Psaradakis and Spagnolo (2006) found that information criteria can accurately identify the appropriate number of states for a Markov-switching model. Awirothananon and Cheung (2009) argued for the use of the BIC to select the number of states based on results of Monte Carlo experiments. In the following application, we follow Awirothananon and Cheung (2009) and use the BIC for model selection with respect to the number of states, the lag length and switching behaviour of the drift terms as well as elements of the variance-covariance matrix.

2.5 Empirical analysis

2.5.1 Data

For this study, we observe crude oil price data at weekly frequency from May 1987 until October 2015. All crude oil prices are free on board (FOB) spot prices[4], observed at each Monday and denominated in US dollars per barrel. The time series are obtained from DATASTREAM[5] and the original observations were transformed by taking natural logarithms.

First, the time series are tested for their order of integration. The results of ADF and KPSS unit root tests are reported in Table 2.1. Furthermore, we apply the Lee-Strazicich (LS) unit root test which accounts for two structural breaks in the null and alternative (Lee and Strazicich, 2003). The null hypothesis of the ADF and LS tests cannot be rejected at the 1% significance level for all prices while the null hypothesis of the KPSS test is rejected at all conventional significance levels. We obtain opposite results for the returns. The tests support the hypothesis that all prices follow a unit root process and are integrated of order one.

Table 2.1
Unit root tests of the logarithmized crude oil prices.

Variables	ADF	LS	KPSS	Variables	ADF	LS	KPSS
WTI	-2.635	-2.846^{*}	0.668^{***}	Δ WTI	-22.153^{***}	-37.433^{***}	0.064
Brent	-2.901	-3.087^{**}	0.738^{***}	Δ Brent	-20.234^{***}	-34.436^{***}	0.068
Dubai	-2.794	-3.459^{**}	0.742^{***}	Δ Dubai	-19.773^{***}	-41.847^{***}	0.071
Bonny Light	-2.520	-3.014^{*}	0.742^{***}	Δ Bonny Light	-20.167^{***}	-31.848^{***}	0.069
Tapis	-2.575	-3.485^{**}	0.725^{***}	Δ Tapis	-18.630^{***}	-42.166^{***}	0.072

Note: The ADF, LS and KPSS test equations are estimated including an intercept and trend for the variables in levels. The test equations for the first differences include an intercept. Lag selection is based on the Bayesian Information Criterion (BIC).
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

2.5.2 Linear cointegration analysis

To test for cointegration, we rely on the Johansen rank test which is based on the VECM specified in Equation (2.1). The cointegrating rank r is determined by the number of estimated eigenvalues of the

[4]Pertains to a transaction whereby the seller makes the product available within an agreed on period at a given port at a given price; it is the responsibility of the buyer to arrange for the transportation and insurance. (US Energy Information Administration)

[5]The data can be found using Mnemonic (Code): OILTPMY (S214WT), OILDUBI (T15609), OILBRNP (S04107), CRUDWTC (S369VW), OILAFRB (S00112).

estimated adjustment coefficient matrix Π that are significantly greater than zero. Johansen (1988, 1991) proposed likelihood ratio type tests of which we use the trace test variant[6]. The trace test examines the null hypothesis, $\text{rank}(\Pi) = r_0$, against the alternative hypothesis, $r_0 < \text{rank}(\Pi) \leq k-1$.

The results of the cointegration test are presented in panel (a) of Table 2.2. Since the null hypothesis $r_0 = 3$ can soundly be rejected, we assume the maximum number of cointegrating vectors of four. The normalized cointegrating vectors are displayed in Panel (b) of Table 2.2. We find that the price differentials between WTI and the four remaining crudes are relevant long-run equilibria. The trade-off between a parsimonious specification and sufficiently capturing the short-run dynamics of the system leads to two additional lagged differences ($K = 2$).

Table 2.2
Cointegration tests and linear VECM.

	$N-r$	r	Eig.value	Trace	5% Crit. val.	p-Value
Panel (a): I(1)-analysis						
	5	0	.1084	361.75	76.07	.000
	4	1	.0651	192.16	53.12	.000
	3	2	.0374	92.61	34.91	.000
	2	3	.0292	36.22	19.96	.000
	1	4	.0013	1.95	9.24	.783
	WTI	Brent	Bonny	Dubai	Tapis	μ
Panel (b): Cointegration vectors						
β_1	-1.087				1	.276
β_2	-1.136			1		.584
β_3	-1.097		1			.355
β_4	-1.094	1				.363
Panel (c): Adjustment coefficients						
α_1	.066*	.104***	.110***	.049	−.162***	
	(1.879)	(2.974)	(3.159)	(1.514)	(−6.198)	
α_2	.028	.064**	.063**	−.016	.061***	
	(1.005)	(2.270)	(2.238)	(−.606)	(2.877)	
α_3	−.214*	−.229**	−.432***	−.272***	−.001	
	(−1.933)	(−2.062)	(−3.915)	(−2.638)	(−.018)	
α_4	.198*	.053	.257**	.242**	.090	
	(1.701)	(.451)	(2.216)	(2.237)	(1.028)	
Panel (d): Weak exogeneity						
LR(4)	16.47***	22.87***	33.38***	7.54	43.07***	
Lag	1	2	3	4	5	
Panel (e): Test for residual autocorrelation						
	3.398	9.366	66.174***	148.79***	196.38***	
Panel (f): Test for ARCH effects						
	2081.5***	2937.7***	3790.5***	4971.8***	5529.4***	

Note: Panel (a) reports Johansen (1988) cointegration tests. The critical values are taken from Osterwald-Lenum (1992). p-values are computed using a simulation study with 10,000 replications. Panel (b) displays the estimates of the cointegrating vectors. Insignificant variables have been excluded from the cointegrating vector. Panel (c) reports the estimates of the adjustment coefficients with t-statistics in parentheses. Estimates of the short-run dynamics, drift terms and variance-covariance matrix are not shown to conserve space. Panel (d) reports weak exogeneity tests. The likelihood ratio (LR) statistics are χ^2 distributed with degrees of freedom in parentheses. Panel (e) shows the results of vector portmanteau tests of the residuals. Panel (f) shows the results of tests for ARCH effects.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

We now briefly turn to the results of the linear VECM to obtain a useful summary of the 'average' adjustment dynamics provided by a linear specification. The adjustment coefficients of the linear VECM are reported in panel (c) of Table 2.2. A surprising feature of the results is the adjustment of the cointegrated system to the WTI-Brent price differential. Neither WTI, nor Brent adjust strongly to the

[6]The maximum eigenvalue test reaches the same conclusion: The null hypothesis of at most three cointegration vectors is rejected.

deviations from their long-run equilibrium. By contrast, Bonny Light and Dubai react to deviations from the WTI-Brent price differential in the previous period. Tests for weak exogeneity of particular crude oil prices are presented in panel (d). The tests suggest weak exogeneity of Dubai, although it adjusts significantly to the WTI-Brent and WTI-Bonny Light price differentials. This discrepancy can be attributed to a generally lower power of Wald-type statistics. WTI is found to adjust to all price differentials except WTI-Dubai. Hence, WTI does not seem to be an exogenous price setter although it is the most closely watched benchmark crude oil price in the US.

2.5.3 Markov-switching error correction models

Given the evolution of the market conditions, described in Section 2.2, we suspect that the adjustment coefficients among crude oil prices do not remain constant over time and therefore consider a MSVECM which allows the model parameters to change between different regimes. As noted previously, the model specification of the MSVECM in terms of number of states is typically not clear a priori. Therefore, we consider both a two-state and a three-state specification and choose the final model specification based on the BIC[7]. Further, in line with the principle of parsimony, we reduce the number of parameters to estimate by testing whether allowing a switching behavior in a parameter matrix improves the model with regard to the BIC. More specifically, in the two-state MSVECM with two lags, henceforth MS(2)VECM(2), the vector of drift terms is restricted to be constant over both states and the variance-covariance matrix Σ is allowed to switch over states. In the three-state MSVECM with two additional lags, henceforth MS(3)VECM(2), we impose constancy of the drift terms and allow for a switching scale of the variance-covariance matrix. A comparison between the MS(2)VECM(2)[8] and MS(3)VECM(2) based on the BIC suggest that the increased goodness-of-fit of a three-state MSVECM indeed outweighs the increasing number of parameters. The regime-specific adjustment parameters for the MS(3)VECM(2) are reported in Table 2.3. We have excluded the short-run dynamics to conserve space and focus on the adjustment to the long-run equilibria. We find evidence for distinct regime-switching, reflected by non-zero transition probabilities and a state variable that assumes state 1 in 17%, state 2 in 15% and state 3 in 68% of the sample period. We refer to those points in time in which the model is confident of being in state 1 as regime 1 (R1), in state 2 as regime 2 (R2) and in state 3 as regime 3 (R3). Smoothed probabilities reflect the estimated probabilities of occurrence of each state at each point in time. This allows us to gain insights into the evolution of the adjustment dynamics over time. The smoothed probabilites are depicted in Figure 2.3.

The cointegrated system seems to be predominantly in state 1 at the beginning of the observational period. The first regime, thus, comprises almost exclusively of the first part of the sample, reaching from 1987 to 1994 and we refer to this as the 'early regime'[9]. High probabilities of state 2 can be linked to exogenous global events and volatile economic environments. Probabilities close to one coincide with, among others, the period around the events of September 11, 2001, the period after the invasion of Iraq in 2003, and the Financial Crisis beginning in 2008. The second regime can therefore be associated with volatile economic and geopolitical times, hence we call it the 'crisis regime'. The remaining regime

[7]Higher order MSVECM ($q > 3$) are not in line with a parsimonious model specification

[8]The results for the MS(2)VECM(2) specification are reported in Table 2.A–1 in the appendix.

[9]Please note that the labelling of the regimes primarily serves the purpose of illustration. The transition probabilities are estimated to be nonzero. Hence, it is, for example, possible that the state variable takes value one at a later point in time and the system switches to the 'early regime' again.

associated with state 3 is referred to as the 'tranquil regime' and reflects behavior of the system in periods of relative calm. We investigate the role of each crude oil price in all three regimes. The regime-specific dynamics help us to obtain new insights regarding the changing roles of regional crudes in the world crude oil market.

We report the results of regime-specific and overall weak exogeneity test in panel (c) of Table 2.3. We find no evidence against the null hypothesis of weak exogeneity of WTI in the 'early regime' and in the 'tranquil regime' during the later parts of the sample period. However, WTI adjusts significantly to the WTI/Bonny Light and WTI/Brent price differential in the 'crisis regime'. The hypothesis of overall weak exogeneity is rejected which can be attributed to the significant adjustment in the 'crisis regime'. In other words, WTI seems to react to other crude oil prices primarily in times of uncertainty about future supply and demand. Brent is a weakly exogenous variable in the 'early regime' and the 'crisis regime'. However, Brent adjusts to the WTI/Tapis and WTI/Dubai price differentials in the 'tranquil regime'. Bonny Light is weakly exogenous in the 'early regime', adjusts to WTI/Bonny Light and WTI/Brent price differentials in the 'crisis regime' and to the WTI/Tapis and WTI/Dubai price differentials in the 'tranquil regime'. These findings suggest that WTI and Brent are important signals of world crude oil market news for Bonny Light in crisis periods whereas the price differentials with the Arabian Dubai and the Asian Pacific Tapis are constant factors in the price determination of Bonny Light. This can in parts be explained by the fact that Dubai is a close regionally competitor to the Nigerian Bonny Light. A reaction to its WTI price differential is attributed to the fact that the US is the largest importer of Nigerian crude oil so that US crude oil demand shocks are transmitted to the price of Bonny Light.

Dubai is the only weakly exogenous variable in all regimes. The results of the overall weak exogeneity test for Dubai in the three-state model is in line with the findings for the two-state MSVECM and the linear model (see panel (d) in Table 2.2 and panel (c) in Table 2.A–1). Also, an alternative normalization in which Dubai is allowed to be an exogenous variable in each equation left the results virtually unchanged. The results of this model are reported in Table 2.A–3 in the appendix. Economically, the result implies that Dubai acts as a price setter in this set of benchmark crude oil prices. Finally, Tapis is a price taker in all three states.

The orthogonalized impulse response functions[10] are displayed in Figure 2.4. We find that shocks to one variable in the 'early regime' do not evoke strong responses from the other variables. In contrast, shocks in the 'crisis regime' lead to visible reactions of the system. Adjustment to shocks is relatively fast whereas it takes the system more time to adjust to shocks in the 'tranquil regime'. These findings are in line with Ji and Fan (2015) who document stronger market integration if global exogenous shocks occur.

[10]The ordering of the variables which is used for the Cholesky decomposition is given as follows: Dubai → WTI → Brent → Bonny Light → Tapis.

Table 2.3

Markov-switching error correction model for major crude oil prices (three-state model).

	WTI			Brent			Bonny Light			Dubai			Tapis		
Panel (a): Switching adjustment coefficients															
	R1	R2	R3	R1	R2	R3	R1	R2	R3	R1	R2	R3	R1	R2	R3
$\alpha_1(s_t)$	−.030	.083	.073*	−.059	.136	.115***	.016	.090	.127***	−.067	.065	.057	−.099***	−.312***	−.108***
	(−.620)	(.673)	(1.740)	(−1.080)	(1.080)	(2.630)	(.314)	(.725)	(2.920)	(−1.300)	(.562)	(1.520)	(−4.660)	(−3.430)	(−3.240)
$\alpha_2(s_t)$	.019	.031	.039	.004	.052	.095**	−.037	.073	.096**	−.076	−.007	.002	−.009	.137**	.039
	(.368)	(.383)	(1.020)	(.077)	(.641)	(2.390)	(−.727)	(.930)	(2.390)	(−1.460)	(−.087)	(.072)	(−.369)	(2.330)	(1.300)
$\alpha_3(s_t)$	.094	−.944***	.052	.114	−.740**	.006	−.301*	−.953***	−.060	−.053	−.735**	−.012	.119*	−.402	.197*
	(.615)	(2.780)	(.347)	(.686)	(−2.190)	(.039)	(−1.950)	(−2.880)	(−.393)	(−.335)	(−2.320)	(−.091)	(1.780)	(−1.630)	(1.660)
$\alpha_4(s_t)$	.045	1.236***	−.150	−.060	.714*	−.254	.321**	.969**	−.200	.255	.821**	−.073	−.000	.661	−.162
	(.287)	(3.180)	(.951)	(−.350)	(1.830)	(−1.560)	(2.000)	(2.510)	(−1.230)	(1.570)	(2.260)	(−.515)	(−.003)	(2.370)	(−1.290)
Panel (b): Weak exogeneity															
LR(4)	5.770	17.128***	5.858	1.251	6.980	19.261***	4.428	10.661**	20.992***	7.088	5.970	5.029	23.571***	21.208***	12.919**
LR(12)		29.252***			30.275***			38.908***			17.759			58.702***	

Lag	1	2	3	4	5	6	7	8	9	10
Panel (c): Test for residual autocorrelation										
	4.615	13.313	51.284	89.917	121.29	142.67	177.06	212.82	250.97	275.12
	(.999)	(.999)	(.984)	(.755)	(.577)	(.652)	(.442)	(.254)	(.113)	(.132)
Panel (d): Test for ARCH effects										
	1.631	1.352	1.314	1.316	1.303					
	(.025)	(.050)	(.037)	(.020)	(.014)					

	R1	R2	R3
Panel (e): Transition probabilities			
R1	.952	.200	.021
R2	.042	.761	.039
R3	.005	.043	.940

Note: R1 refers to the 'early regime', R2 to the 'crisis regime' and R3 to the 'tranquil regime', respectively. Panel (a) reports the estimates of the adjustment coefficients for three regimes with t-statistics in parentheses. The estimated cointegrating vectors are identical to panel (a) in Table 2.2. Estimates of the short-run dynamics, drift terms and variance-covariance matrix are not shown to conserve space. Panel (b) reports weak exogeneity tests for each regime (first row) and over all three regimes (second row). The likelihood ratio (LR) statistics are χ^2 distributed with degrees of freedom in parentheses. Panel (c) shows the results of vector portmanteau tests of the residuals with p-values in parentheses. Panel (d) shows the results of tests for ARCH effects with p-values in parentheses. Panel (e) displays the estimated transition probabilities.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure 2.3
Smoothed probabilities MS(3)VECM(2).

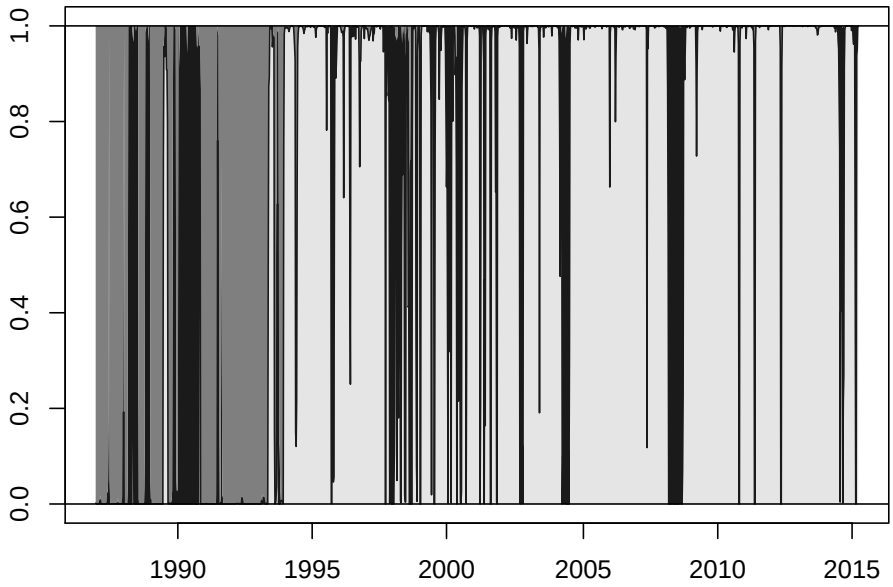

Notes: This figure shows the probabilities for the cointegrated system being in the 'early regime' (grey), probabilities of being in the 'crisis regime' (black) and probabilities of being in the 'tranquil regime' (light-grey). The probabilities sum up to one in each period.

2.6 Discussion

Overall, the results are in line with the findings of Lu et al. (2014) and Ji and Fan (2015), indicating a stronger market integration in turbulent times. While a globally stable oil market promotes the use of nearby oil fields with lower transportation costs, extreme economic conditions create incentives to re-evaluate the attractiveness of different crude oil sources. Therefore, crude oil prices have to incorporate global information beyond the regional supply and demand changes.

Furthermore, the allocation of regime 1 to the earlier part of our sample, helps to emphasize the evolution of the world crude oil market. With the exception of Tapis, we do not reject weak exogeneity for any crude oil in the 'early regime'. The later part of the sample is partitioned into the 'tranquil regime' and the 'crisis regime', so that either Brent and Bonny Light adjust to long-term equilibria in tranquil times or WTI adjusts to its WTI/Brent and WTI/Bonny Light price differentials to maintain a long-run equilibrium relationship under extreme economic conditions. Dubai's price setting role supports the hypothesis in Bentzen (2007) which states that OPEC prices are gaining influence in the world crude oil market.

Similar to our results, Guelen (1999) finds that crude oil market integration is not stable and is especially strengthened during tight market conditions. His results, however, rely on a pre-specified structural break (the full sample is divided into two subperiods 1991-1993 and 1994-1996). Our study, following a more flexible approach, reveals that focusing only on the magnitude of prices does not seem to provide a more comprehensive picture of the crude oil market dynamics. Specifically, the application of a Markov-switching model to a longer and more varied sample period shows that crude oil market inte-

gration is strengthened in periods following geopolitical and economic events. The prices of benchmark crude oil reflect changing market conditions and, for example, tend to increase if supply is uncertain, but we document faster adjustment primarily in high volatility periods.

Moreover, the extent of market integration seems to coincide with the level of macroeconomic and financial uncertainty. To illustrate our notion, we compare the occurrence of the 'crisis regime' with two measures for financial and economic uncertainty. First, we contrast the evolution of the state indicator variable with the CBOE Volatility Index (VXO) which is based on 30-day S&P 100 index at-the-money options. It is a widely used measure for uncertainty in the financial market and has the advantage over other uncertainty measures that it spans the full sample period and is available at weekly frequency. The VXO, however, primarily measures uncertainty in the financial markets while *economic* uncertainty may also be influenced by fluctuations in real activity.

Second, we therefore also compare the occurrence of 'crisis' episodes in the crude oil market with a measure for macroeconomic uncertainty, recently developed by Jurado, Ludvigson, and Ng (2015). This new measure for macroeconomic uncertainty essentially is an index based on various indicators including real output and income, unemployment, consumer spending and foreign exchange measures. The smoothed probabilities for the 'crisis regime' and our uncertainty measures are depicted graphically in Figure 2.5. It is obvious that the occurrence of the 'crisis regime' matches various peaks in the VXO, particularly, after the stock market crash in 1987, during the Persian Gulf crisis 1990-1991, the September 11, 2001 attack in the US, the 2003 Iraq war and the Financial Crisis starting late 2007. Likewise, peaks in macroeconomic uncertainty match 'crisis' episodes in the crude oil market. Compared to the VXO, Jurado et al. (2015)'s measure for macroeconomic uncertainty, however, is much smoother and its relation with the 'crisis regime' appears to be generally less pronounced. Finally, we consider the linear relation between the VXO and the 'crisis regime' indicator.[11] The contemporary correlation of the two time series is 0.277.

In essence, these findings provide descriptive evidence for a link between global economic uncertainty and world crude oil market integration. While they support our notion they do not enable an inferential analysis which we leave for future research.

[11] Computing correlations between our state indicator variables and the measure for macroeconomic uncertainty is not possible due to different data frequencies.

Figure 2.4

Regime-specific orthogonalized impulse response functions for one standard deviation shock in Dubai, WTI, Brent, Bonny Light and Tapis. The dotted, dashed and solid lines represent the OIRF in the 'early regime', the 'crisis regime' and the 'tranquil regime', respectively.

Panel (a): Response by Dubai

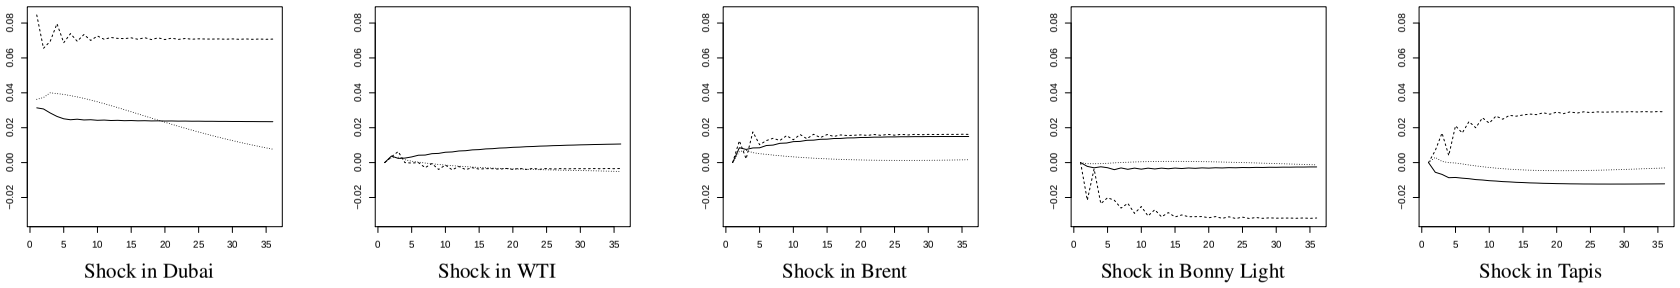

Panel (b): Response by WTI

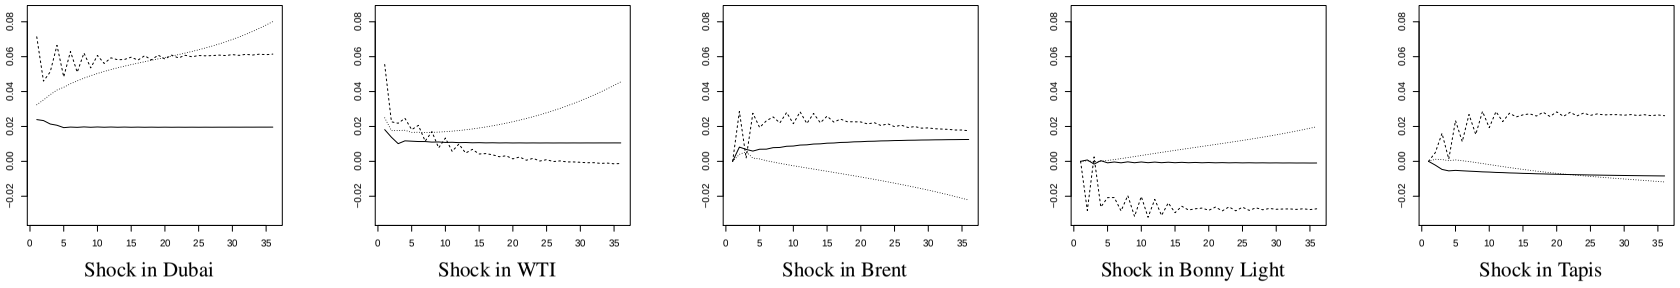

Panel (c): Response by Brent

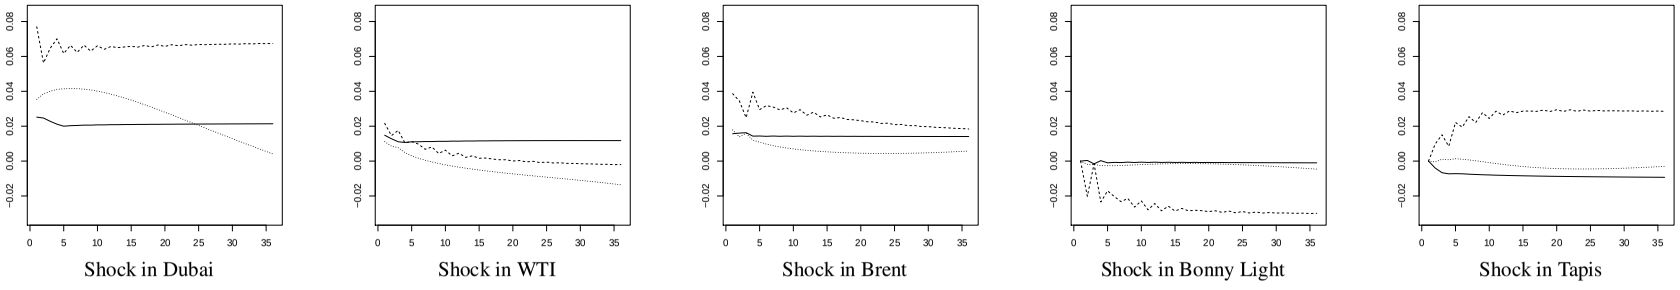

Figure 2.4 *(continued)*

Regime-specific orthogonalized impulse response functions for one standard deviation shock in Dubai, WTI, Brent, Bonny Light and Tapis. The dotted, dashed and solid lines represent the OIRF in the 'early regime', the 'crisis regime' and the 'tranquil regime', respectively.

Panel (d): Response by Bonny Light

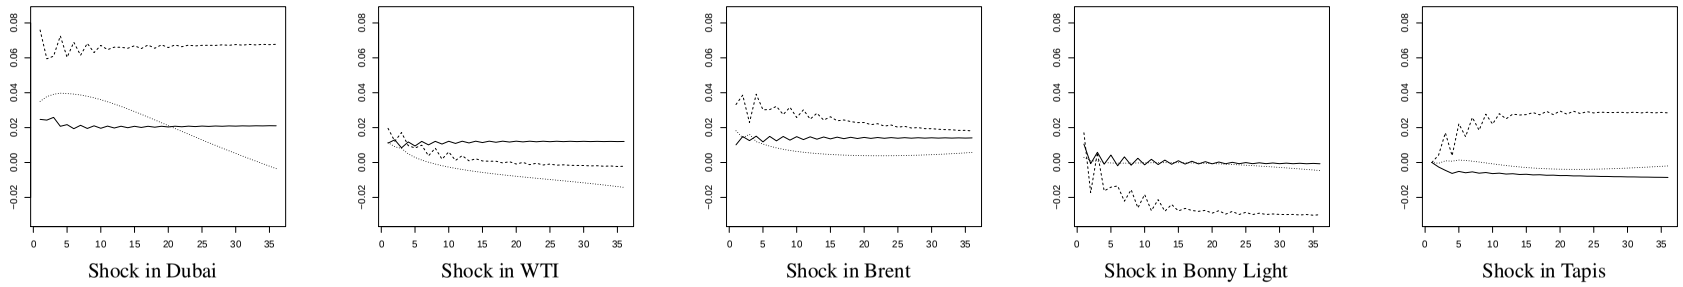

Panel (e): Response by Tapis

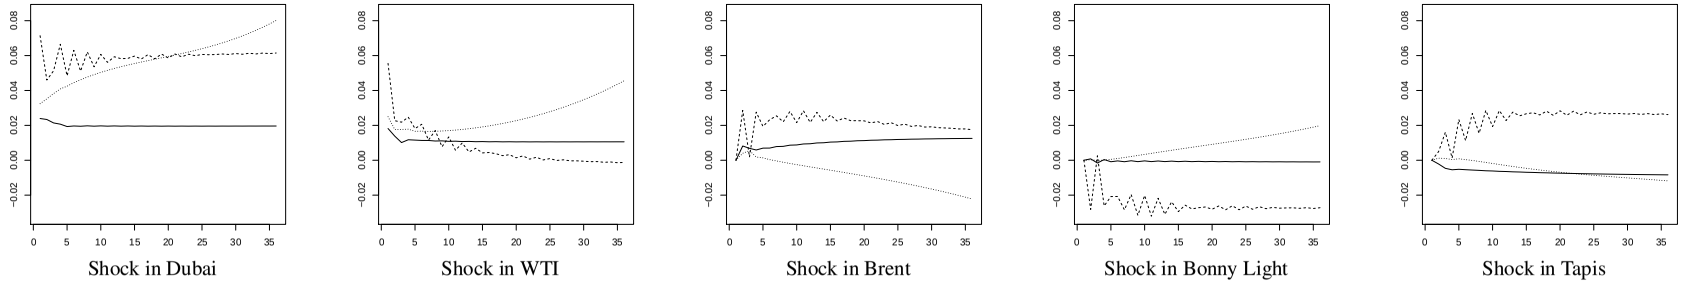

Figure 2.5
Smoothed probabilities of the 'crisis regime' and uncertainty measures.

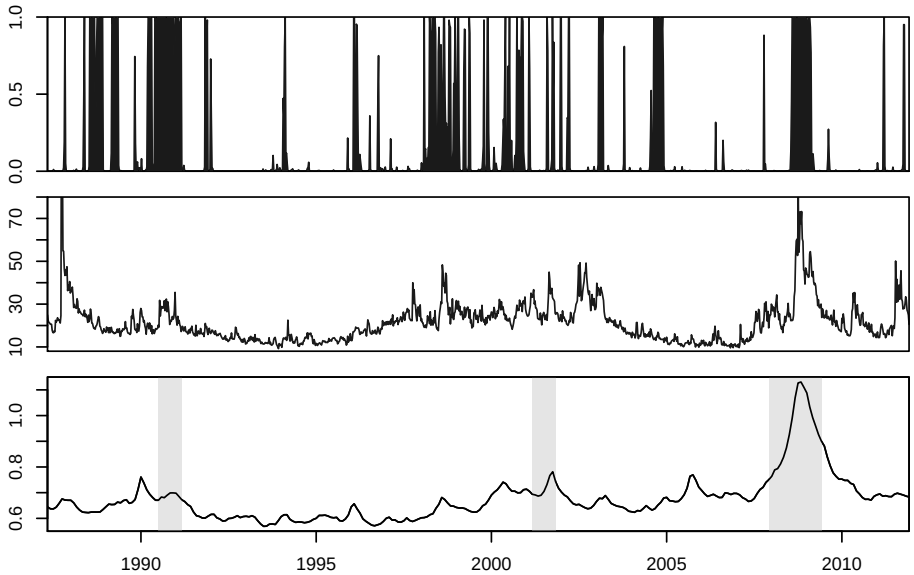

Source: [dataset] Jurado, K. and Ludvigson, S. C. and Ng, S. (2015), [dataset] NBER (2017)

Notes: This figure compares the smoothed probabilities of the cointegrated system being in the 'crisis regime' (row one) with the CBOE Volatility Index (row two) and the measure for macroeconomic uncertainty (grey shaded area: NBER recession dates) by Jurado et al. (2015) (row three).

2.7 Conclusion

This study provides a dynamic perspective on crude oil market integration. We employ a Markov regime-switching model based on the vector error correction model to study regime-switching adjustment behavior to constant long-run equilibria. Thereby, we identify three regimes to describe the adjustment behavior in different market conditions. The results highlight the changing landscape of the world crude oil markets. While the crude oil prices did not seem to maintain a long-run equilibrium from 1987 to 1994, the degree of crude oil market integration has strengthened in the later part of the sample. However, the roles of price setter and price taker can change drastically depending on the state of the global economy. Moreover, the results reveal the important role of Dubai as a price setter. Understanding crude oil market dynamics should therefore not be confined to a precise monitoring of WTI and Brent prices but should include Dubai as a third important benchmark price. Although the relationship between crude oil benchmark prices is changing over time, we do not find evidence for a decoupling of the WTI benchmark after the introduction of hydraulic fracturing to the shale oil fields of the US. It seems, that instead global events trigger adjustment to other regional benchmarks, thereby increasing world crude oil market integration.

Appendix 2.A Figures and Tables

Figure 2.A–1
Smoothed probabilities MS(2)VECM(2).

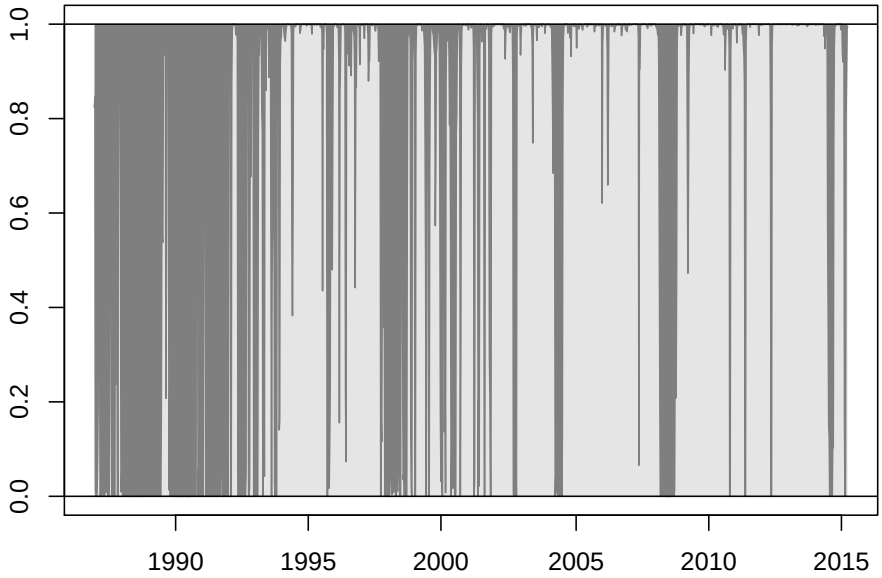

Notes: This figures depicts the probabilities for the cointegrated system being in regime 1 (grey) and probabilities of being in regime 2 (light-grey). The probabilities sum up to one in each period.

Figure 2.A–2
Smoothed probabilities MS(3)VECM(2), Dubai normalization.

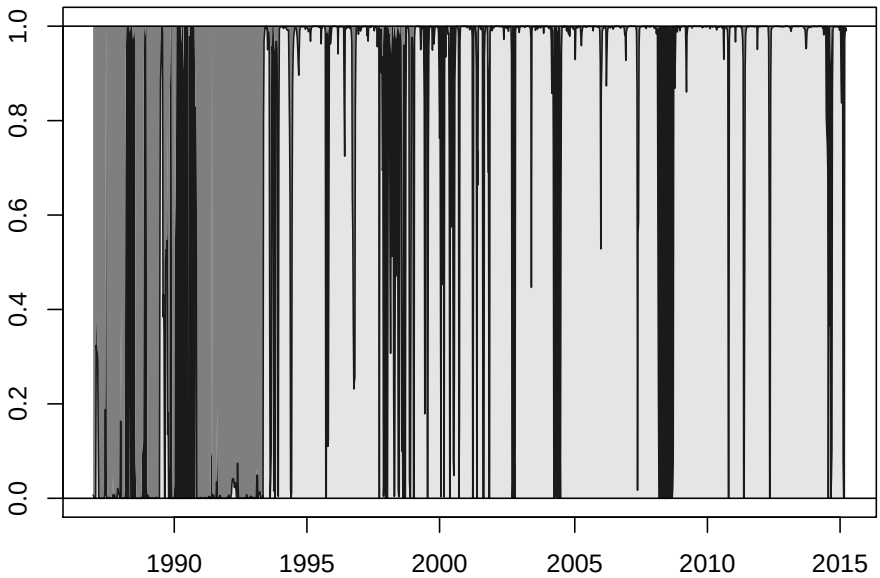

Notes: Probabilities for the cointegrated system being in the 'early regime' (medium-grey), probabilities of being in the 'crisis regime' (dark-grey) and probabilities of being in the 'tranquil regime' (light-grey). The probabilities sum up to one in each period.

Table 2.A–1

Markov-switching error correction model for major crude oil prices (two-state model).

	WTI		Brent		Bonny Light		Dubai		Tapis	
Panel (a): Switching adjustment coefficients										
	R1	R2	R1	R2	R1	R2	R1	R2	R1	R2
$\alpha_1(s_t)$	.039	.092**	.064	.142***	.052	.155***	.031	.063*	−.227***	−.097***
	(.553)	(2.340)	(.884)	(3.550)	(.731)	(3.870)	(.454)	(1.780)	(−4.560)	(−3.110)
$\alpha_2(s_t)$	.022	.055	.032	.122***	.036	.124***	−.024	.007	.093**	.038
	(.422)	(1.550)	(.600)	(3.350)	(.683)	(3.400)	(−.468)	(.224)	(2.500)	(1.330)
$\alpha_3(s_t)$	−.538***	.044	−.431**	−.015	−.709***	−.082	−.497**	−.040	−.154	.174
	(−2.630)	(.319)	(−2.070)	(−.108)	(−3.470)	(−.598)	(−2.480)	(−.326)	(−1.070)	(1.580)
$\alpha_4(s_t)$	.772***	−.178	.436*	−.289**	.735***	−.235	.589***	−.053	.321**	−.146
	(3.410)	(−1.230)	(1.890)	(−1.970)	(3.240)	(−1.610)	(2.670)	(−.414)	(2.010)	(−1.260)
Panel (b): Weak exogeneity										
LR(4)	22.131***	10.941**	6.328	34.432***	14.024***	37.712***	7.319	5.571	28.873***	11.757**
LR(8)	35.555***		42.029***		52.948***		12.586		41.605***	

Lag	1	2	3	4	5	6	7	8	9	10
Panel (c): Test for residual autocorrelation										
	3.090	7.793	44.582	94.171	118.78	148.50	189.55	228.23	275.35**	297.66**
Panel (d): Test for ARCH effects										
	4.028***	3.241***	2.704***	3.616***	3.413***					

	R1	R2
Panel (e): Transition probabilities		
R1	.938	.140
R2	.062	.860

Note: Panel (a) reports the estimates of the adjustment coefficients for two regimes (R1 and R2) with t-statistics in parentheses. The estimated cointegrating vectors are identical to panel (a) in Table 2.2. Estimates of the short-run dynamics, drift terms and variance-covariance matrix are not shown to conserve space. Panel (b) reports weak exogeneity tests for each regime (first row) and over both regimes (second row). The likelihood ratio (LR) statistics are χ^2 distributed with degrees of freedom in parentheses. Panel (c) shows the results of vector portmanteau tests of the residuals. Panel (d) shows the results of tests for ARCH effects. Panel (e) displays the estimated transition probabilities.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.A–2
Cointegration tests and linear VECM (Dubai normalization).

$N-r$	r	Eig.value	Trace	5% Crit. val.	p-Value
Panel (a): I(1)-analysis					
5	0	.1084	361.75	76.07	.000
4	1	.0651	192.16	53.12	.000
3	2	.0374	92.61	34.91	.000
2	3	.0292	36.22	19.96	.000
1	4	.0013	1.95	9.24	.783

	WTI	Brent	Bonny	Dubai	Tapis	μ
Panel (b): Cointegration vectors						
β_1				−.958	1	−.284
β_2			1	−.966		−.209
β_3		1		−.963		−.200
β_4	1			−.881		−.515
Panel (c): Adjustment coefficients						
α_1	.066*	.104***	.110***	.049	−.162***	
	(1.879)	(2.974)	(3.159)	(1.514)	(−6.198)	
α_2	−.214*	−.229**	.432***	−.272***	−.001	
	(−1.933)	(−2.062)	(−3.915)	(−2.638)	(−.018)	
α_3	.198*	.053	.257**	.242**	.090	
	(1.701)	(.451)	(2.216)	(2.237)	(1.028)	
α_4	−.086***	.007	.001	−.003	.011	
	(−3.254)	(.277)	(.050)	(−.110)	(0.570)	
Panel (d): Weak exogeneity						
LR(4)	16.47***	22.87***	33.38***	7.54	43.07***	

Lag	1	2	3	4	5
Panel (e): Test for residual autocorrelation					
	3.398	9.366	66.174***	148.79***	196.38***
Panel (f): Test for ARCH effects					
	2081.5***	2937.7***	3790.5***	4971.8***	5529.4***

Note: Panel (a) reports Johansen (1988) cointegration tests. The critical values are taken from Osterwald-Lenum (1992). p-values are computed using a simulation study with 10,000 replications. Panel (b) displays the estimates of the cointegrating vectors. Insignificant variables have been excluded from the cointegrating vector. Panel (c) reports the estimates of the adjustment coefficients with t-statistics in parentheses. Estimates of the short-run dynamics, drift terms and variance-covariance matrix are not shown to conserve space. Panel (d) reports weak exogeneity tests. The likelihood ratio (LR) statistics are χ^2 distributed with degrees of freedom in parentheses. Panel (e) shows the results of vector portmanteau tests of the residuals. Panel (f) shows the results of tests for ARCH effects.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.A–3

Markov-switching error correction model for major crude oil prices (three-state model, Dubai normalization).

	WTI			Brent			Bonny Light			Dubai			Tapis		
Panel (a): Switching adjustment coefficients															
	R1	R2	R3	R1	R2	R3	R1	R2	R3	R1	R2	R3	R1	R2	R3
$\alpha_1(s_t)$	−.018	.080	.068	−.036	.102	.122***	.016	.063	.134***	−.030	.025	.055	−.144***	−.312***	−.105***
	(−.393)	(.535)	(1.590)	(−.719)	(.666)	(2.650)	(.305)	(.405)	(2.910)	(−.658)	(.175)	(1.440)	(−3.660)	(−2.700)	(−3.280)
$\alpha_2(s_t)$	−.073	−.826**	.072	−.053	−.645*	.009	−.479***	−.835**	−.055	−.217	−.615*	−.015	.035	−.360	.202*
	(−.453)	(−2.220)	(.495)	(−.313)	(−1.740)	(.062)	(−2.660)	(−2.220)	(−.369)	(−1.390)	(−1.750)	(−.115)	(.256)	(−1.270)	(1.860)
$\alpha_3(s_t)$	.166	1.124***	−.176	.094	.627	−.284*	.513***	.853**	−.232	.350**	.720*	−.079	.074	.611*	−.172
	(.968)	(2.670)	(−1.140)	(.524)	(1.470)	(−1.750)	(2.680)	(1.970)	(−1.430)	(2.160)	(1.800)	(−.552)	(.511)	(1.890)	(−1.480)
$\alpha_4(s_t)$	−.098	−.426***	−.018	.018	−.142	.036	.023	−.160	.035	−.022	−.127	.027	.027	−.084	.033*
	(−1.450)	(−3.320)	(−.738)	(.256)	(−1.120)	(1.450)	(.315)	(−1.240)	(1.420)	(−.338)	(−1.050)	(1.210)	(.488)	(−.852)	(1.740)
Panel (b): Weak exogeneity															
LR(4)	3.277	13.250**	5.962	.962	4.173	18.255***	7.744	6.271	19.733***	6.013	3.491	4.816	14.596***	12.992**	13.814***
LR(12)		23.029**			26.316***			36.293***			13.450			42.970***	

Lag	1	2	3	4	5	6	7	8	9	10
Panel (c): Test for residual autocorrelation										
	5.615	12.342	43.175	82.409	103.99	126.05	138.06	196.39	229.13	250.39
	(.999)	(.999)	(.999)	(.999)	(.914)	(.923)	(.816)	(.559)	(.411)	(.481)
Panel (d): Test for ARCH effects										
	2.583	2.382	2.317	2.048	1.964					
	(.000)	(.000)	(.000)	(.000)	(.000)					

	R1	R2	R3
Panel (e): Transition probabilities			
R1	.952	.182	.047
R2	.037	.770	.033
R3	.012	.048	.919

Note: R1 refers to the 'early regime', R2 to the 'crisis regime' and R3 to the 'tranquil regime', respectively. Panel (a) reports the estimates of the adjustment coefficients for three regimes with t-statistics in parentheses. The estimated cointegrating vectors are identical to panel (a) in Table 2.A–2. Estimates of the short-run dynamics, drift terms and variance-covariance matrix are not shown to conserve space. Panel (b) reports weak exogeneity tests for each regime (first row) and over all three regimes (second row). The likelihood ratio (LR) statistics are χ^2 distributed with degrees of freedom in parentheses. Panel (c) shows the results of vector portmanteau tests of the residuals with p-values are given in brackets. Panel (d) shows the results of tests for ARCH effects with p-values are given in brackets. Panel (e) displays the estimated transition probabilities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Chapter 3

The Quantile-Heterogeneous Autoregressive Model of Realized Volatility: New Evidence from Commodity Markets[†]

3.1 Introduction

This chapter presents new insights into the dynamics of gold-, silver and crude oil market volatility. Specifically, we start by using a heterogeneous autoregressive model of realized volatility (HAR-RV) to obtain insights into the temporal dependence of today's volatility with respect to past daily, weekly, and monthly volatility aggregates. This allows us to explicitly capture the possibly different contributions of market participants with diverse trading motives and investment horizons. The interaction of heterogeneous traders seems particularly evident in commodity markets. On the one hand there are producers and commercial traders that aim to hedge against (future) price changes (long-term horizon). On the other hand, there are investors that have no 'real' interest in the underlying asset but consider the future itself as a financial asset (short- and medium-term horizon). The resulting linear description of the average volatility dynamics and relative importance of past volatility components might, however, provide an incomplete view if the volatility process, in fact, is non-linear (i.e. depends on the level of volatility itself).

Therefore, we additionally consider a so-called quantile-heterogeneous autoregressive model of realized volatility (Q-HAR-RV).[1] This model is particularly well suited in this context since it allows to capture time-varying volatility dynamics and hence to explicitly consider different 'states' of the financial markets. Upper quantiles of the realized volatility indicate high-volatility days. These days are particularly interesting from a risk management perspective as they are typically associated with increased

[†]This chapter is joint work with Robert Maderitsch and has been originally published as KUCK, K. AND R. MADERITSCH (2019). "The Quantile-Heterogeneous Autoregressive Model of Realized Volatility: New Evidence from Commodity Markets," in *Financial Mathematics, Volatility and Covariance Modelling*, Volume 1, ed. by Julien Chevallier, Stéphane Goutte, David Guerreiro, Sophie Saglio and Bilel Sanhaji, New York, NY: Routledge. DOI 10.4324/9781315162737

[1]Similar models have been used in slightly different forms by Bonaccolto and Caporin (2014) and by Žikeš and Baruník (2016).

uncertainty and high trading intensity in a market. Intermediate (and lower) quantiles of the volatility distribution, by contrast, represent medium- and low-volatility days that are typically related to phases of calm market conditions and relatively low trading volumes.

Interest in commodity market volatility has increased steadily over the last years. Particularly through index investment, commodity futures have become part of more and more investment portfolios due to expected diversification benefits (Tang and Xiong, 2012). An increasing share of investors thus regards commodity futures as an asset (Adams and Glück, 2015). Typically, trading commodity futures is now also embedded into complex risk management processes ('dynamic risk management'), while the perceived complexity of commodity futures markets seems to have increased for commodity producers and consumers alike (see, e.g., Aepli, Füss, Henriksen, and Paraschiv, 2017).[2]

The remainder of this chapter is structured as follows: we start with a brief introduction into realized volatility estimation. Then, we describe the heterogeneous autoregressive model of realized volatility (HAR-RV) which provides the basis for the quantile-heterogeneous autoregressive model of realized volatility that we present in the subsequent section. We discuss basic econometric features of this new model. Also, we explain why we consider this model particularly interesting for the analysis of commodity market volatility and how it contributes to answering commodity market-specific economic questions. To illustrate our theoretical considerations (different traders, trading motives), we present an empirical application for selected commodity futures (gold, silver and crude oil) together with a detailed economic interpretation. We conclude with a short summary and outlook.

3.2 Realized volatility and quantile regression

3.2.1 Realized volatility

Return volatility is typically perceived as the *risk* associated with a particular asset. Alternatively, it is (often) interpreted as a measure for *information flow* in financial markets. It is central in finance, both in theory and practice. Option pricing and asset allocation, portfolio selection and re-balancing are only some applications to which *volatility* is a highly important input parameter. Various popular risk measurement and risk management approaches such as Value-at-Risk rely on it either. Volatility estimation hence is an important topic. As most of these applications are typically embedded in a dynamic context, understanding the behavior of volatility over time is clearly of great relevance as well.

A fundamental problem with volatility is the fact that real world financial time series are observable at certain points in time (frequencies) only while the true underlying *price processes* are continuous. Therefore, volatility, in fact, is latent and needs to be estimated from data, often observed at lower frequencies. The literature proposes various different solutions. For example, historical volatility, range-based volatility, implied volatility, and the popular ARCH and GARCH models belong to the methods widely used to estimate volatility from daily, weekly and monthly return data. Recently, the use of *realized volatilities* has become more and more popular. On the one hand, this can be explained by the increasing availability of high-frequency data; on the other hand this may be due to the realized volatility's attractive empirical features.

[2]Moreover, speculative trading with commodities, particularly with agricultural commodities, is perceived controversially in both media and academia. See for example Gilbert (2010) and the literature mentioned therein on the heated discussions about the drivers of agricultural futures prices.

The theoretical foundation for the realized volatility relies on the assumption of a continuous (stochastic) price process and the theory of quadratic variation. Andersen, Bollerslev, Diebold, and Ebens (2001a) have shown that the volatility of the latent price process can be obtained consistently, if the sampling frequency tends to an infinitesimal time interval. The daily volatility, for instance, can then simply be estimated from the square root of the sum of squared intra-day returns observed at ultra-high frequencies.[3] The resulting quantities can then be treated as observed realizations of daily volatility, which means that they can directly be used in conventional econometric models.

In practice, however, an important problem with high-frequency financial data is the presence of so-called market microstructure noise. While the theory suggests to sample at the finest time intervals possible, in reality, price observations at ultra-high frequencies are typically characterized by various distorting effects such as discreteness, jumps, bid-ask bouncing (occurring when transactions are priced between the bid and ask prices), or non-synchronous trading. Presence of either of these microstructure effects leads to a bias in the resulting volatility estimates. The benefits of frequent sampling are thus traded off against the issues caused by cumulating noise.

Empirical applications typically avoid these problems by using a sampling frequency that balances the trade-off between bias and variance. There are various different tools that can help to pick the optimal sampling frequency. For example, volatility signature plots, proposed by Andersen and Bollerslev (1998), help in visualizing this bias-frequency relationship and hence can provide guidance in finding the appropriate sampling frequency.[4] Typically, the findings suggest to sample at moderate intra-daily frequencies such as every 5, 10, or 30 minutes and to discard the information from within those time intervals. Overall, however, the 5-minute frequency has become most common in the literature. It has been shown to be most adequate in solving the trade-off between bias and variance in the realized volatility estimator (Andersen, Bollerslev, Frederiksen, and Nielsen, 2010; Liu, Patton, and Sheppard, 2015). Alternatively, bias correction techniques or market microstructure robust volatility estimators can be used for data sampled at ultra-high frequencies instead. Andersen et al. (2001a), for instance, propose a moving-average filter for raw high-frequency returns. Martens and van Dijk (2007) propose the realized range estimator, which is robust to bid-ask bounce, while Andersen, Dobrev, and Schaumburg (2012) suggest the jump-robust median realized volatility.

Compared to traditional volatility proxies such as squared daily returns, realized volatilities are characterized by a much higher precision since they incorporate more information about the price process (see Andersen and Bollerslev, 1998). Additionally, and in contrast to volatilities estimated from the popular ARCH and GARCH models, realized volatilities are *model-free ex-post* estimates of the underlying asset return volatilities in the sense that they do not rely on the assumption(s) of a parametric model. That is, no potentially restrictive assumptions on the dynamics of the underlying conditional distribution are necessary.

Until recently, studies on realized volatilities focused on the stock and foreign exchange markets (see, for example, Andersen et al., 2001a,b; Andersen, Bollerslev, Diebold, and Labys, 2003; Corsi, 2009; Bubák, Kočenda, and Žikeš, 2011). With growing interest in commodity markets (index investment, risk

[3] See also Andersen and Bollerslev (1998), Andersen et al. (2001a), Andersen, Bollerslev, Diebold, and Labys (2001b), and Barndorff-Nielsen and Shephard (2002) for fundamental literature.

[4] If microstructure noise is present, it will distort the average realized volatility if the sampling frequency is increased. Plotting the average realized volatility against the return frequency hence may help to find the appropriate sampling frequency, i.e., where the average realized volatility stabilizes.

management) and an increasing availability of commodity market high-frequency data, however, some authors started to consider these markets as well (see e.g. Martens and Zein, 2004; Wang, Wu, and Yang, 2008; Souček and Todorova, 2013; Todorova, Worthington, and Souček, 2014; Todorova, 2015).

3.2.2 The heterogeneous autoregressive model of realized volatility (HAR-RV)

One model for the daily realized volatility that has attracted considerable attention in the literature recently is the heterogeneous autoregressive model of realized volatility (HAR-RV), initially proposed by Corsi (2009). In essence, this model is an autoregressive-type cascade model for volatility components aggregated over different time horizons (daily, weekly, monthly). From a theoretical point of view, the HAR-RV model builds upon the considerations of the heterogeneous market hypothesis according to Müller, Dacorogna, Davé, Olsen, Pictet, and von Weizsäcker (1997). It basically assumes that volatilities measured over different time resolutions reflect the perceptions and activities of market agents with heterogeneous investment time horizons. The HAR-RV model identifies three main time resolutions: (1) *daily* to capture the activities of short-term traders with a daily or even higher trading frequency, (2) *weekly* for medium-term investors typically re-balancing their positions about once per week, and (3) *monthly* for longer-term traders with a time horizon of one month or more. This choice corresponds also broadly to behavior observed in real financial markets (Corsi, 2009).

The popularity of the HAR-RV model can be explained mainly by its remarkable empirical performance. Despite its simplicity, it is well suited to reproduce the pronounced volatility persistence as well as many of the other features of financial data. Put differently, it captures the persistence and summarizes the dynamics in the mean of realized volatility with only a few parameters. Further, the parameters in the HAR-RV model have a clear economic interpretation which is a considerable advantage compared to many other approaches that are often characterized by a lack of economic intuition (see Corsi, 2009, and the literature mentioned therein).

Following Corsi (2009) in the first step of our analysis, we consider a basic univariate specification of the HAR-RV model for each of our different markets. More precisely, we express the conditional mean of the realized volatility of today as a linear function of past realized volatilities measured over daily, weekly and monthly time horizons:

$$RV_t^{(d)} = \alpha + \beta^{(d)} RV_{t-1}^{(d)} + \beta^{(w)} RV_{t-1}^{(w)} + \beta^{(m)} RV_{t-1}^{(m)} + \varepsilon_t, \tag{3.1}$$

with $RV_{t-1}^{(d)}$, $RV_{t-1}^{(w)}$ and $RV_{t-1}^{(m)}$ corresponding to the above-mentioned daily, weekly and monthly volatility aggregates and ε_t being a serially uncorrelated zero mean innovation term. Specifically, $RV_{t-1}^{(d)}$ corresponds to the lagged daily realized volatility, whereas the weekly and monthly volatility components, $RV_{t-1}^{(w)}$ and $RV_{t-1}^{(m)}$, are averages over the past 5 and 22 trading days, computed as $(1/5) \times (RV_{t-1}^{(d)} + RV_{t-2}^{(d)} + \ldots + RV_{t-5}^{(d)})$ and $(1/22) \times (RV_{t-1}^{(d)} + RV_{t-2}^{(d)} + \ldots + RV_{t-22}^{(d)})$, respectively. This specification corresponds to a parsimonious autoregressive-type model for the realized volatility, just like an AR(22) process with restrictions.

Using conventional ordinary least squares (OLS) regression techniques, Corsi (2009) estimates the conditional mean of the (daily) realized volatility, given the three right-hand-side volatility components, realized over different time horizons. The estimated dependence parameters allow to obtain insights with respect to the relative importance of past volatilities on the daily volatility. For the S&P 500 In-

dex Future, for example, Corsi (2009) finds a decreasing dependence from the daily, to the weekly, to the monthly volatility component (i.e., $\beta^{(d)} > \beta^{(w)} > \beta^{(m)}$). Differences between dependencies across different volatility aggregates might indicate an asymmetric information flow between short-, medium-, and long-term traders and different proportions of such traders being present in the markets.

3.2.3 The quantile-heterogeneous autoregressive model of realized volatility (Q-HAR-RV)

Being confined to the analysis of the conditional mean of realized volatility, however, the HAR-RV model according to Corsi (2009) is not able to capture any details about state dependence in the volatility dynamics. Put differently, the HAR-RV model cannot describe the potentially time-varying relative importance of past volatility components. Since volatility is typically used in dynamic contexts, it would be very interesting to obtain insights into the temporal dependence with respect to different volatility levels.

Inspired by the work of Baur (2013) and Žikeš and Baruník (2016), we therefore extend the basic HAR-RV model by using the quantile (auto-)regression framework according to Koenker and Xiao (2006). Conventional OLS based autoregressive models focus on the effects of a variable's own lags on its conditional mean. Quantile autoregression instead can be used to estimate the dependence of specific conditional quantiles of the dependent variable on its own lags. That is, the analysis is not confined to study location shifts, but allows to obtain a detailed description of the tails of the dependent variable's conditional distribution. Quantile regression basically allows to describe the complete conditional distribution of the dependent variable. Further, it is robust to conditional heteroskedasticity, skewness, and leptokurtosis which are common features of realized volatility as reflected also in the descriptive statistics of our realized volatility time series (see Table 3.1). Using quantile regressions, we thus essentially relax the assumption of a (constant) linear impact of the different volatility aggregates onto the conditional volatility. More specifically, we can uncover non-linearities and asymmetries in the temporal dependence across specific conditional quantiles of the realized volatility, given the volatility components, measured over different time horizons. Compared to other non-linear models such as the threshold autoregressive models (Tong, 1983) and Markov-switching models (Hamilton, 1989, 1994), quantile regression is flexible in the sense that also complex non-linear dynamics can be captured accurately with a few a priori restrictions. Particularly, it is not necessary to predetermine the number of 'states' or thresholds.

Until now, there are various general applications of quantile regression in empirical finance (e.g., Koenker and Zhao, 1996; Chernozhukov and Umantsev, 2001; Engle and Manganelli, 2004). An application in the context of the HAR-RV model, however, has to the best of our knowledge, so far, only been considered by Bonaccolto and Caporin (2014) and Žikeš and Baruník (2016). Similarly to us, Žikeš and Baruník (2016) employ quantile regression techniques to estimate the dependence parameters for univariate specifications of the HAR-RV model. Their focus, however, is on forecasting specific quantiles of the conditional distribution of daily realized volatility. Our aim, by contrast, is to provide a detailed description of the volatility dynamics with respect to different states of the financial markets. Further, their motivation is rather of a pure econometric nature, whilst our goal is to draw concrete economic conclusions which seem to be absent in this particular literature so far. Apart from that, we apply our Q-HAR-RV in order to study commodity futures markets which is to the best of our knowledge a unique feature as well.

Based on the flexible semi-parametric quantile regression framework, Equation (3.1) now takes the following form of a quantile-heterogeneous autoregressive model of realized volatility (Q-HAR-RV):

$$Q_{RV_t^{(d)}}(\tau \mid \mathscr{F}_{t-1}) = \alpha(\tau) + \beta^{(d)}(\tau)RV_{t-1}^{(d)} + \beta^{(w)}(\tau)RV_{t-1}^{(w)} + \beta^{(m)}(\tau)RV_{t-1}^{(m)}, \tag{3.2}$$

where $\mathscr{F}_{t-1}$ denotes the information available at day $t-1$ and $Q_{RV_t^{(d)}}(\tau \mid \mathscr{F}_{t-1})$ is τth quantile of the realized volatility, conditional on $\mathscr{F}_{t-1}$. Hence, $\beta^{(d)}(\tau)$, $\beta^{(w)}(\tau)$ and $\beta^{(m)}(\tau)$ are the quantile-specific volatility-dependence parameters which are of central interest to us. We interpret them as quantile-specific daily, weekly, and monthly persistence parameters. Estimating Equation (3.2) over the range of all quantiles $\tau \in [.01, .02, \ldots, .99]$ allows us to characterize volatility persistence and dynamics with respect to different states of the commodity markets. Primarily, we are interested in comparing high and low volatility, i.e., upper and lower quantiles of the volatility distribution. Medium volatility around the central quantiles might be of interest as well. According to Baur (2013), the sequence of temporal dependence parameters across all quantiles describes the *structure of dependence*, in our case, however, with respect to the level of today's volatility. In the presentation of our results, we contrast these quantile sequences to the corresponding coefficient estimates from OLS regressions as a benchmark and measure for the *degree* of dependence.

Besides the graphical presentation of the results, we perform a Kolmogorov-Smirnov-type as well as a Wald-type test to formally test for non-linearities and assess the presence of asymmetries in the temporal dependence across low- and high-volatility days. First, to test for the presence of overall quantile effects with respect to the corresponding (linear) dependence parameter (estimated by OLS), we use a Kolmogorov-Smirnov type test according to Bera, Galvao, and Wang (2014). This test allows us to check the general presence of quantile-specific dependencies that differ from the OLS benchmark. In particular, we produce a sequence of Wald statistics $(W(\tau = L), \ldots, W(\tau = U))$ out of individual Wald tests for specific ranges of quantiles ($\mathscr{T} = [L, U]$). The null hypothesis for the presence of an overall quantile effect is then $H_0 : \beta(\tau = j) = \beta^{OLS}$ for all $j \in \mathscr{T}$. We evaluate the supremum of the sequence of Wald statistics:

$$W_n := \sup_{\tau \in \mathscr{T}} W(\tau), \tag{3.3}$$

where W_n does not follow a standard χ_p^2-distribution. According to Bera et al. (2014), we approximate the p-values by using an upper boundary, taking the form

$$\Pr(W_n > u) \leq \Pr(\chi_p^2 > u) + \frac{u^{\frac{p-1}{2}}}{2^{\frac{p}{2}} \exp(u/2)\Gamma(p/2)} \int_{\mathscr{T}} E \left| \frac{\partial W_n^{\frac{1}{2}}(\tau)}{\partial \tau} \right| d\tau. \tag{3.4}$$

We estimate $\int_{\mathscr{T}} E \left| \frac{\partial W_n^{\frac{1}{2}}(\tau)}{\partial \tau} \right| d\tau$ from the total variation $V = \left| W_n^{\frac{1}{2}}(\tau_1) - W_n^{\frac{1}{2}}(\tau_L) \right| + \left| W_n^{\frac{1}{2}}(\tau_2) - W_n^{\frac{1}{2}}(\tau_1) \right| + \cdots + \left| W_n^{\frac{1}{2}}(\tau_U) - W_n^{\frac{1}{2}}(\tau_k) \right|$, where τ_1, τ_2,...,τ_k are the turning points of $W_n^{\frac{1}{2}}(\tau)$ and L and U are the lower and upper bounds of τ.

Further, to test for the presence of differences in the volatility dynamics with respect to different volatility levels prevailing, we use Wald tests again. The null hypothesis to compare a range of depen-

dence coefficients across quantiles is $H_0 : \beta(\tau = .01) + \cdots + \beta(\tau = d) = \beta(\tau = u) + \cdots + \beta(\tau = .99)$, where the Wald statistic takes the form

$$W_{d,u} = (\mathbf{R}\boldsymbol{\beta}(\boldsymbol{\tau}) - r)'(\mathbf{R}\hat{\Omega}_B\mathbf{R}')^{-1}(\mathbf{R}\boldsymbol{\beta}(\boldsymbol{\tau}) - r), \tag{3.5}$$

with $\mathbf{R} = (1, \ldots, 1, -1, \ldots, -1)$, $\boldsymbol{\beta}(\boldsymbol{\tau}) = (\beta(\tau = .01), \ldots, \beta(\tau = d), \beta(\tau = u), \ldots, \beta(\tau = .99))'$, $r = 0$ and $\hat{\Omega}_B$ as the covariance matrix.[5] For one linear restriction, as in our case, the Wald statistic is approximately χ^2-distributed with one degree of freedom.

The model described in Equation (3.2) is linear in parameters (for a given quantile, τ) and all parameters can be estimated according to the standard optimization routine presented in Koenker and Bassett (1978) and implemented, for example, in R. Further, we estimate asymptotic standard errors based on bootstrap methods.

While quantile regression is not well-suited for (multiple step ahead) forecasting applications (since future quantiles are unknown in advance), it can be used to obtain a precise description of the volatility dynamics. In essence, the Q-HAR-RV model allows us to broaden the perspective and to study differences and changes in the volatility dynamics along two dimensions: first, motivated by economic theory and like the pure HAR-RV model, it captures the relative importance of volatility measured over different time resolutions which we assume to reflect the activities of heterogeneous market participants. Second, the use of quantile regression allows us to capture the changing relative importance of volatility components, depending on the state of the volatility or put differently, the respective quantile of the (conditional) volatility distribution.

3.3 Motivation for an application to commodity markets

3.3.1 Volatility persistence and differences across markets: HAR-RV

The heterogeneous market hypothesis proposed by Müller et al. (1997) states that volatility measured over different time resolutions reflects the perceptions and activities of market agents characterized by heterogeneous investment horizons. While traders may differ along several dimensions, a distinction between different types of traders with respect to their investment time horizon seems economically plausible and useful (see, e.g., Müller et al., 1997; Corsi, 2009; Pennings and Garcia, 2010).

In our point of view, the interplay between market agents with different trading motives and time horizons might be particularly evident in commodity markets. On the one hand, there are producers and commercial traders, having a 'real' and rather longer term interest behind their transactions. These traders will typically trade infrequently. Often, the primary objective behind their transactions is to transfer their exposure to future price movements to speculators taking the corresponding long and short positions. On the other hand, there is a large fraction of speculators who have no interest in the delivery of the underlying commodity but consider the futures contract itself a (financial) asset. Among them, there are index investors and participants trading a lot among each other and executing transactions also at very high frequencies, constantly monitoring and re-evaluating the markets. Obviously, these short-term traders might react quickly to changes in short-term volatility. Short-term traders, and thus the short-term

[5] Note that we estimate the variance-covariance matrix Ω for the quantile-dependent autoregressive coefficients by using bootstrap techniques.

(daily) volatility, however, may also be affected by medium- and long-term volatility. By contrast, for investors with rather long time horizons, the short-term volatility might be relatively unimportant. Their decisions can be assumed to be mostly impacted by long-term volatility and 'fundamentals' (Müller et al., 1997, p. 217).

In addition, volatility dynamics may vary across commodity markets (agriculture, energy, precious metals, rare-earth element(s)) due to differences in certain characteristics of the underlying commodities. Precious metals, such as gold and silver, are durable and provide a store of value. Hence, precious metals are used as *investment asset*.[6] For gold, this might be particularly true. It is widely used for portfolio diversification and is kept as a reserve by virtually every central bank in the world. By the nature of silver and gold, we would hence expect a large share of speculative investment traders to be active in these markets. Since they are used as investment assets they might in many respects behave very similar to financial assets, like for example, equities. Beyond that, precious metals are also used in industry (Schweikert, 2018). For instance, gold is used in restorative dentistry and for high-quality electrical connectors. Silver is the most reflective metal and is used in photography, optics and the solar-energy industry. By contrast, other commodities have industrial use only, are non-durable and hence share the properties of a *consumption asset*. Crude oil, for instance, is primarily used for transportation and heating fuel. It is also utilized for electricity generation and as an input for the production of asphalt and road oil as well as the making of chemicals, plastics, and synthetic materials. Crude oil hence is an important economic input factor. This may also manifest in the presence of a more pronounced 'real' delivery aspect. The cost of carry, the lease rate (compensation for lending) and the convenience yield are more important than in case of silver and gold. Moreover, the price of crude oil reacts more sensitively to the supply side, since the stockpile of crude oil is substantially less than the annual extraction. In contrast, the prices of precious metals are rather driven by the demand-side since they are durable assets and stockpile outweighs annual production.

We believe that the above-mentioned considerations can provide valuable economic intuition for the interpretation of our estimation results. We translate the above considerations into two summarizing statements on commodity market realized volatility dynamics:

1. Measured volatility dependence might differ across (commodity) markets. The dependence patterns identified might provide valuable indications about investors' time horizons and the proportion between short- and long-term traders being active in the markets.

2. For non-durable assets (*consumption assets*) like crude oil, the volatility of today might depend more strongly on the long-term and less strongly on the short-term volatility component in comparison to *investment assets* such as silver and gold.

3.3.2 State-dependent volatility dynamics: Q-HAR-RV

The use of the Q-HAR-RV model is mainly motivated by the potential presence of state dependence in the volatility dynamics. Traders' heterogeneity in time horizons might not only result in heterogeneous reactions to price fluctuations captured over different time intervals, but also in different reactions or attention shifts depending on high, moderate and low volatility, and thus state dependence. For financial

[6] Note that in addition to bars and coins, jewelry provides a store of value and may be purchased for the purpose of portfolio diversification as well.

returns, for example, it is a well-known phenomenon that temporal dependencies can be characterized by asymmetries under certain circumstances (see, e.g., Barberis, Shleifer, and Vishny, 1998; Veronesi, 1999; Zhang, 2006; Baur, 2012; Kuck, Maderitsch, and Schweikert, 2015).

Theoretically, asymmetries in realized volatility dependence might be explained by shifts in traders' attention or by changes in the shares of certain types of traders being active in the markets. In presence of very high volatility (today), for instance, uncertainty prevails and speculators might react more strongly to past volatility than in the presence of moderate volatility. If prices fluctuate strongly (in the short run), the potential for short-term gains is high. Traders might therefore open and close their positions more quickly than during episodes of moderate volatility due to an increased attractiveness of short-term speculation. This might even be self-reinforcing, by first increasing the share of speculators and second, causing further increases in volatility on high volatility days. At the same time, episodes of uncertainty and crises are typically characterized by extreme negative returns. The need to close-out positions in order to avoid losses or meet margin calls might be stronger then. That is, tensions on the financial markets and commodity markets might also force traders to react and rebalance their portfolios more quickly. This effect might be reflected in a particularly pronounced increase in the dependence on the past short-term volatility component.

Finally, drawing attention to the differences between consumption and investment assets, we would again expect the share of short-term speculators to be particularly high in investment markets which might result in stronger differences in volatility patterns between high and low volatility days in these markets. By contrast, for consumption assets, state dependence of the volatility dynamics could be less pronounced due to a smaller share of short-term speculators and a tendency of market participants to focus on the long term.

In essence, there are various different reasons to assume that volatility dependence might be characterized by asymmetry. They translate into the following summarizing statements:

3. Realized volatility dependence might differ statistically significantly across the high and low volatility state.

4. The asymmetries in dependence across the high and low volatility state might be more pronounced in case of investment assets (gold, silver, S&P 500) than in case of consumption assets (crude oil).

3.4 Empirical evidence from major commodity markets

3.4.1 Data

The data that we use for the empirical illustrations consists of WTI light sweet crude oil, gold and silver futures bid and ask quotes from the Thomson Reuters TickHistory database.The sample covers the period from January 2007 to September 2014 for gold and silver and from January 2000 to September 2014 in case of WTI light sweet crude oil.

While multiple futures contract maturities are traded simultaneously, it is convention to consider only the price of the next-to-maturity contract (front contract). To get a continuous series of futures prices over an extended period it is thus necessary to 'switch' from the front contract to the next one (the second next-to-maturity) at some point. Here, the expiry date of the front-contract is chosen as the so-called

rollover date, that is, the point at which we switch from one contract to another.[7] Assuming that gold-, silver and WTI futures contracts are traded actively, i.e. are liquid, the choice of the rollover date should not affect the resulting time series in any systematic way. Carchano and Pardo (2009), for example, are not able find any evidence for a significant impact of the choice of the rollover date on the return series.

For the construction of intra-day returns underlying our realized measures, we follow Scharth and Medeiros (2009), Giot, Laurent, and Petitjean (2010), Andersen et al. (2012) and others and use the midpoint of the bid and ask quotes as a price indication. An appealing feature of proceeding this way is that we can rule out the presence of bid-ask bounce effects which might distort volatility estimation when transaction prices are used instead.[8]

As common in the literature, we sample our intra-day returns at the 5-minute interval. Then we compute the daily realized volatilities by

$$RV_t^{(d)} = \sqrt{\sum_{j=1}^{M} r_{t,j}^2},$$

where $r_{t,j}$ is the jth 5-minute return on day t and M is the number of intra-day returns. That is, we construct the daily realized volatilities for each of our markets by summing up the squared intra-daily log-returns ($r_{t,j}^2$) for each single trading day in the sample. Following e.g. Bubák et al. (2011), we define a trading day to last from 21:00 GMT to 20:59 GMT the following day. We delete weekends as well as public holidays such as Christmas Eve and New Year in order to avoid periods of low trading volume. For model estimation, we log-transform the realized volatilities, in order to improve the statistical properties of the time series.

The (daily) mid-prices and the corresponding log-realized volatilities of our three futures are depicted in Figure 3.1. All series show typical characteristics of financial time series. Whereas the mid-prices resemble non-stationary random series, the realized volatilities appear stationary and show volatility clustering (strong persistence).

Table 3.1
Descriptive statistics for the daily realized volatilities.

	Mean	SD	Skew	Kurt	Min	Max	Obs
Gold RV	1.2334	0.8187	1.90	7.67	0.0971	6.0314	2648
log(Gold RV)	0.0182	0.6323	−0.28	3.58	−2.3324	1.7970	2648
Silver RV	2.1122	1.3131	1.62	6.64	0.1458	10.9597	2410
log(Silver RV)	0.5622	0.6394	−0.56	3.87	−1.9253	2.3942	2410
WTI RV	1.6759	1.1194	2.11	9.89	0.1165	10.0306	2714
log(WTI RV)	0.3157	0.6606	−0.48	3.62	−2.1500	2.3056	2714
S&P500 RV	1.0480	0.6755	2.96	18.03	0.1030	8.1343	3807
log(S&P500 RV)	−0.1027	0.5306	0.27	3.68	−2.2731	2.0961	3087

[7] In our case this results in 12 rollover dates each year due to monthly expiry cycle.

[8] As argued in Hansen and Lunde (2006), using mid-quotes instead of transaction data reduces the spurious serial correlation in the high-frequency returns due to bid-ask bounce and non-synchronous trading effects.

Figure 3.1
Time series of prices and realized volatilities.

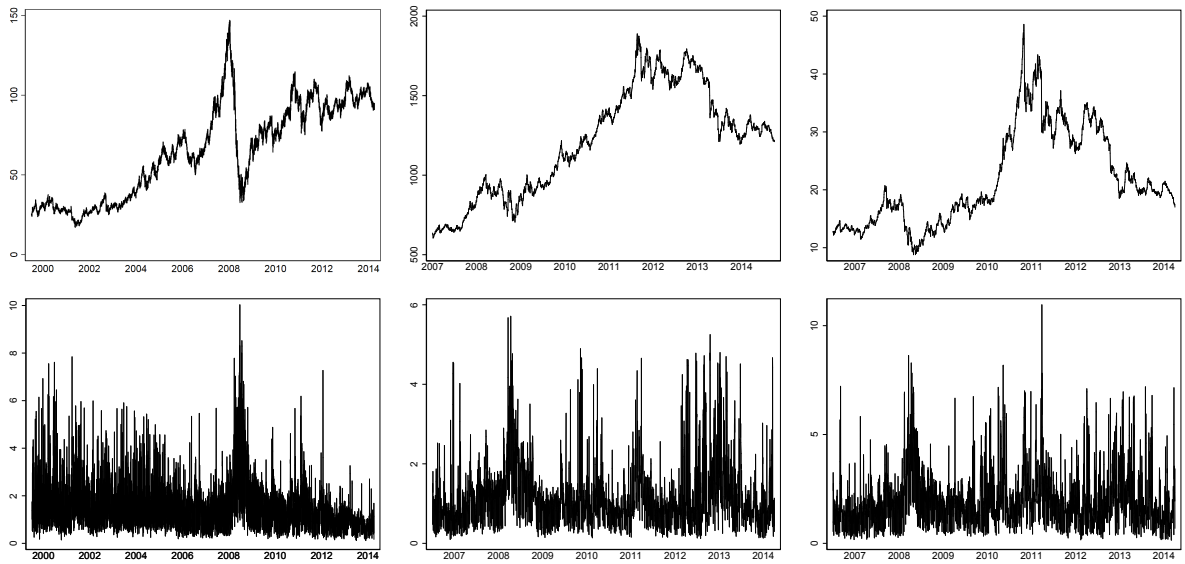

Light Sweet Crude Oil Gold Silver

Notes: Daily mid-prices are depicted in the upper panel. The lower panel shows daily realized volatilities.

3.4.2 Estimation Results

HAR-RV Results

The estimation results for the pure HAR-RV model for the three commodities are reported in Table 3.2. In order to compare the magnitude of the estimated coefficients, all variables have been standardized before estimation (Verbeek, 2017). To put our findings into perspective, we include also the estimation results for a HAR-RV model for the S&P 500 futures.[9] First, the temporal dependence coefficients for the past daily, weekly and monthly volatility components are, generally, statistically significant across all four series, indicating long memory, i.e., persistence in the volatility series. The only exceptions exist for gold and silver where the coefficient for the past weekly realized volatility components are not statistically significantly different from zero. Regarding the magnitudes of the coefficients, gold and silver appear to be very similar. Substantial differences are, however, evident between precious metals and both crude oil and the equity market.

Following Corsi (2009), temporal dependence on past realized volatilities aggregated over different time scales might represent the activities of market participants with different time horizons. The coefficients on the volatility aggregates might hence provide an indication if a market is composed rather by homogeneous or heterogeneous market participants. The differences across coefficients that are evident in Table 3.2, indeed support the presence of heterogeneous traders across the four markets (Statement 1). As a general pattern, we find the coefficients for the monthly component to be (slightly) more pronounced than the daily and weekly components in case of commodities. This might indicate the relative impor-

[9] Specifically, we consider the E-mini S&P 500 Futures contracts.

tance of market participants with long-term time horizons in the commodities markets, particularly in the crude oil market.[10]

Comparing now in detail gold, silver, crude oil and S&P 500 volatility dynamics, it is apparent that the coefficient on monthly volatility for precious metals is similar in magnitude to that of the S&P 500's monthly component.[11] In contrast to precious metals, the dependence on past weekly volatility appears to be relatively more important than both the daily and monthly volatility component for equity. Like for equities (S&P 500), the volatilities of precious metals exhibit marked dependence on past daily volatility, although somewhat less pronounced. This result supports our idea that precious metals are also used similar to traditional investment assets. In case of crude oil, by contrast, we find a strikingly low dependence on the daily volatility component, while the dependence on the monthly volatility component is rather strong (in absolute terms), also relative to that of the other assets.[12] Moreover, we find a negative relation between past weekly volatility and daily realized volatility in case of crude oil futures. This may indicate some form of cyclicality.

Overall and in line with Statement 2, the impact of long-term volatility is strongest for crude oil while we find it to be substantially weaker for gold, silver and S&P 500 (in absolute terms). This supports the idea that the crude oil futures market might be characterized by a comparably larger share of traders with a long-term investment horizon whereas the gold and silver futures markets appear to be composed of both short-term and long-term oriented traders. This might also reflect that precious metals, particularly gold, are investment assets but also have industrial use, while crude oil is of industrial use only.

Table 3.2
HAR-RV estimation results.

	Gold	Silver	WTI Crude Oil	S&P500
Constant	0.0012	0.0026	−0.0017	−0.0016
	(.07)	(.14)	(−.13)	(−.18)
Daily	0.2265***	0.1886***	0.0918***	0.2725***
	(10.40)	(8.14)	(5.58)	(14.03)
Weekly	0.0238	0.0414	−0.1830***	0.3965***
	(.78)	(1.29)	(−6.81)	(13.80)
Monthly	0.2464***	0.2476***	0.5339***	0.1949***
	(8.90)	(8.51)	(21.06)	(8.69)

t-statistics in parantheses.
*** (**, *) denotes significance at a 1% (5%, 10%) significance level.

[10] If investors have limited attention and short-term traders care about short-, medium-, and long-term volatility whereas long-term traders only care about long-term volatility, then this should be reflected in a stronger dependence on the long-term volatility component than the short-term (medium-term) component.

[11] Considering the coefficient's dimensions and their precision (large t-statistics) allows to uncover differences in the parameters across different models. Approximately, if the confidence intervals of two coefficients do not overlap this points to their statistical diversity. Note that it is not possible to formally test parameter equality across equations (different models) using conventional Wald-type tests.

[12] Given the high precision of the estimates, it can well be assumed that the confidence intervals on the coefficient of the monthly component of crude oil do not overlap with that of gold and silver, respectively.

Q-HAR-RV results

For ease of understanding, we present the estimation results of the Q-HAR-RV models graphically.[13] For each market, Figure 3.1 shows the sequence of the estimated $\hat{\beta}(\tau)$s over all percentiles τ (black solid curve) together with the corresponding 99% confidence band (gray-shaded area). In addition, we show the benchmark OLS coefficient with its 99% confidence interval (dashed line). The lower and upper bounds of the confidence bands for the quantile regression coefficients are based on bootstrapping techniques.

On the one hand, as evident in Figure 3.1, there are substantial differences in the volatility dependence parameters across quantiles for commodities, indicating that the volatility dynamics are time-varying. For the S&P 500, by contrast, the quantile specific coefficients are very close to their OLS benchmark. Though the asymmetry is not very pronounced for the S&P 500, the overall pattern in the temporal dependence structure, in particular with respect to past daily and weekly volatility, appears similar to that of silver and gold. On the other hand, there are also some general patterns and similarities in dependencies across our commodity futures: first, the dependence tends to be stable (the quantile-specific dependence parameters evolve relatively horizontal) across quantiles for the daily volatility component, except for very extreme quantiles. Second, the impact of the weekly volatility component is significantly negative for lower and positive for upper (conditional) quantiles of the daily realized volatility. Moreover, the dependence on past weekly volatility seems to become more important when today's volatility increases from moderate to high levels. Third, in contrast, the temporal dependence tends to decrease from lower to upper quantiles for the monthly volatility aggregate. More specifically, the dependence on the monthly volatility component becomes insignificant for extreme upper quantiles. In essence, this suggests that information generated over the medium term (five days) gains relative importance in phases of increased uncertainty. To sum up, the estimated impact of past realized volatility is significantly positive across quantiles except for extreme quantiles. Only for the weekly volatility component in the commodities markets we find a significant negative impact on today's volatility at lower quantiles.

Specifically, our results indicate that the daily volatility component tends to be a stronger volatility driver (or amplifier) on low volatility days. On high volatility days, however, it tends to have a relatively weaker (positive) impact. The weekly component, by contrast, tends to have a dampening impact on low volatility days and a positive, amplifying effect on high volatility days. For the monthly volatility component, our findings suggest the presence of an amplifying impact that tends to be particularly strong when today's volatility is low. By contrast, when volatility is high, the dependence on past monthly volatility becomes insignificant. In sum, volatility on low volatility days tends to be driven strongly by past daily and monthly volatility, whereas volatility on high volatility days tends to be driven strongly by the weekly volatility component.

From an economic perspective, these findings suggest the following conclusions: first, when volatility is low (moderate) it seems to be mainly caused by the activities of traders with short- and long-term investment horizon. Second, high volatility is driven particularly by the activities of traders with medium-term time horizon. Moreover, the observed negative dependence on past weekly volatility in commodity markets when volatility is low might be related to some sort of cyclicality. It seems that for the weekly component, a change in traders' composition might be apparent. Alternatively, traders might tend to

[13] Again, we standardize the dependent and independent variables before estimation in order to be able to compare the magnitudes of coefficients.

focus more strongly on ex-post volatility measured over the medium-term in case of high (extreme) volatility days. Simultaneously, their attention to the long-run volatility component tends to decrease.

Finally, we use Wald-type tests according to Bera et al. (2014) in order to formally assess the statistical significance of the non-linearities in the volatility dynamics. First, a sup-Wald test is used to study the presence of significant overall non-linearities (quantile effects). Secondly, we apply a Wald test to investigate the significance of potential asymmetries in dependencies for different levels of the daily realized volatility. Specifically, we test the presence of differences in the temporal dependence on past volatilities across the 1%–10% and the 90%–99% quantiles (see Statement 3). The sup-Wald tests confirm the presence of non-linearities in the temporal dependence on past volatilities aggregated over daily, weekly and monthly time horizons, see Table 3.1, Panel (a). Only in case of crude oil and S&P 500, the null hypothesis of linear temporal dependencies cannot be generally rejected at conventional significance levels. Specifically, the dependence on past daily volatility is not characterized by non-linearities in case of crude oil. For the S&P 500, the dependence on past weekly and monthly volatilities seems to be linear. Likewise, the tests confirm the presence of asymmetric dependence particularly between low and high volatility, see Table 3.1, Panel (b). Only for past daily and past monthly volatility the null hypothesis of symmetric temporal dependence across quantiles cannot be rejected in case of crude oil and S&P 500, respectively.

3.5 Summary and Outlook

This chapter revisited the HAR-RV and Q-HAR-RV models and demonstrated their usefulness to obtain interesting new insights into volatility dynamics of commodity markets. Specifically, we studied the temporal dependencies of gold, silver and crude oil Futures volatility in detail.

First, the differences across the dependence on the daily, weekly, and monthly volatility components that we find are theoretically consistent with the presence of traders with heterogeneous time horizons. Specifically, the impact of long-term volatility appears to be more dominant in our analyzed commodity markets than in the equity market. In addition, the differences across different commodity markets that we find might indicate different proportions between short-term and long-term traders across investment and consumption asset markets. In the case of crude oil as a consumption asset, we find the dependence on the longer term (monthly) volatility component to be substantially stronger than in case of the investment assets gold, silver, and the S&P 500.

Second, using the Q-HAR-RV model, we find volatility dynamics in the markets for precious metals and crude oil to be characterized by substantial state dependence. The non-linear dependence patterns that we reveal appear to be theoretically consistent with the idea of shifts in investor attention or changes in investor composition depending on the level of current volatility. Interestingly, we find strong similarities in the dependence patterns across our commodity futures. Specifically, for the weekly and monthly volatility component there are substantial differences in the dependence on past volatility aggregates between low- and high-volatility states. In presence of high volatility, the dependence on the long-term (monthly) volatility component tends to decrease generally while the dependence on the medium-term (weekly) volatility aggregate tends to increase notably.

It will be of interest in future work to compare these findings based on a small selection of commodity markets with a broader set of commodities as well as other financial markets (e.g., foreign exchange).

Further studies could also investigate the dependence structure across indexed and non-indexed commodities.

Table 3.1
Tests for non-linearities and asymmetries between low and high volatility.

	Daily	Weekly	Monthly
Panel (a): Tests for non-linearities (quantile effects)			
Gold	30.06	84.30	18.46
	(0.00)	(0.00)	(0.00)
Silver	66.31	144.00	49.71
	(0.00)	(0.00)	(0.00)
Oil	11.51	63.57	22.73
	(0.03)	(0.00)	(0.00)
S&P 500	21.48	14.21	6.65
	(0.00)	(0.01)	(0.34)
Panel (b): Tests for asymmetries			
Gold	24.72	77.47	25.24
	(0.00)	(0.00)	(0.00)
Silver	71.27	92.96	31.37
	(0.00)	(0.00)	(0.00)
Oil	2.92	46.93	22.57
	(0.09)	(0.00)	(0.00)
S&P 500	14.97	8.70	1.71
	(0.00)	(0.00)	(0.21)

p-values in parantheses.

Figure 3.1
Q-HAR-RV estimation results.

Panel (a): Gold futures

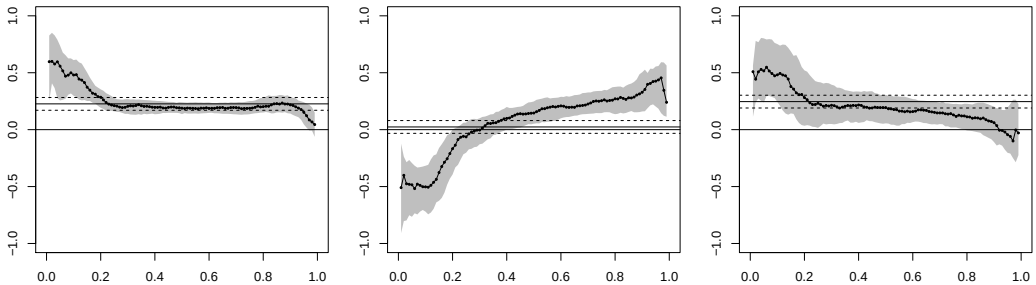

Panel (b): Silver futures

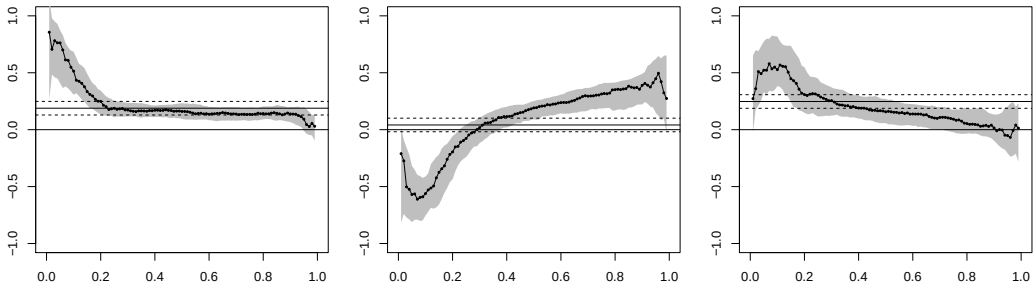

Panel (c): Light sweet crude oil futures

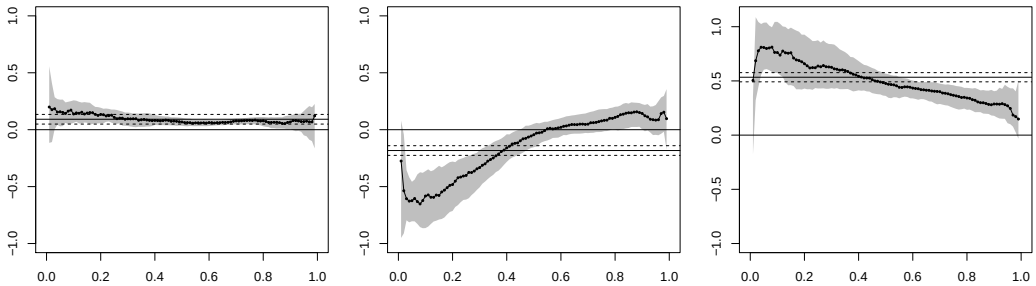

Panel (d): S&P 500 futures

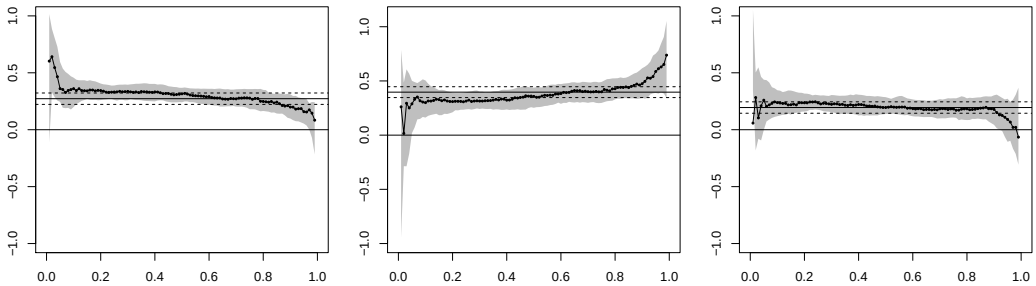

Notes: Dependence on past daily (left), weekly (middle) and monthly (right) volatility. Solid horizontal line: OLS coefficient; dashed horizontal lines: corresponding 99% confidence bands. Solid curve: quantile-specific coefficients; Gray-shaded: area within the corresponding 99% confidence band; Confidence bands are based on standard errors estimated via bootstrapping techniques.

Chapter 4

The Timing of the Flight to Gold: An Intra-day Analysis of Gold and the S&P500†

4.1 Introduction

The role of gold as a safe haven asset is well established since the work by Baur and Lucey (2010) and Baur and McDermott (2010). Although the safe haven property of gold has been assessed relative to different asset markets and currencies, there is no study that analyses the safe haven effect of gold at the intra-day level.[1] In other words, it is not clear when and how fast investors react to large negative shocks in the stock market over short time horizons with purchases of gold, that is, a "flight to gold".[2]

This paper aims to fill this gap in the literature. We use high-frequency intra-day returns of the S&P500 and gold spot and futures, sampled at 5-min intervals, and study how the prices of those two assets evolve over the trading day on average and in the presence of extreme losses in the equity market. Our sample spans the 11-year period from 2007 until 2018 covering episodes of sustained uncertainty (e.g. the Global Financial Crisis as well as the (Eurozone) sovereign debt crisis) associated with strongly declining equity prices. Hence, the dataset is well suited for the analysis of the safe haven effect of gold from an intra-day perspective. Since the gold futures market might be more liquid than gold spot prices, we analyze the dynamic relationships with the S&P500 for both types of gold.

The remainder of this paper contains an Empirical Analysis (Section 4.2) which presents our dataset and provides a descriptive analysis of extreme events and an analysis of the average response of gold to large stock market shocks. The paper concludes with a summary of the main findings and their implications (Section 4.3).

†This chapter is joint work with Dirk G. Baur and has been originally published as BAUR, D. G. AND K. KUCK (2019). "The Timing of the Flight to Gold: An Intra-day Analysis of Gold and the S&P500," *Finance Research Letters*. DOI 10.1016/j.frl.2019.05.005

[1]Ranaldo and Söderlind (2010) analyze the safe haven property of major currencies using intra-day data and Batten, Lucey, McGroarty, Peat, and Urquhart (2017) study general statistical patterns of intra-day precious metals returns.

[2]A "flight to gold" implies portfolio re-balancing and a *strong safe haven effect* (negative correlation of gold price changes with equity price changes in extreme market conditions. In contrast, a *weak safe haven effect* (zero correlation) does not automatically imply a "flight to gold" since it is also possible that investors do not buy gold and re-balance their portfolios but simply do not sell their gold holdings when all other assets (asset classes) suffer extreme losses.

4.2 Empirical Analysis

4.2.1 Data

We use intra-daily prices for the S&P500, gold spot and futures sampled at the 5-min frequency, covering the period from January 03, 2007 until July 13, 2018. All time series have been downloaded from Thomson Reuters Tick History (DataScope).[3] Regarding the futures prices, we focus on the front contract, as common in the literature.

Table 4.1 presents the descriptive statistics for 5-min intra-day returns and daily returns for the full sample. The returns are based on the midpoint of the bid-ask quotes. The average intra-day gold spot and futures returns are positive when the full trading day is considered and negative during the business hours of the S&P500 (S&P500 open-to-close). In addition, the standard deviations of gold returns are larger when the S&P500 is traded compared to the full day suggesting that equity market trading is generating gold-relevant information and thus increasing the volatility of gold.

Table 4.1
The table presents descriptive statistics of 5-min intra-day and daily returns for the S&P 500 (.SPX), gold spot (XAU=) and gold futures (GCc1) for the full sample period (January 03, 2007 until July 13, 2018).

	Mean	Std.Dev.	Skew	Kurt	Min	Max	Obs
Intra-day							
S&P 500	$0.212 \cdot 10^{-5}$	0.1069	0.35	38.69	−2.6796	3.9496	223377
Gold	$5.753 \cdot 10^{-5}$	0.0714	−0.16	44.92	−2.9520	2.9324	692447
Gold futures	$4.895 \cdot 10^{-5}$	0.0714	−0.14	41.51	−2.8215	2.3018	688825
Intra-day (S&P 500 open-to-close)							
Gold	$-6.451 \cdot 10^{-5}$	0.0875	−0.19	31.95	−2.9520	2.1864	223554
Gold futures	$-6.116 \cdot 10^{-5}$	0.0886	−0.14	33.41	−2.8215	2.3018	221256
Daily (open-to-close)							
S&P 500	$-0.5748 \cdot 10^{-2}$	0.9919	−0.67	15.00	−8.8461	7.2399	2861
Gold	$1.4483 \cdot 10^{-2}$	1.0603	−0.14	10.32	−7.5400	10.1064	2846
Gold futures	$2.7277 \cdot 10^{-2}$	1.0755	−0.09	10.32	−7.6572	10.2744	2803
Daily (S&P 500 open-to-close)							
S&P 500	$-0.5748 \cdot 10^{-2}$	0.9919	−0.67	15.00	−8.8461	7.2399	2861
Gold	$-0.2075 \cdot 10^{-2}$	0.7682	0.37	21.19	−7.7230	10.1064	2950
Gold futures	$-0.2051 \cdot 10^{-2}$	0.7710	0.34	20.77	−7.8344	10.0147	2950
Daily (close-to-close)							
S&P 500	$2.1119 \cdot 10^{-2}$	1.2148	−0.42	11.64	−9.5741	10.2671	2743
Gold	$1.3943 \cdot 10^{-2}$	1.1192	−0.52	8.31	−9.0599	6.8685	2719
Gold futures	$2.3441 \cdot 10^{-2}$	1.1754	−0.22	10.34	−9.0445	10.7649	2532

Note, the statistics for the intra-day returns are based on 5-min intra-day returns over the complete trading day or during the business hours for the S&P 500 (9:30–16:00 local time).

The average correlations between S&P500 and gold intra-day returns are displayed in Table 4.2 and are all positive. The correlation between the gold spot and gold futures returns is 0.93 and the correlation

[3]The RICs are `.SPX`, `XAU=` and `GCc1`, for S&P500, gold spot and gold futures, respectively.

between the S&P 500 and gold returns is 0.1. Moreover, the correlation estimates between the gold spot and futures appear to be unaffected by different trading times throughout the day.

Table 4.2
Correlations between intra-day returns (full sample).

	S&P 500	Gold	Gold futures
Full day			
S&P 500	1		
Gold	0.0991***	1	
Gold futures	0.0999***	0.9291***	1
S&P 500 open-to-close			
S&P 500	1		
Gold	0.0991***	1	
Gold futures	0.0999***	0.9320***	1
S&P 500 close-to-open (overnight)			
Gold		1	0.926***

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.2.2 Descriptive Analysis

This section compares the average performance of the gold price and the S&P500 over the full sample period with their performance on days when the S&P500 exhibits strong declines, i.e. extreme negative returns. We define "extreme" days by daily open-to-close equity returns that fall below -4.5 percent. Using this threshold, our subset of "extreme" days contains the top 10 days with the most severe open-to-close price declines of the S&P500. In this subset, open-to-close S&P500 returns are only negative and range from -8.85 to -4.56 percent with an average of -6.43 percent.

These extreme days and their returns are displayed in Table 4.3; Detailed descriptive statistics for the returns of the two assets on extreme days are presented in Table 4.4. Interestingly, the standard deviations of the returns are significantly larger on extreme return days compared to the full sample period for both assets, i.e. despite moving in opposite directions, both assets experience significantly larger fluctuations on extreme return days than on average return days. It appears that the increased volatility in the stock market spills over to the gold market. The result is also consistent with a different asymmetric volatility effect of gold compared with stocks as documented in the literature (Baur, 2012).[4]

On extreme days (see Table 4.5), the correlation between the S&P500 and gold spot is positive and marginally negative with gold futures. Moreover the correlation between gold spot and futures returns tends to be affected by different trading times throughout the trading day.

Figure 4.1 depicts the aggregate intra-day returns for two key dates of the global financial crisis (GFC). September 15, 2008 can be considered the start of the GFC with the Lehman bankruptcy and October 15, 2008 marks the largest negative return of the S&P500 (-8.3%) in the sample. The graphs illustrate the negative relationship between the two series once the S&P500 starts to display extreme negative price changes.

[4]Whilst negative shocks increase the volatility of stock returns by more than positive shocks, positive shocks increase the volatility of gold returns by more than negative shocks.

Table 4.3
Top 10 extreme dates in the equity market (sorted descending by S&P 500 loss).

Date	S&P 500	Gold	Gold futures
09-May-11	−8.85	0.32	0.30
09-Oct-08	−8.48	3.25	3.34
15-Oct-08	−7.22	1.41	1.41
07-Oct-08	−6.39	0.04	0.25
29-Sep-08	−6.05	2.98	3.00
19-Nov-08	−6.03	−1.55	−1.60
20-Nov-08	−6.02	0.44	0.41
01-Dec-08	−5.63	−1.64	−1.63
08-Aug-11	−5.08	1.24	1.21
06-Nov-08	−4.56	−2.89	−2.96

Note, the table displays open-to-close returns computed during the business hours for the S&P 500 (9:30–16:00 local time).

Table 4.4
The table presents descriptive statistics of 5-min intra-day and daily returns for the S&P 500 (.SPX), gold spot (XAU=) and gold futures (GCc1) for the top 10 most extreme days in the equity market during the period January 03, 2007 until July 13, 2018.

	Mean	Std.Dev.	Skew	Kurt	Min	Max	Obs
Intra-day							
S&P 500	$-7.0838 \cdot 10^{-2}$	0.3727	−0.17	5.59	−2.1728	1.5540	778
Gold	$0.2297 \cdot 10^{-2}$	0.1672	0.03	9.78	−1.0617	1.2513	2679
Gold futures	$0.3598 \cdot 10^{-2}$	0.1733	−0.12	11.06	−1.4130	1.2388	2756
Intra-day (S&P 500 open-to-close)							
Gold	$0.0926 \cdot 10^{-2}$	0.2141	0.31	7.96	−1.0284	1.2513	741
Gold futures	$0.5196 \cdot 10^{-2}$	0.2233	0.12	8.27	−1.2199	1.2388	786
Daily (open-to-close)							
S&P 500	−6.43	1.3792	−0.60	2.33	−8.8461	−4.5600	10
Gold	1.05	2.0458	−0.06	1.98	−2.4992	3.7160	10
Gold futures	1.05	2.0903	−0.06	2.02	−2.5963	3.7962	10
Daily (S&P 500 open-to-close)							
S&P 500	−6.43	1.3792	−0.60	2.33	−8.8461	−4.5600	10
Gold	0.36	1.9845	−0.08	2.07	−2.8935	3.2504	10
Gold futures	0.37	2.0139	−0.10	2.11	−2.9572	3.3417	10

Note, the statistics for the intra-day returns are based on 5-min intra-day returns over the complete trading day or during the business hours for the S&P 500 (9:30–16:00 local time).

Table 4.5
Correlations between intra-day returns (extreme days).

	S&P 500	Gold	Gold futures
Full day			
S&P 500	1		
Gold	0.0993***	1	
Gold futures	−.0135***	0.8937***	1
S&P 500 open-to-close			
S&P 500	1		
Gold	0.0993***	1	
Gold futures	−.0135***	0.916***	1
S&P 500 close-to-open (overnight)			
Gold		1	0.875***

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 4.2 displays two days when material news was released when the stock market was closed but the gold market was open. On May 06, 2010 US stock markets fell sharply in response to news related to the Eurozone sovereign debt crisis. The other day is June 23, 2016 which marks the date of the Brexit referendum with important geopolitical implications. The graphs illustrate that gold acts as safe haven even before the news arrives in the stock market.

Figure 4.1
The figure illustrates the negative relationship between the equity market and the gold spot on two key dates during the Global Financial Crisis (GFC).

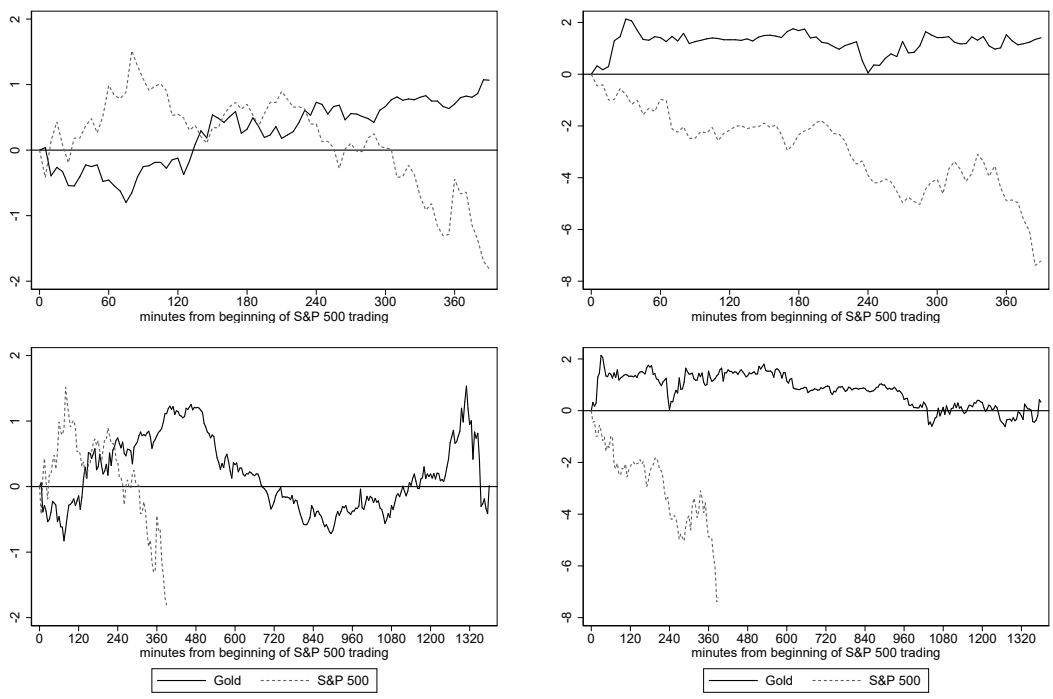

(a) September 15, 2008 **(b)** October 15, 2008

Figure 4.2
The figure illustrates the negative relationship between the equity market and the gold price for two extreme events during the sample period.

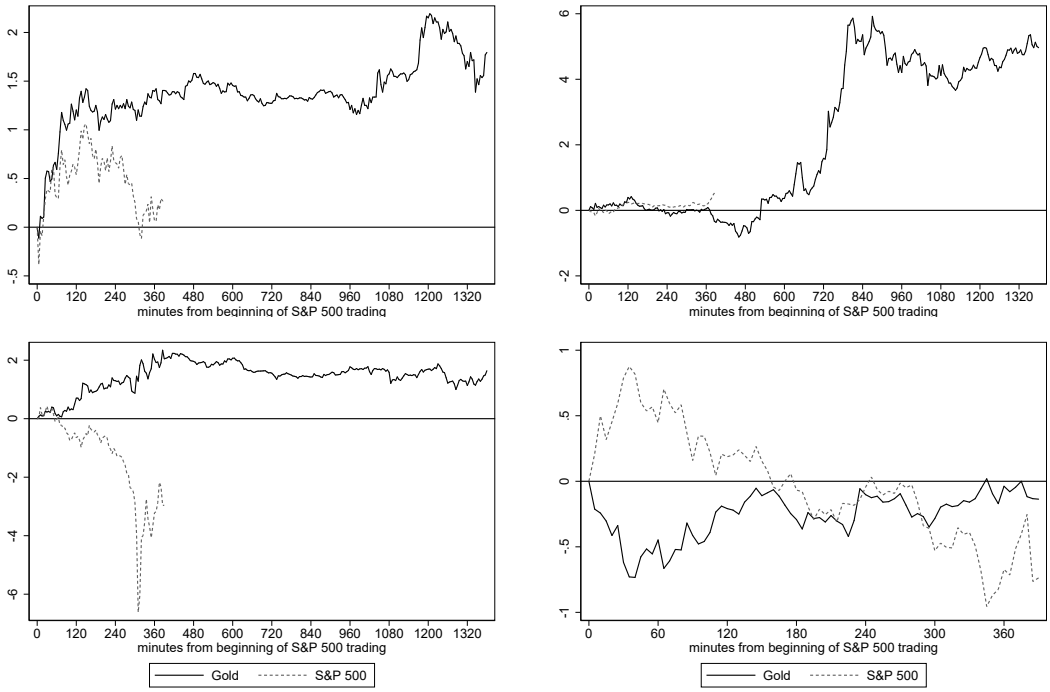

(a) Eurozone sovereign debt shock –
May 05, 2010 (top) and May 6, 2010 (bottom)

(b) Brexit shock –
June 23, 2016 (top) and June 24, 2016 (bottom)

There are also (extreme) days where both the S&P500 and gold decline (e.g. November 19, 2008 or December 1, 2008). This finding may be interpreted as a contradiction to the safe haven property of gold documented in the literature (see e.g. Erb and Harvey, 2013). However, the safe haven property of an asset must be assessed in a broader and dynamic context of the financial market conditions, i.e. extreme days must be seen in the sequence of events and should not be analyzed in isolation. Gold fulfilled its safe haven role at the beginning of the GFC but only during the initial phase (see e.g. Baur and McDermott, 2016). Extreme return days in late October, November and December 2008 were potentially associated with the sale of gold to limit losses incurred in the stock market. Therefore, the safe-haven property is short-lived (Baur and Lucey, 2010) and negative gold returns during a financial crisis do not automatically and necessarily indicate that gold lost its safe haven property.

4.2.3 Average Cumulative Returns

This section analyzes the average evolution of S&P500 and gold prices over all days and over a subset of extreme negative returns of the S&P500. Therefore, we first construct the aggregate intra-day returns for each 5-minute interval $m = 1, 2, \ldots, M$ from the beginning of S&P500 trading on each day (14:30h/13:30h UTC, 09:30h local) until its opening the next day (14:25h/13:25h UTC, 09:25h local).

Then, for each of these intervals, we compute the average over all days and, alternatively, over a subset of extreme days:

$$\bar{r}_m = \frac{1}{T}\sum_{t=1}^{T} r_{m,t},$$

where $r_{m,t} = \sum_{j=1}^{m} r_{j,t}$ denotes the cumulative intra-day return from the S&P500 market opening to the interval m on day t.

The average cumulative intra-day returns of the S&P 500 and gold futures for the full sample are depicted in Figure 4.3. Figure 4.4 presents a similar graph but extends the trading hours until the opening of trading in the equity market the day (open-to-open). The graphs illustrate the low average correlation of the two assets and their volatility.

Figure 4.3
The graph shows the average cumulative 5-min returns for the S&P500 and gold futures for the full sample period (S&P open-to-close).

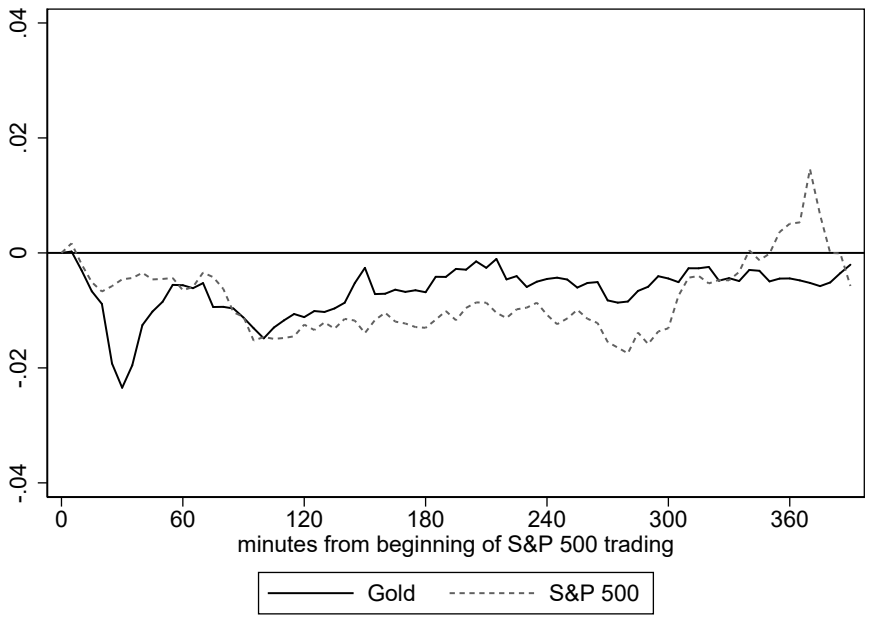

Figures 4.5 and 4.6 present the average intra-day returns for the subset of extreme days in the stock market and show a relatively stable and positive evolution of gold returns contrasted by the continuously decreasing cumulative S&P500 returns. The relative stability of gold prices on average confirms the safe haven property of gold and clearly illustrates that gold does (on average) not co-move with the stock market in extreme conditions. The Figures also show that cumulative average gold returns are not consistently positive over all days of extreme negative stock market returns but turn and remain positive if stock market returns fall by more than 3% within the first 4 hours of stock trading. Figure 4.6 also reveals that the gold price continues to rise during the overnight period after the end of S&P500 trading on extreme days. The cumulative gold returns are clearly positive on such days but the average cumulative gold return is only about 1%. The reason for the rather small increase of gold prices relative to the fall in equity valuations could be that the buying pressure on gold due to a flight from stocks to gold is compensated by investors who sell gold to cover losses incurred in the stock market. Another

Figure 4.4
The graph shows the average cumulative 5-min returns for the S&P500 and gold futures for the full sample period (S&P open-to-open).

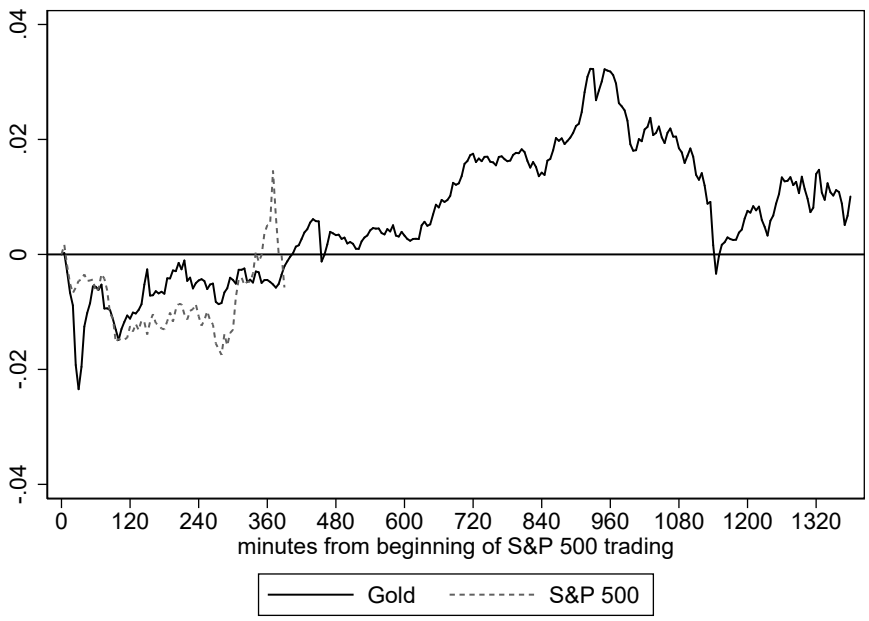

reason may be that most investors already hold gold and do not buy additional gold in reaction to extreme negative stock market returns.

Finally, Figure 4.7 illustrates how the correlation changes throughout the trading day and turns negative on extreme days towards the end of the trading day. The graph presents the correlation estimates based on forward-rolling 30-min gold and S&P500 returns across all (extreme) days.[5]

4.2.4 Econometric Analysis

To investigate the cross-market dynamics of gold and equity in adverse conditions in the equity market, we estimate a safe haven regression model to test if large negative shocks of the S&P500 are associated with immediate[6] (positive) reactions of the gold price. Specifically, to account for conditional heteroscedasticity in the time series of intra-day gold and equity returns we estimate the following asymmetric Glosten, Jagannathan, and Runkle (1993) GARCH model:

$$r_{\text{gold},t} = a + b_t r_{\text{S\&P500},t} + \psi r_{\text{S\&P500},t-1} + \phi r_{\text{gold},t-1} + e_t \tag{4.1}$$

$$\begin{aligned} b_t = {} & c_0 + c_1 D(r_{\text{S\&P500},t} < -.30) + c_2 D(r_{\text{S\&P500},t} < -.50) + \\ & c_3 D(r_{\text{S\&P500},t} < -.75) + c_4 D(r_{\text{S\&P500},t} < -1.00) \end{aligned} \tag{4.2}$$

$$h_t = \omega + \alpha e_{t-1}^2 + \gamma D\left(e_{t-1} > 0\right) e_{t-1}^2 + \beta h_{t-1} \tag{4.3}$$

$$e_t \sim N(0, h_t), \tag{4.4}$$

[5]Graphs for higher frequency (5min, 10min) return correlations can be obtained from the authors.
[6]We interpret a reaction within 5 minutes as "immediate".

Figure 4.5
The figure shows the average cumulative 5-min returns for the S&P500 and gold futures on "extreme" days (S&P open-to-close).

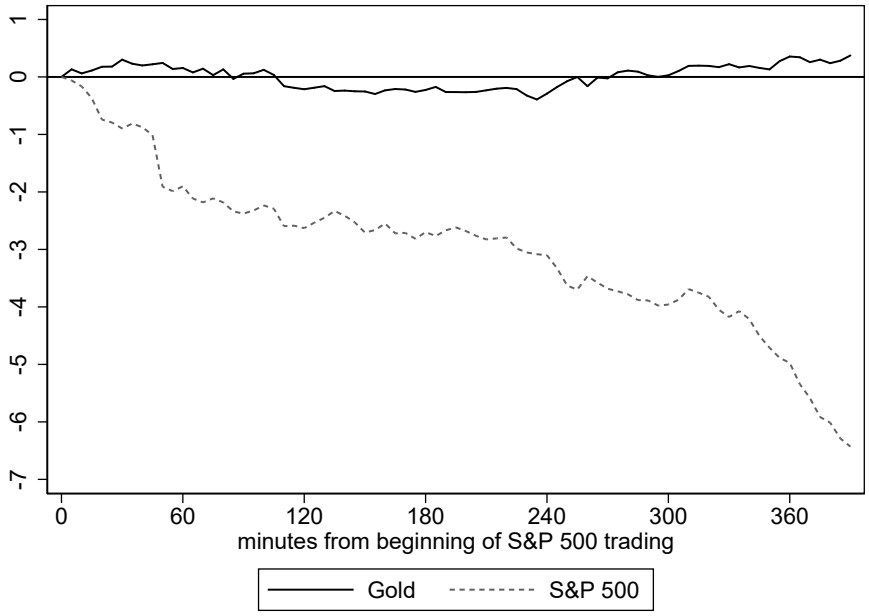

where $D(\cdot)$ is an indicator function that takes on the value of 1 if its argument is true, and 0 otherwise.[7] This model is similar to the one proposed in Baur and Lucey (2010) for daily data but uses intra-day 5-minute returns and enables us to study the presence of short-term reactions of gold to negative stock market *shocks* of different magnitudes.

Table 4.6 reports the estimation results and shows that only extreme negative S&P500 returns that materialize over a short intra-day time horizon exhibit negative and statistically significant coefficients implying a positive reaction of gold. More specifically, in the rare case of a negative shock exceeding -1%, gold acts as a weak safe haven over short intra-day time horizon(s). That is, in the 58 cases of negative 5-min returns smaller than -1% in our sample, gold and the S&P500 appear contemporaneously uncorrelated. Less pronounced negative S&P500 returns, by contrast, do not imply a positive reaction of gold. The fact that gold only positively reacts to extreme negative shocks is fully consistent with the notion of a safe haven asset. If gold returns were positive at all times gold would be a perfect "safe asset" in the sense of Gorton (2017) but not a typical "safe haven" (Baur and McDermott, 2016).

The additional quantile regression estimates reported in Table 4.7 provide an alternative to the GARCH-framework and a robustness analysis. The results confirm the findings reported above and also show that there are significant differences between the lower quantile and the upper quantile estimates. The strongest safe haven effects are found for the upper quantiles when gold exhibits large positive returns, the results are weaker for lower quantiles when gold exhibits negative returns. Another interesting observation is the increased persistence of gold returns in the extreme tails of the distribution. The AR(1) estimates are positive and highly significant in both the 1% and the 99% quantiles compared to a negative AR(1) coefficient at the 50% quantile and thus in the center of the distribution.

[7]Note that this specification of the Glosten et al. (1993)-GARCH model is specific to STATA. A negative and significant coefficient estimate for γ implies that positive shocks increase the variance by less than negative shocks.

Table 4.6
Safe haven regression model based on 5-min intra-day gold spot and futures returns (GJR-GARCH).

	Gold		Gold		Gold futures		Gold futures	
Mean equation								
Constant	0.0001	(0.91)	0.0001	(1.04)	0.0001	(0.82)	0.0001	(0.76)
S&P 500 return$_t$	0.0457***	(59.43)	0.0407***	(30.11)	0.0498***	(62.18)	0.0424***	(31.78)
S&P 500 return$_t < -.30$	0.0394***	(16.91)	0.0029	(0.64)	0.0289***	(12.74)	0.0053	(1.17)
S&P 500 return$_t < -.50$	−0.0119***	(−3.37)	0.0227**	(2.81)	−0.0009	(−0.23)	0.0201*	(2.54)
S&P 500 return$_t < -.75$	0.0006	(0.09)	0.0310*	(2.12)	0.0355***	(6.73)	0.0381**	(2.83)
S&P 500 return$_t < -1.0$	−0.1350***	(−19.74)	−0.0899***	(−5.33)	−0.1680***	(−29.98)	−0.0869***	(−5.10)
AR(1)	−0.0285***	(−26.63)	−0.0696***	(−34.31)	−0.0157***	(−13.25)	−0.0599***	(−29.41)
S&P 500 return$_{t-1}$	0.0049***	(4.85)	0.0066***	(5.11)	0.0037***	(3.53)	0.0060***	(4.72)
Variance equation								
ω	0.0000***	(75.98)	0.0000***	(23.63)	0.0000***	(76.48)	0.0000***	(23.05)
α	0.1300***	(186.97)	0.1240***	(58.22)	0.1280***	(185.31)	0.1250***	(58.22)
β	0.8660***	(2078.60)	0.8730***	(645.89)	0.8660***	(2064.48)	0.8730***	(655.03)
γ	−0.0065***	(−7.34)	−0.0154***	(−6.09)	−0.0060***	(−6.73)	−0.0162***	(−6.42)
ν			4.274				4.2075	
Observations	218,048		218,048		216,053		216,053	

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: ν denotes the degrees of freedom of the t-distribution

Table 4.7
Quantile safe haven regression based on 5-min returns for different quantiles.

	$\tau = .01$		$\tau = .05$		$\tau = .10$		$\tau = .50$	
Constant	-0.2450^{***}	(-126.83)	-0.1190^{***}	(-189.56)	-0.0774^{***}	(-211.43)	0.0001	(1.19)
S&P 500 return$_t$	-0.0126	(-0.61)	0.0297^{***}	(4.39)	0.0420^{***}	(10.63)	0.0573^{***}	(45.43)
S&P 500 return$_t < -.30$	1.1040^{***}	(17.70)	0.6160^{***}	(30.49)	0.4449^{***}	(37.52)	0.0118^{**}	(3.12)
S&P 500 return$_t < -.50$	-0.3830^{***}	(-4.09)	-0.0388	(-1.28)	-0.0508^{**}	(-2.86)	0.0209^{***}	(3.69)
S&P 500 return$_t < -.75$	-0.0857	(-0.65)	-0.0846^{*}	(-1.98)	-0.0278	(-1.11)	-0.0128	(-1.60)
S&P 500 return$_t < -1.00$	-0.3510^{*}	(-2.45)	-0.2970^{***}	(-6.38)	-0.2290^{***}	(-8.40)	-0.0816^{***}	(-9.39)
AR(1)	0.0764^{***}	(3.46)	0.0036	(0.50)	-0.0278^{***}	(-6.63)	-0.0722^{***}	(-53.97)
S&P 500 return$_{t-1}$	0.0106	(0.58)	0.0115	(1.95)	0.0109^{**}	(3.15)	0.0054^{***}	(4.88)
Observations	218,048		218,048		218,048		218,048	

	$\tau = .90$		$\tau = .95$		$\tau = .99$	
Constant	0.00305	(0.91)	0.0043	(0.77)	0.0180	(1.00)
S&P 500 return$_t$	0.0786^{***}	(222.33)	0.1180^{***}	(198.95)	0.2340^{***}	(123.39)
S&P 500 return$_t < -.30$	0.143^{***}	(37.49)	0.1790^{***}	(27.81)	0.2420^{***}	(11.83)
S&P 500 return$_t < -.50$	-0.429^{***}	(-37.59)	-0.5650^{***}	(-29.36)	-0.8000^{***}	(-13.08)
S&P 500 return$_t < -.75$	0.0657^{***}	(3.83)	0.0027	(0.09)	-0.0821	(-0.89)
S&P 500 return$_t < -1.00$	-0.0800^{***}	(-3.31)	-0.0038	(-0.09)	-0.1600	(-1.23)
AR(1)	0.189^{***}	(7.18)	0.0527	(1.19)	0.4400^{**}	(3.12)
S&P 500 return$_{t-1}$	-0.0493^{***}	(-12.19)	-0.0318^{***}	(-4.67)	-0.0040	(-0.19)
Observations	218,048		218,048		218,048	

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 4.6
The graph shows the average cumulative 5-min returns for the S&P500 and gold futures on "extreme" days (S&P open-to-open).

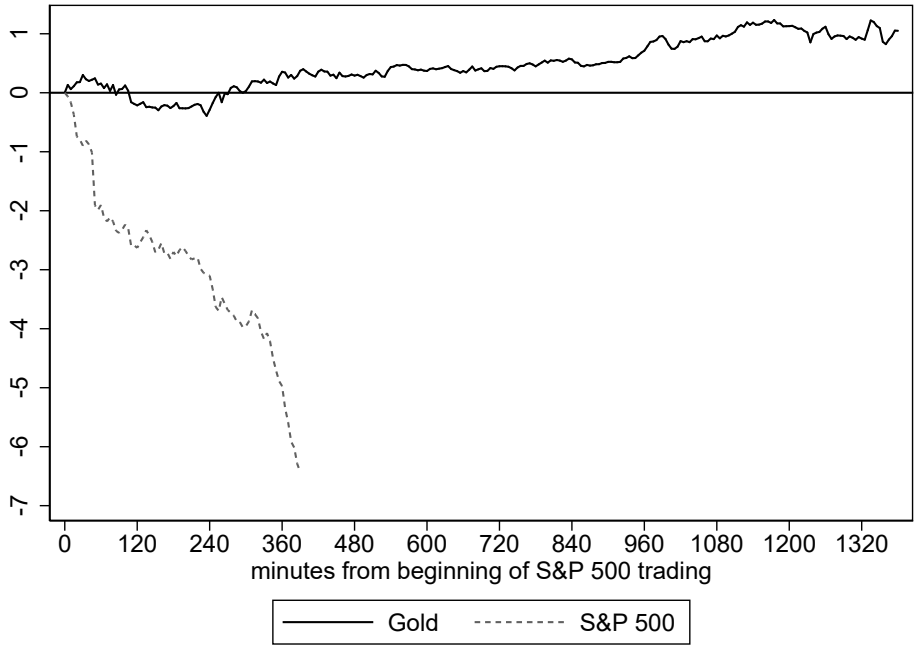

Figure 4.7
The figure shows the correlation between 30-min intra-day gold and S&P500 returns in the full sample and for "extreme" days.

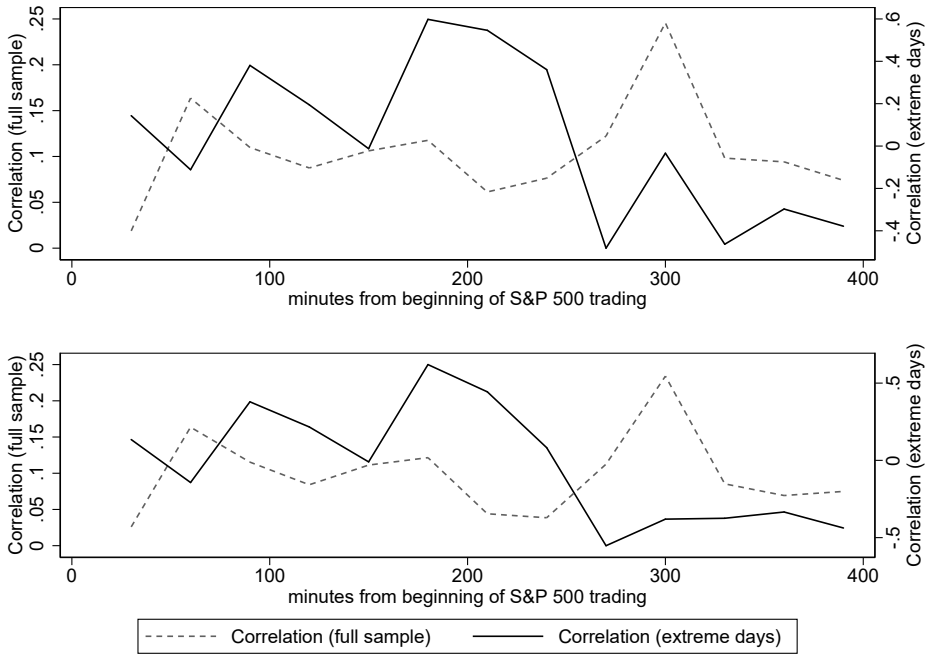

Notes: Upper panel: Gold spot and S&P500 returns. Lower panel: Gold futures and S&P500 returns.

4.3 Summary and Concluding Remarks

This paper extended existing research on the safe haven property of gold by demonstrating that the safe haven property of gold can also be identified at short 5-min time intervals. The use of intra-day data instead of lower frequency daily data allowed us to better understand the timing of extreme events, flight to quality and the safe haven effect. We find that gold returns react fast to extreme negative changes in the S&P500 but that extreme returns over an entire trading day evolve relatively slowly. The fact that extreme stock market movements do not occur over a 5-min period but generally extend over several hours also means that the full extent of flight to quality builds over several hours. The analysis in this paper points to the importance of using high-frequency intra-day data to better understand the dynamics of crashes and the role of gold.

Chapter 5

Gold volatility and the safe haven effect

5.1 Introduction

We aim to study the relationship of gold returns and volatility with equity returns. Dependencies between these two assets attract special attention, since gold has been found to act as 'safe haven' for equity. More specifically, empirical evidence suggests that gold holds its value or exhibits positive returns when equity prices decline severely (Baur and McDermott, 2010; Baur and Kuck, 2019). Naturally, this state-dependent correlation between the returns of the two assets is of high interest for investors with respect to portfolio diversification and risk management since it suggests that gold may have the ability to limit losses due to negative shocks in the equity market. However, apart from the return relationship, the effectiveness of gold as a safe haven depends on its volatility behaviour in presence of strongly declining stock prices. Specifically, for a portfolio composed of the two assets, the positive effect of negative or no correlation between gold and stock returns may be compromised by an increase of the gold volatility (Baur, 2012). In particular, the risk of a portfolio may increase if the volatility of gold is higher when the equity market is in distress and gold and stocks are conditionally uncorrelated (or positively correlated).[1] A comprehensive analysis of the safe haven effect, therefore, should consider the volatility of gold as well.

Although the relationship between gold and equity returns, as well as gold and other financial returns, has been investigated extensively in the literature (see e.g. Baur and McDermott, 2010; Reboredo, 2013; Baur and McDermott, 2016; Miyazaki, 2019), the gold volatility dynamics and, specifically their behaviour conditionally on extremely negative equity returns, have attracted much less attention in the literature. So far, only Baur (2012) investigates the influence of changes in stock prices on the gold volatility. Using asymmetric GARCH-X models, he documents an inverted asymmetric gold volatility response to past shocks in gold returns, which means that past positive shocks increase the volatility by more than negative shocks of the same magnitude. However, he cannot find evidence for a response in gold volatility to (past) stock returns. More recently, Todorova (2017) studies the presence of asymmetric reactions of realized volatility of gold with respect to past positive and negative gold return semi-variance. She finds that past positive long-term semi-variance has relatively more explanatory power

[1] Baur (2012) simulates the effect of increasing volatility of gold in a stock-gold portfolio for various combinations of weights and asset correlations. The simulation results suggest that in presence of negative correlation between stock and gold, the portfolio variance may even decrease despite higher gold volatility.

for future volatility. Her study, however, does not incorporate equity returns and, hence, cannot present insights regarding the gold and equity return volatility relationship.

This paper aims to contribute to the existing literature and provides a comprehensive analysis of the relationship between both gold returns and volatilities and the equity market. We base our analysis on synchronized returns and realized measures for gold and the S&P500 which we obtain from intra-day data. More specifically, we measure the returns of gold over the stock trading hours and compare them to open-to-close returns of equity. In contrast to studies based on conventional daily returns which are typically measured over different time intervals (e.g. AM fixing returns for gold and close-to-close returns of equity), this allows for a more accurate analysis of the contemporaneous relationship between gold and equity. Further, whilst the gold futures market might be more liquid than the spot market, futures only represent a claim for the delivery of the underlying asset at a later time. However, (immediate) tangibility might be an important aspect of the safe haven effect. Therefore, our study consider both, spot and futures market. The sample spans an 11-year period from January 2007 to July 2018 and covers episodes of sustained uncertainty such as the Global Financial Crisis, associated with strongly declining equity prices. Therefore, our dataset is well suited for a comprehensive analysis of the gold-stock return and volatility relationship under different market conditions.

Naturally, the idea of a safe haven asset involves non-linear (temporal) dependencies between two assets, which might not be adequately described in the context of linear models. Therefore, we use quantile regression techniques which allows us to provide a precise description of the contemporaneous dependence structure between gold and equity in both returns and volatility. Quantile regression is widely used in financial econometrics (see e.g., Koenker and Zhao, 1996; Chernozhukov and Umantsev, 2001; Baur, Dimpfl, and Jung, 2012; Žikeš and Baruník, 2016), and has recently been applied by Miyazaki (2019) and Liu (2018) to study the safe haven property of gold and other financial assets. However, employing quantile regression models for financial returns to describe their variance dynamics is an innovative approach, introduced only recently by Baur and Dimpfl (2018). To the best of our knowledge, this approach has not been used to investigate the relationship between gold return variance and stock returns. In addition to quantile regression models for gold returns, we apply the heterogeneous autoregressive model of realized volatility (HAR-RV) by Corsi (2009) to model the volatility dynamics based on realized measures. More specifically, to account for the influence of negative shocks in the equity market, we include additional regressors in the HAR-RV model (HAR-RV-X).

We find that gold returns are contemporaneously uncorrelated with extreme negative equity returns on average. This supports the idea that gold acts as a weak safe haven for equity in the sense of Baur and McDermott (2010). However, and in contrast to the results of Baur (2012), we also find that the volatility of gold is substantially higher on days with negative shocks in equity prices, reflecting increased uncertainty in the gold market. Consequently, the risk of a portfolio composed of the two assets may be affected due to the relationship between the gold volatility and extremely negative stock returns.

The remainder of the paper is structured as follows: Section 5.2 presents the dataset and describes our econometric framework. In Section Section 5.3 we present and discuss our empirical results. Section 5.4 provides a summary of the key findings.

5.2 Empirical framework

5.2.1 Data

We base our analysis on contemporaneous returns and volatilities of gold and the S&P 500. Further, regarding gold, our studies considers both the gold futures and spot market.[2] Whilst the gold futures market might be more liquid than the spot market, futures only represent a claim for the delivery of gold at a later time. However, in presence of financial uncertainty the (immediate) tangibility of gold might be an important aspect of the safe haven effect. Therefore, investigating both types gold may provide further insights regarding the safe haven effect of gold.

From an initial data set of 5-minute intra-day bid/ask quotes, we obtain open-to-close returns of S&P 500 and returns of gold spot and futures measured over the S&P500 trading hours (S&P500 open-to-close) and the full trading day.[3] Further, we use the 5 minute intra-day returns to construct various realized measures of gold. Specifically, we compute separately realized variances of gold during the S&P500 trading hours and for the overnight period until the beginning of stock trading on the following day. Using synchronized returns and realized variances allows us to analyze the contemporaneous behaviour of returns and volatility under different market conditions and is a unique feature of our study.

All time series have been downloaded from Thomson Reuters TickHistory (DataScope).[4] The final dataset of daily returns and realized measures covers the period from January 03, 2007 until July 13, 2018.

Table 5.1 provides descriptive statistics of S&P 500 returns as well as returns and realized measures for gold spot and futures. The mean returns of equity and gold spot and futures returns are close to zero but negative during the S&P500 trading hours. Further, equity returns have smaller minimum return (-8.85% versus -7.8) whereas gold exhibits higher positive returns (10% versus 7%). Their variability, when measured by their standard deviations is comparable, although slightly higher for S&P500. When the full trading day is considered, average returns are positive but close to zero.[5] The realized variances of gold measured over the trading hours of S&P500 and the full day are characterized by the substantial excess kurtosis and autocorrelation. Further, logarithmic realized variances are found to be much closer to a normal distribution. Interestingly, for the realized variances captured during the overnight period, autocorrelation appears to be absent.

5.2.2 Quantile regression models

Conventional linear (auto)regressive models based on ordinary least squares (OLS) focus on the effects of exogenous or lagged own variable(s) on the conditional mean of the dependent variable. Quantile (auto)regression according to Koenker and Bassett (1978) and Koenker and Xiao (2006) instead allows to estimate the dependence of specific quantiles of the dependent variable conditional on some information set, $\mathscr{F}_t$. Specifically, for each (conditional) quantile, a linear model is specified and estimated separately. Compared to OLS, quantile regression enhances the research possibilities by allowing for an accurate

[2] Regarding the futures prices, we focus on the front contract, as common in the literature.

[3] Whilst trading of the S&P 500 takes place only between 14:30h/13:30h and 21:00h/20:00h UTC (09:30h to 16:00 local), gold is traded over the full 24 hours of a day.

[4] The RICs are `.SPX`, `XAU=` and `GCc1`, for S&P500, gold spot and gold futures, respectively.

[5] In case of S&P500 close-to-close returns are considered as full day returns whilst for gold the period 09:00h to 09:25h on the following day (local) is considered.

Table 5.1
Descriptive statistics.

	Mean	SD	Skew	Kurt	Min	Max	$Q_r(1)$	$Q_{\hat{\varepsilon}}(1)$	$Q_{\hat{\varepsilon}^2}(1)$	N
Panel (a): Returns										
S&P 500 trading hours (open-to-close)										
S&P 500	−0.0057	0.9919	−0.6740	15.0047	−8.8461	7.2399	0.0000	0.7399	0.0000	2861
Gold	−0.0049	0.7739	0.3303	20.9117	−7.7230	10.1064	0.0740	0.9482	0.0000	2860
Gold futures	−0.0043	0.7774	0.2953	20.4800	−7.8344	10.0147	0.1219	0.9476	0.0000	2857
Full day (close-to-close)										
S&P 500	0.0208	1.2108	−0.4285	11.7204	−9.5741	10.2671	0.0000	0.9010	0.0000	2740
Gold	0.0123	1.1201	−0.5240	8.3280	−9.0599	6.8685	0.1407	0.9964	0.0000	2693
Gold futures	0.0219	1.1751	−0.2200	10.4157	−9.0445	10.7649	0.7571	0.9772	0.0000	2508
Panel (b): Realized Variance										
S&P 500 trading hours (open-to-close)										
Gold	0.5727	0.9183	8.6471	135.6862	0.0241	21.1788	0.0000	—	—	2857
log(Gold)	−1.0581	0.9337	0.4024	3.3128	−3.7250	3.0530	0.0000	—	—	2857
Gold futures	0.5822	1.0375	9.8384	153.3029	0.0265	21.9448	0.0000	—	—	2841
log(Gold futures)	−1.0672	0.9449	0.4453	3.4294	−3.6300	3.0885	0.0000	—	—	2841
Overnight (close-to-open)										
Gold	0.6086	0.9479	5.8604	56.4398	0.0000	13.9446	0.0000	—	—	2857
log(Gold)	−1.4493	1.8219	−1.3864	5.2881	−11.3587	2.6351	0.1335	—	—	2857
Gold futures	0.6065	0.9608	6.0791	60.2998	0.0003	14.8873	0.0000	—	—	2842
log(Gold futures)	−1.3744	1.6315	−1.0284	3.7397	−8.1165	2.7005	0.9303	—	—	2842
Full day (close-to-close)										
Gold	1.1813	1.6715	6.2943	67.1902	0.0250	27.9763	0.0000	—	—	2857
log(Gold)	−0.2890	0.9094	0.2171	3.5867	−3.6902	3.3314	0.0000	—	—	2857
Gold futures	1.1885	1.7924	7.2653	86.9296	0.0345	31.2215	0.0000	—	—	2842
log(Gold futures)	−0.2910	0.9041	0.3018	3.6593	−3.3674	3.4411	0.0000	—	—	2842

Notes: The table provides descriptive statistics of S&P500 returns and returns and realized variances of gold spot and futures. Means, Min and Max are in percentage terms for returns. The last three columns hold p-values for Ljung–Box tests for autocorrelation up to the lag indicated in parentheses. The null hypothesis is independent distribution. $Q_r(1)$ tests the return or realized variance time series and $Q_{\hat{\varepsilon}}(1)$ the residuals from a AR(1) model. $Q_{\hat{\varepsilon}^2}(1)$ uses the squared residuals of an AR(1) model, that is, testing for heteroscedasticity in the returns.

description of the complete distribution, conditional on $\mathscr{F}_t$. Furthermore, it is robust to conditional heteroskedasticity, skewness and leptokurtosis which are common features of financial return time series.

The linear quantile (auto)regression (LQ(A)R) model reads as follows:

$$Q_{y_t}(\tau|\mathscr{F}_t) = \alpha(\tau) + \beta(\tau)\mathbf{x}_t, \tag{5.1}$$

where $Q_{y_t}(\tau|\mathscr{F}_t)$ is the τth quantile of y_t conditional on $\mathbf{x}_t$, containing, for example, contemporaneous returns of other assets, or lags of the dependent variable.

Using quantile regression to investigate the dependence of gold returns and their variability on the state of the stock market is inspired by Baur (2013) and the methodology presented in Baur and Dimpfl (2018). First, Baur (2013) demonstrates that quantile regression, by estimating separate models for each quantile of interest, proves useful to study potentially non-linear dependencies among financial returns. Second, Baur and Dimpfl (2018) propose to exploit this feature of the quantile regression model to estimate the response in the return variability conditionally on lags of the dependent variable. The basic idea is that the estimated (extreme) lower and upper quantile regression lines span a range of possible future return realizations which capture their variability.

In the following, we build on these ideas to study the link between stock returns and the returns of gold and their variability. We start our analysis by estimating a safe haven regression using quantile regression techniques and focus on the effect of extreme negative shocks in the stock market, defined as equity returns smaller than the unconditional 1% percentile of the S&P500 return distribution:

$$\begin{aligned} Q_{r_t^{\text{Gold}}}(\tau|\mathscr{F}_t) &= \alpha(\tau) + \beta_t(\tau) r_t^{\text{S\&P500}} \\ \beta_t(\tau) &= \delta_1(\tau)\cdot \mathbb{1}\left(r_t^{S\&P500} \geq 0\right) + \delta_2(\tau)\cdot \mathbb{1}\left(r_t^{S\&P500} < 0\right) + \\ &\quad \gamma(\tau)\cdot \mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right), \end{aligned} \tag{5.2}$$

where $Q_{r_t^{Gold}}(\tau|\mathscr{F}_t)$ is the τth quantile of the gold return conditional on the information set, $\mathscr{F}_t$, containing contemporaneous S&P500 returns. $\mathbb{1}(\cdot)$ is a Heaviside indicator function equal to one if the argument in parentheses is true, and zero otherwise; and $q_{1\%}^{S\&P500}$ denotes the 1% quantile of the unconditional S&P500 return distribution. Disentangling stock returns into positive and (moderate) negative returns enables us to study asymmetric relations between the two assets with respect to the sign of stock returns. The two parameters $\{\delta_1(\tau), \delta_2(\tau)\}$ capture the strength and direction of positive and moderate negative stock returns on the τth conditional gold return quantile. Further, $\gamma(\tau)$ captures how the τth conditional quantile is affected by an extreme negative S&P500 return. Therefore, the relationship between gold and equity in presence of extremely negative equity returns is given by $\delta_2(\tau) + \gamma(\tau)$. The sequence of quantile-specific regression coefficients over the range of quantiles $\tau \in [.01, .02, \ldots, .99]$ characterize the *structure of the dependence* in the spirit of Baur (2013).

Additionally, we use Equation 5.2 to study how stock returns, and particularly shocks in the equity market affect the gold return variability. We follow the ideas of Baur and Dimpfl (2018) and use the range of possible dependent variable realizations spanned by the quantile regression lines of opposing extreme quantiles, $\tau^\star = .02$ and $(1-\tau^\star) = .98$, to investigate the gold return variability conditional on contemporaneous returns in the stock market.[6] More specifically, for a simple quantile (auto)regression

[6] Baur and Dimpfl (2018) use a first order quantile autoregressive (QAR(1)) model for financial returns to investigate the presence of asymmetric volatility responses to positive and negative past returns.

model with only one explanatory variable, the relationship between x_t and the inter-quantile range of y_t is described by:

$$|[\alpha(1-\tau^\star)+\beta(1-\tau^\star)\cdot x_t)]|-|[\alpha(\tau^\star)+\beta(\tau^\star)\cdot x_t]|. \tag{5.3}$$

From Equation 5.3 follows that the (conditional) range of the dependent variable is determined by the location, $\alpha(\tau^\star)$ and $\alpha(1-\tau^\star)$, and the slopes of the quantile regression lines, $\beta(\tau^\star)$ and $\beta(1-\tau^\star)$. In addition, the sign and relative steepness in the slopes of the quantile regression lines determine if and to what extent the volatility response is asymmetric with respect to positive and negative realizations of x_t. The implications of the signs of $\beta(\tau^\star)$ and $\beta(1-\tau^\star)$ and their relative absolute size for the volatility asymmetry is summarized in Table 5.2. For instance, if $\beta(\tau^\star)$ and $\beta(1-\tau^\star)$ have opposing signs, the volatility is (inverted) asymmetric if the slope is positive (negative) for the $Q_{y_t}(\tau^\star|\mathscr{F}_t)$-quantile regression line and negative (positive) for the $Q_{y_t}(1-\tau^\star|\mathscr{F}_t)$-regression line.[7] This is depicted by the black solid lines in Panel (a) of Figure 5.1, whilst the dashed lines illustrate inverted asymmetry. Furthermore, if the two slope coefficients $\beta(\tau^\star)$ and $\beta(1-\tau^\star)$ have the same sign, the potential volatility asymmetry is determined by their relative (absolute) size. For instance, if the slopes for $\tau^\star$ and $(1-\tau^\star)$ are negative and the slope coefficient in the lower quantile exceeds that in the upper quantile in absolute terms, i.e. $|\beta(\tau^\star)|>|\beta(1-\tau^\star)|$, the volatility is inverted asymmetric in the sense that positive realizations of x_t induce a larger volatility response than negative realizations. This situation is illustrated in Panel (b) of Figure 5.1. The inter-quantile range is larger for positive realizations of x_t than for negative realizations of the same magnitude. Only when the upper and lower quantile regression lines are perfectly parallel, conditional variability is symmetric and not influenced by the sign of x_t.

The quantile regression models are estimated using the standard optimization routine presented in Koenker and Bassett (1978), and implemented, for example, in R. Asymptotic standard errors are estimated using bootstrap methods. As we use covariance stationary time series data, we perform a block bootstrap with average block length equal to 36 and 2000 replications (see, for example, Chernick, 2008, chap. 5).

Table 5.2
Coefficient signs and relative magnitude and the form asymmetric volatility response.

positive signs: $\theta(\tau)>0,\ \theta(1-\tau)>0$		*negative signs:* $\theta(\tau)<0, \theta(1-\tau)<0$	
$\lvert\theta(\tau)\rvert<\lvert\theta(1-\tau)\rvert$:	inverted asymmetry	$\lvert\theta(\tau)\rvert<\lvert\theta(1-\tau)\rvert$:	asymmetry
$\lvert\theta(\tau)\rvert>\lvert\theta(1-\tau)\rvert$:	asymmetry	$\lvert\theta(\tau)\rvert>\lvert\theta(1-\tau)\rvert$:	inverted asymmetry
opposing signs:			
$\theta(\tau)>0,\ \theta(1-\tau)<0$:	asymmetry		
$\theta(\tau)<0,\ \theta(1-\tau)>0$:	inverted asymmetry		

Notes: This table summarizes how the signs and relative slopes of upper- and lower-tail regression lines affect the volatility response.

[7] By *inverted* volatility asymmetry we mean that positive returns increase the volatility by more than negative returns of the same magnitude.

Figure 5.1
Quantile regression and asymmetric volatility.

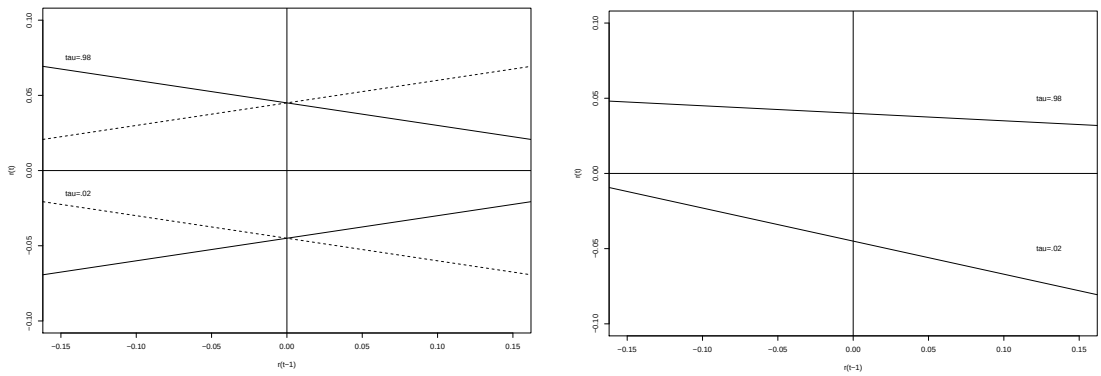

(a) Asymmetric and inverted asymmetric volatility for opposing signs between upper- and lower-tail regression lines

(b) Inverted asymmetric volatility for equal signs between upper- and lower-tail regression lines

Notes: Panel (a) shows (inverted) asymmetric volatility described by the inter-quantile range for opposing signs between the upper-tail and lower-tail regression lines. The black solid lines depict asymmetric volatility, whilst the dashed lines illustrate inverted asymmetry. Panel (b) depicts inverted asymmetric volatility for equal signs between upper- and lower-tail regression lines.

5.2.3 Heterogenous autoregressive models for realized volatility (HAR-RV)

In addition to an analysis of the inter-quantile range-based price variability, captured from contemporaneous gold returns over the S&P500 trading hours using quantile regression, we study the behaviour of gold volatility based on various realized measures, estimated from intra-day data. The motivation is that realized measures are more precise proxies for the latent volatility in the sense that any intra-daily price variation, i.e. the within day volatility, is taken into account. In contrast, volatility measures based on daily returns or high-low price ranges, by construction, ignore this information, since they are based on the magnitude of price changes at only two different time points. The within trading day price variation, however, reflects the uncertainty in the gold market which seems of particular interest in the context of our analysis as well.

To measure the daily quadratic variation, we build on Andersen and Bollerslev (1998), who originally proposed to estimate the daily realized volatility by the sum of squared intra-day returns, sampled at ultra-high frequencies. The corresponding RV estimator is defined as:

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2, \tag{5.4}$$

where $r_{t,j}$ denotes the jth intra-day return on day t, and M is the number of intra-day returns. Whilst the theory suggests to sample at the finest time intervals possible, in reality, price observations at ultra-high frequencies are typically contaminated by microstructure noise. This renders the simple realized variance estimator in Equation 5.4 biased and inconsistent. Therefore, we resort to the moving-average based estimator of Hansen, Large, and Lunde (2008) to obtain realized measures of gold. Another nice

feature of realized measures is that they can be treated as observed realizations of the latent volatility and, therefore, can be directly used in conventional econometric models.

Given the time series of realized variances, we apply the popular heterogeneous autoregressive model for realized volatility (HAR-RV) of Corsi (2009) to describe their dynamics and, particularly, to investigate if the volatility is higher on days when stock prices decline strongly. Since we are interested in the volatility behaviour of gold during and after the stock trading hours, we consider two different specification of the HAR-RV model. The equations corresponding to the gold volatility during the S&P500 trading hours (open-to-close, OC), $RV_{t,OC}$ and the overnight period post stock trading hours until the opening of the stock market on the following day (close-to-open, CO), $RV_{t,CO}$, are as follows:

$$RV_{t,OC} = \alpha_{OC} + \beta_{OC}^{(d)} RV_{t-1}^{(d)} + \beta_{OC}^{(w)} RV_{t-1}^{(w)} + \beta_{OC}^{(m)} RV_{t-1}^{(m)} + \gamma_{OC} \cdot \mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right) + \varepsilon_{t,OC}, \tag{5.5}$$

$$\begin{aligned} RV_{t,CO} = {} & \alpha_{CO} + \beta_{CO}^{(OC)} RV_{t,OC} + \beta_{CO}^{(d)} RV_{t-1}^{(d)} + \beta_{CO}^{(w)} RV_{t-1}^{(w)} + \beta_{CO}^{(m)} RV_{t-1}^{(m)} + \\ & \gamma_{0,CO} \cdot \mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right) + \gamma_{1,CO} \cdot \mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right) RV_{t,OC} + \varepsilon_{t,CO}, \end{aligned} \tag{5.6}$$

with $RV_{t-1}^{(d)}$, $RV_{t-1}^{(w)}$ and $RV_{t-1}^{(m)}$ corresponding to daily, weekly and monthly volatility aggregates and $\varepsilon_{t,(.)}$ being a serially uncorrelated zero mean innovation term. Specifically, $RV_{t-1}^{(d)}$ corresponds to the lagged daily realized volatility, whereas the weekly and monthly volatility components, $RV_{t-1}^{(w)}$ and $RV_{t-1}^{(m)}$, are averages over the past 5 and 22 trading days, computed as $(1/5) \times (RV_{t-1}^{(d)} + RV_{t-2}^{(d)} + ... + RV_{t-5}^{(d)})$ and $(1/22) \times (RV_{t-1}^{(d)} + RV_{t-2}^{(d)} + ... + RV_{t-22}^{(d)})$, respectively. This specification corresponds to a parsimonious autoregressive-type model specification, just like an AR(22) process with restrictions. While $\beta^{(OC)}$, $\beta^{(d)}$, $\beta^{(w)}$ and $\beta^{(m)}$ measure the effect of past volatility, γ_{OC}, $\gamma_{0,CO}$ and $\gamma_{1,CO}$ describe how the volatility (realized variance) of gold is influenced by negative shocks in the stock market, and thus are of central interest in our analysis. In particular, γ_{OC} captures the effect of extreme equity price declines on the realized variance measured during the stock trading hours. If $\gamma_{OC} > 0$, the volatility in the gold market is contemporaneously higher when the stock market is in distress. Further, if $\gamma_{0,CO} + \gamma_{1,CO} > 0$, the gold volatility during the overnight period is higher when there was a negative shock during the preceding equity trading hours.

For comparison, we also estimate a HAR model for the gold realized variance of the full trading day (09:30 to 09:25 on the following day, local time), $RV_t^{(d)}$:

$$RV_t^{(d)} = \alpha + \beta^{(d)} RV_{t-1}^{(d)} + \beta^{(w)} RV_{t-1}^{(w)} + \beta^{(m)} RV_{t-1}^{(m)} + \gamma \cdot \mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right) + \varepsilon_t. \tag{5.7}$$

The results from this model provide insights with regard to the average volatility behaviour over the full gold trading day when the equity market is in distress. Again, γ measures the impact of a negative shock in the equity market on the (realized) variance of gold.

All models are estimated by OLS using logarithmic realized variances instead of the realized variance itself. This has the advantage that no parameter restrictions are necessary to ensure non-negativity of the realized variance. Further, the distribution of the logarithmic realized variance is substantially closer to a normal distribution which is attractive from an econometric point of view.

5.3 Empirical Results

5.3.1 Returns

The estimated coefficients of the safe haven regression based on gold returns synchronized with the stock trading hours are depicted graphically in Figure 5.1. In order to make the magnitude of the estimated coefficients comparable, all variables have been standardized before estimation (Verbeek, 2017). Panel (a) of Figure 5.1 displays the results for the gold spot market whereas Panel (b) shows the results for the futures market. The gray-shaded area is the 99% confidence interval obtained from a block bootstrap. For comparison, the corresponding coefficients from the mean equation of a Glosten et al. (1993) asymmetric GARCH-X model are displayed as well.[8] Figure 5.1 reveals several interesting features of the relationship between gold and S&P 500 returns.

First, the GARCH-X estimates are close to zero and statistically insignificant, implying that gold and the S&P500 do not move in tandem on average in both normal times and during a crisis in the stock market. Consequently, gold is both a weak hedge and a weak safe haven in the sense of Baur and McDermott (2010). Second, the quantile-specific dependence parameters deviate substantially from the corresponding GARCH-X estimates for both positive and moderate negative S&P500 returns, and are significantly different from zero in the tails of the gold return distribution. By contrast, the quantile specific coefficients for extremely negative returns in the S&P500 do not significantly deviate from the GARCH-X estimates in any discernible pattern. More precisely, lower quantiles of gold returns are negatively related to positive S&P500 returns whereas upper quantiles are positively related. In contrast, the association of gold returns with moderate negative S&P500 returns is positive at lower quantiles and becomes negative at upper quantiles. This suggests that under extreme conditions, reflected by lower and upper conditional quantiles, changes in the price of gold are amplified by contemporaneous stock returns. Interestingly, gold return quantiles are not affected by a negative shock in equity prices.

Our findings on the relationship in the returns of the two assets support the notion that gold is a weak safe haven for equity and do not present a contradiction thereof, since on average, gold holds its value in presence of distress in the equity market.

For comparison, we re-estimate the model in Equation 5.1 using gold returns of the full trading day. The results are presented in Figure 5.2 and confirm our findings obtained from contemporaneous returns.

Beyond an analysis of the return-relationship, we employ the model in Equation 5.1 to study the range spanned by the quantile regression lines for $\tau^\star = .02$ and $(1 - \tau^\star) = .98$. Again, all variables have been standardized prior to estimation in order to be able to compare the estimated coefficients. The results are displayed in Figure 5.3 for synchronized (upper panel) and full day returns (lower panel) in the spot (left) and futures (right) market of gold. The figure depicts the regression lines and shows the inter-quantile range spanned by extreme upper and lower quantiles. Since we explicitly disentangle the influence of positive and negative stock market returns on the gold return quantiles in the estimation, the regression lines exhibit a kink in the origin. The gray dashed vertical line marks the 1% S&P500 return quantile.

As can be seen in Figure 5.3, the state of the equity market has an influence on the inter-quantile range spanned by the regression lines for upper and lower conditional quantiles of gold returns. Interestingly, the inter-quantile range of gold spot and futures returns measured during the stock trading hours is

[8] The full GARCH estimations results are not reported to save space, but are available from the authors.

Figure 5.1
Quantile safe haven regression using gold returns during the S&P500 trading hours.

Panel (a): Gold spot market

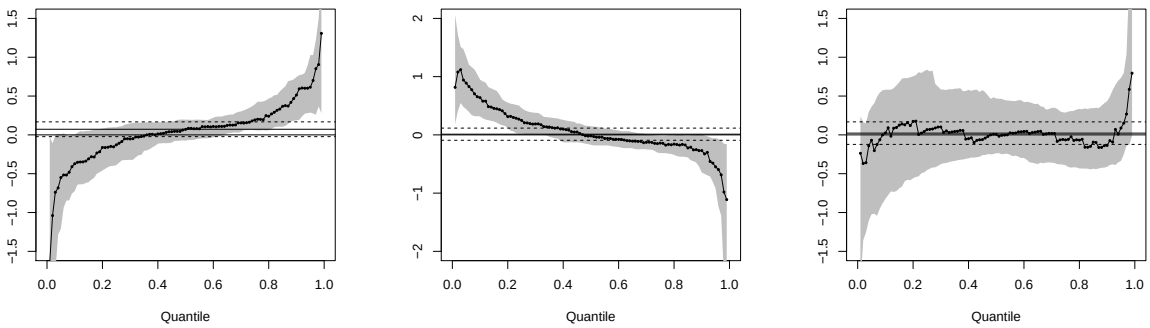

Panel (b): Gold futures market

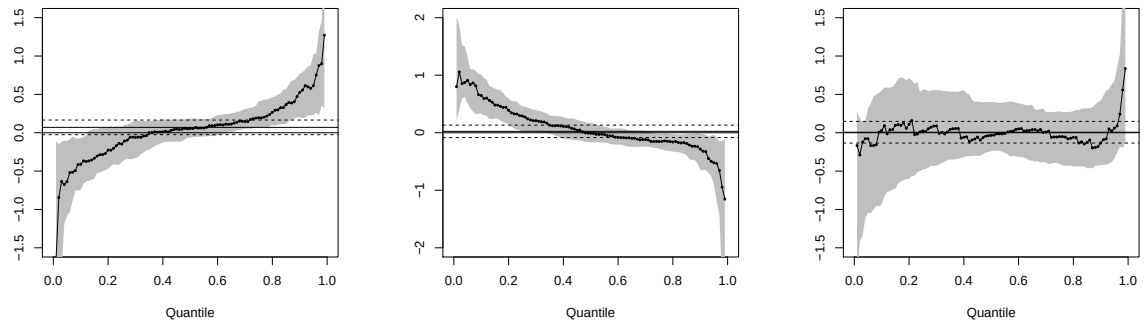

Notes: Relationship between S&P500 returns and gold spot (row one) and gold futures (row two). Relationship between gold returns and positive (left), moderate negative (middle) and extremely negative (right) equity returns. Solid curve: quantile-specific coefficients (structure of relationship); gray shaded area is the 99% confidence interval obtained from bootstrap. Solid horizontal line: coefficients of mean equation of GARCH model; dashed horizontal line: corresponding 99% confidence band.

smaller when stock returns are below their 1% quantile. That is, the variability of gold returns measured over the stock trading hours is found to be lower when the stock market is in distress. In other words, the magnitude of synchronized gold price changes is lower in this situation. However, when gold returns over the full trading day are considered instead, different volatility characteristics between spot and futures returns are found. Whilst there is no effect in the lower tail regression line for spot returns, the upper tail regression line appears to be slightly shifted downwards. That is, the inter-quantile range seems to become smaller due to a lower conditional magnitude of positive gold returns. By contrast, the inter-quantile range of gold futures returns over the full day is affected by a negative shock in the equity market.

5.3.2 Realized Variance

In addition to an analysis of gold volatility based on daily return data within the quantile regression framework in the previous section, we also study the volatility behaviour using realized measures as proxies for the latent volatility. Realized measures also account for short-lived gold price changes over

Figure 5.2
Quantile safe haven regression using gold returns over the full day.

Panel (a): Gold spot market

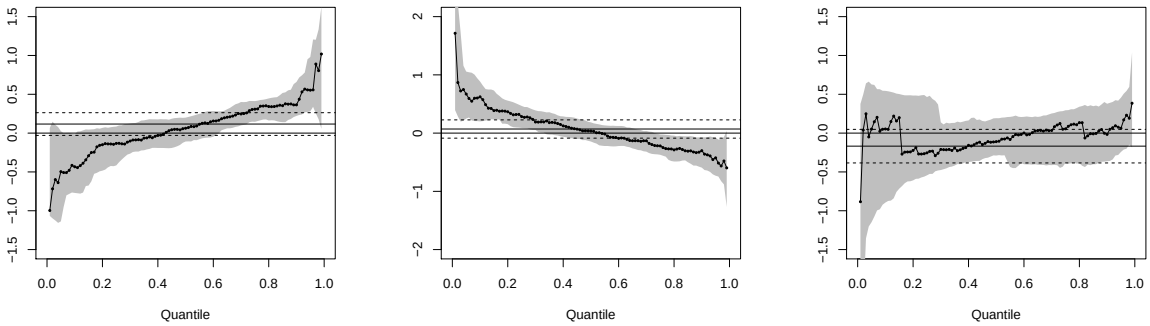

Panel (b): Gold futures market

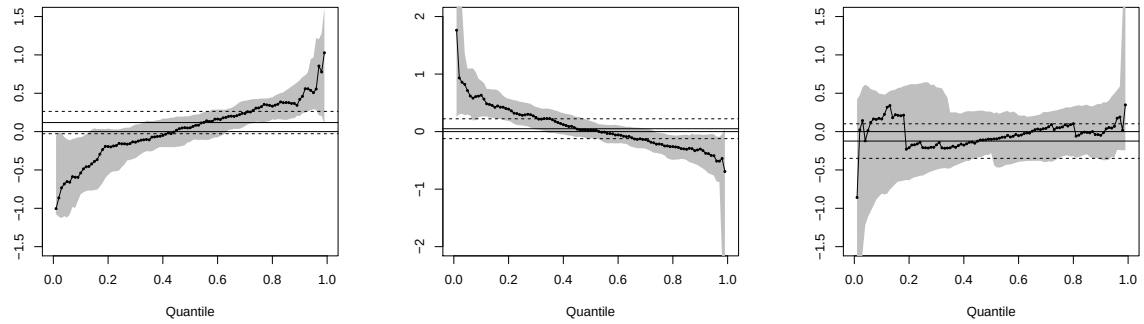

Notes: Relationship between S&P500 returns and gold spot (row one) and gold futures (row two). Relationship between gold returns and positive (left), moderate negative (middle) and extremely negative (right) equity returns. Solid curve: quantile-specific coefficients (structure of relationship); gray shaded area is the 99% confidence interval obtained from bootstrap. Solid horizontal line: coefficients of mean equation of GARCH model; dashed horizontal line: corresponding 99% confidence band.

Figure 5.3
Influence of extremely negative S&P500 returns on the inter-quantile range of gold spot and futures returns.

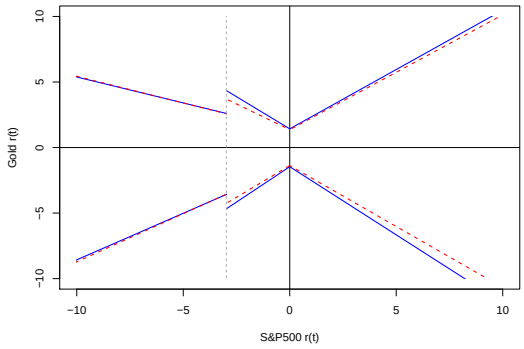

Gold spot, S&P500 trading hours

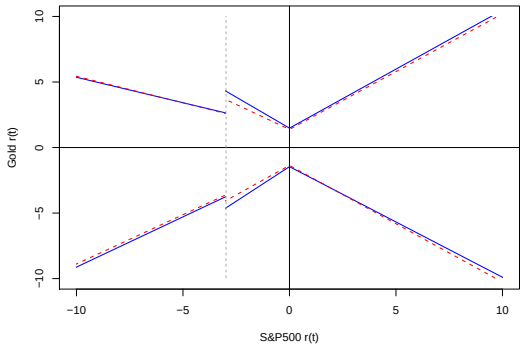

Gold futures, S&P500 trading hours

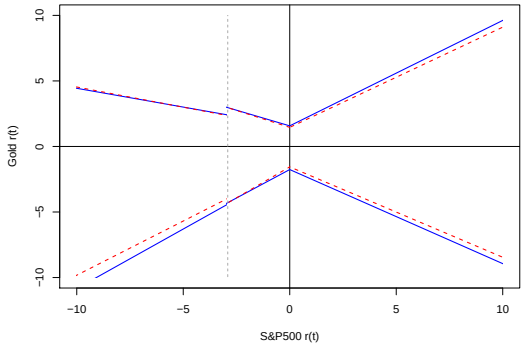

Gold spot, full day

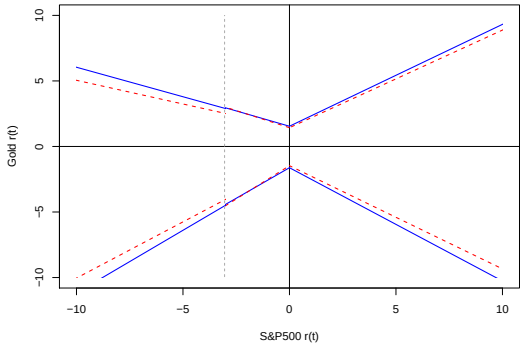

Gold futures, full day

Notes: The blue solid lines are the regression lines for $\tau^{\star} = .02$ and $(1-\tau^{\star}) = .98$; the red dashed lines are the 'average' regression lines for extreme lower and upper quantiles, $\tau' \in \{.01, \ldots, .05; .95, \ldots, .99\}$.

the trading day (intra-daily price variation) and therefore can precisely reflect the uncertainty in the gold market. If the uncertainty in the gold market is higher when the equity market is in distress, we expect an increase in the volatility measured by the realized variance but not necessarily in the range spanned by the lower and upper quantile regression lines based on synchronized or daily returns. More precisely, the change in gold prices measured from the opening until the close of the stock market has found to be limited in presence of a crisis in the equity market.

The results of our HAR models for gold spot and futures are presented in Panel (a) and (b) of Table 5.1, respectively. First, it can be noted the gold volatility exhibits a diurnal pattern and is higher when the S&P500 is traded compared to the overnight period when the stock exchange is closed. The gold volatility measured during the trading hours of the S&P500, RV_t^{OC}, depends only positively on past volatility, reflecting the high persistence in volatility. Further, the contemporaneous volatility of gold spot and futures is substantially higher when the equity market is in distress. The volatility during the post trading hours and overnight period is inversely related to the volatility of the same day but captured when the S&P was traded. Interestingly, the volatility of gold during the overnight period is not affected by negative shocks in the stock market. For comparison, we report in Table 5.1 the results of a HAR model fitted to the volatility of gold measured over the full trading day in the gold market. We find the daily volatility of daily gold volatility to be highly persistent and particularly sensitive to lagged weekly volatility. Moreover, the increase in the volatility due to a negative shock in the equity market documented for the contemporaneous volatility is visible in the full day volatility as well. Finally, the behavior of the spot and futures market appears to be very similar.

Overall, our findings suggest that the uncertainty in the gold market is higher in presence of a negative shock to stock prices. An alternative interpretation of our results is that the uncertainty in gold spot and futures markets is affected by the same factors that initially caused stock prices to decline strongly.

Together with the findings for the gold-stock return relationship, this has implications for the role of gold as safe haven. First, we find no correlation between gold and moderate negative equity returns on average, but a substantial positive relationship in the lower quantiles of gold returns. In addition, gold returns do not seem to be contemporaneously affected by negative shocks in the S&P500. Therefore, the variance of a portfolio composed of the two assets is likely to increase.

Table 5.1
HAR-RV estimation results for gold spot and futures.

	S&P trading hours	S&P500 non-trading hours	Full trading day
Panel (a): Gold spot			
Constant	-0.01	-0.00	-0.01
	(0.01)	(0.02)	(0.01)
RV_t^d	—	-0.51^{***}	—
		(0.03)	
RV_{t-1}^d	0.30^{***}	-0.26^{***}	0.08^{***}
	(0.02)	(0.03)	(0.02)
RV_{t-1}^w	0.24^{***}	0.51^{***}	0.46^{***}
	(0.03)	(0.05)	(0.04)
RV_{t-1}^m	0.28^{***}	-0.00	0.20^{***}
	(0.03)	(0.04)	(0.03)
$\mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right)$	0.34^{**}	0.27	0.44^{***}
	(0.12)	(0.37)	(0.13)
$\mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right) \times RV_t^d$	—	0.04	—
		(0.18)	
Panel (b): Gold futures			
Constant	-0.01	-0.00	0.26^{***}
	(0.01)	(0.02)	(0.02)
RV_t^d	—	-0.53^{***}	—
		(0.03)	
RV_{t-1}^d	0.33^{***}	-0.26^{***}	0.09^{***}
	(0.02)	(0.03)	(0.02)
RV_{t-1}^w	0.21^{***}	0.52^{***}	0.62^{***}
	(0.03)	(0.05)	(0.05)
RV_{t-1}^m	0.27^{***}	-0.05	0.28^{***}
	(0.03)	(0.04)	(0.05)
$\mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right)$	0.39^{**}	0.19	0.52^{**}
	(0.14)	(0.38)	(0.16)
$\mathbb{1}\left(r_t^{S\&P500} < q_{1\%}^{S\&P500}\right) \times RV_t^d$	—	0.15	—
		(0.21)	

Notes: Standard errors in parantheses. "*(**/***)" means statistical significance at the 5%(1%/.1%) level.

5.4 Conclusion

This paper explored the relationship between gold and equity, using gold returns and realized measures synchronized with the trading hours of the stock market. Using gold returns measured during the S&P500 trading hours, we show that two assets are uncorrelated on average and that the price of gold is not contemporaneously affected by negative shocks in the equity market. In other words, we confirm that gold acts as a weak safe haven for equity in the sense that it does not move in tandem with stocks when the equity market is in distress. However, we also show that the volatility of gold spot and futures markets is influenced by the state of the stock market. More specifically, comparing the findings for contemporaneous returns and realized variance of gold, we reveal that the higher volatility reflects increased uncertainty and is not caused by the magnitude of the change in the value of gold in presence of a negative shock in the equity market. Therefore, the implications for the safe haven effect of gold appear to be limited for investors with daily and longer trading horizon since the variation in daily returns appears not to be affected.

Chapter 6

Intra-day dynamics of exchange rates: New evidence from quantile regression†

6.1 Introduction

The aim of this paper is to investigate the potentially non-linear intra-day dynamics in foreign exchange (FX) markets over the full 24 hours of the trading day. The FX market is the largest and most liquid financial market in the world. The average daily turnover in the spot currency market amounts to 2,046 billion US-Dollar, with the US-Dollar involved in the majority of all FX transactions (Bank for International Settlements, 2016). Hence, understanding the short-term dynamics of US-Dollar exchange rates is relevant for market makers and currency traders in order to provide liquidity and ensure reasonable investment decisions. This knowledge is relevant for international day traders as well, since transactions in stock and bond markets might require currency trades as a by-product. Moreover, and in contrast to other financial markets, a distinct feature of the FX market is its fundamental 'two-sided' nature. That is, any movement in an exchange rate has opposite interpretations depending on the perspective of the market. Domestic currency appreciations (depreciations) correspond to foreign currency depreciations (appreciations) and, consequently, news are not universally perceived as 'positive' or 'negative'.

Using quantile autoregression according to Koenker and Xiao (2006), we examine the presence of nonlinear *autocorrelation*, or *temporal dependence*, in FX returns sampled at intra-daily time intervals from ten minutes to three hours. In line with the ideas of Fama (1965) and Samuelson (1965), financial returns are typically thought to be uncorrelated since prices should accurately reflect all past information and adjust to news *instantaneously*. Despite that the idea of 'instantaneous' price adjustment is related to price moves at very short time horizons, daily, weekly and monthly data were typically used to analyze FX return dynamics so far (see e.g. Hsieh, 1988; Liu and He, 1991; Lobato, Nankervis, and Savin, 2001; Charles and Darné, 2009; Escanciano and Lobato, 2009a,b; Chortareas, Jiang, and Nankervis, 2011a). Although there are episodes where FX returns are found to exhibit temporal dependence, the FX market generally seems to be very well-described by weak-form efficiency (Charles et al., 2012). In highly active and liquid markets – like the FX market – the return dynamics may, however, not be well-captured in data observed at low frequencies, such as daily or weekly. Therefore, autocorrelation could be present

†This chapter is joint work with Robert Maderitsch and has been originally published as KUCK, K. AND R. MADERITSCH (2019). "Intra-day dynamics of exchange rates: New evidence from quantile regression," *The Quarterly Review of Economics and Finance*, **71**, 247–257. DOI 10.1016/j.qref.2018.09.001

in FX returns sampled at fine time intervals while no temporal dependence might be observable in daily returns due to the aggregation of information. The number of studies investigating statistical features of intra-daily FX data is, however, rather limited. Wasserfallen and Zimmermann (1985), Feinstone (1987) and Ito and Roley (1987) are among the first studying intra-day FX returns, but they confine their analyses to certain trading locations and their business hours. Only Müller, Dacorogna, Olsen, Pictet, Schwarz, and Morgenegg (1990), Guillaume, Dacorogna, Davé, Müller, Olsen, and Pictet (1997) and, more recently, Neely and Weller (2003) consider intra-day FX returns collected over the full 24 hours of the trading day. Moreover, while studies investigating return dynamics commonly focus on the conditional mean of the return distribution to analyze the effect of past return(s) on the current return, we describe the complete conditional return distribution. Put differently and in contrast to measures of linear temporal dependence (like the conventional autocorrelation coefficient), quantile autoregression enables us to capture non-linear temporal dependencies. We are thus able to investigate if the autocorrelation depends on both, the sign and the magnitude of a return. Baur et al. (2012) use a similar approach to explore stock market dynamics at daily and weekly frequencies. For the FX market, there is, to the best of our knowledge, only the study of Kuck et al. (2015) that investigates non-linearities in the dynamics of various spot exchange rate returns at the daily frequency. They find that large US dollar appreciations tend to exhibit positive dependence on past returns, whereas large US dollar depreciations tend to exhibit negative dependence on past returns. A complete description of the (conditional) return distribution is also interesting in the context of intra-day risk management, where typically other moments than the mean of the (conditional) return distribution are relevant. For instance, the intra-day value-at-risk (VaR), proposed by Giot (2005) and employed by Liu and Tse (2015), can be computed straightforwardly from our quantile autoregression model without the need to rely on a distributional assumption for the intra-day returns. Due to the autoregressive nature of the model, this would also account for the dynamics in the tails of the conditional return distribution in the VaR.[1]

We consider intra-day FX return data of the Euro (EUR), the British Pound (GBP) and the Japanese Yen (JPY) against the US-Dollar (USD), to explore the return autocorrelation of the most actively traded currency pairs (Bank for International Settlements, 2016). Specifically, we 'zoom in' and employ an 11-year high-frequency dataset and estimate quantile autoregressive models for FX returns sampled at frequencies from ten minutes up to three hours. For comparison, daily FX returns are considered as well. A special feature of our study is the use of non-intermittent return time series over the whole sample period. This is an important aspect because the FX market is a decentral OTC market, operating the full 24 hours a day, with dealers around the globe continuously publishing price quotes.[2] As a general feature, we find the autocorrelation of intra-daily FX returns to be symmetrically U-shaped. Specifically, we observe pronounced negative autocorrelation for moderate FX returns, i.e. moderate USD appreciations and depreciations. For extreme appreciations and depreciations, by contrast, a tendency for positive but statistically insignificant autocorrelation is observed. Moreover, this symmetric non-linear form of autocorrelation is remarkably stable regardless of the trading intensity (intra-day realized volatility) and level of uncertainty (daily realized volatility) in the FX market.

The remainder of this paper is structured as follows: Section 6.2 provides an introduction to the

[1] Modeling and monitoring market risk at the intra-daily level seems of importance for highly active market participants with short trading horizons such as market makers and high-frequency traders.

[2] Specifically, trading in the FX market proceeds in three major markets: London (Europe), New York (North America) and Tokyo (Asia).

econometric framework. In Section 6.3 we describe the construction of the FX return series and provide descriptive statistics while in Section 6.4 we present and discuss the empirical results. Finally, Section 6.5 concludes.

6.2 The quantile autoregression framework

Conventional autoregressive models based on ordinary least squares (OLS) focus on the effects of a variables' own lags on its conditional mean. Quantile autoregression according to Koenker and Xiao (2006) can instead be used to estimate the dependence of specific conditional quantiles of the dependent variable on its own lags. This allows for an accurate description also of the tails of the dependent variable's conditional distribution. Further, it is robust to conditional heteroskedasticity, skewness and leptokurtosis which are common features of financial return time series.

As a baseline specification, we use a simple quantile autoregressive model of order one, QAR(1):

$$Q_{r_t}(\tau|\mathscr{F}_{t-1}) = \alpha(\tau) + \beta(\tau) r_{t-1}, \tag{6.1}$$

where $Q_{r_t}(\tau|\mathscr{F}_{t-1})$ denotes τth quantile of the exchange rate return conditional on the information set, $\mathscr{F}_{t-1}$, containing the first lag of the return. Specifically, the model described in Equation (6.1) is estimated separately for all currency pairs (USD/EUR, USD/GBP and USD/JPY) and frequencies (ten minutes up to three hours) we consider. The parameter $\beta(\tau)$ is the quantile-specific autocorrelation coefficient. It measures the strength and the direction of the effect of the return's first lag on the τth conditional return quantile and is of central interest in our study. The sequence of quantile-specific autocorrelation coefficients over the range of quantiles $\tau \in [.01, .02, \ldots, .99]$ characterizes the *structure of the dependence* in the spirit of Baur (2013), in our case however with respect to currency appreciations and depreciations. Beyond that, our model can be used for constructing (one-period ahead) forecasts of the intra-day return distribution, conditional on some current return.[3] It is thus possible to forecast the magnitude of appreciations and depreciations within short intra-day time horizons conditional on past return(s). This can be of use, for instance, in the context of risk management, where the dynamics in the tails of the return distribution are of relevance. Put differently, an intra-daily VaR can be computed straightforwardly from the intra-day data without relying on a particular distributional assumption for the intra-day returns. Modeling and monitoring market risk also in (near) real-time is of importance in markets composed of highly active market participants with short trading horizons such as market makers and high-frequency traders (Liu and Tse, 2015).

Furthermore, the quantile-specific first-order autoregressive coefficient can be compared to the corresponding simple OLS-AR(1) coefficient, denoted by β^{OLS}, resulting from conventional linear autoregression:

$$r_t = \alpha^{OLS} + \beta^{OLS} r_{t-1} + u_t,$$

which we consider as a benchmark. Significant deviations of the quantile-specific autocorrelation coefficient(s) from this benchmark point to the presence of non-linear dependence and potentially asymmetric short-term over- and undershooting of intra-day exchange rate returns.

[3] Note that, quantile autoregression is, however, not suited to construct multiple-period ahead forecasts in the conventional sense since future return quantiles (future states) are always unknown in advance.

In addition to the baseline model presented in Equation (6.1), we consider two extensions thereof to investigate whether different *market conditions* of the FX market affect the intra-day return dynamics. Specifically, we consider the impact of (i) different intra-day volatility levels caused by changes in trading volume and information processing intensity over the trading-day, and, (ii) daily volatility as a measure for uncertainty in the FX market (inter-daily comparison). First, accounting for the intra-day volatility is motivated by Dacorogna, Gencay, Müller, Olsen, and Pictet (2001) who find that trading activity is positively related to intra-day volatility. Controlling for the intra-day volatility hence allows, at least to some extent, to investigate the influence of trading intensity on the dynamics. Another motivation is provided by LeBaron (1992) who documents a negative relationship between volatility and stock-return serial correlation at daily and weekly frequencies. Second, we consider the daily (realized) volatility of an exchange rate as a proxy for uncertainty in the price discovery process since return dynamics might vary with the level of uncertainty. For instance, Baur and Dimpfl (2012) find that an extreme absolute return on the previous day affects the structure of dependence of daily equity returns. As an extended model specification, we use:

$$Q_{r_t}\left(\tau|\mathscr{F}_{t-1}\right) = \alpha(\tau) + \beta(\tau) r_{t-1} + \gamma(\tau) r_{t-1} \mathbb{1}\left(\sigma > \sigma^*\right), \tag{6.2}$$

where $\mathbb{1}\left(\sigma > \sigma^*\right)$ is a Heaviside indicator function equal to one if the argument in parenthesis is true and zero otherwise. Specifically, σ either is (i) the intra-day realized volatility (for some 30-min time interval k), $RV_{t,k}^{intraday}$, or (ii) denotes the daily realized volatility, RV_t^{daily}, on day t. The respective threshold value is given by σ^*. First, when controlling for the intra-day volatility seasonality, we assume the limit σ^* to be exceeded when the 30-min intra-day realized volatility is relatively 'high' due to enhanced trading activity during the 'main' business hours. Since the FX market is active 24 hours without a particular geographical location, there is no clear a priori choice for the 'main' trading hours. The most active trading hours, hence, are chosen individually for the USD/EUR, USD/GBP and USD/JPY, based on a visual inspection of their (average) 30-min intra-day realized volatilities across all trading days (see Figure 6.A–1). For the USD/EUR and USD/GBP our choice closely coincides with the business hours reported in Neely and Weller (2003). Second, the threshold σ^* is chosen to be the 95% quantile of the unconditional distribution of the daily realized volatility, when we take into account the level of uncertainty in the price-formation process of a currency pair in our empirical application. Since we aim to study serial correlation in presence of higher trading intensity and increased uncertainty in the FX market, we are mainly interested in the sum of the two coefficients $\beta(\tau) + \gamma(\tau)$ from Equation (6.2).

All model parameters are estimated using the standard optimization routine presented in Koenker and Bassett (1978) and implemented, for example, in R. Asymptotic standard errors are estimated using bootstrap methods. As we use covariance-stationary time series data, we perform a block bootstrap with block length equal to five days of intra-day data observations and 2000 replications (see, for example, Chernick, 2008, ch. 5).[4]

[4]The block length equals 40, 60, 120, 240, 480 and 720 observations for 3/2/1-hour, 30/15/10-min FX returns, respectively. In each case, the (average) block-length thus equals 5 days of intra-day data.

6.3 Data

Our initial data-set covers 5-min intra-day closing bid- and ask prices of the USD/EUR, USD/GBP and USD/JPY exchange rates in direct quotation from the viewpoint of the US market.[5] The data for the USD/EUR has been provided by Olsen Financial Technologies, while the data for the other exchange rates has been downloaded from Thomson Reuters TickHistory. The sample period spans almost 11 years from January 3, 2000, to September 30, 2011. Since the FX market is a worldwide OTC market and participants are not required to register their transactions with any central agency, transaction prices are usually not available. Market makers, however, publish (non-binding) quotes which are shared globally by data providers. It is, common to regard the bid/ask midpoint (average of bid and ask price) as a reasonable proxy for the transaction price (Dacorogna et al., 2001).

To construct the exchange rate returns, we first compute bid/ask midpoints. Specifically, to avoid the need to model periods of low trading volume due to limited trading activity, we follow Andersen et al. (2001b, 2003), Bubák et al. (2011) or Chortareas, Jiang, and Nankervis (2011b) and discard weekend periods from Friday 21:00 UTC until Sunday 21:00 UTC as well as major public holidays.[6] We moreover follow Bubák et al. (2011) and Andersen et al. (2001b) and define the trading day as the interval from 21:00 UTC to 20:59 UTC of the following day.[7] After these adjustments the data-set covers 3,029 trading days of intra-day data in total. Then, we construct the intra-day returns of the bid/ask midpoint as percentage log-differences for each currency against the US-Dollar:

$$r_t = 100 \times [\log(s_t) - \log(s_{t-1})], \tag{6.3}$$

where s_t denotes the spot bid/ask midprice of a particular currency pair. Specifically, we consider intra-day FX returns at 10-min, 15-min, 30-min, hourly, two-hour and three-hour time-horizons. Since OTC trading in the foreign exchange market takes place 24 hours a day, we obtain sequences of non-intermittent high frequency returns over the whole sample period.[8] Using midprices in direct quotation from the perspective of the US-market means that positive FX returns reflect US-Dollar depreciations, whereas negative returns coincide with US-Dollar appreciations. The time series for the midprices and returns rates are depicted in Figure 6.1 for all three exchange rates.

The statistical properties of the returns are presented in Table 6.1. As characteristic for financial returns, all time series exhibit means close to zero, a relatively pronounced skewness and excess kurtosis. Also, the kurtosis tends to increase from lower to higher return frequencies and all return series are characterized by volatility clustering, as apparent in Figure 6.1. Moreover, consistent with, for example, Dacorogna et al. (2001), Andersen et al. (2001b) or Neely and Weller (2003), negative serial correlation is apparent for intra-day FX returns as a typical feature of high-frequency financial return data as well.

[5] Specifically, we use 5-min bid- and ask close quotes, representing the last qualified value in a particular 5-min interval.

[6] Particularly, we discard January 1, December 25 and 26 as well as Easter Monday.

[7] The definition of the the trading day should not have any influence on the results since, except for weekends and public holidays, we study non-intermittent time series of returns.

[8] Price changes over the weekend (weekend returns) are discarded.

Table 6.1
Summary statistics.

USD/EUR

	Mean	Std. Dev.	Skewness	Kurtosis	Min	Max	$\rho(1)$	N
5 min	$0.259 \cdot 10^{-4}$	0.0408	0.51	48.00	-1.38	2.70	-0.02^{***}	865,759
10 min	$0.493 \cdot 10^{-4}$	0.0572	0.39	27.56	-1.28	2.21	-0.03^{***}	432,701
30 min	$0.152 \cdot 10^{-3}$	0.0970	0.09	17.00	-1.97	2.31	-0.01^{***}	144,153
1 Hour	$0.351 \cdot 10^{-3}$	0.1371	0.05	12.77	-2.03	2.24	0.00	72,060
2 Hour	$0.687 \cdot 10^{-3}$	0.1929	0.18	11.27	-1.85	2.49	0.00	36,026
3 Hour	$0.976 \cdot 10^{-3}$	0.2358	0.19	11.13	-1.80	3.15	0.00	24,042
Daily	$0.890 \cdot 10^{-2}$	0.6595	0.05	5.54	-3.85	4.62	0.02	2,407

USD/GBP

	Mean	Std. Dev.	Skewness	Kurtosis	Min	Max	$\rho(1)$	N
5 min	$-0.358 \cdot 10^{-5}$	0.0388	-0.18	27.04	-1.39	1.07	-0.04^{***}	872,316
10 min	$-0.716 \cdot 10^{-5}$	0.0536	-0.11	20.36	-1.33	1.28	-0.03^{***}	436,158
30 min	$-0.215 \cdot 10^{-4}$	0.0909	-0.12	17.74	-1.45	2.00	-0.02^{***}	145,386
1 Hour	$-0.429 \cdot 10^{-4}$	0.1278	-0.22	15.55	-2.22	1.93	-0.01^{***}	72,693
2 Hour	$-0.854 \cdot 10^{-4}$	0.1788	-0.15	12.79	-2.47	1.94	0.00	36,346
3 Hour	$-0.129 \cdot 10^{-3}$	0.2193	-0.20	13.31	-3.22	2.23	-0.01^{**}	24,231
Daily	$0.578 \cdot 10^{-2}$	0.6084	-0.03	7.60	-4.20	4.58	0.05^{*}	2,408

USD/JPY

	Mean	Std. Dev.	Skewness	Kurtosis	Min	Max	$\rho(1)$	N
5 min	$0.216 \cdot 10^{-4}$	0.0425	0.40	40.69	-1.28	1.82	-0.05^{***}	872,316
10 min	$0.431 \cdot 10^{-4}$	0.0583	0.36	30.82	-1.38	2.14	-0.04^{***}	436,158
30 min	$0.129 \cdot 10^{-3}$	0.0975	0.42	23.49	-1.72	2.56	-0.01^{***}	145,386
1 Hour	$0.259 \cdot 10^{-3}$	0.1370	0.39	19.62	-1.71	3.07	-0.01^{***}	72,693
2 Hour	$0.514 \cdot 10^{-3}$	0.1918	0.27	15.28	-2.41	3.12	-0.02^{***}	36,346
3 Hour	$0.776 \cdot 10^{-3}$	0.2336	0.06	11.64	-2.48	2.51	-0.02^{***}	24,231
Daily	$0.867 \cdot 10^{-2}$	0.6586	0.32	5.98	-3.09	4.42	-0.02	2,408

Note: Means are in percentage terms, $\rho(1)$ is the first order autocorrelation.
"*(**/***)" means statistical significance at the 10%(5%/1%) level.

Figure 6.1
Time series plots for the exchange rate series.

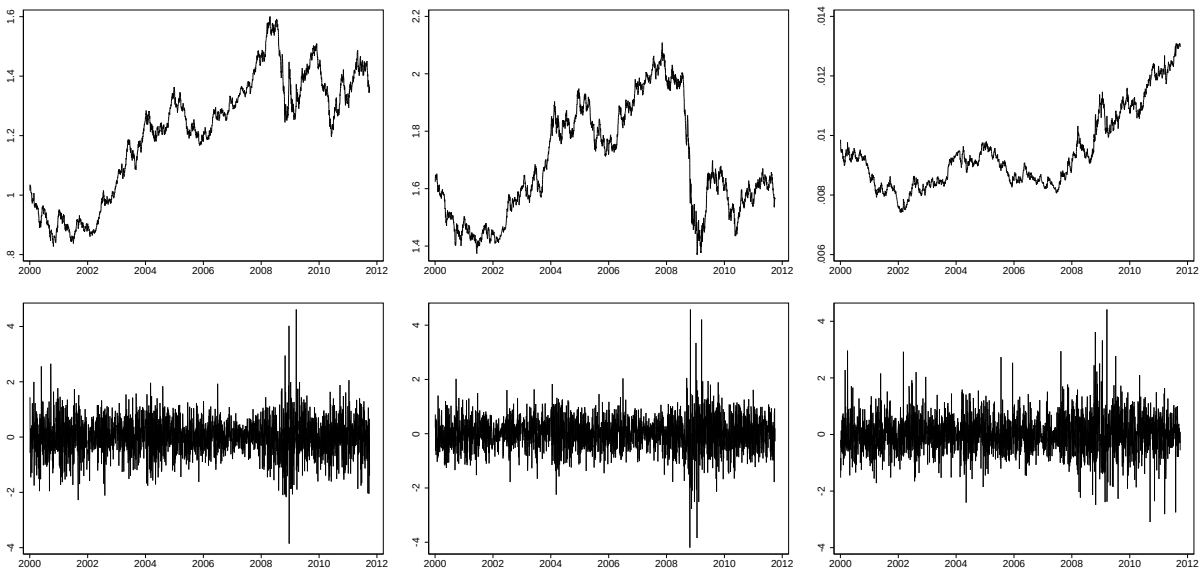

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

6.4 Empirical results

6.4.1 Baseline specification

We start the empirical analysis with the baseline QAR(1) as given in Equation (6.1). For ease of understanding, we present the results graphically. For each currency pair, we depict the sequence of the estimated $\hat{\beta}(\tau)$s over all quantiles τ (black solid curve) together with the corresponding 99.9% confidence band (grey-shaded area). In addition, we show the benchmark OLS-AR(1)-coefficient with its 99.9% confidence interval (dashed line). The results for the baseline model for 10-min, 15-min and 30-min returns are depicted in Figure 6.1.

Most notably, the intra-day FX dynamics are remarkably similar across the three currencies. Specifically, the 'structure' of temporal dependence in FX returns is found to be mostly symmetrically U-shaped. Statistically significant negative autocorrelation can be observed for the majority of return quantiles. However, for extreme US-Dollar appreciations or depreciations (extreme quantiles), there is a tendency for positive, although generally insignificant serial correlation. Significant positive autocorrelation is observed only in a few cases and only for returns sampled at time intervals up to one hour. For instance, in the case of extreme depreciations of the US-Dollar against Euro and Yen, significant positive autocorrelation can be observed. By contrast, for the British Pound, only extreme US-Dollar appreciations seem to exhibit significant positive autocorrelation at return frequencies higher than 60 minutes. At lower intra-day frequencies, the U-shapedness of the dependence structure is still apparent but seems to be less pronounced (see Figure 6.2).

Positive return serial correlation indicates that the market does not instantaneously incorporate news such that prices adjust only partially. In such a situation, news – although unambiguous – potentially do not arrive simultaneously or are processed with different speed by different market participants. This

might be particularly true for extreme news (lower and upper quantiles) and hence provide an explanation for the symmetry observed in the dependence structure. A distinct feature of the FX market, compared to other financial markets is its two-sidedness. Returns have opposite interpretation depending on the perspective of the market, such that returns will not universally be interpreted as 'positive' or 'negative'.

Negative first-order autocorrelation in intra-day returns has already been documented in the literature and often is attributed to the presence of microstructure effects that might disappear after completion of the price-formation. For instance, Goodhart (1990), Goodhart and Figliuoli (1991), Dacorogna et al. (2001) and Neely and Weller (2003) report substantial negative autocorrelation for intra-day FX returns. Moreover, Bianco and Renò (2006) find negative autocorrelation for Italian stock index futures. Among the explanations offered in the literature we found non-synchronous trading (Baillie and Bollerslev, 1991)[9], erroneous data or "screen-fighting" (Zhou, 1996). In our case however, none of the existing explanations appear fully convincing.

Microstructure noise is typically present at ultra high frequencies such as tick-by-tick or minute-to-minute, but disappears at lower intra-day frequencies such as 15- or 30-minutes. Moreover, it has been shown that returns at 5-minute time intervals typically no longer exhibit microstructure noise (Andersen et al., 2010).[10] Bubák et al. (2011) study a very similar dataset as we do (also provided by Olsen Financial Technologies) and do not identify any microstructure noise in the exchange rate returns (for the USD/EUR) for the 5-minute frequency. Since the smallest time interval that we consider is 10 minutes, we are convinced that our findings are not overshadowed by any such microstructure effects. In line with this notion, we also do not observe substantial differences in the structure of dependence between the 10-minute to 30-minute frequencies. Moreover, Dacorogna et al. (2001) provide another explanation for the presence of negative autocorrelation in high-frequent financial returns. They point out that negative first-order autocorrelation could be caused by *diverging opinions* on the effect of new information on the price of a financial asset. Market participants may observe different news or interpret the same information differently. In case of the FX market, this argument is particularly convincing, given its two-sidedness and decentral, global nature. The evaluation of the overall effect of news might thus be harder than in case of stocks, for example.

Further, strong similarities in the intra-day FX return autocorrelations of our currency pairs are apparent. First, the quantile specific autocorrelation differs significantly from the OLS benchmark at high intra-day return frequencies. More specifically, for the vast majority of quantiles, the quantile specific autocorrelation turns out to be more pronounced than by linear autocorrelation. Second, the magnitude of autocorrelation declines with increasing time intervals. From an economic perspective, the autocorrelation tends to become insignificant from the 180-min frequency onwards. Finally, the pronounced symmetric U-form of the dependence structure is a distinct feature of intra-day FX returns. It is not apparent at daily return frequencies (see Figure 6.3). As a robustness check, we also considered intra-day returns for two other assets in order to see whether the symmetric U-shaped dependence structure is a specific feature of the FX market. Particularly, we investigated intra-day returns for S&P500 spot index and Gold returns, but were not able to find evidence for similar patterns in the temporal dependence.[11]

[9] While non-synchronous trading may induce spurious positive first-order serial correlation for an index or portfolio, it can also induce artificial negative autocorrelation in the return series on a single asset (Campbell, Lo, and MacKinlay, 1997).

[10] Returns computed over 5-minute intervals are commonly used in the literature, and realized volatilities are typically constructed from returns measured at that frequency.

[11] We therefore suspect that the symmetric U-shaped dependence structure is related to the 'two-sidedness' of the FX market. The results for the robustness checks are available on request.

Figure 6.1
QAR(1)-Baseline model estimation results for 30-min to 10-min intra-day returns.

Panel (a): 30-minute returns

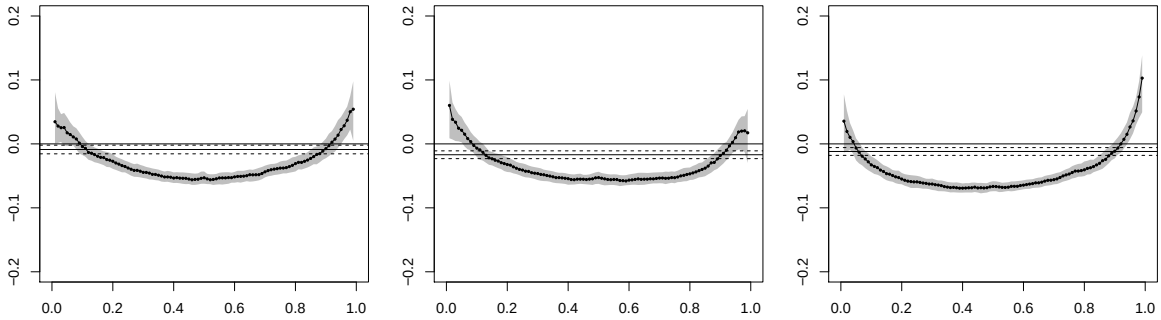

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

Panel (b): 15-minute returns

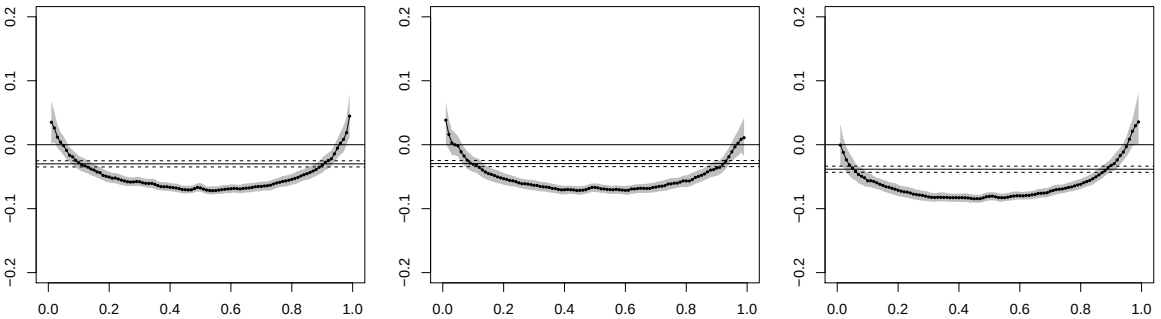

(d) USD/EUR **(e)** USD/GBP **(f)** USD/JPY

Panel (c): 10-minute returns

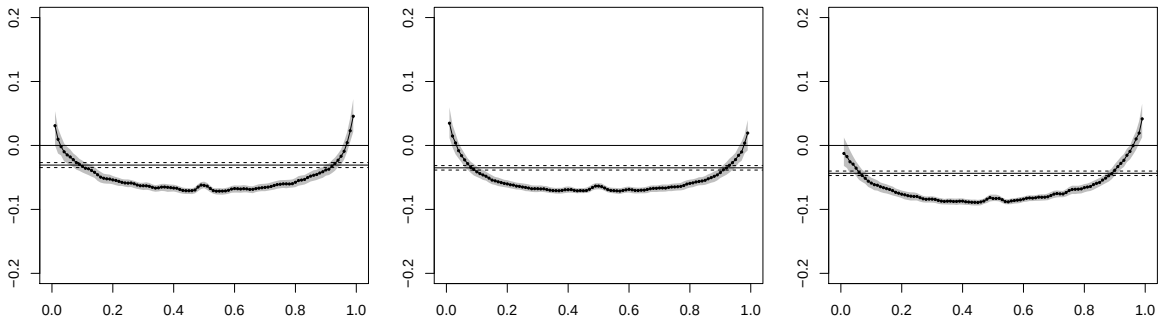

(g) USD/EUR **(h)** USD/GBP **(i)** USD/JPY

Notes: Solid horizontal line: OLS-coefficient; dashed horizontal lines: corresponding 99.9% confidence bands. Solid curve: quantile-specific coefficients; Gray-shaded: area within the corresponding 99.9% confidence band; Confidence bands are based on standard errors estimated via bootstrapping techniques.

Figure 6.2
QAR(1)-Baseline model estimation results for 3 hr to 1 hr intra-day returns.

Panel (a): 180-minute returns

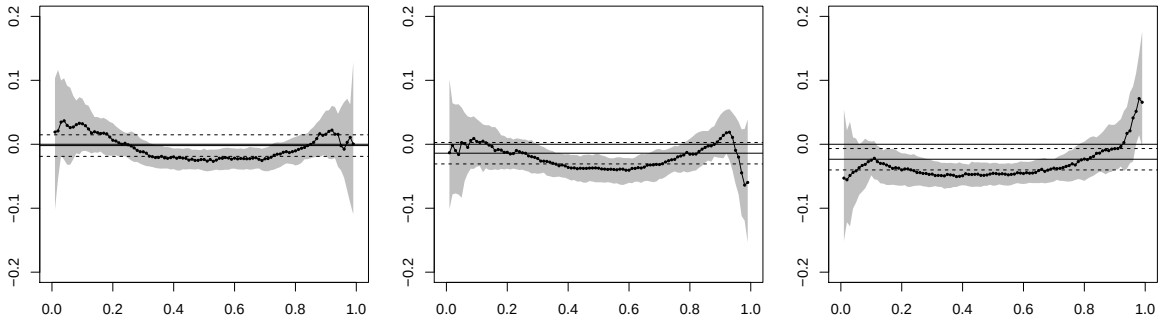

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

Panel (b): 120-minute returns

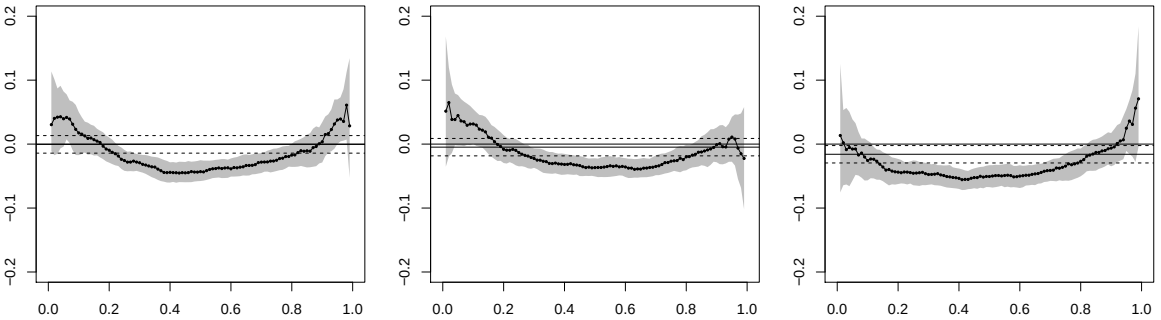

(d) USD/EUR **(e)** USD/GBP **(f)** USD/JPY

Panel (c): 60-minute returns

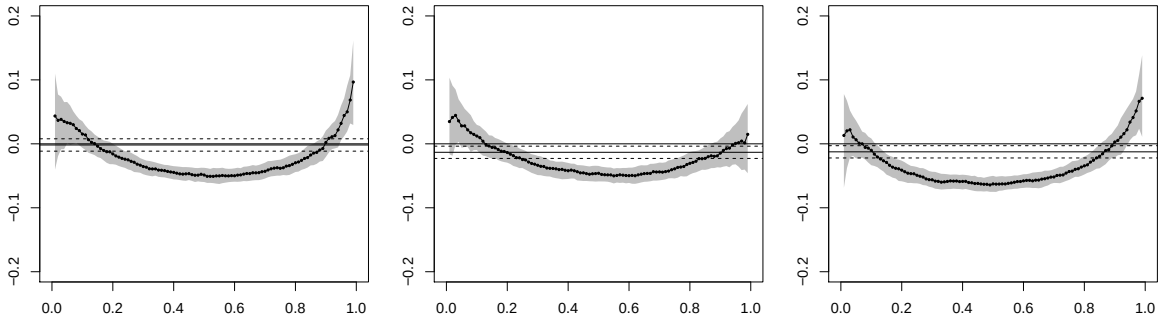

(g) USD/EUR **(h)** USD/GBP **(i)** USD/JPY

Notes: Solid horizontal line: OLS-coefficient; dashed horizontal lines: corresponding 99.9% confidence bands. Solid curve: quantile-specific coefficients; Gray-shaded: area within the corresponding 99.9% confidence band; Confidence bands are based on standard errors estimated via bootstrapping techniques.

Figure 6.3
QAR(1)-Baseline model estimation results for daily returns.

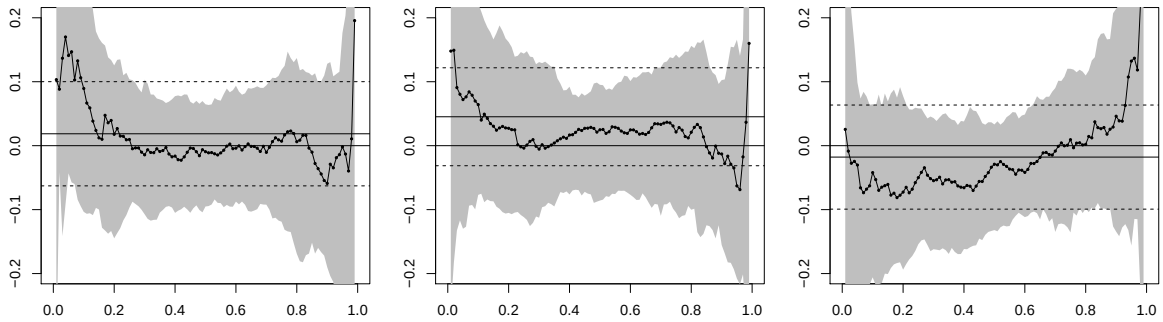

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

Notes: Daily FX returns reflect the change in the midprice between 16:00 GMT on day $t-1$ and 16:00 at day t. Solid horizontal line: OLS-coefficient; dashed horizontal lines: corresponding 99.9% confidence bands. Solid curve: quantile-specific coefficients; Gray-shaded: area within the corresponding 99.9% confidence band; Confidence bands are based on standard errors estimated via bootstrapping techniques.

6.4.2 Extended specification

In this section, we consider the extensions of the baseline model as proposed in Equation (6.2) in order to analyze the impact of diurnal volatility variation and increased uncertainty. The results are depicted in Figures 6.4 and 6.5. The black solid curve represents the (baseline) quantile specific first order autocorrelation while the red solid curve is the autocorrelation when either the 30-min intra-day realized volatility is increased or the daily realized volatility is high.[12]

We start by studying the impact of the intra-daily seasonality of the volatility on the temporal dependence of FX returns. In Figure 6.4 it is obvious that the return autocorrelation structure is hardly affected by higher intra-day realized volatility due to increased trading activity. More precisely, the quantile specific FX return autocorrelation does not differ significantly from the baseline first order autocorrelation. The U-shaped dependence structure described in the previous section, thus, is a feature of intra-day FX returns which is not driven by phases of high-trading intensity in a currency pair during business hours.

Beyond the impact of the volatility's intra-day seasonality, we were interested in the (in-) stability of the dependence structure with respect to the 'condition' of the FX market. We therefore measure FX market uncertainty by the daily realized volatility and study whether the structure of dependence changes when the ex-post daily realized volatility of a currency pair is extremely high. Most notably, we find the dependence structure to be hardly affected by uncertainty. The dynamics do not deviate significantly from those revealed by the baseline specification. However, as Figure 6.5 also shows, the shape of the dependence structure tends to get flatter in the sense that the U-shape becomes less pronounced on high volatility days (red curve). In essence, the structure of dependence of intra-day FX returns turns out to be remarkably stable both with regard to intra-day seasonality of volatility and uncertainty in the FX markets.

[12]To save space, we focus on intra-day FX returns for 30-, 15- and 10-min time intervals. Results for longer intra-day time intervals are similar and available on request.

Figure 6.4
Impact of intra-day volatility seasonality.

Panel (a): 30-minute returns

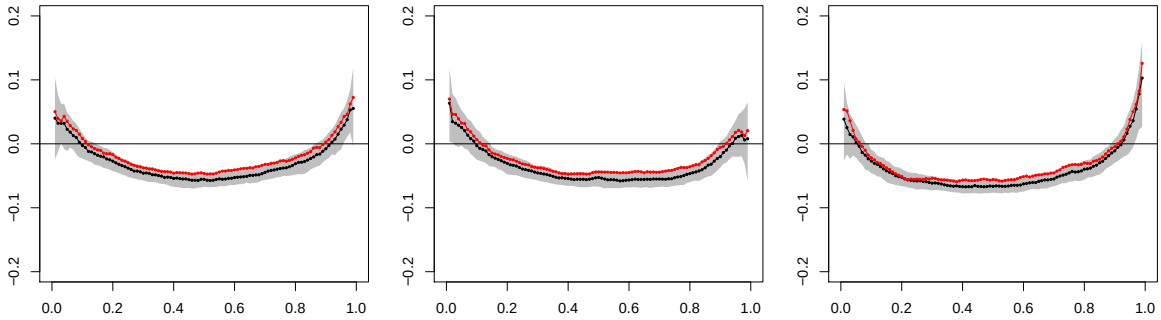

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

Panel (b): 15-minute returns

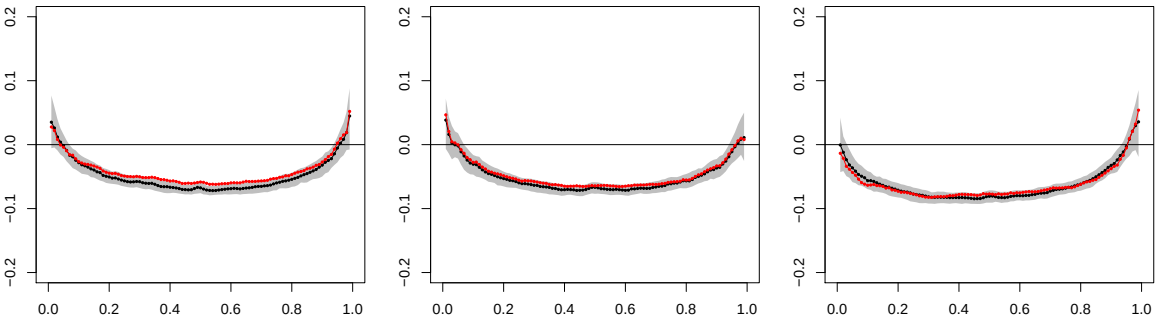

(d) USD/EUR **(e)** USD/GBP **(f)** USD/JPY

Panel (c): 10-minute returns

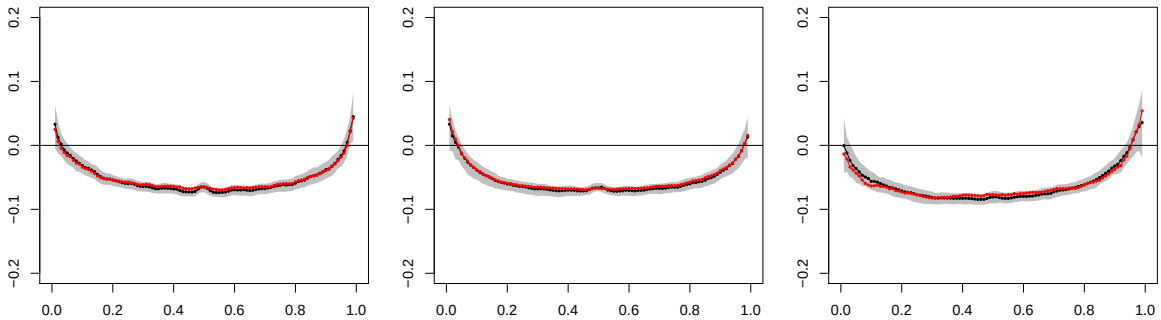

(g) USD/EUR **(h)** USD/GBP **(i)** USD/JPY

Notes: Black solid curve: baseline temporal dependence of FX returns at different intra-day time horizons; Gray-shaded: area within the corresponding 99.9% confidence band; Confidence bands are based on standard errors estimated via bootstrapping techniques. Red solid curve: return autocorrelation when intra-day realized volatility is relatively high during main business hours in a particular market.

Figure 6.5
Impact of increased financial uncertainty.

Panel (a): 30-minute returns

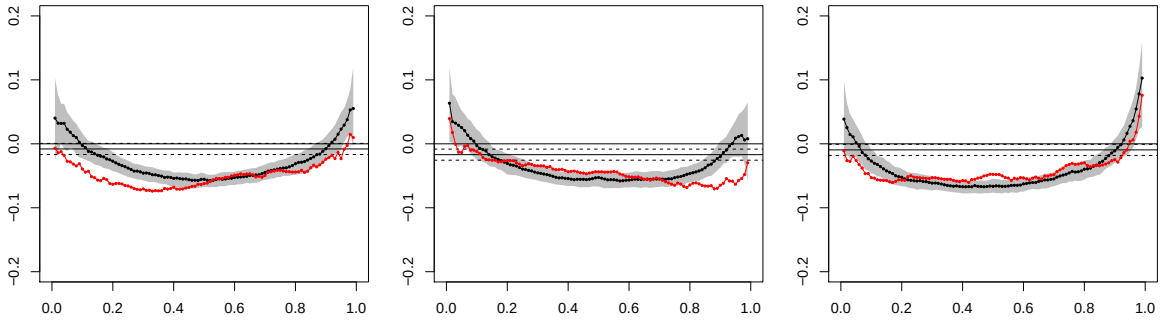

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

Panel (b): 15-minute returns

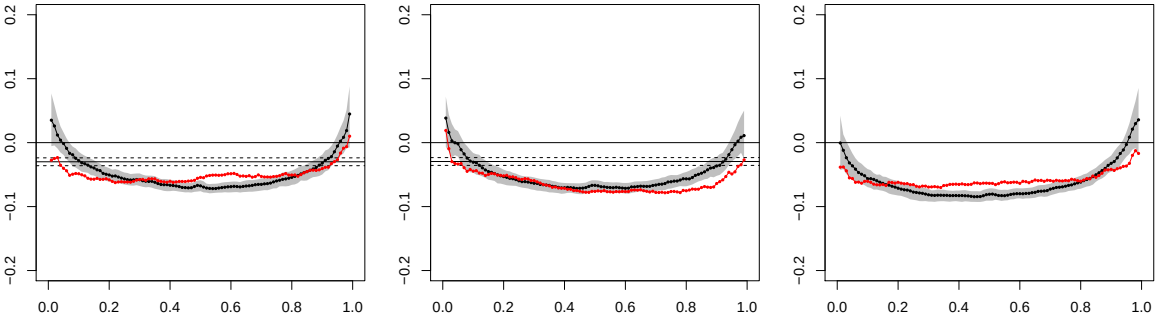

(d) USD/EUR **(e)** USD/GBP **(f)** USD/JPY

Panel (c): 10-minute returns

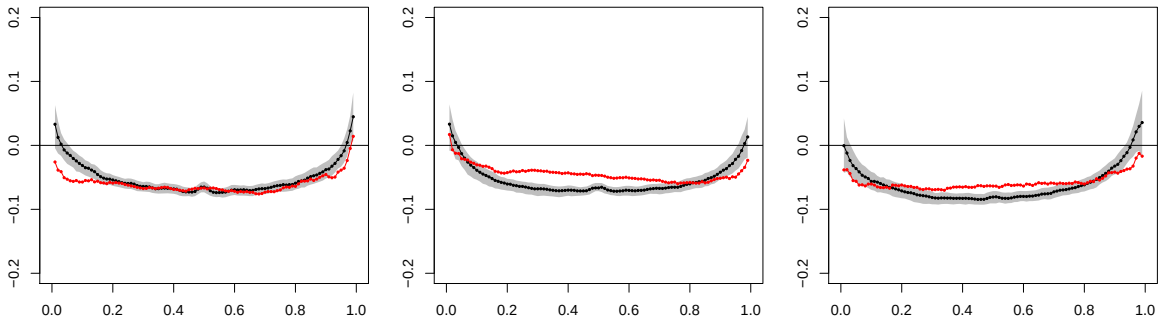

(g) USD/EUR **(h)** USD/GBP **(i)** USD/JPY

Notes: Black solid curve: baseline temporal dependence of FX returns at different intra-day time horizons; Gray-shaded: area within the corresponding 99.9% confidence band; Confidence bands are based on standard errors estimated via bootstrapping techniques. Red solid curve: return autocorrelation when daily realized volatility is relatively high (exceeding the 95% quantile of the unconditional volatility distribution).

6.5 Conclusion

Describing characteristics and establishing stylized facts on the behavior of financial time series is of substantial importance in financial econometrics. Specifically, stylized facts and statistical properties about financial time series may help to enhance the understanding of financial markets and to design and refine appropriate theoretical models and simulations. This paper provides a differentiated picture on the dynamics of US-Dollar spot exchange rate returns at an intra-day level. The use of quantile autoregression techniques, as developed by Koenker and Xiao (2006), allows us to shed more light on intra-day FX return dynamics than conventional linear autoregressive models. This seems to be of interest, for instance, in the context of intra-day riskmanagement where the dynamics in the tails of the return distribution are of greater relevance than the mean.

In contrast to previous studies, we find negative FX return autocorrelation not attributable to common market microstructure effects. Beyond this, we also identify a pronounced U-shaped structure of dependence that is very similar across all exchange rates considered. Specifically, we find that in particular moderate FX rate changes tend to be characterized by significant negative autocorrelation. Theoretically, this finding is consistent with the presence of currency overshooting under calm market conditions. The fact that this negative autocorrelation disappears slowly from higher to lower return frequencies points to the presence of asymmetric information processing and diverging opinions. Strong FX market appreciations or depreciations, by contrast tend to be characterized by the absence of (significant) autocorrelation.

Further, we show that our newly detected U-shaped dependence structure is remarkably stable with regard to different market conditions (trading intensity, uncertainty). The dependence structure is virtually unaffected by diurnal patterns in intra-day volatility and different levels of trading intensity. Moreover, we find quantile specific return autocorrelations to be stable with regard to increased uncertainty in the FX markets as well.

Appendix 6.A Figures and Tables

Figure 6.A–1
Average 30-min intra-day realized volatility over 24 hours.

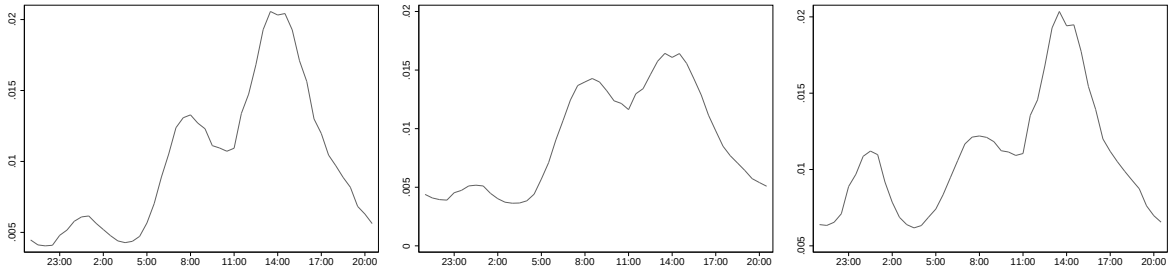

(a) USD/EUR **(b)** USD/GBP **(c)** USD/JPY

Notes: This figure shows the average 30-min realized volatility of the three currency pairs over the trading day (GMT). For the USD/EUR and USD/GBP pairs, phases of relatively high intra-day volatility match the corresponding (most) active business hours of the European trading session (around 08:00 to 16:00 GMT).

Chapter 7

Discussion

This thesis explores state-dependence in the context of financial market dynamics and cross-market linkages. More specifically, the view on *state-dependence* in this thesis is from an *economic perspective* and refers to the time-varying behaviour in and across financial markets, related to the *economic state* of the market. Changing dynamics in response to varying economic conditions or extreme events is a widely observed behaviour of financial markets. The corresponding statistical feature is *non-linear dependence* of financial variables and refers to changes in their temporal dependencies or relationships over time which is associated with observed and latent factors.

The findings regarding financial market behaviour presented in this thesis are of interest for investors and other market agents to ensure reasonable investment decisions and appropriate risk management. Furthermore, accurate descriptions of the statistical properties of financial time series contribute to the understanding of financial markets and help to design and refine theoretical and statistical models. In particular, this thesis addresses the following issues:

1. **Are the dynamics among crude oil prices stable or time-varying? Are the crude oil markets generally integrated or 'regionalized'? Is there a leading benchmark price?**

   Applying a Markov-switching vector error correction model to five crude oil benchmark prices (WTI, Brent, Bonny Light, Dubai, Tapis), we find the world crude oil market to be integrated, although the *degree* of integration is varying over time. Specifically, we identify three distinct regimes characterized by different price dynamics, reflecting the changing landscape of the world crude oil market. First, we find that the world crude oil market was regionalized throughout the 1980s but became integrated in subsequent years. Second, whilst the crude oil market is generally integrated after the 1980s, the degree of integration varies with the level of global economic uncertainty. Further, our empirical results suggest that Dubai acts as a price setter in all regimes, whereas the adjustment behaviour of the other benchmark prices varies over time. This implies that there is no benchmark which universally leads the other crude oil prices at all times. Therefore, instead of a single benchmark price rather the system of benchmark prices should be considered for a precise assessment of the global crude oil price.

   The Markov-switching model is a useful approach to describe state-dependence in the dynamics of a process for a given number of regimes. A specific feature of Markov-switching models, compared to other non-linear models, is that the state variable is latent. This means that the states are estimated form observed realizations of the process which are potentially generated under

different and latent regimes. This also seems of particular usefulness when the state governing variable is unknown or cannot be observed. From an econometric perspective, an important issue of Markov-switching models in the context of vector error correction models, however, is that a large number of parameters needs to be estimated, even when only a small number of different states is assumed. This implies that usually only two or three states can be considered in practical applications. Moreover, the process should not oscillate between states, as this impedes their economic interpretation and the use of the model in forecasting applications. In addition, since the non-linearity in the process dynamics enters entirely data driven, further interpretation requires that the states can be motivated by economic theory or reasoning.

2. **How are the volatility dynamics of crude oil and precious metals affected by the level of volatility? Are there differences between crude oil and precious metals?**

   We find the volatility dynamics in the gold, silver and crude oil futures to be characterized by substantial state-dependence. More specifically, estimating a HAR-RV model in the spirit of Corsi (2009) by quantile regression techniques, we document that the commodity volatilities are time-varying with the level of current daily volatility. Specifically, for the weekly and monthly volatility aggregates, there are substantial changes in the dependence between high and low volatility states: When volatility is low, past monthly volatility has a strong amplifying effect on the current volatility, whilst there is an inverse dependence on past weekly volatility. On high volatility days, by contrast, we find a strong positive effect of the weekly volatility aggregate whereas the dependence on the long-term volatility becomes insignificant. This implies that information generated over the medium-term becomes more important when uncertainty is high. Our findings are theoretically consistent with the idea of shifts in investor attention depending on the uncertainty and risk in the commodity market. An alternative interpretation of our findings is that the composition of traders changes with the level of uncertainty and risk in the commodity markets. Further, our results suggest that the (average) volatility dynamics differ between commodities: In particular, for crude oil which is a non-durable 'consumption' good, associated with high transportation costs, we find a pronounced dependence of the daily volatility on past monthly volatility, whereas the impact of past daily volatility is limited. From an economic viewpoint, this may reflect the 'real' interest in the underlying commodity and the rather long time horizon of traders active in the crude oil market. By contrast, in the case of gold and silver, past daily and monthly volatility appear to be equally important for current volatility. This may reflect that gold and silver have industrial use but are important 'investment' assets as well. More precisely, it suggests the presence of investors with rather long time horizon who might be interested in precious metals as a input factor or as store of value as well as investors with short time horizon who trade at high frequencies.

   Beyond describing non-linearities in the volatility dynamics, our Q-HAR-RV models can be employed to predict quantiles of the (conditional) distribution of future volatility (Žikeš and Baruník, 2016). That is, our models allow to assess the degree of future uncertainty, which seems of immediate use for market agents exposed to commodity volatility and investors in the context of portfolio diversification. However, a certain limitation of Q-HAR-RV models with regard to forecasting is that future states of the process, reflected by quantiles, are always unknown and unpredictable.

3. **What can be learned from intra-day data regarding the safe-haven property of gold? How**

fast do investors react to negative shocks in the equity market? Do negative shocks in the equity market affect the volatility of gold and what are the implications for the role of gold as a safe haven?

Using intra-day instead of daily data allows to provide a more accurate description of the state-dependence in the relationship between gold and stock prices. Studying intra-day data enhances the understanding of the time dimension of the safe haven effect of gold. Our descriptive results show that extreme equity price movements, ex-post observed in daily returns, do not occur within short 5 minute periods but materialize over the trading day and that gold prices do not decline on these days. We find that on average, gold prices do not move in tandem with stocks when the equity market is in distress. Further, we show that gold holds its value also in the overnight period following a negative shock in the equity market. Although ex-post observed negative shocks in equity prices accumulate over the trading day, extremely strong negative 5 minute returns lead to a positive reaction of the gold price. These findings suggest a fast reaction of the gold price to extreme equity price declines and are in line with the results reported in Baur and McDermott (2010) and Baur and Lucey (2010) for daily data. Based on a quantile regression model for gold returns measured during the stock trading hours, we document that the absolute size of returns does not increase in response to a negative shock in the equity market. However, the results of a heterogeneous autoregressive model of realized volatility (HAR-RV) in the spirit of Corsi (2009) show that the volatility of gold as measured by the realized variance during the stock trading hours is higher when the equity market is in distress. Since realized measures are more accurate proxies of the volatility as they also capture the impact of short-term intra-daily returns, we conclude that this finding reflects an increase in the uncertainty in the gold market. The effect of the increased contemporaneous volatility on the safe haven effect of gold, however, seems limited for investors with daily trading horizon since the variation in daily returns appears not to be affected.

An important issue of our studies is that we define "extreme days" based on ex-post observed open-to-close equity returns. Whilst this allows for a precise description of the behaviour of gold when the stock market is in distress, it is not possible to design an active trading strategy that exploits this feature in the relationship between the two assets. However, our findings suggest that investors can compensate extreme losses in the equity market by diversifying their portfolios with gold. Generally increased gold holdings by investors, however, could reduce the effectiveness of gold as a safe haven since investors might sell gold to cover losses, reduce risk or to maintain their investment preferences in response to negative shocks in the equity market (Baur and Glover, 2012). From an econometric perspective, an important concern with our studies is the risk of structural breaks in the relationship due to the length of our sample period. However, studying the safe haven effect involves investigating asset dynamics in extreme and rare situations, implying that data over extended periods is necessary to ensure that enough observations generated under extreme conditions are available.

4. **What can be learned from intra-day data about temporal dependencies and information processing in the foreign exchange (FX) market?**

   Using quantile autoregression techniques, we show that the intra-day FX market dynamics are characterized by state-dependence. More precisely, non-linear temporal dependence is present

in FX returns observed at various intra-day sampling frequencies from 10 minutes to 3 hours. Moderate changes in an FX rate exhibit statistically significant negative serial correlation, whereas strong intra-day appreciations and depreciations over a short time interval tend to be characterized by the absence of autocorrelation. This finding may be attributed to *diverging opinions* of traders regarding the impact of news on the direction of the FX rate, but also suggests a tendency for an immediate adjustment when unambiguous news – reflected by extreme FX rate returns – arrive. Further, we show that the dependence structure is not affected by diurnal volatility variation due to the level of trading activity and is stable with regard to increased daily FX return volatility.

In addition to a description of the non-linear dynamics, our quantile autoregression models can be used to estimate the entire (conditional) distribution of future returns. This allows to assess the magnitude of future FX rate appreciations and depreciations at different intra-day time horizons which could find immediate use in intra-day risk management applications since moves in the FX rate might affect the profitability of intra-daily trades. However, the use of predictions based on time series quantile regression models is limited in the sense that future quantiles are always unknown. Our approach, therefore, cannot be used to design an active trading strategy that aims to generate risk free profits and consequently, the presence of non-linear temporal dependence in FX rate returns should be interpreted carefully in the context of weak-form efficiency.

The central message behind all studies is identical: The dynamics of financial markets are characterized by state-dependence which must be taken into account in risk management as well as the design of forecasting models. From a statistical perspective, our findings suggest that results based on linear models may be highly inaccurate.

This thesis provides insights regarding the dynamics and interdependencies of different financial markets, but several issues still remain for future research: For instance, our findings suggest that the adjustment behaviour of crude oil prices is time-varying and that there is no benchmark price which universally acts as price setter. Therefore, investigating the individual relevance of different crude oil prices within the pricing process, using information shares in the sense of Hasbrouck (1995), would allow a more accurate assessment of the world crude oil price. Moreover, our investigation of commodity volatility could be extended to a multivariate setting and a broader set of commodities from different sectors (agriculture, energy, precious metals, rare earths). This would allow to study non-linear interdependencies of commodity volatility which might be of direct use in the context of portfolio diversification using commodities. Our studies on the relationship between the gold and equity market confirm the role of gold as a safe haven asset. Further research on the safe haven effect could focus on transaction and sales data to investigate which types of investors buy and which sell gold in response to a negative shock in the equity market. In this context, it seems also interesting to investigate how investors purchase gold. Finally, our research on the FX market could be extended to an analysis of the short-run dynamics in FX forward and futures markets, which is of interest for investors exposed to currency risk.

References

ADAMS, Z. AND T. GLÜCK (2015). "Financialization in commodity markets: A passing trend or the new normal?" *Journal of Banking and Finance*, **60**(Supplement C): 93–111.

ADELMAN, M. A. (1984). "International oil agreements," *The Energy Journal*, **5**(3): 1–9.

AEPLI, M. D., R. FÜSS, T. E. S. HENRIKSEN, AND F. PARASCHIV (2017). "Modeling the multivariate dynamic dependence structure of commodity futures portfolios," *Journal of Commodity Markets*, **6**: 66–87.

ANDERSEN, T. G. AND T. BOLLERSLEV (1998). "Answering the skeptics: yes, standard volatility models do provide accurate forecasts," *International Economic Review*, **39**(4): 885–905.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND H. EBENS (2001a). "The distribution of realized stock return volatility," *Journal of Financial Economics*, **61**(1): 43–76.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND P. LABYS (2001b). "The distribution of realized exchange rate volatility," *Journal of the American Statistical Association*, **96**(453): 42–55.

——— (2003). "Modeling and Forecasting Realized Volatility," *Econometrica*, **71**(2): 579–625.

ANDERSEN, T. G., T. BOLLERSLEV, P. FREDERIKSEN, AND M. Ø. NIELSEN (2010). "Continuous-time models, realized volatilities, and testable distributional implications for daily stock returns," *Journal of Applied Econometrics*, **25**(2): 233–261.

ANDERSEN, T. G., D. DOBREV, AND E. SCHAUMBURG (2012). "Jump-robust volatility estimation using nearest neighbor truncation," *Journal of Econometrics*, **169**(1): 75–93.

AWIROTHANANON, T. AND W.-K. A. CHEUNG (2009). "On Joint Determination of the Number of States and the Number of Variables in Markov-Switching Models: A Monte Carlo Study," *Communications in Statistics - Simulation and Computation*, **38**(8): 1757–1788.

BAILLIE, R. T. AND T. BOLLERSLEV (1991). "Intra-Day and Inter-Market Volatility in Foreign Exchange Rates," *Review of Economic Studies*, **58**(3): 565–585.

BANK FOR INTERNATIONAL SETTLEMENTS (2016). "Triennial Central Bank Survey of foreign exchange and OTC derivatives markets in 2016," *Technical Report*.

BARBERIS, N., A. SHLEIFER, AND R. VISHNY (1998). "A model of investor sentiment," *Journal of Financial Economics*, **49**(3): 307–343.

BARNDORFF-NIELSEN, O. E. AND N. SHEPHARD (2002). "Estimating quadratic variation using realized variance," *Journal of Applied Econometrics*, **17**: 457–477.

BATTEN, J., B. LUCEY, F. MCGROARTY, M. PEAT, AND A. URQUHART (2017). "Stylized facts of intraday precious metals," *PLoS ONE*, **12**(4).

BAUR, D. AND K. KUCK (2019). "The Timing of the Flight to Gold: An Intra-day Analysis of Gold and the S&P500," *Finance Research Letters*.

BAUR, D. G. (2012). "Asymmetric Volatility in the Gold Market," *The Journal of Alternative Investments*, **14**(4): 26–38.

——— (2013). "The structure and degree of dependence: A quantile regression approach," *Journal of Banking & Finance*, **37**(3): 786–798.

BAUR, D. G. AND T. DIMPFL (2012). "State-dependent momentum in international stock markets," Available at SSRN.

——— (2018). "A Quantile Regression Approach to Estimate the Variance of Financial Returns," *Journal of Financial Econometrics*.

BAUR, D. G., T. DIMPFL, AND R. C. JUNG (2012). "Stock return autocorrelations revisited: A quantile regression approach," *Journal of Empirical Finance*, **19**(2): 254–265.

BAUR, D. G. AND K. J. GLOVER (2012). "The Destruction of a Safe Haven Asset?" *Applied Finance Letters*, **1**(1): 8–15.

BAUR, D. G. AND B. M. LUCEY (2010). "Is Gold a Hedge or a Safe Haven? An Analysis of Stocks, Bonds and Gold," *Financial Review*, **45**(2): 217–229.

BAUR, D. G. AND T. K. MCDERMOTT (2010). "Is gold a safe haven? International evidence," *Journal of Banking & Finance*, **34**(8): 1886–1898.

——— (2016). "Why is gold a safe haven?" *Journal of Behavioral and Experimental Finance*, **10**: 63–71.

BENTZEN, J. (2007). "Does OPEC influence crude oil prices? Testing for co-movements and causality between regional crude oil prices," *Applied Economics*, **39**(11): 1375–1385.

BERA, A. K., A. GALVAO, AND L. WANG (2014). "On Testing the Equality of Mean and Quantile Effects," *Journal of Econometric Methods*, **3**(1): 47–62.

BIANCO, S. AND R. RENÒ (2006). "Dynamics of intraday serial correlation in the Italian futures market," *Journal of Futures Markets*, **26**(1): 61–84.

BONACCOLTO, G. AND M. CAPORIN (2014). "Modelling and Forecasting the Realized Range Conditional Quantiles," *Available at SSRN.*

BUBÁK, V., E. KOČENDA, AND F. ŽIKEŠ (2011). "Volatility transmission in emerging European foreign exchange markets," *Journal of Banking & Finance*, **35**(11): 2829–2841.

BUYUKSAHIN, B. AND J. HARRIS (2011). "Do Speculators Drive Crude Oil Futures Prices?" *The Energy Journal*, **32**(2): 167–202.

CAMPBELL, J. Y., A. W. LO, AND A. C. MACKINLAY (1997). *The Econometrics of Financial Markets*, Princeton University Press.

CANDELON, B., M. JOËTS, AND S. TOKPAVI (2013). "Testing for Granger causality in distribution tails: An application to oil markets integration," *Economic Modelling*, **31**: 276–285.

CARCHANO, O. AND A. PARDO (2009). "Rolling over stock index futures contracts," *Journal of Futures Markets*, **29**(7): 684–694.

CHARLES, A. AND O. DARNÉ (2009). "Testing for random walk behavior in Euro exchange rates," *Economie Internationale*, (119): 22–45.

CHARLES, A., O. DARNÉ, AND J. H. KIM (2012). "Exchange-rate return predictability and the adaptive markets hypothesis: Evidence from major foreign exchange rates," *Journal of International Money and Finance*, **31**(6): 1607–1626.

CHENG, I.-H. AND W. XIONG (2014). "Financialization of Commodity Markets," *Annual Review of Financial Economics*, **6**(1): 419–441.

CHERNICK, M. R. (2008). *Bootstrap methods: A guide for practitioners and researchers*, vol. 619, Hoboke, NJ: John Wiley & Sons.

CHERNOZHUKOV, V. AND L. UMANTSEV (2001). "Conditional value-at-risk: Aspects of modeling and estimation," *Empirical Economics*, **26**(1): 271–292.

CHORTAREAS, G., Y. JIANG, AND J. C. NANKERVIS (2011a). "The random-walk behavior of the Euro exchange rate," *Finance Research Letters*, **8**(3): 158–162.

——— (2011b). "Forecasting exchange rate volatility using high-frequency data: Is the euro different?" *International Journal of Forecasting*, **27**(4): 1089–1107.

CORSI, F. (2009). "A simple approximate long-memory model of realized volatility," *Journal of Financial Econometrics*, **7**: 174–196.

DACOROGNA, M. M., R. GENCAY, U. MÜLLER, R. B. OLSEN, AND O. V. PICTET (2001). *An Introduction to High-Frequency Finance*, San Diego, CA: Academic Pr Inc.

[DATASET] JURADO, K. AND LUDVIGSON, S. C. AND NG, S. (2015). "Aggregate Uncertainty," "Measuring Uncertainty," *American Economic Review*, Additional Material.

[DATASET] NBER (2017). "US Business Cycle Expansions and Contractions (table). *International Energy Statistics*," `http://www.nber.org/cycles.html`, [Online; accessed 08-February-2017].

[DATASET] US ENERGY INFORMATION AGENCY (no date). "Total Petroleum and Other Liquids Production (table). *International Energy Statistics*," `http://www.eia.gov/beta/international/data/browser/`, [Online; accessed 08-February-2017].

DOMANSKI, D. AND A. HEATH (2007). "Financial Investors and Commodity Markets," *BIS Quarterly Review*, March 2007: 53–67.

ENGLE, R. AND S. MANGANELLI (2004). "CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles," *Journal of Business & Economic Statistics*, **22**: 367–381.

ERB, C. B. AND C. R. HARVEY (2013). "The Golden Dilemma," *Financial Analysts Journal*, **69**(4): 10–42.

ESCANCIANO, J. C. AND I. N. LOBATO (2009a). "An automatic Portmanteau test for serial correlation," *Journal of Econometrics*, **151**(2): 140–149.

——— (2009b). "Testing the Martingale Hypothesis," in *Palgrave Handbook of Econometrics: Volume 2: Applied Econometrics*, ed. by T. C. Mills and K. Patterson, London: Palgrave Macmillan UK, chap. 20, 972–1003.

EUROPEAN COMMISSION (2016). "Renewable Energy Directive," `https://ec.europa.eu/energy/en/topics/renewable-energy/renewable-energy-directive`, [Online; accessed 2016-08-05].

FAMA, E. F. (1965). "The Behavior of Stock-Market Prices," *The Journal of Business*, **38**(1): 34–105.

FATTOUH, B. (2006). "The origins and evolution of the current international oil pricing system: A critical assessment," in *Oil in the 21st Century - Issues, Challenges and Opportunities*, ed. by R. Mabro, Oxford: Oxford University Press, chap. 3, 41–101.

——— (2007). "WTI Benchmark Temporarily Breaks Down: Is It Really a Big Deal?" *Oxford Energy Comment*, 1–7.

——— (2010). "The dynamics of crude oil price differentials," *Energy Economics*, **32**(2): 334–342.

FEINSTONE, L. J. (1987). "Minute by minute: Efficiency, normality, and randomness in intra-daily asset prices," *Journal of Applied Econometrics*, **2**(3): 193–214.

GILBERT, C. L. (2010). "How to Understand High Food Prices," *Journal of Agricultural Economics*, **61**(2): 398–425.

GIOT, P. (2005). "Market risk models for intraday data," *The European Journal of Finance*, **11**(4): 309–324.

GIOT, P., S. LAURENT, AND M. PETITJEAN (2010). "Trading activity, realized volatility and jumps," *Journal of Empirical Finance*, **17**(1): 168–175.

GLOSTEN, L. R., R. JAGANNATHAN, AND D. E. RUNKLE (1993). "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *The Journal of Finance*, **48**(5): 1779–1801.

GOODHART, C. (1990). "News and the Foreign Exchange Market," FMG Discussion Papers dp71, Financial Markets Group.

GOODHART, C. A. E. AND L. FIGLIUOLI (1991). "Every minute counts in financial markets," *Journal of International Money and Finance*, **10**(1): 23–52.

GORTON, G. (2017). "The History and Economics of Safe Assets," *Annual Review of Economics*, **9**(1): 547–586.

GUELEN, G. (1999). "Regionalization in the World Crude Oil Market: Further Evidence," *The Energy Journal*, **20**(1): 125–140.

GUILLAUME, D. M., M. M. DACOROGNA, R. R. DAVÉ, U. A. MÜLLER, R. B. OLSEN, AND O. V. PICTET (1997). "From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets," *Finance and Stochastics*, **1**(2): 95–129.

HAMILTON, J. D. (1989). "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, **57**(2): 357–384.

——— (1994). *Time Series Analysis*, Princeton, NJ: Princeton University Press.

HAMMOUDEH, S. M., B. T. EWING, AND M. A. THOMPSON (2008). "Threshold cointegration analysis of crude oil benchmarks," *The Energy Journal*, **29**(4): 79–95.

HANSEN, P. R., J. LARGE, AND A. LUNDE (2008). "Moving Average-Based Estimators of Integrated Variance," *Econometric Reviews*, **27**(1-3): 79–111.

HANSEN, P. R. AND A. LUNDE (2006). "Consistent ranking of volatility models," *Journal of Econometrics*, **131**(1-2): 97–121.

HASBROUCK, J. (1995). "One Security, Many Markets: Determining the Contributions to Price Discovery," *The Journal of Finance*, **50**(4): 1175–1199.

HÖÖK, M., R. HIRSCH, AND K. ALEKLETT (2009). "Giant oil field decline rates and their influence on world oil production," *Energy Policy*, **37**(6): 2262–2272.

HSIEH, D. A. (1988). "The statistical properties of daily foreign exchange rates: 1974-1983," *Journal of International Economics*, **24**(1): 129–145.

INTERCONTINENTAL EXCHANGE (2016). "ICE Crude Oil," `https://www.theice.com/publicdocs/ICE_Crude_Oil.pdf`, [Online; accessed 2016-08-05].

ITO, T. AND V. V. ROLEY (1987). "News from the U.S. and Japan," *Journal of Monetary Economics*, **19**(2): 255–277.

JI, Q. AND Y. FAN (2015). "Dynamic integration of world oil prices: A reinvestigation of globalisation vs. regionalisation," *Applied Energy*, **155**: 171–180.

JOHANSEN, S. (1988). "Statistical Analysis of Cointegration Vectors," *Journal of Economic Dynamics and Control*, **12**: 231–254.

——— (1991). "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models," *Econometrica*, **59**(6): 1551–1580.

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015). "Measuring Uncertainty," *American Economic Review*, **105**(3): 1177–1216.

KOENKER, R. AND G. J. BASSETT (1978). "Regression Quantiles," *Econometrica*, **46**(1): 33–50.

KOENKER, R. AND Z. XIAO (2006). "Quantile Autoregression," *Journal of the American Statistical Association*, **101**(475): 980–990.

KOENKER, R. AND Q. ZHAO (1996). "Conditional Quantile Estimation and Inference for ARCH Models," *Econometric Theory*, **12**(05): 793–813.

KROLZIG, H. M. (1997). *Markov-Switching Vector Autoregressions: Modelling, Statistical Inference, and Application to Business Cycle Analysis*, New York: Springer.

KUCK, K., R. MADERITSCH, AND K. SCHWEIKERT (2015). "Asymmetric over- and undershooting of major exchange rates: Evidence from quantile regressions," *Economics Letters*, **126**(C): 114–118.

LEBARON, B. (1992). "Some Relations Between Volatility and Serial Correlations in Stock Market Returns," *The Journal of Business*, **65**(2): 199–219.

LEE, J. AND M. C. STRAZICICH (2003). "Minimum Lagrange Multiplier Unit Root Test with Two Structural Breaks," *The Review of Economics and Statistics*, **85**(4): 1082–1089.

LIU, C. Y. AND J. HE (1991). "A Variance-Ratio Test of Random Walks in Foreign Exchange Rates," *The Journal of Finance*, **46**(2): 773–785.

LIU, L., C.-C. CHEN, AND J. WAN (2013). "Is world oil market 'one great pool'?: An example from China's and international oil markets," *Economic Modelling*, **35**: 364–373.

LIU, L. Y., A. J. PATTON, AND K. SHEPPARD (2015). "Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes," *Journal of Econometrics*, **187**(1): 293–311.

LIU, S. AND Y.-K. TSE (2015). "Intraday Value-at-Risk: An asymmetric autoregressive conditional duration approach," *Journal of Econometrics*, **189**(2): 437–446.

LIU, W. (2018). "Are Gold and Government Bond Safe-Haven Assets? An Extremal Quantile Regression Analysis," *International Review of Finance*, **0**(0).

LOBATO, I., J. C. NANKERVIS, AND N. E. SAVIN (2001). "Testing for Autocorrelation Using a Modified Box-Pierce Q Test," *International Economic Review*, **42**(1): 187–205.

LU, F., Y. HONG, S. WANG, K. LAI, AND J. LIU (2014). "Time-varying Granger causality tests for applications in global crude oil markets," *Energy Economics*, **42**: 289–298.

MANN, J. AND P. SEPHTON (2016). "Global relationships across crude oil benchmarks," *Journal of Commodity Markets*, **2**(1): 1–5.

MARTENS, M. AND D. VAN DIJK (2007). "Measuring volatility with the realized range," *Journal of Econometrics*, **138**(1): 181–207.

MARTENS, M. AND J. ZEIN (2004). "Predicting financial volatility: High-frequency time-series forecasts vis-à-vis implied volatility," *Journal of Futures Markets*, **24**(11): 1005–1028.

MIYAZAKI, T. (2019). "Clarifying the Response of Gold Return to Financial Indicators: An Empirical Comparative Analysis Using Ordinary Least Squares, Robust and Quantile Regressions," *Journal of Risk and Financial Management*, **12**(1).

MÜLLER, U. A., M. M. DACOROGNA, R. D. DAVÉ, R. B. OLSEN, O. V. PICTET, AND J. E. VON WEIZSÄCKER (1997). "Volatilities of different time resolutions – Analyzing the dynamics of market components," *Journal of Empirical Finance*, **4**(2-3): 213–239.

MÜLLER, U. A., M. M. DACOROGNA, R. B. OLSEN, O. V. PICTET, M. SCHWARZ, AND C. MORGENEGG (1990). "Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Scaling Law, and Intraday Analysis," *Journal of Banking & Finance*, **14**: 1189–1208.

NEELY, C. AND P. WELLER (2003). "Intraday technical trading in the foreign exchange market," *Journal of International Money and Finance*, **22**(2): 223–237.

OSTERWALD-LENUM, M. (1992). "A Note with Quantiles of the Asymptotic Distribution of the Maximum Likelihood Cointegration Rank Test Statistics," *Oxford Bulletin of Economics and Statistics*, **55**(3): 461–472.

PENNINGS, J. M. E. AND P. GARCIA (2010). "Risk and hedging behavior: The role and determinants of latent heterogeneity," *Journal of Financial Research*, **33**(4): 373–401.

PSARADAKIS, Z. AND N. SPAGNOLO (2006). "Joint Determination of the State Dimension and Autoregressive Order for Models with Markov Regime Switching," *Journal of Time Series Analysis*, **27**(5): 753–766.

RANALDO, A. AND P. SÖDERLIND (2010). "Safe Haven Currencies," *Review of Finance*, **14**(3): 385–407.

REBOREDO, J. C. (2011). "How do crude oil prices co-move? A copula approach," *Energy Economics*, **33**(5): 948–955.

——— (2013). "Is gold a hedge or safe haven against oil price movements?" *Resources Policy*, **38**(2): 130–137.

RODRIGUEZ, A. E. AND M. D. WILLIAMS (1993). "Is the World Oil Market 'One Great Pool'? A test," *Energy Studies Review*, **5**(2): 121–130.

SAMUELSON, P. A. (1965). "Proof that properly anticipated prices fluctuate randomly," *Industrial Management Review*, **6**(2): 41–49.

SCHARTH, M. AND M. C. MEDEIROS (2009). "Asymmetric effects and long memory in the volatility of Dow Jones stocks," *International Journal of Forecasting*, **25**(2): 304–327.

SCHWEIKERT, K. (2018). "Are gold and silver cointegrated? New evidence from quantile cointegrating regressions," *Journal of Banking & Finance*, **88**: 44–51.

SOUČEK, M. AND N. TODOROVA (2013). "Realized volatility transmission between crude oil and equity futures markets: A multivariate HAR approach," *Energy Economics*, **40**(C): 586–597.

TANG, K. AND W. XIONG (2012). "Index Investment and the Financialization of Commodities," *Financial Analysts Journal*, **68**(6): 54–74.

TODOROVA, N. (2015). "The course of realized volatility in the LME non-ferrous metal market," *Economic Modelling*, **51**(Supplement C): 1–12.

——— (2017). "The asymmetric volatility in the gold market revisited," *Economics Letters*, **150**: 138 – 141.

TODOROVA, N., A. WORTHINGTON, AND M. SOUČEK (2014). "Realized volatility spillovers in the non-ferrous metal futures market," *Resources Policy*, **39**(Supplement C): 21–31.

TONG, H. (1983). *Threshold models in Nonlinear Time Series Analysis*, New York, NY: Springer.

US ENERGY INFORMATION AGENCY (2016). "Spot Prices," `http://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm`, [Online; accessed 2016-08-05].

VERBEEK, M. (2017). *A guide to modern econometrics*, West Sussex, England: John Wiley & Sons, 5 ed.

VERONESI, P. (1999). "Stock market overreactions to bad news in good times: a rational expectations equilibrium model," *Review of Financial Studies*, **12**(5): 975–1007.

ŽIKEŠ, F. AND J. BARUNÍK (2016). "Semi-parametric Conditional Quantile Models for Financial Returns and Realized Volatility," *Journal of Financial Econometrics*, **14**(1): 185–226.

WANG, T., J. WU, AND J. YANG (2008). "Realized volatility and correlation in energy futures markets," *Journal of Futures Markets*, **28**(10): 993–1011.

WASSERFALLEN, W. AND H. ZIMMERMANN (1985). "The behavior of intra-daily exchange rates," *Journal of Banking & Finance*, **9**(1): 55–72.

WEINER, R. J. (1991). "Is the World Oil Market 'One Great Pool'?" *The Energy Journal*, **12**(3): 95–107.

——— (1993). "The World Oil Market is not 'One Great Pool': A Reply to Rodriguez and Williams," *Energy Studies Review*, **5**(3): 225–230.

WILMOT, N. A. (2013). "Cointegration in the Oil Market among Regional Blends," *International Journal of Energy Economics and Policy*, **3**(4): 424–433.

WLAZLOWSKI, S., B. HAGSTRÖMER, AND M. GIULIETTI (2011). "Causality in crude oil prices," *Applied Economics*, **43**(24): 3337–3347.

ZHANG, X. F. (2006). "Information Uncertainty and Stock Returns," *Journal of Finance*, **61**(1): 105–137.

ZHANG, Y.-J. AND L. ZHANG (2015). "Interpreting the crude oil price movements: Evidence from the Markov regime switching model," *Applied Energy*, **143**: 96–109.

ZHOU, B. (1996). "High-Frequency Data and Volatility in Foreign-Exchange Rates," *Journal of Business & Economic Statistics*, **14**(1): 45–52.