IMPERIAL COLLEGE LONDON
IMPERIAL COLLEGE BUSINESS SCHOOL

RETURN PREDICTABILITY AND OPTIMAL PORTFOLIO CHOICE: EVIDENCE FROM COMMODITY AND GLOBAL FUTURES MARKETS

JAN DANILO AHMERKAMP

SUBMITTED IN PART FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY OF IMPERIAL COLLEGE LONDON
AND THE DIPLOMA OF IMPERIAL COLLEGE LONDON

LONDON, UK
OCTOBER 2013

DECLARATION

I herewith certify that this thesis constitutes my own work and that all material, which is not my own work, has been properly acknowledged. The last chapter is a collaborative work with James B. Grant.

Jan Danilo Ahmerkamp

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

to my parents, Katha, Soeren and Lilli

in memory Oma

ABSTRACT

Time variation in expected returns is understood to be a common feature across aggregate asset classes as diverse as equities, currencies and bonds. Less is known on aggregate commodity return predictability. Findings of this thesis provide evidence of time variation of commodity returns and assesses its economic value within portfolio allocation strategies. Further findings demonstrate the effect aggregate hedge fund capital has on the profitability of carry and momentum strategies across global futures markets. Besides the financial implications, the findings of this thesis also affect policies with respect to international macroeconomics, poverty alleviation, energy and climate policy, and commodity market regulation.

First, I study variation of expected commodity return portfolios. I run regressions of 1 to 9 months portfolio holding returns on lagged average futures discounts (AFD). I find that the AFD predicts commodity portfolio returns with 9 months R^2 values of 10 percent. Most predictable variation is a result of spot premia variation, while term premia are only significantly time-varying on short term horizons. Variation in the AFD is procyclically related to macroeconomic conditions. The procyclical relation leads to strong return predictability: the AFD and US industrial production growth rates forecast up to 16 percent of the commodity portfolio holding return variation at the 9 months horizon.

Second, I study the statistical and economic value of macroeconomic, financial and commodity market specific variables in predicting commodity returns. I estimate the models within a data-rich Bayesian model averaging (BMA) framework. I find that commodity portfolio returns and volatility are predictable across commodity sectors. Posterior model probabilities reveal that most of the predictable variation in commodity returns is due to macroeconomic variables of industrial production growth and the variation of the aggregate commodity basis. Portfolio volatility is related to lagged dividend yield, default spread, and inflation growth. I further find that an investor will pay a high performance fee to switch from a dynamic portfolio strategy based on a simple autoregressive benchmark models to a BMA model. In contrast a conditional volatility strategy does not generate significant economic gains.

Third, in collaboration with James B. Grant, we provide evidence that hedge funds capital is a key determinant for the profitability of carry and momentum strategies in futures markets

across asset classes. We parameterize carry and momentum portfolios from the perspective of a utility maximizing risk averse investor. We find that the returns to optimal carry and momentum strategies yield high Sharpe ratios (above 1.2), which are not a compensation for traditional risk exposure or time-varying risk due to macroeconomic cycles or funding liquidity, however they are related to pro-cyclical hedge fund capital flows. Larger capital flows lead to higher carry and momentum returns, implying that expected returns decrease with the total amount of assets under management by hedge funds. We argue that these findings are consistent with the notion of limits to arbitrage.

ACKNOWLEDGMENTS

First, I would like to thank my supervisors, Dr. Robert Kosowski, Associate Professor at Imperial College Business School and Dr. Walter Distaso, Professor at Imperial College Business School for their guidance, continuous encouragement and patience throughout these years. It was not the easiest journey, but you always offered me your assistance. Thank you. Additionally, I am very grateful for the full scholarship that the Grantham Institute of Climate Change offered me to conduct my research.

Furthermore, I am grateful to have met so many great people during my time in London. I would like to thank my fellow PhD cohort, including Julia Arnold, Saad Badaoui, Mai Daher, Saeed Iqbal, Farouk Jivraj, Aleksejs Krecetovs, Lisa Majmin, Komkrit Owen Ovararin, Lorenzo Pitotti, Worrawat Gof Sritrakul, and Paul Whelan, for exchanging ideas, helping each other and making the long hours we spent together an fun task. I must in particular thank Andrea Carnelli, Jamie Grant, Andres Reibel, and Simon Schillebeeck for being great friends. We spend almost every free minute in London together. Thank you for sharing research and trading ideas, having drinks together (always too many), collaborating on articles and having some of the greatest time in my life. I will never forget it. Outside of London I would like to thank Stephan Gsell, Bastian Koenings, Dennis Ott and Caspar Schmelzer. I feel honored to have met you guys in high school and that you are still my best friends, even if we do not see each other a lot. Every hour we spent together is as if we always lived in one place. Thanks again to everyone I met, also if I did not mention him/her in this short paragraph.

I would also like to express my loving thanks to my family. I owe you everything. Thanks to my parents, Johanna and Juergen, for being the best parents in the world, supporting me during my whole life and helping me when I needed it. Thanks to Katha and Soeren for not only being my siblings but also my best friends.

Last I would like to take this moment to thank the most important person in my live: Lilli. Thank you for your love, support, encouragement, patience and countless hours of Skype and phone calls.

TABLE OF CONTENTS

DECLARATION ii

ABSTRACT vi

ACKNOWLEDGMENTS viii

LIST OF TABLES xi

LIST OF FIGURES xiii

1 INTRODUCTION AND LITERATURE REVIEW 1

2 TIME-VARYING EXPECTED RETURNS IN COMMODITY FUTURES PORTFOLIOS 14
- 2.1 Introduction 14
- 2.2 Definitions and Empirical Implementation 18
  - 2.2.1 Holding Returns 18
  - 2.2.2 Data 20
- 2.3 Time Variation in Holding Returns 26
  - 2.3.1 Predictability Results 26
  - 2.3.2 Average Futures Discount and Individual Futures Contracts 32
  - 2.3.3 Robustness 34
- 2.4 Time Variation in Spot and Term Premia 44
- 2.5 Macroeconomic Influences and Market Participants' Positions 51
  - 2.5.1 Properties of the Average Futures Discount 52
  - 2.5.2 Macroeconomic Factors, Average Futures Discount and Commodity Return Predictability 54
  - 2.5.3 Trading Position, Average Futures Discount and Commodity Returns 57
- 2.6 Conclusion 60

3 PREDICTING COMMODITY RETURNS AND VOLATILITY: A BAYESIAN APPROACH 61
- 3.1 Introduction 61
- 3.2 Data and Descriptive Statistics 65
  - 3.2.1 Commodity Data 65
  - 3.2.2 Predictors Data 69
- 3.3 Bayesian and Economic Framework 73
  - 3.3.1 Bayesian Setup 73
  - 3.3.2 Portfolio Choice and Economic Value 79
- 3.4 Empirical Results 83
  - 3.4.1 Bayesian Model Averaging: In-Sample Analysis 83
  - 3.4.2 Bayesian Model Averaging: Out-of-sample Analysis 95
- 3.5 Conclusion 109

4 OPTIMAL FUTURES PORTFOLIOS AND HEDGE FUND CAPITAL 110
- 4.1 Introduction 110
- 4.2 Portfolio Construction and Data 115
  - 4.2.1 Optimal Futures Portfolios with Predictable Returns 115

4.2.2 Portfolio Estimation . . . 118
4.2.3 Construction of Tradable Futures Return Series . . . 119
4.2.4 Transaction Costs . . . 119
4.3 Data . . . 120
4.3.1 Price Data . . . 121
4.3.2 Macroeconomic and Liquidity Data . . . 124
4.3.3 Hedge Fund Data . . . 125
4.4 Characterising Optimal Futures Portfolios . . . 126
4.4.1 Carry and Momentum Portfolio Returns Across Asset Classes . . . 126
4.4.2 Effects of Changing Risk Aversion . . . 134
4.4.3 Diversified Investor with Optimal Futures Portfolio . . . 136
4.5 Business Cycles, Limits to Arbitrage and Hedge Fund Capital . . . 140
4.5.1 Optimal Futures Returns with Transaction Costs . . . 141
4.5.2 Risk Factor Exposure . . . 143
4.5.3 Macroeconomic Influences and Funding Liquidity . . . 146
4.5.4 Hedge Fund Activity and Capital Flows . . . 150
4.6 Conclusion . . . 157

5 CONCLUSIONS, POLICY IMPLICATIONS AND FUTURE RESEARCH . . . 158
5.1 Appendix A . . . 164
5.1.1 Statistical Inference of optimal future portfolio coefficients . . . 164

LIST OF TABLES

2.1 Summary Statistics for Individual Commodities . . . 21
2.2 Summary Statistics for Portfolios of Commodity Futures . . . 24
2.3 Predicting Portfolio Front Roll and Excess Holding Returns with Average Futures Discount 31
2.4 Predictability of Individual Holding Returns Using Univariate and Average Futures Discount . . . 33
2.5 Predicting Commodity Portfolio Holding and Excess Holding Returns for Various Sub Samples . . . 35
2.6 Predicting Commodity Portfolio Holding and Excess Holding Returns with Alternative Construction Methods of the Aggregate Futures Discount . . . 37
2.7 Predicting Commodity Portfolio Holding and Excess Holding Returns with Average Futures Discount (AFD), Common Systematic and Commodity Specific Variables . . . 39
2.8 Economic and Out-of-Sample Significance of Forecasting Regressions . . . 43
2.9 Predicting Commodity Holding Returns and Spot Changes of Futures Portfolios with the Average Futures Discount (AFD) . . . 50
2.10 Average Futures Discount and Systematic Variables . . . 53
2.11 Predicting Commodity Portfolio Holding, Spot, Front Roll and Excess Holding Returns with Average Futures Discount and Macroeconomic Variables . . . 56
2.12 Commodity Trader Long and Short Positions and the Average Futures Discount . . . 59

3.1 Summary statistics for Commodity Portfolio Returns and Portfolio Volatility . . . 68
3.2 Summary statistics for predictive variables . . . 70
3.3 Predictive Regressions for Global Commodity Portfolio Returns and Portfolio Volatility using BMA estimation . . . 85
3.4 Top 5 Models, Variable Inclusions and Statistical Inference of Bayesian Estimation . . . 90
3.5 Predictive Regressions for Sector Commodity Portfolio Returns and Portfolio Volatility using BMA estimation . . . 92
3.6 Predictive Regressions for Characteristics Based Commodity Portfolio Returns and Portfolio Volatility using BMA estimation . . . 94
3.7 Out-of-sample Statistical Criteria for BMA Combination Forecasts of Global Portfolio Returns and Portfolio Volatility . . . 96
3.8 Out-of-sample Statistical Criteria for BMA Combination Forecasts of Sector and Characteristics Based Portfolio Returns and Portfolio Volatility . . . 98
3.9 Economic Value of Sector Portfolio Return and Portfolio Volatility Predictability using BMA and BBM Combination Forecasts . . . 100
3.10 Economic Value of Global Portfolio Return and Portfolio Volatility Predictability using BMA Combination Forecasts for Specific Predictor Classes . . . 104
3.11 Economic Value of Sector Portfolio Return and Portfolio Volatility Predictability in Subsample Analysis . . . 106
3.12 Economic Value of Global Portfolio Return and Portfolio Volatility Predictability using BMA Combination Forecasts . . . 108

4.1 Summary Statistics for Individual Futures Contracts . . . 122
4.2 Optimal Futures Strategy Portfolio Returns . . . 128
4.3 Varying Risk Aversion . . . 135
4.4 Diversification Benefit of the Optimal Futures Portfolio . . . 138
4.5 Optimal Futures Returns with Transaction Costs . . . 142

4.6 Risk Factors and Optimal Futures Portfolios 145
4.7 Macroeconomic and Funding Liquidity Influences on Optimal Carry and Momentum Portfolios 148
4.8 Return Decomposition of Hedge Fund Indices 152
4.9 Hedge Funds, Speculative Capital and the Optimal Carry and Momentum Strategy . . . 154

LIST OF FIGURES

2.1 Individual long run average holding returns and individual futures discounts 22
2.2 Nine months average futures discount and 9 months holding returns 25

3.1 Cumulative Returns of Global Portfolio Allocation Using BMA and BBM Forecasts and Comparisons 102
3.2 Dynamics of Portfolio Weights of Global Portfolio Allocation Using BMA and BBM Forecasts and Weights of Autoregressive Benchmark 103

4.1 Number of Available Futures Contracts over Time 123
4.2 Cross Sectional Summary Statistics of Carry and Momentum Signal 124
4.3 Out of Sample Parameter Estimates over Time (θ estimates) 131
4.4 Historical Performance of Carry, Momentum and Carry and Momentum 132
4.5 Rolling Average Returns and Draw-down Dynamics 133
4.6 Diversified Investor 139

CHAPTER 1
INTRODUCTION AND LITERATURE REVIEW

Understanding whether returns of financial assets are predictable lies at the heart of financial economic theory. In this thesis, I focus on studying predictive relations in commodity and global futures markets and test whether predictability can be exploited to generate economic profits within novel econometric and portfolio choice frameworks. I provide evidence on the time-varying behaviour of expected returns of commodity futures portfolios across futures maturities and the relation to macroeconomic fundamentals. I further investigate the relative statistical and economic value of a multitude of financial, macroeconomic and commodity-specific variables in predicting the returns and volatility of commodities within a Bayesian model averaging framework. Finally, I study the performance of carry and momentum strategies in global futures markets in a parametric portfolio choice framework to show that their excess returns are linked to the aggregate capital available to the hedge fund industry.

While early evidence of return predictability was interpreted as evidence against the **efficient market hypothesis**[1], predictability is by now understood to be directly linked to time-varying expected excess returns[2] (excellent summaries of the exhaustive literature are provided by Fama (1991) or more recently by Cochrane (2011)). Empirical evidence shows that movements in price-based valuation ratios and yields translate to expected excess return variation. Fama and French (1988), for example, show that dividend yields forecast returns, however not dividend growth. Cochrane (2008) and Binsbergen and Koijen (2010), motivated by the approximate present value identity of Campbell and Shiller (1988), find that cash flow growth and returns are predictable by the dividend yield, however most of the predictability is attributed to expected return variation. Subsequently, evidence on return predictability using yields and valuation ratios has been gathered in a variety of asset classes such as currencies, bonds, treasuries, sovereign debt, the housing

1. The efficient market hypothesis relates asset prices to information sets. In particular, it states that financial markets are efficient when prices aggregate all available information (see Fama (1965), Fama (1970); Jensen (1978); Shiller (1978); Summers (1986)). Testing the efficient market hypothesis requires a *market model* that specifies how information gets incorporated into asset prices.

2. In the following exposition expected excess returns, discount rates and risk premia are used synonymously.

market, and hedge funds[3]. The empirical in-sample evidence consistently reports statistically and economically significant coefficient estimates in predictive regressions and R^2 values that increase with forecasting horizon.

Several theories have been proposed to explain time variation in expected returns. A vast literature has generalized consumption-based frameworks to introduce time variation in discount rates[4]. These **macroeconomic theories** of discount rate variation predict that expected returns are high and prices are low at times when consumption, output, and investment are low, unemployment is high and industrial production is decreasing, and vice versa. Consequently, macroeconomic models link expected returns to macroeconomic data. Cooper and Priestley (2009) empirically confirm the counter-cyclical behaviour of expected excess returns for equity markets, showing that the industrial production gap predicts global stock market returns. Ludvigson and Ng (2009) show that the first principal components of a large set of macroeconomic data predict excess bond returns and Lustig et al. (2013) show that industrial production growth predicts returns of currency portfolios.

Theories that include financial frictions have become increasingly popular in understanding return predictability, especially since the financial crisis in 2007-2008. These **friction-based theories** try to understand the effects segmented markets, intermediaries and market liquidity have on discount rates. Theories of segmented and intermediated markets typically imply that the limited risk sharing ability of market participants leads to downward sloping demand or supply curves[5]. Consistent with these theories, Adrian et al. (2011) show that measures of intermediary balance sheet aggregates contain strong predictive power for excess returns on equities. These theoretical and empirical concepts have been extended to currency markets (Adrian et al. (2011)) and trea-

3. E.g., Fama and Bliss (1987) and Campbell and Shiller (1988) show that a rising yield curve of treasuries signals better 1-year returns for long-term bonds and not higher interest rates, Fama (1986) show that credit spreads over time and across firms predicts returns and not default probabilities, Hansen and Hodrick (1980) and Fama (1984) show that international interest rate spreads predict returns and not exchange depreciation, an effect underlying the famous carry trade, and Avramov et al. (2011) show that various macroeconomic variables predict hedge fund returns.

4. Campbell and Cochrane (1999) introduce models with time-varying risk aversion, Bansal and Yaron (2004) and Bansal et al. (2010) advocate time-varying aggregate consumption risk, Gabaix (2008) and Gabaix (2012) introduce time-varying consumption disasters, Lustig and Van Nieuwerburgh (2005) model time variation in risk sharing opportunities among heterogeneous agents and Timmermann (1993) introduces time variation in believes.

5. He and Krishnamurthy (2012) show that in asset markets where a financial intermediary is the marginal investor, tightening capital constraints of the intermediaries lead to increasing risk premia. Similarly, Duffie and Strulovici (2012) show that asset markets that differ only by the committed capital earn lower risk premia if less capital is available. Intermediaries engage in these markets, against a fee, to move capital from one market with the lower to the market with the higher returns and amount of capital.

suries (Vayanos and Vila (2009), Krishnamurthy and Vissing-Jorgensen (2012) as well as different classes of intermediaries (Jylhä and Suominen (2011)). Furthermore, aggregate market liquidity is a priced risk factor in equity and currency markets (Pastor and Stambaugh (2002) and Mancini et al. (2013)) and consequently aggregated measures of market liquidity predict returns. The effect of market liquidity on asset prices is tightly linked to segmented and intermediated market theories. Brunnermeier and Pedersen (2009), for example, show that if financial market participants or intermediaries are increasingly capital constraint, the unwinding of trading positions leads to higher trading costs, affecting asset returns, and subsequently even tighter capital constraints; an effect termed *liquidity spirals*. Brunnermeier et al. (2008) document liquidity spirals empirically in currency markets. In summary, macroeconomic and financial friction-based theories can explain predictability in asset returns within a rational expectations model.

An alternative view to explain return predictability are **behavioural theories** of asset pricing. Behavioural theories explain return predictability by psychological biases (Barberis and Thaler (2003)). Hong and Stein (1999), for example, develop a theory of underreaction and delayed overreaction where information are being included in prices only gradually leading to momentum and reversal effects. Hong et al. (2007) show empirically that information contained in industry returns diffuses only gradually into aggregate stock markets. Baker and Wurgler (2006) construct a market wide sentiment index to show that shocks to aggregate sentiment impact stocks that are difficult to arbitrage. Limits to arbitrage lead then to cross-sectional return differences. Stambaugh et al. (2011) show that equity anomalies, such as momentum returns, are stronger when following high levels of sentiment and vice versa. The differentiation between behavioural and discount rate theories are, however, blurred. Behavioural theories can be embedded within discount rate theories by assuming that market participants systematically have "wrong" expectations (Cochrane (2011)). Systematically "wrong" expectations distort the probabilities to discount future payoffs and are equivalent to psychological biases. In particular, Cochrane (2011) states that it is pointless to argue rational versus behavioural theories, because saying that "The market went up, risk aversion must have declined" is as vacuous as saying "the market went up, sentiment must have increased".

Return predictability in commodity markets, embedded within the framework of financial eco-

nomic theory, has received less attention in the academic literature[6]. Interestingly, from a theoretical perspective, commodity return predictability and time-varying risk premia predates the literature in other asset classes. Keynes (1930) and Hicks (1939) argue that futures prices should embed a risk premium to entice speculators to offset the spot price risk that commodity producers (hedgers) are facing. This Theory of Normal Backwardation postulates that aggregate futures positions of hedgers are net short and the risk premium increases or decreases with the amount of net short positions. Early empirical evidence such as by Working (1949), Telser (1958), and Dusak (1973) do not find evidence of time-varying risk premia in commodity futures, while, more recently, De Roon et al. (2000) find some evidence on the "hedging pressure effect". Fama and French (1987) are the first to link time variation in expected returns to variation in the futures discount or yield (difference between log spot price and log futures price), extending the evidence from other asset classes to commodity futures. They find that the individual futures discount has some power in explaining expected return variation for a variety of individual futures contract. However, they also note that most of futures discount variation on an individual futures level is related to idiosyncratic sources. In the context of storage models the commodity futures basis is affected by expected changes in the spot price of the underlying commodity and therefore, idiosyncratic factors play a role in explaining the basis variation, e.g., Deaton and Laroque (1996) and Routledge et al. (2002). The resulting high idiosyncratic volatility makes the statistical results difficult to interpret. To reduce the idiosyncratic sources of time variation, Szymanowska et al. (2013) and Gorton et al. (2007) study cross-sectional commodity sorts based on the futures discount and find statistically significant return premia, confirming the presence of time-varying risk in commodity markets.

While the futures discount is a direct ex-ante measure of expected return variation in commodity markets, various authors suggest that price-based measures of expected returns from other asset classes should have predictive power for expected commodity returns, e.g., Bessembinder and Chan (1992) and Bailey and Chan (1993). From an asset-pricing perspective, it is systematic risk in the economy that should lead to higher expected returns and therefore should be common across

6. Much of the literature on commodity markets has focussed on the effect inventory decisions of commodity producers have on futures and spot prices (Deaton and Laroque (1996) and Routledge et al. (2002)), on the link between exhaustible resources, production decisions and commodity prices (Carlson et al (2007), the effect of supply and demand imbalances on commodity prices (Kilian (2009)), and the role of commodity price shocks on the macroeconomy (Kilian (2007) and Kilian and Lewis (2011)).

asset classes. Bailey and Chan find that basis variation of commodities is indeed related to common systematic variables, while it is not clear how the variation is reflected in expected returns. Bessembinder and Chen show that expected returns of some futures contracts can be predicted by the dividend yield, term spread, default spread, and short rate. Bakshi et al. (2011) show that changes in the baltic dry index predict the returns to commodity indices. Lastly, Gargano and Timmermann (2013) show that various macroeconomic aggregates have statistical in- and out-of-sample predictive value for spot price return indices.

Recently, friction based theories have also been applied to commodity markets. Acharya et al. (2013) and Etula (2014) show theoretically that at times the limited risk sharing ability of speculators in commodity markets cannot offset the hedging demand by commodity produces. The authors show empirically, that measures of aggregate speculative capital predicts expected commodity returns, consistent with the theoretical predictions. In a related vain, Hong and Yogo (2013) show that open interest predicts commodity returns in the presence of hedging pressure and downward-sloping demand in futures markets.

In **Chapter 2** of this thesis I contribute to the literature on commodity return predictability and study the time variation in expected returns for commodity futures portfolios. Studying return predictability on a portfolio level helps to diversify the idiosyncratic noise of return variation and focuses on the aggregate and systematic sources of risk that are central to the modern notion of discount rate variation. I, therefore, complement and extend the recent literature that focusses on the behaviour of commodity returns from an investment perspective (see Rouwenhorst and Tang (2012) for a review). Gorton and Rouwenhorst (2004) and Erb and Harvey (2006) report that unconditional risk premia for diversified commodity portfolios are similar to those in equity markets. I focus on conditional or time-varying risk premia by decomposing the expected return of a commodity futures portfolio, that includes energy, livestock, agricultural, metal and precious metal commodities, into the average futures discount (AFD) (mean of futures discounts across futures contracts in the portfolio) and expected spot price changes. The AFD is a model-free component of the expected returns that can be observed ex-ante. Time variation of AFD does, however, not immediately imply time-varying risk premia. The AFD can change even if expected returns are constant when the change is offset by an equal but opposite-signed expected spot price

appreciation.

To test the relation between AFD and expected returns, I follow Fama and French (1987), Fama (1984) and Lustig et al. (2013) and run complimentary regression of 1 to 9 months excess portfolio returns and average price changes on the AFD. I find that the AFD predicts commodity portfolio excess returns with R^2 values of more than 9 percent at a 9 months horizon. The results for increasing horizons reveal interesting dynamics of the coefficient estimates: Coefficients for short-term horizons are significantly smaller than 1, implying a negative relation between futures discount and spot changes, while coefficient estimates for longer horizons are 1 or larger than 1, implying a one-to-one relation between expected return variation and futures discount variation. The negative relation between futures discount and future spot price changes at short horizons is a consequence of The Theory of Storage. Low inventory levels lead to high futures discounts and high contemporaneous spot prices and consequently low expected spot changes. These results are robust to various time horizons, alternative construction methods of the predictive variable, and out-of-sample regressions. In addition, investors profit from significant economic gains when conditioning commodity allocation on the futures discount.

In the next section, following Szymanowska et al. (2013), I decompose portfolio holding returns into spot premia and term premia by taking long positions in short maturity futures contracts and combining long and short positions in contracts with different maturities, respectively. Spot premia are associated with the risk in the underlying spot price and term premia are associated with risk in the convenience yield analog to risk premia in the term structure of bonds, see Cochrane and Piazzesi (2002). I find that time variation in expected returns is mostly associated with time-varying spot premia, while term premia satisfy a version of the expectation hypothesis with significant unconditional risk premia, however no time-variation.

These empirical results imply that commodity portfolio returns vary over time with movements in the AFD. Given the lack of theoretical and empirical evidence, it is not clear ex-ante, whether the time variation can be explained by macroeconomic theories or commodity market specific sources of risk (e.g., hedging pressure (De Roon et al. (2000)), inventory levels (Gorton et al. (2007)) and open interest (Hong and Yogo (2013))) explain this time variation. More recently, Baker and Routledge (2012) and Casassus et al. (2005) developed equilibrium models of oil price and risk premia dynamics showing that time variation in expected oil risk premia is related to macroeconomic conditions.

Motivated by these finding, I further test, whether macroeconomic and systematic variables are also related to the variability of the AFD in my sample and consequently predict expected holding returns. I find that commodity portfolio returns are pro-cyclically related to macroeconomic activity because the AFD is statistically significantly correlated with macroeconomic and common systematic variables. Furthermore, macroeconomic variables contain information about the cyclical behaviour above and beyond the information contained in the AFD, and therefore, predict expected spot price growth. Including US industrial production growth (IP growth) and the AFD in the predictive regression increases the adjusted R^2 values for a 9 months horizon to 16 percent. An interesting observation is that term premia are counter-cyclical, however statistically insignificant. When analysing the role of trading activity on commodity portfolio returns, I do not find any indication that hedgers' or speculators' positions, as declared by the Commodity Futures Trading Commission (CFTC), predict expected returns.

Despite the evidence for time-varying expected returns in commodity futures portfolios in Chapter 2 and the reviewed evidence on time variation in other asset classes, return predictability remains a subtle feature of the data. The forecasting relationship between financial ratios and future stock returns exhibits several concerning statistical features. Stambaugh (1999) and Nelson and Kim (1993) argue that the near-unit root behaviour of the lagged regressors and its innovations, that are highly correlated with returns of the endogenous variable, lead to special finite sample properties of the ordinary least square estimator. These properties lead to a substantial deviation of the standard errors that depart from the ordinary least square assumptions. To avoid this *small sample bias*, Ang and Bekaert (2007) and Campbell and Yogo (2006) derive asymptotic distributions for coefficients under the assumptions that the forecasting variable follows a local-to-unit root and still find evidence for statistically significant predictability. Furthermore, Bossaerts and Hillion (1999), Goyal and Welch (2003), Welch and Goyal (2008) and Inoue and Kilian (2006) show that the significant in-sample predictive results, do not necessarily produce significant **out-of-sample** test statistics. Out-of-sample tests are, however, important for real-time investors, who make investment decisions based on data that is only available up to this point in time. While several statistical measures have been proposed to assess the predictive out-of-sample value, more recently, various authors suggest the use of portfolio investment strategies to determine the **eco-**

nomic value of asset return predictability. Della Corte et al. (2009) show that futures discount forecasts of currencies add economic value to a portfolio relative to a random walk benchmark, while macroeconomic variables do not. Thornton and Valente (2012) show that neither the Fama and Bliss (1987) forward-spot spread nor the Cochrane and Piazzesi (2002) forecast utilizing the full term structure of forward rates adds economic value to portfolio investment strategies relative to the random walk benchmark. This is a surprising result given that Cochrane and Piazzesi (2002) document highly significant in- and out-of-sample test statistics. Recently, McCracken and Valente (2012) find marginal evidence of economic value of predictability using various predictors in equity return premia.

An even more important data issue is that the specification of the "correct" predictive model is not trivial. The large amount of predictors resulting from the many theoretical models, that are mostly not explicit about which variables should enter the regressions, makes the empirical evidence difficult to interpret. Some predictive variables can be statistically significant based on a particular collection of variables, but not on a competing one. Furthermore, in regression frameworks, it is difficult to assess the relative value a particular variable adds relative to another variable. Recently, various authors have employed Bayesian model averaging (BMA) methodologies to overcome these problems (Avramov (2002), Cremers (2002) and Wright (2008)). The BMA framework is attractive because it incorporates model uncertainty and is therefore robust to model misspecifications. In particular, the BMA methodology allows for inclusion of an arbitrary number of variables that are weighted according to the posterior probabilities of inclusion. The posterior probabilities can be further used to assess the relative value a particular variable adds to the overall predictive power.

In **Chapter 3**, I use the BMA methodology to examine the relative statistical importance and economic value of macroeconomic, financial and commodity-specific variables in predicting commodity portfolio returns and volatility. The BMA framework allows to include a rich set of variables in the linear predictive model without specifying, ex-ante, the correct variables or forecasting model. In contrast, the BMA methodology assigns a priori equal probabilities to each potential model and further assigns posterior probabilities of inclusion only after having observed all potential model combinations. This characteristic of the BMA framework is particularly desirable when studying commodity markets with ambiguous economic theory regrading the sources of return and volatility

variation. In particular, it remains difficult to determine whether commodities are integrated within traditional asset classes, influenced by commodity market-specific sources of risk, or a combination of both (see discussion above). Furthermore, I study portfolio volatility predictability, which is a separate task to return predictability. For the choice of potential predictors, I rely on previous research on commodity markets and do not propose any new set of variables.

I rank model specifications and consequently the predictive variables by computing posterior probabilities of each model and estimate the BMA weighted coefficients and corresponding t-ratios adjusted for model uncertainty. The ranking reveals that macroeconomic variables, which proxy for global and US IP growth, are key variables to predict commodity portfolio returns. Furthermore, the commodity-specific futures discount has some additional power in predicting returns according to the cumulative posterior inclusion probability. Given that the futures discount is an ex-ante measure of expected returns, suggests that important sources of return variation have not been captured by the variables included in this study. Commodity portfolio volatility, on the other hand, is predictable by inflation growth, dividend yield and default spread in addition to the lagged portfolio volatility. Macroeconomic variables do not have any significant posterior inclusion probabilities for volatility predictions. This finding suggests that commodity portfolio returns and portfolio volatility are driven by different sources of risk. I find very similar results for commodity sector portfolios. The characteristics based portfolios that are constructed according to the individual momentum and basis signal are not predictable by the variables considered in this study.

The in-sample results do not guarantee significant out-of-sample predictive results for a real-time forecaster or investor. Therefore, I assess the value of real-time predictability using statistical out-of-sample measures as well as portfolio investment strategies. I find that the statistical measures consistently report significant out-of-sample predictability for commodity portfolio returns for global and sector portfolios. Predicting portfolio volatility does not result in significant out-of-sample results relative to an autoregressive benchmark model. For the economic assessment, I employ mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility, which allows me to quantify how risk aversion affects the economic value of predictability, building on empirical studies of volatility timing in stock returns by Fleming et al. (2001) and Marquering and Verbeek (2004). I find that the BMA portfolio return predictions generate high economic value relative to an investor, who uses autoregressive models to forecast mean returns for

varying risk aversion and volatility levels. This predictive power is robust to reasonable transaction costs. The volatility predictions, on the other hand, consistent with the statistical out-of-sample results, do not add any economic value relative to the autoregressive benchmark model.

In summary, I find that commodity portfolio returns and volatility is predictable, however with different sources of risk attribution. The in-sample results extend to out-of-sample significance for commodity return predictability but not to volatility predictability. Furthermore, the return predictability generates significant economic value.

In the previous sections, I described aggregate sources of return variation of commodity futures portfolios and other asset classes. A large complimentary stream of literature suggests that characteristics of individual assets lead to significant predictable variation in cross-sectional sorts. Fama and French (1993) and Fama and French (1996) show that sorting stocks by their book-to-market ratio and size leads to significant return premia. Jegadeesh and Titman (1992) and Jegadeesh and Titman (2002) show that stocks that recently outperformed their peers will continue to do so. A characteristic that has been termed *momentum*. The characteristics-based approach has been extended to other asset classes. Lustig and Verdelhan (2008) show that individual interest rate differentials, which are equivalent to the futures discounts or carry in currency markets, predict currency excess returns. Menkhoff et al. (2012) find evidence for momentum in a large set of different currencies. Szymanowska et al. (2013) and Gorton et al. (2007) show that carry and momentum sorts also work in commodity markets[7]. More recently, Moskowitz et al. (2011), Asness et al. (2013) and Koijen et al. (20012) extend the evidence of carry and momentum strategies across asset classes. While the literature on characteristics based strategies and return premia is exhaustive, the literature lacks theoretical explanation[8].

In **Chapter 4**, which was written in collaboration with James B. Grant, we study optimal carry and momentum strategies and their relation to hedge fund capital. We extend the analysis

7. Further evidence is provided by Miffre and Rallis (2007), Fuertes et al. (2010) and Fuentes (2013).

8. Theoretical explanations for the return premia, generally fall into the categories, rational risk-based (e.g. Fama and French (1996), Jegadeesh and Titman (2002), Sarno et al. (2012)), characteristics-based models (e.g. Moskowitz and Grinblatt (1999)), and explanations based on limits to arbitrage (e.g. Jylhä and Suominen (2011), Lesmond et al. (2004)).

from commodity to global futures markets, including equity, bond and currency futures. We do so because the power of looking at various asset classes and strategies together greatly improves the ability to identify common factor structures (see Asness et al. (2013)). Furthermore, futures contracts across various asset classes represent an excellent laboratory for implementing dynamic strategies. In contrast to stock markets, most futures markets are very liquid, feature high transaction volumes, and have very low transaction costs. By the very nature of futures contracts, taking long or short positions is a natural process for investors, who only maintain a margin account that makes a dynamic use of leverage easily attainable. Furthermore, futures markets are mainly populated by sophisticated institutional investors such as hedge funds and investment banks, who do not face short selling constraints, which would prevent market participants from fully exploiting pricing anomalies, as investigated in Stambaugh et al. (2011).

Here, we study carry and momentum strategies within the framework of parametric portfolio choice problems by Brandt et al. (2009). Most research on characteristic-based asset allocation strategies resort to equally- or volatility-weighted portfolios. Even though simple trading rules tend to outperform more complex ones (DeMiguel et al. (2009) and Jacobs et al (2010)), they also abstract from real investors challenges. Investors have to choose the weight they apply to each signal, include transaction costs in their asset allocation process, and be able to adjust their level of risk aversion in terms of their utility function. The parametric portfolio choice framework accommodates all these uncertainties in a unifying framework. Furthermore, in contrast to Garleanu and Pedersen (2009), who also introduce portfolio allocation strategies with transaction costs and predictable returns, the framework of Brandt et al. (2009) does not require to model or estimate the many moments of the return generating process. In contrast, we model the portfolio weights of each asset as a function of the asset's momentum or carry signal and take the perspective of a risk averse utility maximising investor with CRRA preference structure. The advantage of CRRA preferences, in contrast to mean/variance portfolios, is that it penalizes higher moments of the portfolio return process; a desirable property when studying carry and momentum strategies (Brunnermeier et al. (2008) and Daniel (2011)).

By characterizing the optimal carry and momentum strategy, we find that the out-of-sample parametric portfolio policy generates a Sharpe ratio of 1.2, which corresponds to a certainty-equivalent (CE) gain relative to the risk-free rate of an annualized 15 percent. The CE and Sharpe

ratio are 8 percent and 0.5 higher, respectively, than the performance of a volatility-weighted long-only portfolio. We find that the optimal out-of-sample portfolio weight of the momentum signal is 1.5 times the carry signal. Nevertheless, both signals contribute statistically significantly to the performance of optimal carry and momentum strategies. Increasing the level of risk aversion reduces the allocation to carry and momentum strategies, however, the relative importance of carry and momentum signal in the portfolio remains constant. To analyze transaction costs, we modify the return equation and re-run the portfolio optimisations. For both, a conservative flat one-way trading cost of 0.05 percent and the market based bid-ask spread cost, we find no material change in neither the optimal portfolio nor its returns. Furthermore, the significant return premia are not explained by standard (Fama and French (1992) and Pastor and Stambaugh (2002)) or non-standard (Asness et al. (2013)) risk factors. Motivated by these findings, we study the diversification benefit of including optimal carry and momentum strategies in portfolios of stock market factors and bonds, as in Barroso and Santa-Clara (2013) and Burnside et al. (2006). We find that including optimal carry and momentum strategies in an optimized portfolio increases the Sharpe ratio on average by 0.7, out-of-sample. Even more important, including these strategies reduces the skewness and kurtosis of the original portfolios and therefore minimizes the overall risk of the portfolios.

We test two possible explanations for this puzzling performance of the carry and momentum strategies and find that capital frictions rather than business cycle risk is a key determinant. First, we investigate if the returns to the strategies capture fluctuations in macroeconomic and liquidity risk. Bessembinder and Chan (1992) find some evidence that macroeconomic variables explain return fluctuations of long-only futures strategies. Asness et al. (2013) find that value and momentum strategies, which include futures markets, are related to liquidity risk. Regressing lagged variables of macroeconomic activity and liquidity proxies (Chordia and Shivakumar (2002)) on the strategy returns, we find that macroeconomic variables are pro-cyclically related to future returns, while liquidity risk does not show any significant exposure. The pro-cyclical behaviour, however, contrasts the theories on discount rate variation.

The second approach to explain our results relies on capital market frictions, limits to arbitrage and hedge fund activity. First, we find that the returns to optimal carry and momentum strategies and a aggregate market factor (MSCI) explain more than 40 percent of aggregate hedge fund index returns. Furthermore, our strategies explain a larger fraction of hedge fund returns, and further

reduce the statistical significance of the residual alpha, relative to factors by Fung and Hsieh (2004). Our results have important applications in the literature determining alpha contribution of hedge funds (Avramov et al. (2012), Kosowski et al. (2007), Buraschi et al. (2013) and references therein). Second, having determined that hedge funds indeed follow carry and momentum strategies, we follow the approach of Jylha et al. (2010), Jylhä and Suominen (2011) and Baltas and Kosowski (2013), and regress optimal carry and momentum returns on proxies of aggregate hedge fund capital and capital flow into the hedge fund industry. The theoretical literature on capital frictions and market segmentation implies that capital flow to hedge funds and consequently to carry and momentum strategies should contemporaneously lead to increasing returns due to downward sloping demand curves and lead to lower expected returns relative to the aggregate capital available to the hedge fund industry. Our empirical investigation confirms the theoretical predictions. We find that hedge fund capital flows are contemporaneously positive related to the optimal strategy returns, while the predictive coefficient of aggregate hedge fund capital is negatively related to the carry and momentum strategies. For robustness we run complimentary regressions with macroeconomic and liquidity proxies as control variables. We find that the hedge fund capital variables completely drive out the predictability of the macroeconomic and liquidity variables. These observations suggest that the pro-cyclical behaviour of the carry and momentum returns is related to hedge fund capital and not macroeconomic risk. If we regress long-only strategies on hedge fund capital and macroeconomic variables, we find that macroeconomic risk is counter-cyclically related to the returns, consistent with theoretical predictions, while hedge fund capital does not influence the performance. In summary we find that hedge fund activity is a strong determinant for the performance of optimal carry and momentum strategies.

In chapter 5 of this thesis I conclude, highlight policy implications of this research, and outline areas of further research.

CHAPTER 2
TIME-VARYING EXPECTED RETURNS IN COMMODITY FUTURES PORTFOLIOS

2.1 Introduction

Do expected commodity portfolio returns vary over time in a statistical and economical meaningful way? If so, is the return variation related to macroeconomic and commodity market specific state variables? Recent literature in financial economics suggests that predictable time series variation of aggregate portfolio returns is a common feature across asset classes as diverse as equities, bonds, currencies or the housing market (Cochrane (2011)). The main objective of this chapter is to systematically investigate predictable variation of commodity futures portfolios and test if commodity futures markets feature the same characteristics of common asset classes.

This study complements the work by Fama and French (1987) who investigate time-variation in expected returns for individual commodities. They find that the individual futures discount, the difference between logarithmic spot and futures price, has only limited power in predicting expected returns for a variety of individual futures contract, therefore not being able to reject the null hypothesis of constant or zero risk premia. However, they also note that most of futures discount variation on an individual futures level is related to idiosyncratic factors. In storage models, for example, the commodity futures discount is affected by expected changes in the spot price of the underlying commodity and thus, idiosyncratic factors such as expected changes in demand and supply play a large role in explaining the basis variation, see e.g. Deaton and Laroque (1996) and Routledge et al. (2002).

I extend the study by Fama and French (1987) in two dimensions. First, I combine individual futures contracts into equally weighted futures portfolios. Studying futures portfolios diversifies the idiosyncratic risk and improves statistical inference by reducing error variance. A casual inspection of the data set reveals that the average return volatility of individual futures contracts is 31 percent while the volatility of the futures portfolio is only 15 percent confirming the diversification benefit. Second, I use futures price data for maturities of 1 to 9 months across 21 futures contracts spanning the commodity classes agricultural, livestock, energy and precious commodities from January 1973 to September 2012. The increased sample length and the use of various maturity improves the sta-

tistical inference when investigate slow moving signals such as futures discounts and macroeconomic variables.

I study time-variation in expected returns by decomposing holding returns of commodity futures portfolios into the average futures discount, denoted as AFD henceforth, and average expected price appreciation. The AFD is defined as the average logarithmic difference of contemporaneous spot and futures prices across 21 commodities [1]. In contrast to spot price changes, the AFD is a model-free component of the expected return that can be observed ex-ante. Time-variation of AFD does, however, not immediately imply time-varying risk premia. The AFD can change even when expected returns are constant if the change is offset by an equal but opposite-signed expected spot price appreciation. To test the relation between AFD and expected returns I follow Fama and French (1987),(1986) and Lustig et al. (2013) and run complimentary regression of excess portfolio returns and average price changes on the AFD.

I report several findings. First, the AFD is statistically significant related to expected commodity portfolio returns across holding periods. The 1 month ahead AFD explains around 1% of the variation in a portfolio of commodity excess returns over the next month. As the horizon increases, the R^2 value increases, because the AFD is a persistent variable with high coefficients of autocorrelation. At the 9 months horizon, the AFD explains up to 10% of the variation in returns. Interestingly, the coefficient estimates of the predictive regression increase with maturity. For short term horizons, the beta coefficient is below 1 implying that some of the predictable variation is offset by expected mean reversion of spot changes. Over longer horizons, expected returns vary close to one-to-one with variation in the AFD implying that spot changes are random walks. The negative relation between futures discount and spot changes at the short horizon is a consequence of The Theory of Storage. The Theory of Storage by Kaldor (1939), Working (1949), and Brennan (1958) explains the futures discount in terms of storage and warehousing cost, and a convenience yield. In particular, the theory predicts that a high futures discount, which is related to low inventory levels and a resulting convenience yield that is larger than interest and storage costs, leads to a decrease in expected spot price changes. Second, for precious metals I find even stronger predictive results with R^2 values above 13 percent for holding returns and significantly time-varying term

1. In the commodities literature the futures discount is typically defined as the roll return.

premia for long horizons. The presented findings are robust to subperiod analysis, out-of-sample, and alternative construction methods for the AFD and economically significant. Simple trading strategies that exploit the return predictability of commodity portfolios increase the Sharpe ratio by approximately 0.5 in contrast to the random walk benchmark.

Next, following Szymanowska et al. (2013), I motivate theoretically and implement empirically portfolio investment strategies that decompose holding returns to extract information about the risks present in the underlying spot price and the term structure; called spot and term premia, respectively. Decomposing the holding return into the two components helps to identify the fundamental sources of time variation. In particular, I will be able to distinguish between risk variation due to spot risk and variation due to convenience yield risk. Spot premia can be identified by taking long positions in short maturity futures contracts, while detecting term premia requires the combination of long and short positions in contracts with different maturities. Term premia in futures contracts underly the same concept to risk premia in the term structure of bonds that have been widely studied, e.g., see Cochrane and Piazzesi (2002). By running regressions of term and spot premia on the AFD, I find that spot premia are significantly predictable, while term premia are only predictable at a two month horizon. In summary, the empirical results imply that expected holding returns of commodity portfolios vary over time and most of the time variation is related to variation in the underlying spot price.

A voluminous literature in financial economics suggests that investors in equity and bond markets must be compensated for risk associated with recessions. These theories imply that risk premia are counter-cyclically related to macroeconomic activity, see e.g. Campbell and Cochrane (1999) and Bansal and Yaron (2004). While various studies show empirically that risk premia in currencies (Lustig et al. (2013)), equity markets (Cooper and Priestley (2009)) and bonds (Ludvigson and Ng (2009)) are counter-cyclically related to macroeconomic activity, it is not clear theoretically nor empirically whether or how risk premia in commodity markets are related to business cycles. This uncertainty is further complicated by the fact that various authors find that commodity risk premia are related to commodity market specific sources of risk or trading activity such as hedging pressure (De Roon et al. (2000)), inventory levels (Gorton et al. (2007)) and open interest (Hong and Yogo (2013)).

I find that commodity portfolios are pro-cyclically related to expected returns in that the

AFD is statistically significant correlated with macroeconomic and common systematic variables. Furthermore, macroeconomic variables contain information about the cyclical behaviour above and beyond the information contained in the AFD and therefore, predict expected spot price growth. Including US industrial production growth (IP growth) and the AFD in the predictive regression increases the adjusted R^2 values for a 9 months horizon to 16 percent. Interestingly, term premia are counter-cyclical, however statistically insignificant. A precious portfolio, on the other hand, displays counter-cyclical risk premia consistent with the literature in financial economics confirming the role of gold as a pure financial investment. When analysing the role of trading activity on commodity portfolio returns, I find that speculators and hedgers, as declared by the Commodity Futures Trading Commission (CFTC), adjust their long and short positions in futures, according to expected future returns. In particular, hedgers increase their relative short positions, acting as contrarian investors, and speculators increase their long positions. I do not find, however, any indication that hedgers or speculators positions are leading expected returns.

The presented study is closely related to the literature of commodity futures risk premia. Hirshleifer (1988) and De Roon et al. (2000) argue that limited participation, due to informational and capital constraints of commodity producers in equity markets cause the returns on futures to depend on their covariance with non-marketed endowments, the hedging pressure, and their covariance with excess returns on the market portfolio, which is consistent with modern asset pricing theory. Empirically, Hirshleifer (1988) and De Roon et al. (2000) indicate that futures markets are at least partially segmented from other asset classes. Bessembinder and Chan (1992) and Bailey and Chan (1993) show that futures markets and in particular commodity futures markets can be predicted by common macroeconomic variables that have also been shown to predict stock and bond markets, suggesting systematic sources of risk. More recently, Szymanowska et al. (2013) show that the cross section of spot premia can be explained by a high-minus-low portfolio from futures discount sorts and term premia can be explained by two additional basis factors. These results indicate a common factor structure across commodity futures contracts. In contrast, Daskalaki et al. (2014) reject the null hypothesis of a common factor structure, by providing evidence that commodity markets are segmented from equity markets and show that individual commodity futures are considerably heterogenous. The authors use various asset pricing test, but do not find evidence for common factors across individual commodities.

This chapter proceeds as follows: Section 2.2 defines the portfolio construction and data used. Section 2.3 presents the predictive regression results that test for time variation in expected returns of holding returns and are the main results of the presented study. This section also presents several desirable robustness checks. Section 2.4 tests for time variation in spot and term premia. Section 2.5 shows that systematic macroeconomic variables have incremental explanatory power for future commodity portfolio returns and investigates the role of trading activity. Lastly, section 2.6 concludes.

2.2 Definitions and Empirical Implementation

In this section, I will formalize the definitions and construction of returns to aggregate commodity investment portfolios. Furthermore, I present the data used and discuss first summary statistics.

2.2.1 Holding Returns

Definitions *Commodity holding returns* correspond to a simple investment strategy. At time t an investor buys a futures contract guaranteeing delivery of the underlying commodity in one period and sells the commodity in the spot market at expiry. In practice most financial investors do not hold the futures contract until maturity but sell it just prior to maturity to avoid the low liquidity, the physical delivery and irregular price behaviour of many spot commodity markets. In the empirical implementation I sell the futures contract at the end of the month just prior to delivery, but use the spot price for notational convenience in the theoretical motivation. I use s_{t+1} to denote the logarithm of the spot price of the underlying commodity at time $t+1$ and use $f_t^{(1)}$ for the log price of the futures contract at time t maturing in one period. The one period log holding return, $rx^{(1)}$, on buying a futures contract and selling it in the spot market in one period is simply defined as $rx_{t+1}^{(1)} := s_{t+1} - f_t^{(1)}$. Investment strategies using futures contracts are zero cost investment strategies and the holding return is equal to the excess returns of a fully collateralized futures position.

The holding return $rx_{t+1}^{(1)}$ can, equivalently, be stated as the sum between log spot change and log futures discount: $rx_{t+1}^{(1)} = s_{t+1} - s_t + (s_t - f_t^{(1)})$. In the commodities literature the decomposition into log spot changes and log futures discount is typically denoted as spot and roll

return, respectively (see e.g. Erb and Harvey (2006)). I refrain from the terminology and continue to use the same terminology used in the bond and exchange rate literature (Lustig et al. (2013) and Cochrane and Piazzesi (2002)) to highlight the similarities to other asset classes. In equilibrium, a positive futures discount or backwardated term structure implies that immediate ownership of the physical commodity entails a benefit which deferred ownership does not, e.g., see Routledge et al. (2002). This benefit, expressed as a rate, is termed the *convenience yield*. The futures discount in commodity markets, also called carry or basis, is equal to the difference between convenience yield in excess of storage costs and risk-free interest rate. The convenience yield is similar to the dividend yield on equities or the foreign interest rate in currencies. The relation between futures discount, risk-free rate and convenience yield is a direct consequence of The Cost-of-Carry model (e.g. Fama and French (1987)):

$$F_t^{(1)} = S_t e^{(r_t - \delta_t)} \tag{2.1}$$

where r_t is the risk-free interest rate and δ_t the convenience yield. The futures discount is therefore equal to $s_t - f_t^{(1)} = (\delta_t - r_t)$. Hence, the log commodity holding return for one period equals the spot rate changes minus the yield or convenience yield/intrest rate differential: $rx_{t+1}^{(1)} := s_{t+1} - s_t - (\delta_t - r_t)$.

Horizons and Maturities Futures contracts are available at different maturities. However, maturities for different underlyings are not consistently spaced. While energy futures contracts are available that expire in every month of the year, agricultural futures have irregularly spaced maturities. The irregular spacing leads to a spot contract for agriculturals deviating from the "real" spot price by several months. For the study at hand the irregular spacing is not of particular concern. First, financial investors use the same investment procedure I use. Second, using "real" spot prices would introduce too much measurement error due to the low liquidity and would limit my analysis even more.

I use n-month maturity futures contracts to calculate n-month horizon returns (where n=1, 2, 3, 6 and 9)[2]. The log holding return on the n-month commodity contract, i, is defined as

2. I use these maturities to guarantee sufficient liquidity in implementing the commodity investment strategies. I measure liquidity in terms of open interest and trading volume.

$rx^{(n)}_{i,t\to t+n} := s_{i,t+n} - f^{(n)}_{i,t\to t+n} = \Delta s_{i,t\to t+n} + (s_{i,t} - f^{(n)}_{i,t\to t+n})$, where $f^{(n)}_{i,t\to t+n}$ is the log futures price at time t maturing in n periods and $\Delta s_{i,t\to t+n}$ is the n-month change in spot price.

Portfolios of Commodities and Futures Discounts I construct three portfolios of commodity futures. The first portfolio is the *Global with Precious* portfolio including all futures contracts across commodity classes. The second portfolio, denoted as *Global*, contains the futures contracts of all commodity classes excluding precious metals. The final portfolio, denoted *Precious*, contains only gold and silver futures contracts. I distinguish between non-precious and precious commodities because gold and silver are mostly used as financial assets to store wealth, while other commodities are consumption goods.

Average log holding return on commodities in portfolio j over horizon n is defined as $\overline{rx}^{j}_{t\to t+n} = \frac{1}{N^j_t}\sum_{i=1}^{N^j_t} rx^{i}_{t\to t+n}$, where N^j_t denotes the number of commodities in portfolio j at time t. Table 2.1 shows that the start dates of the individual futures contracts vary across assets. A time-varying number of assets explains the t subscript in N^j_t. I use an equally weighted average of individual commodity returns, similar to Hong and Yogo (2013) or Szymanowska et al. (2013). Consequently, the AFD for maturity, n, and portfolio, j, is $\overline{s}_{j,t} - \overline{f}^{(n)}_{j,t\to t+n} = \frac{1}{N^j_t}\sum_{i=1}^{N^j_t} s_{i,t} - f^{(n)}_{i,t\to t+n}$.

2.2.2 Data

I start from price data of daily commodity futures contracts covering a cross section of 21 futures markets that resemble the data set of Szymanowska et al. (2013). These data are collected from a variety of futures exchanges via Bloomberg over the sample period from January 1973 to September 2012. From the daily price data I build end-of-month series for each maturity (n=1,2,3,6,9)[3]. I also collect data on open interest and volume for each of the futures contracts. The total futures sample consists of energy, agricultural, livestock and precious commodities presented in Table 2.1.

3. Due to the irregular spacing of maturities for some contracts I use the closest to maturity contract if a particular date is not available.

Table 2.1: **Summary Statistics for Individual Commodities:** This table reports the exchange, start date, annualized mean and annualized standard deviations of one period holding returns ($rx_{i,t+1}$), spot changes ($\Delta s_{i,t+1}$) and futures discounts ($s_{i,t} - f_{i,t}$) for the sample of individual commodity futures contracts i. Holding returns are constructed by entering a long futures contract at time t holding it for one periods until maturity and selling it at the spot price (front futures contract). The spot change is constructed as the log difference of front month futures contracts and the futures discount is the log difference between spot (front futures contract) and next to maturity futures price. Summary statistics are based on monthly return series that span the cross-section of 19 futures contracts across the commodity classes Energy, Agriculturals, Livestock, Precious Metals and Metals and are reported for the closest to maturity contracts (n=1) only. The sample is from January 1973 to September 2012.

Instrument	Exchange	Start	Holding Returns		Spot Returns		Futures Discount	
			Mean	Std	Mean	Std	Mean	Std
Energy								
Crude Oil	NYMEX	Mar-83	6.7	33.3	4.2	33.4	2.4	6.7
Heating Oil	NYMEX	Apr-86	13.3	36.6	9.6	37.7	3.7	13.1
RBOB Gas	NYMEX	Oct-05	12.0	37.8	10.3	41.1	1.7	11.5
Natural Gas	NYMEX	Apr-90	-17.6	50.6	4.9	56.2	-22.3	21.8
Agriculturals								
Corn	CBOT	Jan-70	-3.9	26.8	4.5	27.8	-8.4	8.8
Soybean	CBOT	Jan-70	2.2	28.9	4.5	29.3	-2.2	6.5
Chicago Wheat	CBOT	Jan-70	-3.8	27.5	4.3	29.2	-8.1	9.1
Soybean Oil	CBOT	Jan-70	4.0	32.7	4.2	32.1	-0.2	7.9
Kansas Wheat	CBOT	Jan-70	1.8	25.9	4.5	26.9	-2.6	7.5
Coffee	NYMEX	Aug-72	1.2	38.2	2.4	38.7	-1.2	10.1
Cotton	NYMEX	Jan-70	1.0	26.9	2.5	30.0	-1.5	15.9
Sugar	NYMEX	Jan-70	-3.5	43.5	4.3	47.0	-7.8	15.9
Cocoa	NYMEX	Jan-70	2.2	32.8	3.1	33.2	-0.8	8.4
Livestock								
Feeder Cattle	CME	Dec-71	1.6	16.8	3.5	17.7	-1.9	6.2
Live Cattle	CME	Jan-70	5.3	18.3	3.5	20.0	2.0	10.1
Lean Hogs	CME	Apr-86	-1.5	26.0	1.0	35.3	-2.5	25.0
Precious Metals								
Gold	COMEX	Jan-75	0.6	19.6	6.3	19.7	-5.6	2.2
Silver	COMEX	Jan-75	-0.3	34.6	5.8	34.7	-6.0	4.5
Metals								
Copper	LME	Jun-97	8.6	27.8	7.7	27.6	0.2	2.3
Lead	LME	Jul-97	6.7	31.3	8.4	31.1	-1.7	3.3
Nickel	LME	Jul-97	7.1	37.5	5.6	37.3	1.4	2.5
Zinc	LME	Jul-97	-3.1	28.5	1.8	28.8	-4.9	1.9

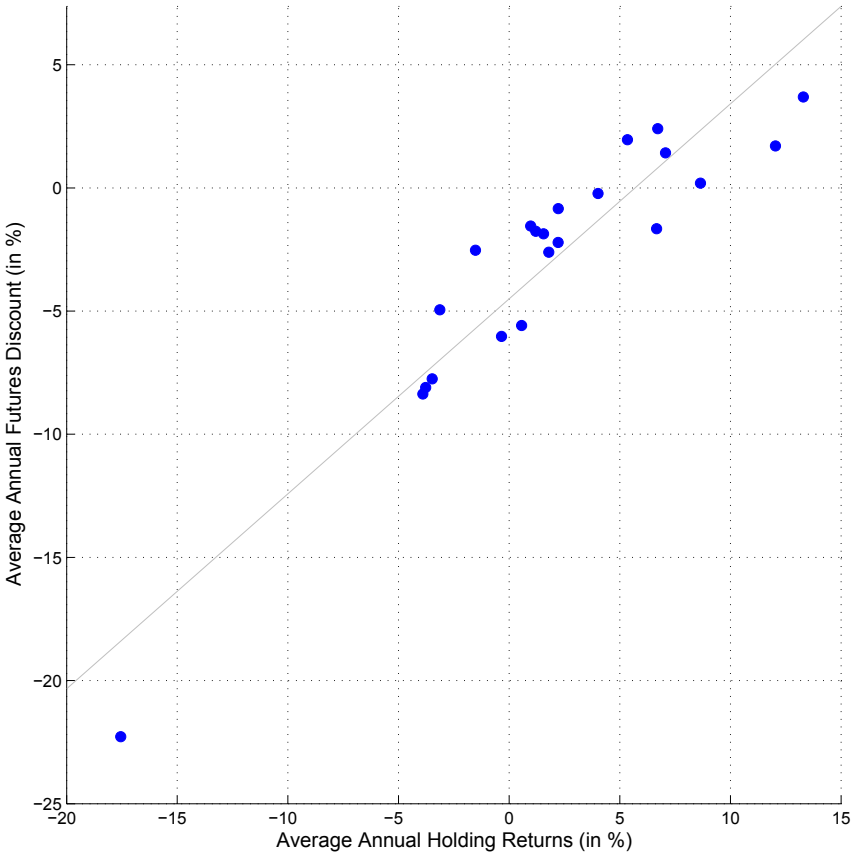


Figure 2.1: **Individual long run average holding returns and individual futures discounts:** This figure displays the average annualized mean holding returns on the y-axes and the local futures discounts on the x-axes. Mean returns are computed for individual commodity contracts over the sample period from January 1973 to September 2012.

Individual Commodity Summary Statistics Table 2.1 reports the summary statistics of the one period holding returns $rx_{t+1}^{(1)}$, the spot changes Δs_{t+1} and the futures discounts $s_t - f_t^{(1)}$ for each individual commodity within the asset classes, energy, agriculturals, livestock, precious metals and metals. One period holding returns display large cross sectional variations. While Heating Oil futures contracts have annualized mean returns of 13.5 percent, Natural Gas contracts have a mean return of -17.6 percent. Over all, 15 contracts have positive mean returns and 7 mean returns are negative. Most commodities also display high volatilities comparable to the volatilities of individual stock returns. Annualized volatilities range from 50.6 percent for Natural Gas to 16.8 percent for Feeder Cattle. The annualized mean of spot return changes also vary cross-sectionally but with a smaller range than the holding returns and are all positive. The annualized volatilities are consistently higher than the holding returns. A declining term structure of commodity futures volatility is known as the "Samuelson (1965) effect" and a consequence of the convenience yield and inventory decisions (e.g. Routledge et al. (2002)). The annual mean of the futures discounts or roll returns displays large cross sectional variation and considerable volatility. Figure 2.1 displays the average annual holding returns against average annual yields for each commodity, as an indication of the long term relation. The data are highly concentrated along a positive sloping linear regression line indicating that futures discounts are related to the expected returns in the long term.

Portfolio Holding Returns Summary Statistics Table 2.2 Panel A reports the summary statistics for average holding returns $\overline{rx}_{t\to t+n}^{j}$, average spot changes $\overline{\Delta s}_{t\to t+n}^{j}$ and AFDs $\overline{s}_{j,t} - \overline{f}_{j,t\to t+n}^{(n)}$ over the maturity horizon of n=1, 2, 3, 6 and 9 months. I report the summary statistics separately for the Global portfolio, the Global and Precious portfolio and Precious only portfolio. Annualized holding returns for the Global with Precious portfolios are positive and increasing, with values ranging from 0.86 percent for a 1 month horizon to 2.59 percent for a 9 months horizon. The annualized standard deviations are increasing as well and ranging from 14.64 percent to 17.87 percent [4]. These results highlight the benefit of aggregating commodities into average portfolios. While the standard deviation of individual commodities is on average higher than 30 percent, the

4. Even though the unconditional returns are positive they are much lower than equity returns. This is in contrast to studies by Gorton and Rouwenhorst (2004) or Erb and Harvey (2006) who report unconditional returns of commodities that are similar to equity returns. However, the sample periods of there studies differ from the sample period considered in this study which might explain this deviation.

Table 2.2: **Summary Statistics for Portfolios of Commodity Futures:** This table reports the summary statistics for portfolios of commodity holding returns ($\overline{rx}_{t\to t+n}$), spot changes ($\overline{\Delta s}_{t\to t+n}$) and average futures discount ($\overline{s_t} - \overline{f_{t\to t+n}}$) in Panel A and the summary statistics for portfolios of front roll ($\overline{frx}_{t\to t+n}$) and excess holding returns ($\overline{hrx}_{t\to t+n}$) in Panel B. Front roll returns for various maturities are constructed by rolling forward one month horizon holding returns. Excess holding returns are the difference between holding returns and front roll returns. The individual return series are aggregated by constructing equally weighted portfolios averages across all commodities. The average futures discount is constructed as the difference between the average spot prices at time t and the average futures returns at time t with maturity n across all commodities. The table shows the annualized mean, annualized standard deviation and autoregression coefficient of order one (AR(1)) for holding horizons of $n = 1, 2, 3, 6, 9$ months. I report the results when aggregating across all available commodity contracts (Global w. P), when aggregating across energy, agriculturals, metals and livestock (Global) and Precious metals only. The sample is from January 1973 to September 2012.

Panel A:									
Horizon	Holding Returns			Spot Changes			Futures Discount		
n	Mean	Std	AR(1)	Mean	Std	AR(1)	Mean	Std	AR(1)
Global									
1	0.73	15.01	0.05	3.81	15.05	0.00	-3.08	3.83	0.39
2	0.96	15.27		3.87	15.04		-2.92	4.20	0.64
3	1.29	15.82		3.87	15.51		-2.58	4.53	0.78
6	2.07	17.34		3.78	16.40		-1.72	5.11	0.88
9	2.68	19.17		3.77	17.50		-1.09	5.40	0.91
Global w. Precious									
1	0.86	14.64	0.07	4.24	14.64	0.02	-3.38	3.42	0.64
2	1.05	15.10		4.27	14.86		-3.22	3.77	0.64
3	1.34	15.62		4.27	15.31		-2.93	4.09	0.78
6	2.03	17.00		4.17	16.02		-2.15	4.73	0.89
9	2.59	17.87		4.18	16.51		-1.58	5.08	0.92
Precious Metals									
1	0.11	25.31	0.06	6.08	25.17	0.03	-5.95	2.27	0.09
2	-0.07	26.00		5.90	26.21		-5.98	2.03	0.46
3	-0.16	25.85		5.81	26.27		-5.97	2.11	0.65
6	-0.18	30.42		5.77	31.34		-5.95	2.64	0.89
9	-0.04	31.29		5.89	32.56		-5.93	3.08	0.94

Panel B:												
	Front Rolls						Excess Holding					
	Global w. P.		Global		Precious		Global w. P.		Global		Precious	
n	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
1	0.85	14.61	0.73	14.95	0.11	25.03						
2	0.81	15.14	0.69	15.36	-0.06	25.76	0.23	1.25	0.27	1.38	-0.01	1.38
3	0.77	15.76	0.66	16.07	-0.15	25.32	0.56	1.62	0.63	1.81	-0.01	1.63
6	0.63	16.97	0.52	17.29	-0.20	25.69	1.40	2.61	1.56	2.95	0.01	1.70
9	0.66	17.95	0.54	18.19	-0.09	26.10	1.92	3.34	2.13	3.79	0.04	1.68

standard deviation of the portfolio reduces to 15 percent when aggregating the contracts. The high individual return volatility indicates that a major part of variation is driven by idiosyncratic sources of risk. However, it is important to stress that there should exist common elements driving excess returns across difference commodities, which is the focus of this study. This hypothesis is further encouraged by the variation of the AFD which is an important component of expected returns. The mean is significantly different from zero with a value of -3.28 percent and is highly time varying as indicated by the equally high standard deviation of 3.42 percent.

The summary statistics for Global portfolios are very similar to the Global with Precious counterpart. For Precious metals I find significant differences. First, holding returns are very close to zero or even negative and the volatilities are much higher, ranging from 25 to 31 percent, which is obvious given that the portfolio consists of only two assets. More interesting is, however, that the AFD variation of Precious metals is much lower than for the Global commodities. The low volatility will have an effect on predictive regressions, which are sensitive to different volatility levels.

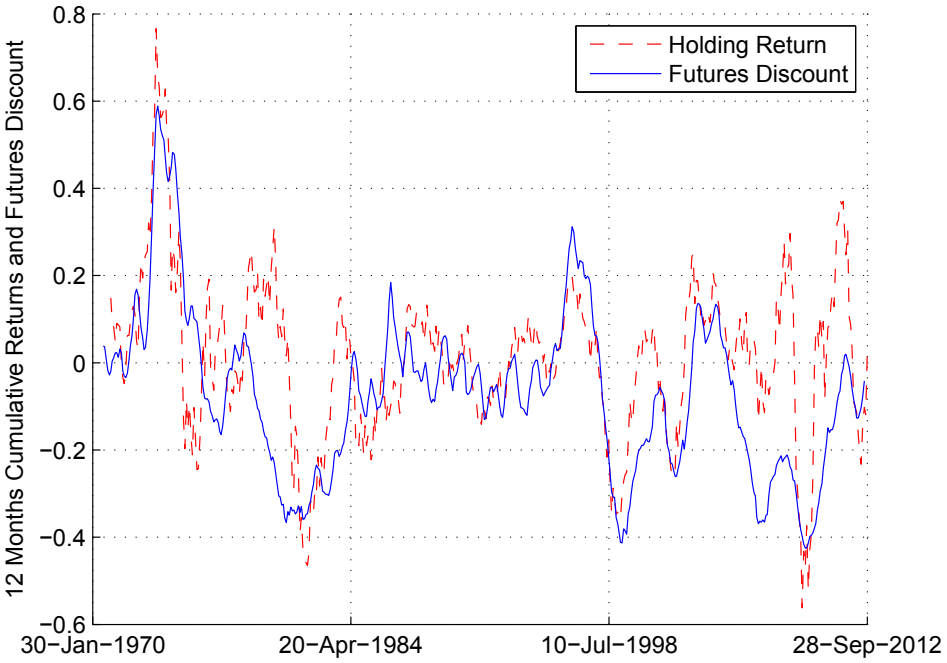


Figure 2.2: **Nine months average futures discount and 9 months holding returns:** This figure displays the average 9 months futures discount and 9 months holding returns.

2.3 Time Variation in Holding Returns

The pattern in the unconditional portfolio return strategies and the variation of the AFD, which is an ex-ante component of expected returns, are indicative of time variation in commodity risk premia. In this section I use predictive regressions to formally test for time variation in expected commodity portfolio returns. For individual commodities predictive regressions have been applied by Fama and French (1987) and are standard in other asset classes, e.g., see Cochrane and Piazzesi (2002) and Lustig et al. (2013). A typical characteristic of individual commodities traded in futures markets is the high uncertainty and therefore, volatility of expected spot price variations. Even if the futures discount variation is large it can be small relative to the spot price variation. At the same time, the large variance of futures discounts means that average risk premia are not sufficient to infer statistically non-zero expected premiums, see e.g. Fama and French (1987). As a consequence it is difficult to obtain reliable and robust results in predictive regressions. Combining commodities into portfolios can increase the power of univariate tests for time-varying expected premium and has been extensively used in various other asset markets, e.g., see Cochrane (2011) for a review. Another important differentiation in this chapter is the time frame and term structure of prices in which I study predictability. I consider a longer sample period and a variety of maturity horizons than most previous studies in commodity predictability. I start the study in 1973 until September 2012 and consider maturities from 1 to 9 months. Variation in expected returns does typically not show up in short term horizons of daily or weekly frequencies but changes within longer frequencies. Obtaining results for longer time horizons and various maturity horizons is therefore crucial.

2.3.1 Predictability Results

Holding Returns Holding returns can be decomposed into the futures discount and spot price change. As a consequence expected portfolio holding returns consist of the AFD, observable ex ante, and model-free and expected portfolio spot price change, resulting in

$$\mathbb{E}_t[\overline{rx}^j_{t\to t+n}] = \mathbb{E}_t[\Delta\overline{s}_{t+1}] + \overline{s}_{j,t} - \overline{f}^{(n)}_{j,t\to t+n} \tag{2.2}$$

for each commodity portfolio j. Equation (2.2) shows that the futures discount contains information on the risk premium and/or the expected spot price change. More specifically, the futures discount is equal to spot and term premia if spot price changes are assumed to follow a random walk. A time-varying futures discount would then immediately imply time variation in risk premia. Unfortunately, it is not clear ex ante if positive futures discounts result in high expected returns. A high futures discount can result in expected spot price depreciation that offsets the gain in expected holding returns. To formally test whether futures discounts predict expected holding returns I follow the regression framework by Fama (1984), Fama and French (1987) and Lustig et al. (2013). Consider the complimentary regression of the holding returns and the spot price change on the AFD:

$$\overline{rx}^j_{t\to t+n} = \psi^j_0 + \psi_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \eta_{t+n} \tag{2.3}$$

and

$$\overline{\Delta s}^j_{t\to t+n} = \zeta^j_0 + \zeta_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \epsilon_{t+n}. \tag{2.4}$$

By construction, equation (2.3) and (2.4) impose restrictions on the coefficient estimates of ψ_0 and ζ_0, as well as, on $\psi_{\mathbf{f}}$ and $\zeta_{\mathbf{f}}$. The difference of the risk premium and the change in spot price is equal to the futures discount, thus the intercepts ψ_0 and ζ_0 must sum to zero and most importantly, the slope coefficients must sum to one. In other words, the regression always allocates all variation of the AFD to the risk premium, the expected change in spot price, or a combination of the two.

The regressions in equation (2.3) and (2.4) test two hypotheses. In regression (2.3), the null hypothesis states that log expected holding returns of commodities are constant. Consequently, a coefficient estimate of $\psi_{\mathbf{f}}$ that is statistically different from zero implies that the futures discount contains information on excess returns or risk premia. Predictable variation in realized risk premiums is evidence of time-varying expected premiums. In regression (2.4), the null states that changes in log spot rates are not predictable. It is sufficient to run one of the two regressions as they are complimentary, but it is instructive to observe both estimates.

While the main task of this article is to determine the magnitude or presence of time variation of risk premia in portfolios of commodities, it is important to understand potential outcomes of the regression (2.3) and (2.4) within the complementary theoretical framework of The Theory of Storage. The Theory of Storage by Kaldor (1939), Working (1949), and Brennan (1958) explains

the futures discount in terms of interest forgone in storing commodity, warehousing cost, and a convenience yield. In particular the theory predicts that a high futures discount, which is related to low inventory levels and a resulting convenience yield that is larger than interest and storage costs, leads to a decrease in expected spot price changes. For individual commodities the relation between futures discounts and expected spot price changes is not controversial[5]. On a portfolio level, however, this relation has not been examined, yet. Negative predictability of spot changes on a portfolio level suggests that common factors across commodities drive variation in the convenience yield. Common drivers could be related to macroeconomic demand and supply cycles that could be potentially systematic across commodities. Fama and French (1988) establish this link on an individual level for metal commodities.

Statistics Inference Statistical inference of the parameters in equation (2.3) and (2.4) and equations (2.23) and (2.24) is complicated by the overlapping data used in the predictive regressions and the autocorrelated AFD. To have confidence in the statistical results, I investigate two variations of statistical inference. I report t-statistics for the slope coefficients $\psi_{\mathbf{f}}$ and $\zeta_{\mathbf{f}}$ as well as $\theta_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$ for two asymptotic tests. If data is overlapping the overlap causes regression errors to be serially correlated even under the null hypothesis of no predictability. Estimates of the covariance matrix of coefficients based on the assumption of homoscedasticity and no autocorrelation overstate the statistical significance of the estimates. First, I use Hansen and Hodrick (1980) methodology to compute asymptotic standard errors (HH) with the number of lags equal to the horizon of the forward contract plus one lag. Second, I use Newey and West (1987) adjusted standard errors (NW) computed with the optimal number of lags computed (Andrews (1991)) to correct for error correlation and conditional heteroscedasticity. Both methodologies can be understood as instances of the Generalized Methods of Moments framework for different choices of the long run covariance matrix of sample moments that estimate the asymptotic covariance matrix of coefficients, taking overlaps and heteroscedasticity into account.

Predictability Results for Holding Returns Table (2.9) reports the results from the regres-

5. In particular Fama and French (1987) or Routledge et al. (2002) show the negative relation between futures discount and expected spot changes.

sions of equation (2.3) and (2.4) over the sample period from 1973 to 2012. The left panel focuses on the Global with Precious portfolio of commodities. There is strong evidence that the AFD is a predictor for holding returns for all horizons considered (n=1,2,3,6 and 9). The estimated slope coefficients, $\psi_{\mathbf{f}}$, in the predictability regression are all highly significant by means of Newey-West or Hanson-Hodrick correction. The adjusted R^2 values increase from 1.2 percent for the monthly horizon to 10.3 percent for a 9 months horizon. The increase in the R^2's is not surprising given the high autocorrelation of the AFD of around 70 percent reported in Table (2.2).

The slope coefficient estimates, $\psi_{\mathbf{f}}$, become larger with holding horizon. For short horizons of 1 month the coefficient is only 0.5 while it increases to 1.16 over a 9 months horizon. The increasing coefficient estimates have implications for spot price changes given the identity $\psi_{\mathbf{f}}=1+\zeta_{\mathbf{f}}$. For very short horizons of one month I find that spot price changes are negatively and significantly related to the ADF. At longer horizons the negative relation decreases and in turn becomes even positive at horizons of 6 and 9 months. This result is not surprising considering the predictability results within The Theory of Storage. As, e.g., Routledge et al. (2002) show, a high futures discount is related to low inventory levels, because low inventory levels lead to high spot prices. As a consequence a high futures discount is related to lower expected returns and explains the negative sign in the predictability regression. Nevertheless, note that for short horizons a common component in the variability of the AFD seems to be common to all commodities considered. Fama and French (1988) show that for metals positive demand shock around business cycle peaks reduce metal inventories, which generates the high AFD.

The economic impact of the AFD on commodity holding returns is increasing with holding periods; for example, the estimated coefficient on the AFD from the 9 month predictability regression for the Global holding returns is 1.16, implying that each 100 basis point increase in the futures discount implies a 116 basis points increase in the expected holding returns. For 1 months horizons, each 100 basis point increase in the futures discount implies only 50 basis points increase in the expected holding returns, implying a 50 basis point decrease in spot changes. The economic significance implies that trading futures contracts with longer maturities conditioned on the AFD can be beneficial for investors as it avoids the strong mean reversion in short term spot prices. I will come back to the economic significance of the predictability in the robustness section.

The mid panel in Table (2.9) reports the predictability results for the Global portfolio selection

that excludes precious metals from the set of commodities. The results are very similar to the results of the Global selection, even though, the adjusted R^2 values are reduced slightly. For the 1 month horizon the R^2 is 1 percent and increases to an R^2 of 9 percent on a nine months horizon. Furthermore, I find the same negative effect of AFD on spot price changes consistent with The Theory of Storage and the economic impact is comparable to the Global with Precious portfolio.

The last two panels in Table (2.9) report the results from regressing holding returns and spot changes of Precious commodities on the AFD for different horizons (n=1,2,3,6 and 9). I find for all horizons statistically significant positive coefficient estimates. In contrast to the Global selection, however, I find significantly higher R^2 values: for the 1 month horizon I find adjusted R^2 of 3 percent that increase to almost 14 percent for the 9 months horizon. Even more pronounced are the magnitudes of the coefficient estimates. Each 100 basis point increase in the futures discount implies a 200 basis points increase in expected holding returns on a 1 month horizon that increases to 294 basis point on a 9 month horizon. The increase of holding returns by more than 100 basis points implies that the AFD predicts spot price changes, too, even though, not statistically significant due to the high volatility. Again, this result is not surprising. While the Global selection contains commodities that are highly dependent on inventory levels, Precious metals are not as affected by cyclical inventory movements.

Table 2.3: **Predicting Portfolio Front Roll and Excess Holding Returns with Average Futures Discount:** This table reports results of forecasting regressions for average front roll and average excess holding returns on average futures discount for portfolios of commodities at horizons of n=1,2,3,6, and 9 months (for excess holding returns I do not report the results for one month horizon as it is by definition zero). The first six columns report the predictability results for Global commodity portfolios. Reported are the estimated slope coefficient $\psi_{\mathbf{f}}$ and $\zeta_{\mathbf{f}}$, the constant coefficients ψ_0 and ζ_0 and the adjusted R^2 values for the time-series regression of the average front roll return $\overline{frx}_{t\to t+n}$ and the average excess holding return $\overline{hrx}_{t\to t+n}$ on the average futures discount, $\overline{s_t} - \overline{f_{t\to t+n}}$. The t-statistics for the slope coefficients in brackets are computed using Newey and West (1987) standard errors (NW) with heteroscedasticity and autocorrelation consistent covariance estimators computed with the number of lags equal to the length of overlap plus one lag and Hansen and Hodrick (1990) standard errors (HH). I repeat the analysis for Non Precious portfolios in columns seven to twelve and Precious portfolios in columns thirteen to eighteen. The sample period is January 1973 to September 2012 for Global and and Global without Precious portfolios and January 1976 to September 2012 for the Precious portfolio.

	Global w. Precious						Global						Precious					
	Front Roll			Excess Holding			Front Roll			Excess Holding			Front Roll			Excess Holding		
	$\overline{frx}_{t\to t+n}$			$\overline{hrx}_{t\to t+n}$			$\overline{frx}_{t\to t+n}$			$\overline{hrx}_{t\to t+n}$			$\overline{frx}_{t\to t+n}$			$\overline{hrx}_{t\to t+n}$		
n	θ_0	$\theta_{\mathbf{f}}$	R^2	ϑ_0	$\vartheta_{\mathbf{f}}$	R^2	θ_0	$\theta_{\mathbf{f}}$	R^2	ϑ_0	$\vartheta_{\mathbf{f}}$	R^2	θ_0	$\theta_{\mathbf{f}}$	R^2	ϑ_0	$\vartheta_{\mathbf{f}}$	R^2
1	0.00	0.49	1.13				0.00	0.44	1.06				0.01	2.00	3.04			
NW	(0.97)	(2.19)					(0.80)	(2.28)					(2.94)	(3.71)				
HH	[0.95]	[2.19]					[0.79]	[2.27]					[3.13]	[3.78]				
2	0.01	0.72	3.00	0.00	0.06	3.45	0.00	0.65	2.97	0.00	0.06	2.67	0.02	2.06	2.47	0.00	0.06	0.61
NW	(1.28)	(2.93)		(2.04)	(2.82)		(1.06)	(3.05)		(1.80)	(2.67)		(2.40)	(1.98)		(2.50)	(1.32)	
HH	[1.16]	[2.77]		[1.86]	[2.58]		[0.95]	[2.85]		[1.65]	[2.46]		[2.46]	[1.96]		[3.14]	[1.45]	
3	0.01	0.75	3.61	0.00	0.07	2.91	0.01	0.69	3.57	0.00	0.06	1.97	0.02	1.65	1.74	0.00	0.10	1.48
NW	(1.23)	(2.73)		(2.72)	(2.26)		(1.00)	(2.78)		(2.55)	(2.00)		(1.50)	(1.34)		(2.45)	(1.43)	
HH	[1.09]	[2.47]		[2.42]	[2.00]		[0.88]	[2.50]		[2.26]	[1.78]		[1.52]	[1.29]		[2.52]	[1.45]	
6	0.01	0.98	7.14	0.01	0.07	1.43	0.01	0.90	6.77	0.01	0.05	0.51	0.07	2.56	7.25	0.00	0.13	4.11
NW	(1.17)	(2.83)		(3.37)	(1.21)		(0.90)	(2.83)		(3.27)	(0.84)		(2.77)	(2.39)		(3.35)	(2.07)	
HH	[0.97]	[2.43]		[2.70]	[0.97]		[0.75]	[2.44]		[2.64]	[0.68]		[2.39]	[2.03]		[2.94]	[1.90]	
9	0.02	1.06	7.08	0.01	0.08	-0.15	0.01	1.07	6.43	0.01	-0.01	-0.20	0.12	2.82	12.15	0.01	0.12	5.20
NW	(0.93)	(2.34)		(2.99)	(0.21)		(0.68)	(2.33)		(3.02)	(-0.07)		(3.15)	(2.68)		(3.03)	(2.06)	
HH	[0.77]	[1.99]		[2.41]	[0.17]		[0.56]	[2.03]		[2.42]	[-0.05]		[2.73]	[2.25]		[2.54]	[1.81]	

2.3.2 Average Futures Discount and Individual Futures Contracts

Most of the literature on commodity return predictability has focused on variation of individual futures contracts. In particular Szymanowska et al. (2013) show that cross-sectional sorts based on individual futures discounts result in statistically significant excess returns[6]. In this section, I test if the AFD is able to forecast individual futures returns in addition to the portfolio level forecasts. In particular, I test if the AFD contains information that are not captured in individual futures discounts, similar to Lustig et al. (2013). Empirically, I set up a pooled panel regression of individual holding returns on the AFD and individual commodity specific local futures discounts;

$$rx_{i,t\to t+n} = \psi_{i,0} + \psi_{\mathbf{f}}(\overline{s}_t^j - \overline{f}_{t\to t+n}^j) + \psi_{\mathbf{f}}(s_{i,t} - f_{i,t\to t+n}) + \eta_{t+n} \tag{2.5}$$

and

$$\Delta s_{i,t\to t+n} = \zeta_{i,0} + \zeta_{\mathbf{f}}(\overline{s}_t^j - \overline{f}_{t\to t+n}^j) + \psi_{\mathbf{f}}(s_{i,t} - f_{i,t\to t+n}) + \epsilon_{t+n}. \tag{2.6}$$

where $\psi_{i,0}$ and $\zeta_{i,0}$ are commodity fixed effects, so that only the slope coefficient are constraint to be the same across commodities. Stacking all futures contracts and dates I run a pooled panel regression and compute t-statistics that account for group-wise clustering by time.

Table (2.4) reports the results of regression (2.5) and (2.6) for the Global with Precious, Global, and Precious portfolios over the horizon of n=2,3,6, and 9. In the rows marked with an (a), I report the regression results including only the individual local futures basis. Rows marked with a (b) report the results when including the AFD and the individual futures discount. First, when observing the (a) rows, I find highly significant coefficient estimates for all maturities and commodity portfolios. This result is consistent with the findings of Szymanowska et al. (2013) and are more robust than the findings by Fama and French (1987). Next, when including the AFD in the panel regression, I find that both the individual and the average futures discount are highly significantly related to individual expected holding returns across maturities. The magnitude of the coefficient estimates is even higher and in most cases more significant for the AFD than for the individual futures discounts. This result extends for the Non Precious and Precious selection.

In summary I find that a single return forecasting variable describes time variation in individual

6. Cross sectional sorts are just an alternative way to regression analysis to test for predictability

Table 2.4: **Predictability of Individual Holding Returns Using Univariate and Average Futures Discount:** This table reports results of pooled panel regressions for holding returns of individual commodities at horizons of n=2,3,6, and 9 months, on both the average futures discount and the commodity specific futures discounts. The pooled panel regression controls for commodity fixed effects to allow for different drifts. For each group of commodities (Global w. Precious, Global, and Precious Metals), I report the slope coefficient on the AFD, $\psi_{\mathbf{f}}$, and on the individual futures discount,ψ_f. I run the regression separately for the individual futures discount in (a) and for both predictors in (b). The t-statistics of the slope coefficients (in brackets) are computed using robust standard errors clustered by month. Data are monthly and the sample runs from January 1973 to September 2012.

		Global w. Precious		Global		Precious Metals	
Horizon		$\overline{rx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$	
n		$\psi_{\mathbf{f}}$	ψ_f	$\psi_{\mathbf{f}}$	ψ_f	$\psi_{\mathbf{f}}$	ψ_f
2	(a)		0.56		0.64		2.16
			(4.01)		(4.90)		(2.86)
	(b)	0.45	0.11	0.55	0.08	2.21	-0.05
		(3.25)	(2.62)	(4.23)	(1.88)	(2.16)	(-0.10)
3	(a)		0.58		0.68		1.79
			(3.60)		(4.53)		(1.72)
	(b)	0.45	0.14	0.57	0.11	2.89	-1.11
		(2.75)	(2.90)	(3.71)	(2.25)	(1.89)	(-1.03)
6	(a)		0.81		0.90		2.69
			(4.03)		(4.71)		(3.05)
	(b)	0.60	0.23	0.73	0.17	4.73	-2.04
		(2.99)	(3.59)	(3.77)	(2.70)	(2.31)	(-1.14)
9	(a)		0.52		0.60		2.03
			(2.87)		(3.39)		(3.14)
	(b)	0.30	0.30	0.39	0.25	2.70	-0.86
		(1.67)	(4.41)	(2.24)	(3.70)	(2.44)	(-0.75)

commodity futures in addition to the commodity specific individual futures discounts. This variable is the AFD of commodity futures portfolios. The result is novel to the commodity literature and is similar to the findings by Lustig et al. (2013) for portfolio of currency excess returns.

2.3.3 Robustness

The results in the previous subsections show that expected holding returns are predictable by the AFD consistent with the presence of time-varying risk premia. In this section, I investigate a number of desirable robustness checks: (1) I show that the results are stable across various sub samples. Only the very recent period of high commodity volatility in 2008 deteriorates the predictive results. (2) I show that the results are robust to alternative construction methods of the AFD. (3) I show that the in-sample results are robust to an extension to out-of-sample statistical measures. (4) I construct various simple "trading rules" based on the predictive results, generating significant economic profits. I report the results for Global and Precious portfolios and holding and excess holding returns only.

Sub-sample Analysis Table (2.5) Panel A addresses the issue of the robustness of predictability by the AFD to a sub-sample analysis by splitting the data into two roughly equal time periods of 20 years. The first period is starting at the beginning of 1973 until 1993 and the second time period ranges from 1993 to 2012. I chose these long sub-samples because the predictability of the AFD is stronger, given the high autocorrelation, at a business cycle frequency than for very short time horizons.

Table 2.5: **Predicting Commodity Portfolio Holding and Excess Holding Returns for Various Sub Samples:** This table presents the results of forecasting regressions for average holding return and average excess holding returns on the AFD for portfolios of commodities at horizons of n=1,2,3,6, and 9 months. Reported are the estimated slope coefficient $\psi_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$ and the adjusted R^2 values for the time series regression of the average holding return $\overline{rx}_{t\to t+n}$ and the average excess holding return $\overline{hrx}_{t\to t+n}$ on the AFD, $\overline{s_t} - \overline{f_{t\to t+n}}$, respectively. I report the results for Global and Precious portfolios for various sub-samples. The first sub-sample is from 1973 to 1993, the second from 1993 to 2012, the third from 1987 to 2007 and the final from 2005 to 2012. The t-statistics for the slope coefficients in brackets are computed using Newey and West (1987) adjusted standard errors with heteroscedasticity and autocorrelation consistent covariance estimators computed with the number of lags equal to the length of overlaps plus one lag.

	1973-1993								**1993-2012**							
	Global				Precious				Global				Precious			
Horizon	$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$	
n	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2
1	0.59	1.79			1.99	4.04			0.38	1.10			0.84	3.37		
	(1.92)				(3.81)				(1.84)				(2.66)			
2	0.76	4.03	0.07	3.65	1.91	1.87	0.08	0.35	0.41	3.24	0.01	-0.24	1.39	4.63	0.02	0.26
	(2.43)		(2.19)		(1.69)		(1.61)		(2.36)		(0.56)		(2.52)		(0.93)	
3	0.77	4.54	0.09	5.45	0.95	0.04	0.13	1.19	0.58	4.15	-0.00	-0.41	1.59	4.25	0.05	0.84
	(2.29)		(2.36)		(0.62)		(1.77)		(2.32)		(-0.04)		(2.29)		(1.16)	
6	1.00	10.85	0.13	4.93	2.97	5.05	0.17	2.43	0.84	5.78	-0.09	0.90	3.07	7.71	0.15	4.17
	(2.61)		(1.69)		(1.93)		(2.27)		(1.95)		(-0.85)		(2.45)		(2.14)	
9	1.11	14.03	0.09	1.71	3.88	10.81	0.10	1.02	0.90	6.92	-0.23	2.77	4.77	12.19	0.21	5.41
	(2.19)		(0.90)		(2.60)		(1.54)		(2.00)		(-1.11)		(2.67)		(2.12)	
	1987-2007								**2005-2012**							
	Global				Precious				Global				Precious			
Horizon	$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$	
n	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\vartheta_{\mathbf{f}}$	R^2
1	0.10	0.25	0.14	-0.20	0.81	3.27	1.83	3.15	0.19	0.14	0.19	0.14	0.84	3.37	0.84	3.37
	(1.04)		(0.69)		(2.29)		(3.37)		(1.04)		(1.04)		(2.66)		(2.66)	
2	0.31	3.05	0.04	0.95	1.28	3.88	0.07	0.68	0.41	1.22	0.01	-0.43	1.39	4.63	0.02	0.26
	(2.17)		(1.31)		(2.06)		(1.46)		(1.24)		(0.68)		(2.52)		(0.93)	
3	0.52	6.52	0.05	1.03	1.34	2.87	0.12	1.76	0.52	1.16	0.01	-0.75	1.59	4.25	0.05	0.84
	(2.66)		(1.21)		(1.74)		(1.61)		(1.11)		(0.24)		(2.29)		(1.16)	
6	0.83	7.61	0.06	0.20	2.61	5.45	0.16	4.87	0.68	0.68	0.01	-0.79	3.07	7.71	0.15	4.17
	(2.28)		(0.55)		(1.94)		(2.32)		(0.92)		(0.12)		(2.45)		(2.14)	
9	0.87	8.71	0.04	-0.26	4.08	9.05	0.14	6.07	0.79	0.57	-0.06	-0.51	4.77	12.19	0.21	5.41
	(1.62)		(0.25)		(2.15)		(2.25)		(0.89)		(-0.41)		(2.67)		(2.12)	

In the early period from 1973 to 1993, I find highly significant coefficient estimates that are in magnitude similar to the full sample results. This result holds for holding and excess holding returns for the Global and Precious commodity portfolios. I furthermore find that the adjusted R^2 values are even higher than for the full time period. For the late period from 1993 to 2012 I find mostly significant coefficients. While for Precious Metals the adjusted R^2 values are even higher than for the late period, I find that the predictability for holding and excess holding returns of the Global selection deteriorated significantly. The magnitude of the coefficient estimates, however, remains relatively stable, suggesting that a change in volatility in the late period can be the source of the deteriorated statistical significance. To examine this effect I test for two additional subsamples in Table (2.5) Panel B. I find that when restricting the sample to the period from 1987 to 2007 coefficient estimates are highly significant with adjusted R^2 values of up to 8 percent. For the very late time period from 2005 to 2012 I cannot obtain significant coefficient estimates anymore. This indeed suggests that the high volatility after 2007 in the commodity markets reduced the significance of the predictive regression.

Alternative Construction Methods of AFD In this passage I evaluate the predictive content of alternative construction methods for the AFD. In particular, I evaluate the predictive content of including futures discounts of different maturities into the predictive regressions. This is motivated by Cochrane and Piazzesi (2002) who show that the whole term structure of bond forward rates contains information about expected bond excess returns. I use two methods to include all maturity futures discounts in the regression framework. First, I use a naive measure that equally weights futures discounts across maturities and across all commodities. This measure is similar to the predictors used in Hong and Yogo (2013). Second, I extract principal components of futures discounts from the panel of all commodities and across futures maturities. The principal components are linear combinations of the individual futures discounts constructed to best explain the total variation in the data. I use the first three principal components as measures for the AFD, which explain more than 40 percent of the total variation of the panel of individual futures discounts. The construction methodology is similar to Gospodinov and Ng (2011).

The upper panel of Table (2.6) presents the results from the forecasting regression of commodity holding returns and excess holding returns on the AFD across all maturities and contracts. For the

Table 2.6: **Predicting Commodity Portfolio Holding and Excess Holding Returns with Alternative Construction Methods of the Aggregate Futures Discount:** This table presents the results of forecasting regressions of commodity portfolios on alternative construction methods for aggregate futures discounts for horizons of n=1,2,3,6, and 9 months. I use two alternative construction methods for the aggregate futures discount: First, I construct the aggregate futures discount as the average of all individual futures discounts across assets and maturities (Avg). Second, I extract the first three principle components of all individual futures discounts across assets and maturities (PC1, PC2 and PC3). The principal components are constructed as linear combinations of the individual futures discounts such that they explain most cross sectional variation. Reported are the estimated slope coefficient ψ and ϑ for the alternative versions of aggregate futures discounts.

	Global				Precious			
Horizon	$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{hrx}_{t\to t+n}$	
n	$\psi_{\mathbf{Avg}}$	R^2	$\vartheta_{\mathbf{Avg}}$	R^2	$\psi_{\mathbf{Avg}}$	R^2	$\vartheta_{\mathbf{Avg}}$	R^2
1	0.23	0.02	0.23	0.02	0.84	0.03	0.84	0.03
	(2.39)		(2.39)		(2.66)		(2.66)	
2	0.52	0.05	0.04	0.04	1.39	0.05	0.02	0.00
	(3.03)		(3.10)		(2.52)		(0.93)	
3	0.76	0.07	0.06	0.04	1.59	0.04	0.05	0.01
	(3.16)		(2.52)		(2.29)		(1.16)	
6	1.36	0.09	0.07	0.01	3.07	0.08	0.15	0.04
	(3.02)		(0.97)		(2.45)		(2.14)	
9	1.72	0.09	-0.03	-0.00	4.77	0.12	0.21	0.05
	(2.50)		(-0.25)		(2.67)		(2.12)	

	Global							
Horizon	$\overline{rx}_{t\to t+n}$				$\overline{hrx}_{t\to t+n}$			
n	$\psi_{\mathbf{PC1}}$	$\psi_{\mathbf{PC2}}$	$\psi_{\mathbf{PC3}}$	R^2	$\psi_{\mathbf{PC1}}$	$\psi_{\mathbf{PC2}}$	$\psi_{\mathbf{PC3}}$	R^2
1	0.00	0.00	-0.00	0.02	0.00	0.00	-0.00	0.02
	(2.81)	(0.76)	(-1.11)		(2.81)	(0.76)	(-1.11)	
2	0.00	0.00	-0.00	0.04	0.00	-0.00	-0.00	0.02
	(3.36)	(0.87)	(-1.05)		(2.15)	(-1.34)	(-1.23)	
3	0.00	0.00	-0.00	0.06	0.00	-0.00	-0.00	0.02
	(3.25)	(1.05)	(-1.05)		(1.27)	(-0.08)	(-1.90)	
6	0.01	0.00	-0.00	0.08	-0.00	0.00	-0.00	0.06
	(2.73)	(1.32)	(-1.55)		(-0.28)	(2.04)	(-2.69)	
9	0.01	0.00	-0.01	0.07	-0.00	0.00	-0.00	0.11
	(2.15)	(1.47)	(-1.69)		(-1.37)	(2.65)	(-2.76)	

Global portfolio I find significant coefficient estimates across all maturities. Similar to the regular AFD I find that the coefficient estimates and the adjusted R^2 values are increasing with holding maturity. The R^2 values are significantly higher than for the regular AFD for short horizons. The R^2 is 2 percent for a 1 month horizon and increases to 7 percent on a 3 months horizon. For 6 and 9 months horizons I find R^2 values of 9 percent. For the excess holding returns I find similar results as for the regular AFD. For the Precious commodity portfolio I find very similar results as in the regular AFD case for both holding and excess holding returns. In the lower panel of Table (2.6), I present the results for the forecasting regressions on the first three principal components. The results are very similar to the results in the upper panel. Noteworthy is that all of the predictive power is related to the first principal component. This suggest that indeed a common factor structure across commodities is responsible for the results in this chapter. In summary, I find that for Global portfolios futures discounts across different maturities indeed contain additional information to the AFD. However, the information content is not as significant as in the bond literature, see e.g. Cochrane and Piazzesi (2002).

AFD, common systematic and commodity specific variables Previous research, such as Bessembinder and Chan (1992), De Roon et al. (2000) and Hong and Yogo (2013), show that common systematic variables and commodity specific variables have some power to predict commodity returns. In Table (2.7) I test whether the common systematic variables also predict portfolios of holding returns and excess holding returns in addition to the AFD. These variables include the Dividend Yield (DivY) and the Cochrane and Piazzesi (2002) bond risk factor (CP), and commodity specific variables, such as the aggregate hedging pressure variable and open interest growth of Hong and Yogo (2013). To test the common predictive power I run the regressions,

$$\overline{rx}^j_{t\to t+n} = \psi^j_0 + \psi_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \beta Z_t + \eta_{t+n} \tag{2.7}$$

and

$$\overline{hrx}^j_{t\to t+n} = \vartheta^j_0 + \vartheta_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \beta Z_t + \epsilon_{t+n} \tag{2.8}$$

where Z_t is a vector of common systematic variables.

Table 2.7: **Predicting Commodity Portfolio Holding and Excess Holding Returns with Average Futures Discount (AFD), Common Systematic and Commodity Specific Variables:** This table presents the results of forecasting regressions of commodity portfolios on the AFD, common systematic and commodity specific variables for horizons of n=2,3,6, and 9 months. I use the Dividend Yield (DivY) and the bond risk factor (CP) by Cochrane and Piazzesi (2002) as common systematic variables (upper panel) and the equally weighted growth in open interest (OI) and the average hedging pressure (HP) across commodities as commodity market specific variables (lower panel). Hedging pressure is defined as the number of commercial short position minus commercial long positions divided by the total open interest of commercials as defined by the Commodity Futures Trading Commission (CFTC). Reported are the estimated slope coefficient $\psi_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$ for the average futures discount, the slope coefficients β for all other variables and the adjusted R^2s.

	Global								Precious							
Horizon	$\overline{rx}_{t\to t+n}$				$\overline{hrx}_{t\to t+n}$				$\overline{rx}_{t\to t+n}$				$\overline{hrx}_{t\to t+n}$			
n	$\psi_{\mathbf{f}}$	β_{DivY}	β_{CP}	R^2	$\vartheta_{\mathbf{f}}$	β_{DivY}	β_{CP}	R^2	$\psi_{\mathbf{f}}$	β_{DivY}	β_{CP}	R^2	$\vartheta_{\mathbf{f}}$	β_{DivY}	β_{CP}	R^2
2	0.77	-0.00	-0.00	4.82	0.06	-0.00	0.00	3.69	1.05	0.02	-0.01	5.54	0.07	-0.00	0.00	0.26
	(3.35)	(-0.08)	(-0.85)		(2.64)	(-0.37)	(1.99)		(1.29)	(2.95)	(-2.75)		(2.24)	(-0.65)	(0.51)	
6	1.10	-0.00	-0.01	13.11	0.05	-0.00	0.00	0.54	3.90	0.06	-0.01	11.74	0.16	-0.00	0.00	4.51
	(2.98)	(-0.01)	(-1.11)		(0.86)	(-0.84)	(0.51)		(1.46)	(2.44)	(-0.68)		(1.58)	(-0.21)	(0.24)	
9	1.16	0.00	-0.01	15.56	-0.00	-0.00	0.00	0.18	7.62	0.11	0.01	22.25	0.13	-0.00	0.00	5.90
	(2.57)	(0.17)	(-1.38)		(-0.04)	(-0.80)	(0.42)		(2.59)	(2.68)	(0.52)		(0.73)	(-0.30)	(0.11)	
	Global								Precious							
Horizon	$\overline{rx}_{t\to t+n}$				$\overline{hrx}_{t\to t+n}$				$\overline{rx}_{t\to t+n}$				$\overline{hrx}_{t\to t+n}$			
n	$\psi_{\mathbf{f}}$	β_{OI}	β_{HP}	R^2	$\vartheta_{\mathbf{f}}$	β_{OI}	β_{HP}	R^2	$\psi_{\mathbf{f}}$	β_{OI}	β_{HP}	R^2	$\vartheta_{\mathbf{f}}$	β_{OI}	β_{HP}	R^2
2	0.74	0.00	0.06	6.89	0.06	-0.00	-0.00	2.42	2.43	0.00	-0.08	2.63	0.11	0.00	-0.00	2.30
	(3.35)	(1.78)	(0.96)		(2.57)	(-0.22)	(-0.41)		(1.81)	(0.66)	(-1.13)		(1.80)	(1.58)	(-0.80)	
6	1.06	0.01	0.04	13.56	0.02	0.00	-0.00	3.08	3.22	0.02	-0.28	11.28	0.21	0.00	0.01	12.20
	(3.15)	(1.30)	(0.44)		(0.31)	(2.00)	(-0.22)		(2.73)	(1.32)	(-2.07)		(2.97)	(0.30)	(1.73)	
9	1.12	0.02	-0.08	17.64	-0.07	0.01	-0.02	6.44	3.39	0.02	-0.39	18.24	0.20	0.00	0.02	16.69
	(2.88)	(1.75)	(-0.95)		(-0.69)	(2.59)	(-1.14)		(3.19)	(1.36)	(-1.98)		(3.16)	(0.09)	(2.11)	

The upper panel in Table (2.7) presents the results for the regressions (2.7) and (2.8) in Global and Precious portfolios on the AFD, the DivY and the CP factor. For the Global portfolios I find that the AFD is still significantly related to holding returns when controlling for the DivY and CP factor. The DivY and CP factor do not seem to add additional information to expected holding returns for any holding periods. For the excess holding returns I find that short horizon returns are statistically significant related to the AFD and the CP factor, while the DivY remains insignificant. It is not surprising that the CP factor is significantly related to expected excess holding returns as those are affected by risk in the convenience yield and interest rates. For the Precious portfolio I find that the DivY is significantly related to holding returns across holding periods. The CP factor is only significant for short maturities. At the same time, the AFD is not significant anymore for horizons of 2 and 6 months but becomes significant for horizons of 9 months. For excess holding returns I do not find any significant relation between the systematic variables and return variation. In summary the DivY and CP factor do not seem to have additional explanatory power for Global portfolios. The results indicate that Precious metals are more related to systematic fluctuation in asset markets than all other commodities.

The lower panel in Table (2.7) presents the results for the regressions of holding returns and excess holding returns in Global and Precious portfolios on the average futures discount, the open interest growth (OI) and hedging pressure (HP). Similar to Hong and Yogo (2013), I find that OI has some additional power to the AFD in predicting holding returns at all maturities for the Global portfolios. The hedging pressure variable is not significant, however. Note that open interest is significantly related to excess holding returns at maturities of 6 and 9 months horizon, when the AFD is not significant anymore. For the Precious portfolios I find that hedging pressure is negatively related to holding returns and positively related to excess holding returns. While a positive sign is consistent with the theory of hedging pressure, see De Roon et al. (2000), a negative sign is troubling. In summary, I find that commodity specific variables do indeed possess some predictive power for holding and excess holding returns but do not subsume the AFD as a predictor, indicating that other sources of risk are responsible for time variation in expected returns.

Given the significant in-sample predictability results, it is of interest for investors to assess the predictive ability of the AFD using only real-time data. From a rational-expectations perspective,

this additional exercise is not necessary given that investors have historical information, and evolved rules of thumb that summarize far longer time series than the data set I use. Therefore, there is no error in using the full-sample forecast results. Nevertheless, I use out-of-sample regression techniques and simple trading strategies to assess the statistical and economic value for an investor, who has access only to real time data.

Out of Sample Results I use two out of sample metrics that have been suggested in the literature: (1) the out-of-sample R^2 by Campbell and Thompson (2008), and (2) the adjusted mean-square prediction error statistic, as introduced by Clark and West (2007). Both tests compare the predictive ability of the AFD to a benchmark model where commodity returns are regressed on a constant, month-by-month, to provide forecasts at each month based on the historical mean. First, the out-of-sample R^2 values which I will denote as OOS R^2, is defined as

$$R^2_{OS} = 1 - \frac{\sum_{t=0}^{T-1}(rx_{i,t\to t+1} - \hat{rx}_{i,t\to t+1})}{\sum_{t=0}^{T-1}(rx_{i,t\to t+1} - \overline{rx}_{i,t\to t+1})} \tag{2.9}$$

where $\hat{rx}_{i,t\to t+1}$ is the prediction for month t+1 from the main regression in equation (2.3) and $\overline{rx}_{i,t\to t+1}$ is the historical mean. I calculate the OOS R^2 values using an expanding window, where I estimate $\hat{rx}_{i,t\to t+1}$ and $\overline{rx}_{i,t\to t+1}$ using all data available up to and including month t. I assume that an investor can utilize all available data when making the regression forecast.

The second test I use is the adjusted mean-square prediction error statistic by Clark and West (2007). In contrast to the OOS R^2 the MSPE-adjusted statistic assesses the statistical significance of out-of-sample predictions of the AFD. Specifically, I investigate whether this test significantly improves on a forecast based solely on the historical average return. As before I use as a benchmark model the random walk specification of the holding returns. I follow Clark and West (2007), who develop an adjusted mean-squared prediction model (MSPE) statistic for evaluating forecasting models, by defining:

$$f_{t+1} = (rx_{i,t\to t+1} - \overline{rx}_{i,t\to t+1})^2 - [(rx_{i,t\to t+1} - \hat{rx}_{i,t\to t+1})^2 - (\overline{rx}_{i,t\to t+1} - \hat{rx}_{i,t\to t+1})^2], \tag{2.10}$$

where $\overline{rx}_{i,t\to t+1}$ and $\hat{rx}_{i,t\to t+1}$ are defined as before. The t-statistic from the regression of f_{t+1} on

a constant is the MSPE-adjusted statistic.

Table (2.8) Panel A presents the results for the OOS R^2 values and the MSPE-adjusted statistic, based on an expanding window with initial length of 60 and 120 months. The OOS R^2 values for the Global holding portfolios are consistently positive across horizons of 1 to 9 months. For a horizon of 1 to 3 months and a 60 months expanding window the OOS R^2 values are around 2 percent and increase to 6 percent for a holding horizon of 6 and 9 months. For the 120 months expanding window the values are even higher for the short horizons. For the Precious portfolios the 1 month holding return does not have positive OOS R^2 values but this changes for horizons from 2 months to 9 months which are all positive. Here, the OOS R^2 values are as high as 7 or 12 percent for the 9 months holding period. The results for the MSPE adjusted statistic are even more encouraging. For all specifications of Global or Precious portfolios the MSPE have statistically significant positive estimates. In summary, the results indicate that the AFD does indeed possess incremental forecasting power for a real time investor, in contrast to a random walk benchmark model.

Economic Significance of Predictability Even though out-of-sample regression tests provide an indication for the predictive strength of the AFD from the perspective of a real-time investor, they do not need to coincide with economic profits when trading on the information content of the regression. For example, Thornton and Valente (2012) show that even though bond risk premia are highly predictable in and out-of-sample the predictions do not provide any economic gains from an investors perspective.

In this section I use two very simple trading strategies to explore the economic significance of the predictive regressions. In the first strategy investors go long the commodity portfolio when the AFD is above its long term mean and short the commodity portfolio otherwise. I call this investment strategy *Simple Trading Strategy*. In the second trading by Cochrane and Piazzesi (2002), I take a position in the holding returns, which size corresponds to the expected value from the forecasting regression, $\hat{rx}_{i,t\to t+1}$. In other words, I calculate the trading rule profits according to

$$rx_{i,t\to t+1} \times \hat{rx}_{i,t\to t+1} = rx_{i,t\to t+1} \times (\psi_{\mathbf{f}}(\overline{s}_t^i - \overline{f}_{t\to t+n}^i)). \tag{2.11}$$

To estimate the regression, I use only data up to time t within an expanding window. To com-

Table 2.8: **Economic and Out-of-Sample Significance of Forecasting Regressions:** This table reports the out-of-sample and economic significance of the predictive regression, where holding returns are regressed on the lagged AFD. In panel A the table presents the out-of-sample R^2 statistic (in %) from Campbell and Thompson (2008) and the one-sided t-statistic for the MSPE-adjusted statistic, developed by Clark and West (2007), for the null hypothesis that using the AFD does not significantly improve on a forecast based solely on the historical average returns. The results are presented for an expanding window with initial length of 60 and 120 months for the Global and Precious futures portfolios. Panel B presents the economic significance of the predictive relation for two trading strategies. The first trading strategy goes long the futures portfolio if the average futures discount is above its long term term and short otherwise. The second investment strategy invests in a futures portfolio with a weight proportional to the estimated regression coefficient times the signal strength. The table reports the mean, standard deviation and Sharpe ratio of these investment strategies. Furthermore, the p-value of the Jobson-Korkie-Memmel test statistic, which tests for equality of Sharpe ratios, is reported. We test for equal Sharpe ratios relative to the buy and hold strategy. The table also reports the results for the long only strategy. For all calculations and regressions I use only the data available to an investor in real time. The sample period is 1973 to 2012.

Panel A:Out of Sample Results								
	60 months				120 months			
	Global		Precious		Global		Precious	
Horizon	OOS $\overline{R}^2$	MSPE	OOS $\overline{R}^2$	MSPE	OOS $\overline{R}^2$	MSPE	OOS $\overline{R}^2$	MSPE
1	1.84	3.25	-1.19	2.20	4.38	2.77	-2.67	2.13
2	2.35	2.96	5.03	2.51	4.84	2.87	-0.23	2.53
3	1.37	2.51	0.09	2.41	3.53	2.50	4.31	2.61
6	6.37	2.09	4.22	2.33	5.77	2.45	9.08	2.72
9	6.95	1.73	7.12	2.65	3.16	2.13	11.69	2.61
Panel B: Economic Significance								
	Simple			Regression			Long only	
	Mean	Std.Dev.	Sharpe	Mean	Std.Dev	Sharpe	Mean	Std.Dev	Sharpe
Non Precious	0.09	0.21	0.44	0.13	0.22	0.58	0.01	0.15	0.04
JKM Test			(0.08)			(0.04)			
Precious	0.06	0.19	0.30	0.09	0.20	0.47	0.00	0.25	0.00
JKM Test			(0.15)			(0.06)			

pare these strategies, I lever the commodity positions ex-post so that all commodity investment strategy returns are as volatile as the long-only commodity holding return. Furthermore, I use the Jobson-Korkie-Memmel test (Memmel (2003) and Jobson and Korkie (1981) to assess the statistical significance of the trading strategy Sharpe ratios relative to the long only benchmark.

Table (2.8) Panel B presents the results for the two trading strategies. For the Global commodity portfolios I find that the simple trading strategy has a mean of 9 percent and a standard deviation of 21 percent, yielding a Sharpe ratio of 0.44. The regression trading strategy is even more profitable with an annualized mean of 13 percent and a standard deviation of 22 percent, yielding a Sharpe ratio of 0.58. Both trading strategies are highly significant and result in economic gains to an investor compared to the long-only strategy that has a Sharpe ratio of only 0.04. This result is also confirmed by the Jobson-Korkie-Memmel test that displays significant p-values for both conditional trading strategies. The same is true for the Precious portfolios. The simple strategy generates a Sharpe ratio of 0.3 and the regression trading strategy of 0.47. Both strategies are significantly better than the long only Sharpe ratio of almost zero.

Of course, real trading profits should also incorporate estimates of the conditional variance of the returns that are heteroscedastic in commodity markets. I leave a more comprehensive assessment of the economic value of predictive regressions in commodity markets for future research.

2.4 Time Variation in Spot and Term Premia

The previous section shows that holding returns of commodity futures portfolios are predictable across various maturities. In this section I go one step further and use the decomposition of holding returns into spot and term premia introduced by Szymanowska et al. (2013) to study the relative contribution of the underlying spot risk and the risk in the convenience yield for the variation in holding returns.

Each commodity futures contract is available with various maturities. Therefore, analogous to bonds, futures contracts possess a term structure of prices. Exposure to investment strategy returns for an n-period horizon can consequently be obtained in at least three dimensions. First, investors can open futures contracts along the term structure, which, e.g., expire in n months, and close the contract just prior to maturity. Second, investors can open futures contracts, e.g.,

the contract closest to maturity, hold it for one period and close the contract, then open the next closest to maturity contract and close it again after one period. Repeating the procedure for n-periods results in an alternative n-period investment strategy return. Third, the investor can open a long and a short position in two futures contracts with different times to maturity and close these contracts after n periods, exposing the investor to movements in the term structure of prices. Because futures contracts are zero cost securities these investment strategies consist of risk premia only. In the following I formalize these thoughts.

In this section I decompose holding returns into spot and term premia that can be identified by taking long positions in short maturity contracts and by combining long and short spreading positions in contracts with different maturities, respectively. Decomposing holding returns into spot and term premia is important for two reasons: (1) Term and spot premia are likely to compensate agents for different sources of risk. While spot premia reflect the price risk of the underlying commodities within the portfolios, term premia compensate agents for the risk present in the term structure or convenience yields of the commodities. (2) Time variation of holding returns, as studied in the previous section might be due to time variation in spot premia, in term premia or a combination of both.

Portfolio Spot Premia The expected log spot price change for an equally-weighted commodity portfolio j is defined as

$$\mathbb{E}_t[\Delta \overline{s}_{j,t+1}] := \mathbb{E}_t[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}\Delta s_{i,t+1}] = \frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(\delta_{i,t} - r_t) + \overline{\pi}_{s,t}, \tag{2.12}$$

where $\overline{\pi}_{s,t}$ is defined as the portfolio spot premium. The spot premium $\overline{\pi}_{s,t}$ can be interpreted as the expected return in excess of the short-term convenience yield and interest rate, similar to stock returns in excess of the short-term interest rate and dividend yield. The portfolio spot risk premium is not only related to changes in spot prices but is equal to the one maturity portfolio holding return, $\overline{rx}^{(1)}$, which follows immediately from the cost of carry formula,

$$\mathbb{E}_t[\overline{rx}^{(1)}_{t+1}] = \mathbb{E}[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(s_{i,t+1} - f^{(1)}_{i,t})] = \mathbb{E}_[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(s_{i,t+1} - s_{i,t} - y_t^1)] = \frac{1}{N_t^j}\sum_{i=1}^{N_t^j}\mathbb{E}_t[s_{i,t+1} - s_{i,t}] - \frac{1}{N_t^j}\sum_{i=1}^{N_t^j} y_t^1 = \overline{\pi}_{s,t} \tag{2.13}$$

The result from equation (2.13) presents an easy method to extract spot premia from trading short-term futures contracts.

Portfolio Term Premia. By defining $y_t^{(n)} := (r_t - \delta_t)$ and taking logarithms, I can rewrite the cost of carry relation in equation (2.1) as,

$$y_t^{(n)} = \frac{f_t^{(n)} - s_t}{n}. \tag{2.14}$$

As in De Roon et al. (1998), I define the forward yield, $h_t^{(n)}$, as the yield or return the investor can earn from simultaneously taking a long position in a futures contract maturing at time $t+n$ and a short position in a futures contract with a shorter maturity. The forward yield can then be defined as

$$h_t^{(n)} := \frac{f_t^{(n)} - f_t^{(1)}}{n-1} = \frac{n y_t^{(n)} - y_t^{(1)}}{n-1} \tag{2.15}$$

where the second equation is a direct consequence of the cost of carry relation in equation (2.14).

In defining the term premium, $\pi_{y,t}^{(n)}$, I decompose the forward yield into the expected futures yield and a risk premium. The term premia in commodity markets are, therefore defined as the expected deviation from the expectations hypothesis of the term structure of the yield, which is similar to term premia in bond markets. More specifically I assume that,

$$h_t^{(n)} = \frac{n y_t^{(n)} - y_t^{(1)}}{n-1} := \mathbb{E}_t[y_{t+1}^{n-1}] - \frac{\pi_{y,t}^{(n)}}{n-1} \tag{2.16}$$

This definition allows me to extract term premia through trading strategies in futures markets. I can rewrite equation (2.16) in terms of futures and spot prices using equation (2.14) as

$$\begin{aligned} \pi_{y,t}^{(n)} = n y_t^{(n)} - y_t^{(1)} - \mathbb{E}_t[y_{t+1}^{(n-1)}] &= f_t^{(n)} - f_t^{(1)} - \mathbb{E}_t[f_{t+1}^{(n-1)} - s_{t+1}] \\ &= \mathbb{E}_t[f_{(t+1)}^{n-1} - f_t^{(n)}] - \mathbb{E}_t[s_{t+1} - f_t^{(1)}] \end{aligned} \tag{2.17}$$

The last term in the equation is the spot premium defined in equation (2.13). Equation (2.17) shows that the term premium equals the expected one period return on a spreading strategy, which involves a long position in a futures contract with n periods to maturity and a short position in a

one period futures contract.

Equivalently, I can define the portfolio term premium for each portfolio j, $\overline{\pi_{y,t}}^{(n)}$, as

$$\overline{\pi_{y,t}}^{(n)} = \mathbb{E}_t[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(f_{(t+1)}^{n-1} - f_t^{(n)})] - \mathbb{E}_t[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(s_{t+1} - f_t^{(1)})]. \tag{2.18}$$

Using this equation it is possible to extract term premia for contracts with different maturities. I will focus on the portfolio term premia definition in equation (2.18) but interested readers can resort to Szymanowska et al. (2013) for alternative constructed methods.

Horizons and Maturities In this section I extend the one period results of the portfolio spot and term premia in the previous sections to the multi period case. First, the conditional expectation or risk premium of the portfolio holding return ,$\overline{rx}_{t\to t+n}^j$, consist of a combination of spot and term premia over time. In particular

$$\begin{aligned}
\mathbb{E}_t[\overline{rx}_{t\to t+n}^j] &= \mathbb{E}_t[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j} rx_{i,t\to t+n}^j] = \mathbb{E}_t[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(s_{i,t+n} - f_t^{(n)})] \\
&= \mathbb{E}_t[\frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(s_{i,t+n} - f_{t+n-1}^{(1)}) + \frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(f_{t+n-1}^{(1)} - f_{t+n-2}^{(2)}) + ... \\
&... + \frac{1}{N_t^j}\sum_{i=1}^{N_t^j}(f_{t+1}^{(n-1)} - f_t^{(n)})] \\
&= \sum_{k=0}^{n-1}\mathbb{E}_t[\overline{\pi}_{s,t+k}] + \sum_{k=0}^{n-1}\mathbb{E}_t[\overline{\pi}_{y,t+j}^{(n-j)}]
\end{aligned} \tag{2.19}$$

Second, I show how to extract portfolio spot premia for longer maturities. In equation (2.13) I show that 1 month horizon portfolio futures returns consist only of spot premia. The returns on those contracts are given by $\overline{rx}_{t+k}^{(1)}$ for k=1,2,...,n. I can therefore extract future spot premia, $\overline{frx}_{t\to t+n}^j$, by investing in one period futures portfolios for n consecutive periods to get

$$\mathbb{E}_t[\overline{frx}_{t\to t+n}^j] = \mathbb{E}_t[\sum_{k=1}^{n}\overline{rx}_{t+k}^{(1)}] = \sum_{k=1}^{n-1}\mathbb{E}_t[\overline{\pi}_{s,t+k}] \tag{2.20}$$

This strategy of rolling forward short maturity futures contracts is the method institutional in-

vestors in commodity markets employ. I will denote the strategy as the short roll.

By comparing the expected returns from the holding returns and the short rolls over varying maturities I can immediately see how to extract long maturity term premia, $\overline{hrx}^j_{t\to t+n}$. I can go long in the holding return strategy and take a short position in the short roll strategy which results in,

$$\mathbb{E}_t[\overline{hrx}^j_{t\to t+n}] = \mathbb{E}_t[\overline{rx}^j_{t\to t+n} - \sum_{k=1}^{n} \overline{rx}^{(1)}_{t+k}] = \sum_{k=0}^{n-1} \mathbb{E}_t[\overline{\pi}^{(n-j)}_{y,t+j}]. \tag{2.21}$$

I will denote the strategy as the *excess holding strategy*. The excess holding strategy is defined for maturities of n=2 or above.

Summary Statistics Front Roll and Excess Holding Returns Table 2.2 Panel B reports the summary statistics of the front rolls $\overline{frx}^j_{t\to t+n}$, which isolate the spot premia, and excess holding returns $\overline{hrx}^j_{t\to t+n}$, which isolate the term premia, for maturity horizons of n=1,2,3,6, and 9 months. The annual mean returns for front rolls for the Global portfolios show a pattern that is decreasing with maturity, from 0.85 percent for a 1 month horizon to 0.66 percent for a 9 month horizon. At the same time the excess holding returns are increasing with maturity from 0.17 percent for a 2 months horizon to 1.43 percent for a nine month horizon. The volatilities of the front roll returns are approximately the same as for the holding returns. The volatilities of the excess holding returns are, however, significantly lower and range from and annual 1.25 percent for a 2 month horizon to 3.34 percent for a nine month horizon. Therefore, the unconditional returns of excess holding strategies are statistically significant. The results for Non-Precious and Precious portfolios show similar patterns, even though, the excess holding returns are not statistically significant.

Figure 2 plots the 12 months rolling average returns to holding and excess holding returns for the commodity portfolios. The profitability of holding returns is time-varying for Global and Precious portfolios. The time variation of excess holding returns is far less pronounced than for the holding returns, however, still observable. More interestingly, the excess holding returns seem to be of exactly the opposite to the holding returns; a result similar to findings in Szymanowska et al. (2013).

Spot and Term Premia In a similar vain as the complimentary regression framework for holding returns I test for time variation in spot and term premia. In equation (2.17) I find that expected

holding returns can be decomposed into spot and term premia:

$$\mathbb{E}_t[\overline{hrx}^j_{t\to t+n}] + \mathbb{E}_t[\overline{frx}^j_{t\to t+n}] = \mathbb{E}_t[\overline{rx}^j_{t\to t+n}] = \mathbb{E}_t[\Delta \overline{s}_{t+1}] + \overline{s}_{j,t} - \overline{f}^{(n)}_{j,t\to t+n}, \tag{2.22}$$

where $\overline{hrx}^j_{t\to t+n}$ are portfolio excess holding returns and $\overline{frx}^j_{t\to t+n}$ are front roll returns. If futures discounts imply time-varying expected holding returns, then the time variation is due to variation in spot or term premia or a combination of both. I can test this relation in the complimentary regression framework given by

$$\overline{frx}^j_{t\to t+n} = \theta^j_0 + \theta_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \eta_{t+n} \tag{2.23}$$

and

$$\overline{hrx}^j_{t\to t+n} = \vartheta^j_0 + \vartheta_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \epsilon_{t+n}. \tag{2.24}$$

By construction equation (2.23) and (2.24) impose restrictions on the coefficient estimates of θ_0 and ϑ_0 as well as on $\theta_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$. The sum of the excess holding and front roll returns is equal to the futures discount, thus the intercepts θ_0 and ϑ_0 must sum up to ψ_0 and the slope coefficients $\theta_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$ sum up to $\psi_{\mathbf{f}}$. In this regression, the variation in the basis is allocated to the spot premia, term premia, or a combination of both. As in the previous regression, because of the adding up constraint it would have been enough to run one of the regressions in equation (2.23) and (2.24).

The regressions in equation (2.23) and (2.24) test the following hypothesis. In both regressions, the null states that the front roll returns and excess holding returns are constant. Coefficient estimates of $\theta_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$ that are statistically different from zero implies that the futures discount contains information on spot and/or term premia.

Predictability Results for Spot and Term Premia Table (2.3) presents the results from the regression of equation (2.23) and (2.24) over the sample period from 1973 to 2012. The left panel reports the results from the Global w. Precious portfolio selection of regressing the front roll returns and excess holding returns on the AFD for horizons of n=1,2,3,6, and 9 months. I find that front roll returns are highly predictable with coefficient estimates similar to the holding returns. The adjusted R^2 values are increasing from around 1.2 percent on the 1 month horizon to 7 percent

Table 2.9: **Predicting Commodity Holding Returns and Spot Changes of Futures Portfolios with the Average Futures Discount (AFD):** This table reports results of forecasting regressions for average holding returns and average spot changes on the average futures discount for commodity portfolios at horizons of n=1,2,3,6, and 9 months. The first four columns report the predictability results for Global w. Precious commodity portfolios. Reported are the estimated slope coefficient $\psi_{\mathbf{f}}$ and $\zeta_{\mathbf{f}}$ and the adjusted R^2 values for the time series regression of the average log holding return $\overline{rx}_{t\to t+n}$ and the log spot return $\overline{\Delta s}_{t\to t+n}$ on the AFD, $\overline{s_t} - \overline{f_{t\to t+n}}$. The t-statistics for the slope coefficients in brackets are computed using Newey and West (1987) standard errors (NW) with heteroscedasticity and autocorrelation consistent covariance estimators computed with the number of lags equal to the length of overlap plus one lag and Hansen and Hodrick (1990) standard errors (HH). I repeat the analysis for Global portfolios in columns five to eight and Precious portfolios in columns nine to thirteen. The sample period is January 1973 to September 2012 for Global and and Global without Precious portfolios and January 1976 to September 2012 for the Precious portfolio.

	Global w. Precious				Global				Precious Metals			
	Holding		Spot		Holding		Spot		Holding		Spot	
Horizon	$\overline{rx}_{t\to t+n}$		$\overline{\Delta s}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{\Delta s}_{t\to t+n}$		$\overline{rx}_{t\to t+n}$		$\overline{\Delta s}_{t\to t+n}$	
n	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2	$\psi_{\mathbf{f}}$	R^2	$\zeta_{\mathbf{f}}$	R^2
1	0.49	1.13	-0.50	1.19	0.44	1.06	-0.56	1.82	2.00	3.04	1.00	0.61
NW	(2.19)		(-2.23)		(2.28)		(-2.89)		(3.71)		(1.85)	
HH	[2.19]		[-2.24]		[2.27]		[-2.88]		[3.78]		[1.89]	
2	0.77	3.52	-0.23	0.13	0.70	3.49	-0.30	0.50	2.12	2.73	1.12	0.62
NW	(3.12)		(-0.92)		(3.23)		(-1.39)		(2.07)		(1.09)	
HH	[2.92]		[-0.86]		[2.99]		[-1.29]		[2.05]		[1.08]	
3	0.82	4.34	-0.19	0.04	0.74	4.26	-0.26	0.37	1.75	2.08	0.75	0.21
NW	(2.89)		(-0.66)		(2.93)		(-1.03)		(1.47)		(0.63)	
HH	[2.58]		[-0.59]		[2.60]		[-0.92]		[1.41]		[0.60]	
6	1.05	9.02	0.05	-0.18	0.96	8.45	-0.05	-0.19	2.69	8.23	1.69	3.29
NW	(3.02)		(0.15)		(2.96)		(-0.14)		(2.56)		(1.61)	
HH	[2.55]		[0.13]		[2.50]		[-0.12]		[2.18]		[1.37]	
9	1.16	10.34	0.15	-0.19	1.07	9.13	-0.07	-0.16	2.94	13.48	1.94	6.26
NW	(2.57)		(0.11)		(2.47)		(-0.19)		(2.82)		(1.86)	
HH	[2.15]		[0.09]		[2.08]		[-0.16]		[2.37]		[1.56]	

for the 9 months horizon. Therefore, the predictability results show that expected aggregate spot premia are time varying.

Next, I present the regression results for Global excess holding returns on the AFD. For the 2 and 3 months horizon I find significant predictable variation of excess holding returns. The predictability, however, deteriorates with maturity. In contrast to the front roll or holding returns, I find significant constant risk premia as evident in a statistically significant ϑ_0 coefficient. This result persists for all maturities considered. Note that for long maturities the excess holding returns satisfy a version of the expectation hypothesis, while for short horizons I find significant time varying term premia in addition to a constant risk premium.

For the Global selection in the two central panels I find almost identical results to the Global selection. The coefficient estimates of the front roll regression and the excess holding regression are close to the Global w. Precious results and the R^2 values are only slightly lower.

I next focus on the final two panels of Table (2.3) that present the regression results for the Precious selection. Again, I find that the front roll returns are highly predictable with respect to the AFD as confirmed by statistically significant coefficient estimates. Furthermore, the adjusted R^2 values increase from 3 percent on a monthly horizon to 12 percent on the nine month horizon. For excess holding returns the results are different than for the Global and Global with Precious selection. I find that the significance of the $\vartheta_{\mathbf{f}}$ coefficients increase with maturity. In particular for horizons of two and three months the coefficient is insignificant and becomes significant at the six and nine months horizon. At the same time for the long maturity horizons I find high adjusted R^2 values of 4 and 5 percent, respectively. The significance of the long maturity predictability indicates that the result is driven by interest rate predictability that is more relevant for long maturities than for short maturities. Similar to the Global and Non Precious selection I find significant constant risk premia at all maturities.

2.5 Macroeconomic Influences and Market Participants' Positions

The empirical results of the previous section imply that expected holding returns of commodity portfolios are time-varying. Having explored *how* expected returns vary over time, I will continue in this section to understand *why* they vary. In this section, I contribute to the understanding of

time-varying commodity risk premia by studying the relation of the AFD and expected holding returns to measures of economic activity.

2.5.1 Properties of the Average Futures Discount

The regressions in the previous chapter imply that the one step ahead forecast of excess holding portfolio returns is given by

$$\hat{\mathbb{E}}_t[rx_{t\to t+n}] = \psi_0 + \psi_{\mathbf{f}}(\overline{s}_t - \overline{f}_{t\to t+n}). \tag{2.25}$$

Therefore, variation in expected holding returns inherit the properties of the AFD. Finding that the AFD is related to systematic macroeconomic variables implies that risk premia in expected returns vary with the business cycles (see Bailey and Chan (1993)). In Table (2.10) I revisit the relation of common macroeconomic variables and the AFD on a portfolio level. I use four financial variables - the dividend yield (DivYield), the short rate (Short), the term spread (TSpread) and the default spread (DefSpread)- and two macroeconomic indicators - the three month US industrial production growth (IP Growth) and the consumer price index/inflation rate (Inflation) to assess the relation to the AFD. I also include the open interest growth that is related to business cycle conditions (see e.g. Hong and Yogo (2013)) as a measure for commodity specific variables. Table (2.10) Panel A reports the results from regressing the AFD of the Global and Precious commodities contemporaneously on the financial variables. The regression results show that the Global AFD is indeed statistically significant related to the DivYield and the DefSpread. The coefficient estimate for the DivYield is positive while it is negative for the DefSpread. Given that the DivYield is pro-cyclical and the DefSpread counter-cyclical, implies that the AFD varies pro-cyclically. For the Precious AFD I find the opposite signed coefficient estimates. The DivYield and the short rate are significantly negative related to the AFD suggesting counter-cyclical variation.

I continue the investigation in Table (2.10) Panel B, where I regress the AFD on the industrial production growth and the inflation rate for various lead and lag combinations. In particular, I run the regression

$$\overline{s}_t - \overline{f}_{t\to t+i+2} = \alpha + \beta Z_{t-i} + \epsilon_t \tag{2.26}$$

Table 2.10: **Average Futures Discount and Systematic Variables:** This table reports the results from regressing the AFD on various sets of systematic variables. Panel A reports the results form regressing the futures discount of Global and Precious portfolios on lagged common systematic and commodity specific variables. These variables include the dividend yield (DivYield), the short rate (Short), the term spread (TSpread), the default spread (DefSpread) and the open interest growth (OI Growth). Panel B reports the results from regressing the AFD on the industrial production growth (IPgrowth) and on changes of the consumer price index (CPI). Regression results are reported for a lead and a lagged AFD by 0, 1, 3, 6 and 9 months. Panel C reports the results from regressing stock and bond returns on the lagged AFD. The left panel reports the results for the Global portfolio and the right panel for the Precious portfolio. For all regressions I report the Newey-West t-statistics. The sample is from 1973 to 2012.

Panel A:							
	α	DivYield	Short	TSpread	DefSpread	OI Growth	R^2
Global	-0.00	0.01	-0.002	0.002	-0.030		0.20
	(-0.14)	(3.73)	(-1.74)	(0.87)	(-8.95)		
Global	-0.01					0.00	0.03
	(-3.55)					(1.70)	
Precious	0.01	-0.01	-0.01	0.00	-0.00		0.83
	(3.39)	(-5.88)	(-9.73)	(1.06)	(-0.56)		
Precious	-0.02					-0.00	0.00
	(-4.97)					(-0.70)	

Panel B:									
	Lead Futures Discout					Lag Futures Discout			
Global									
	9	6	3	1	0	-1	-3	-6	-9
IP Growth	0.62	0.62	0.56	0.40	0.38	0.35	0.23	-0.01	0.00
	(2.24)	(2.00)	(2.43)	(1.93)	(2.32)	(1.80)	(0.91)	(-0.03)	(0.01)
Inflation	-0.15	0.40	0.72	0.98	1.22	0.96	1.02	1.27	1.26
	(-0.17)	(0.48)	(0.97)	(1.72)	(2.56)	(1.71)	(1.46)	(1.45)	(1.36)
Precious									
IP Growth	-0.10	-0.00	-0.00	0.04	0.12	0.18	0.16	0.13	0.31
	(-0.42)	(-0.00)	(-0.00)	(0.31)	(0.85)	(1.03)	(0.96)	(0.66)	(1.54)
Inflation	-2.67	-2.31	-2.23	-2.22	-2.10	-2.13	-1.75	-1.78	-1.48
	(-4.95)	(-5.21)	(-5.50)	(-6.12)	(-6.25)	(-5.69)	(-5.13)	(-4.17)	(-3.91)

Panel C:							
	Global				Precious		
	α	Fut.Disc.	R^2		α	Fut.Disc.	R^2
Stocks	0.01	-0.01	-0.00	Stocks	0.01	-0.06	0.01
	(7.22)	(-0.27)	(-0.00)		(3.60)	(-1.54)	0.01
Bonds	0.621	-6.001	0.004	Bonds	0.60	-5.32	-0.00
	(6.570)	(-1.669)	(0.004)		(3.05)	(-0.65)	-0.00

where Z_{t-i} is either the industrial production growth or the inflation rate for i= -9,-6,-3,-1, 0, 1, 3, 6 and 9 months. The regression results confirm that the industrial production growth is contemporaneously positive and statistically significant related to variations in the Global AFD. Furthermore I find that industrial production growth predicts movements in the AFD for up to 9 months. However, the AFD does not predict industrial production growth for any horizon. These result confirm that the AFD is pro-cyclically related to macroeconomic cycles and the macroeconomic cycles are the drivers of these variations in the AFD. For the Precious AFD I do not find any significant relation. Next, I investigate the role of inflation rates on the AFD. I find that the inflation rate is contemporaneously significant related to the AFD and marginally significant for the lagged futures discount. This suggest that the futures discount has some power in predicting inflation rates, which is consistent with the findings of Gospodinov and Ng (2011). For the Precious AFD the results are striking. I find a highly significant and negative relation between Precious AFD and lagged or lead inflation rates at all horizons. These results imply that high inflation is related to low expected returns, while low inflation is related to high expected returns and is consistent with the view that gold is an inflation hedge.

If the AFD measures systematic risk premia variation, it should also have predictive power for stock and bond returns. Table (2.10) Panel C reports the result from regressing stock and bond returns on the AFD. I find that the Global AFD has marginal predictive power for bonds, while the predictability is insignificant for stocks. However, in both cases the negative coefficient estimates are consistent with counter cyclical variation of stock and bond risk premia, given that the AFD is pro cyclical.

2.5.2 Macroeconomic Factors, Average Futures Discount and Commodity Return Predictability

The predictive power of price based measures such as the AFD is important in order to understand if expected returns are time varying (see e.g. Cochrane (2011). In the previous section I showed that the AFD is indeed related to macroeconomic variables suggesting that expected commodity returns vary with macroeconomic cycles. However, recent literature in financial economics suggests that macroeconomic variables, such as the IP growth, contain information about risk premia that are not captured in price-based information (see Ludvigson and Ng (2009) and Daniel (2011)).

In this section I check if macroeconomic variables contain information about expected commodity holding and excess holding returns above and beyond what is captured by the AFD. I test this conjecture by focusing on the predictive power of industrial production growth, controlling for the AFD, for expected holding returns, as well as expected spot changes.

Using the same framework and notation as in the previous sections Table (2.11) presents the results from the regression,

$$\overline{rx}^j_{t\to t+n} = \psi^j_0 + \psi_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \psi_{\mathbf{IP}}\Delta \log IP_t + \eta_{t+n} \tag{2.27}$$

and

$$\overline{\Delta s}^j_{t\to t+n} = \zeta^j_0 + \zeta_{\mathbf{f}}(\overline{s}^j_t - \overline{f}^j_{t\to t+n}) + \psi_{\mathbf{IP}}\Delta \log IP_t + \epsilon_{t+n}. \tag{2.28}$$

Panel A reports the results for the regressions for (a) when including only the IP growth and (b) when adding the AFD as a control variable. For the individual regressions I find that all the slope coefficients $\psi_{\mathbf{IP}}$ of the Global portfolio and for all holding periods are positive and strongly statistically significant for holding and spot returns. When including the AFD in the regressions as a control variable, I find that the industrial production growth variable remains statistically significant, however the significance decreases slightly. The adjusted R^2 values for the holding returns on a 9 months horizon are around 5 percent for the IP growth and increase to almost 16 percent when including the AFD. Since I are controlling for the AFD the IPgrowth coefficient for the combined regression is the same for holding returns and spot changes, capturing the pure effect of risk premia on expected spot price changes, rather than the return stemming from the convenience yield or interest rates. The adjusted R^2 values for the spot changes are around 5 percent across maturities. The positive coefficients in the estimation confirms the pro-cyclical behaviour of the Global portfolio.

Next, I investigate the regression framework for the Precious portfolio. I find that all slope coefficients $\psi_{\mathbf{IP}}$ are negatively, however for 6 months and above, only marginally significantly related to expected holding returns. When including the AFD the results remain almost unchanged. Also I do not find a significant increase in the R^2 values. Note the negative sign of the coefficient estimates indicating counter-cyclical behaviour of the risk premia.

Table 2.11: **Predicting Commodity Portfolio Holding, Spot, Front Roll and Excess Holding Returns with Average Futures Discount and Macroeconomic Variables:** This table presents the results of forecasting regressions of commodity portfolios on the AFD and macroeconomic variables for horizons of n=2,3,6, and 9 months. I use the IP-Growth as a measure of macroeconomic activity. Reported are the estimated slope coefficient $\psi_{\mathbf{f}}$, $\zeta_{\mathbf{f}}$, $\theta_{\mathbf{f}}$ and $\vartheta_{\mathbf{f}}$ for the AFD as well as the slope coefficients $\psi_{\mathbf{IP}}$, $\zeta_{\mathbf{IP}}$, $\theta_{\mathbf{IP}}$ and $\vartheta_{\mathbf{IP}}$ for the macro variable and the adjusted R^2 values. The numbers in brackets denote the Newey and West (1987) standard errors with heteroscedasticity and autocorrelation consistent covariance estimators computed with the number of lags equal to the length of overlap plus one lag. All results are reported for the Global portfolio and the Precious portfolio from 1973 to 2012.

	Global						Precious Metals					
	Holding			Spot			Holding			Spot		
Horizon	$\overline{rx}_{t\to t+n}$			$\overline{\Delta s}_{t\to t+n}$			$\overline{rx}_{t\to t+n}$			$\overline{\Delta s}_{t\to t+n}$		
n	$\psi_{\mathbf{f}}$	ψ_{IP}	R^2	$\zeta_{\mathbf{f}}$	ψ_{IP}	R^2	$\psi_{\mathbf{f}}$	ψ_{IP}	R^2	$\zeta_{\mathbf{f}}$	ψ_{IP}	R^2
2		2.00	5.90		1.75	4.65		-0.62	-0.04		-0.67	-0.01
		(3.06)			(2.67)			(-0.66)			(-0.70)	
	0.65	1.85	9.00	-0.34	1.84	5.40	1.93	-0.71	2.20	0.93	-0.71	0.34
	(3.06)	(2.85)		(-1.61)	(2.82)		(1.82)	(-0.74)		(0.88)	(-0.74)	
6		3.96	6.12		3.33	4.96		-2.70	0.99		-2.78	1.13
		(2.67)			(2.29)			(-1.65)			(-1.71)	
	1.01	3.44	15.27	0.01	3.32	4.75	2.41	-2.90	7.27	1.41	-2.90	3.20
	(3.12)	(2.29)		(0.03)	(2.34)		(2.26)	(-1.75)		(1.32)	(-1.75)	
9		4.68	5.17		3.88	4.29		-2.92	0.72		-2.99	0.84
		(2.89)			(2.57)			(-1.75)			(-1.88)	
	1.07	3.85	15.65	0.05	3.84	4.10	2.59	-3.12	10.97	1.58	-3.11	4.79
	(2.75)	(2.47)		(0.14)	(2.65)		(2.40)	(-1.89)		(1.46)	(-1.88)	
	Global						Precious Metals					
	Front Roll			Excess Holding			Front Roll			Excess Holding		
Horizon	$\overline{frx}_{t\to t+n}$			$\overline{hrx}_{t\to t+n}$			$\overline{frx}_{t\to t+n}$			$\overline{hrx}_{t\to t+n}$		
n	$\theta_{\mathbf{f}}$	θ_{IP}	R^2	$\vartheta_{\mathbf{f}}$	ϑ_{IP}	R^2	$\theta_{\mathbf{f}}$	θ_{IP}	R^2	$\vartheta_{\mathbf{f}}$	ϑ_{IP}	R^2
2		2.02	5.91		-0.01	-0.19		-0.63	-0.04		0.01	-0.23
		(3.14)			(-0.40)			(-0.66)			(0.36)	
	0.60	1.88	8.52	0.06	-0.03	2.68	1.87	-0.71	1.94	0.07	0.01	0.45
	(2.87)	(2.95)		(2.63)	(-0.79)		(1.73)	(-0.74)		(1.42)	(0.28)	
6		3.97	6.04		-0.08	-0.13		-2.78	1.04		0.09	0.03
		(2.72)			(-0.64)			(-1.69)			(1.18)	
	0.91	3.50	13.34	0.05	-0.11	0.45	2.27	-2.98	6.39	0.15	0.08	4.69
	(3.02)	(2.39)		(0.87)	(-0.79)		(2.08)	(-1.79)		(2.25)	(1.18)	
9		4.62	4.94		-0.09	-0.18		-2.95	0.72		0.02	-0.24
		(2.76)			(-0.37)			(-1.80)			(0.15)	
	0.92	3.91	12.36	-0.00	-0.09	-0.40	2.45	-3.14	9.71	0.14	0.01	5.65
	(2.65)	(2.48)		(-0.00)	(-0.37)		(2.25)	(-1.94)		(2.19)	(0.05)	

Table (2.11) Panel B present the regression results for front roll and excess holding returns on AFD and IPgrowth. For the Global portfolio I find that all slope coefficients are positive for all maturities. For excess holding returns on the other hand I find negative slope coefficients, that are however not significant. These results confirm my earlier finding that most of the time series variation of portfolio holding returns is related to spot premia. For the Precious portfolio the findings are similar. For front rolls I find a marginally significant relation between IP growth and expected returns. For excess holding returns the macro variables does not contain any additional information.

In summary, I find that Global portfolios are pro-cyclically related to expected returns. Furthermore, macroeconomic variables contain information about the cyclical behaviour above and beyond the information contained in the AFD. The Precious portfolio displays counter-cyclical risk premia consistent with the literature in financial economics.

2.5.3 Trading Position, Average Futures Discount and Commodity Returns

A vast majority of the literature in commodity markets has focussed on price discovery in futures returns related to traders' position. The early contribution on "Normal Backwardation" by Keynes (1930) and Hicks (1939) argues that hedgers pay speculators a premium for offsetting their spot price risk in futures markets. This thought-process has been further refined by the "Hedging Pressure" literature of, e.g., De Roon et al. (2000) who show that the commodity risk premium is a combination of hedging positions and systematic risks. In contrast to the early literature that emphasizes the role of hedgers in commodity markets, more recently the high activity of speculators in commodity futures markets has caught the attention of researchers and policy makers alike. In particular, they have asked, whether the increase in speculative interest in commodities has exerted upward pressure on the level of both spot and futures prices. Several authors have found mixed evidence of the role of speculators in commodity futures markets, e.g., see Singleton (2013) and Irwin et al. (2009).

In this section, I contribute to this literature by examining the interaction of trading positions, macroeconomic variables and commodity portfolio returns. In contrast to many other asset classes, the US Commodity Futures Trading Committee (CFTC) requires all individuals and institutions above a certain size to report their positions to the CFTC each week. The CFTC reports aggregates

of these data, broken down into hedging (commercial) and speculative (non-commercial) position. I use these data to construct aggregate measures of relative hedging short positions and relative speculative long positions. Relative hedging short positions are constructed as the average of hedgers short positions minus hedgers long positions divided by all hedgers open interest across all commodities. Relative speculative long positions are constructed as the average of speculators long positions minus speculators short positions divided by all speculative positions.

Table (2.12) Panel A presents the results from regressing aggregate relative speculative long and hedging short positions on the AFD, commodity portfolio holding returns and the IPGrowth for the Global and Precious portfolio. The results are reported when all variables are included contemporaneously and for a predictive framework. The data shows that in the Global and Precious portfolio speculators and hedgers contemporaneously increase their long and short positions, respectively, according to changes in the futures discount and due to higher one period holding returns. While macroeconomic influences do not seem to play a role for the Global portfolio, the data show that speculators also increase long positions according to macroeconomic cycles. I find very similar results for the predictive model. Speculative and hedging positions increase even with a lag of 1 month. In summary, I find that hedgers and speculators seem to adjust their trading positions according to past and contemporaneous AFD and return movements. In particular, I find that speculators increase their long position seemingly knowing of higher expected returns, while hedger act as contrarian traders, consistent with The Theory of Normal Backwardation.

In order to better understand the causality of the relations between trading positions and signals, Table (2.12) Panel B presents the results from regressing the one period commodity holding returns on lagged and contemporaneous values of the relative short hedge position, relative long speculative position, AFD and the IPgrowth. I find that for the Global portfolios neither the short hedge nor long speculative positions are statistically significant related to to holding returns, while the AFD and the IPGrowth still predict holding returns, consistent with the previous sections. These results suggests that traders are not "causing" expected returns as suggested in the hedging pressure literature, but rather adjust their positions according to changes in the price-based signals. I find similar results for the Precious portfolios.

Table 2.12: **Commodity Trader Long and Short Positions and the Average Futures Discount:** This table reports the regression results from the interaction of trading activity, AFD and macroeconomic influences. Panel A reports the results from regression relative speculative long positions and relative hedge short positions on the AFD, the 1 month holding return, and the IPGrowth. Speculators and hedgers long and short positions are from the Commodity Futures Trading Commission (CFTC). Panel B reports the results from regressing the 1 month holding returns on the speculative or hedge positions, the average futures discount and the IPGrowth. All regressions are conduction contemporaneously and for lagged exogenous variables and for Global and Precious commodity portfolios. The numbers in brackets denote the Newey and West (1987) standard errors with heteroscedasticity and autocorrelation consistent covariance estimators. The sample period is from 1987 to 2012.

Panel A:								
	Contemporaneous				Predictive			
Global	Fut.Disc.	$\overline{rx}_{t-1\to t}$	IP_t	R^2	Fut.Disc.	$\overline{rx}_{t-2\to t-1}$	IP_{t-1}	R^2
Spec. Long	2.27	1.41	-1.07	0.17	1.88	1.30	0.53	0.13
	(4.12)	(4.34)	(-0.46)		(3.38)	(4.11)	(0.24)	
Hedge Short	1.15	0.44	0.31	0.21	1.08	0.38	0.72	0.18
	(6.35)	(4.39)	(0.40)		(5.78)	(3.90)	(1.01)	
Precious								
Spec. Long	10.62	1.19	4.44	0.33	11.55	0.67	3.63	0.24
	(8.95)	(6.98)	(2.04)		(9.17)	(4.06)	(1.89)	
Hedge Short	4.64	0.58	1.89	0.22	5.09	0.23	2.25	0.14
	(5.79)	(5.22)	(1.69)		(5.90)	(1.97)	(2.14)	
Panel B:								
	Contemporaneous				Predictive			
Global	Trade Pos.	Fut.Disc.	IP_t	R^2	Trade Pos.	Fut.Disc.	IP_{t-1}	R^2
$\overline{rx}$	0.00	0.36	1.31	0.06	0.00	0.35	2.10	0.10
	(0.43)	(2.62)	(1.99)		(0.15)	(2.60)	(2.23)	
$\overline{rx}$	-0.01	0.39	1.31	0.06	-0.02	0.38	2.12	0.10
	(-0.24)	(2.86)	(1.99)		(-0.66)	(2.89)	(2.24)	
Precious								
$\overline{rx}$	-0.03	2.24	0.82	0.06	-0.02	1.20	1.41	0.04
	(-2.11)	(4.70)	(0.83)		(-1.57)	(3.02)	(1.93)	
$\overline{rx}$	-0.07	2.28	0.85	0.07	-0.04	1.19	1.41	0.05
	(-2.54)	(4.78)	(0.86)		(-1.73)	(3.03)	(1.95)	

2.6 Conclusion

This chapter shows that expected returns of commodity futures portfolios are time-varying and therefore complements the recent cross sectional findings by Szymanowska et al. (2013). I find that most statistically significant variation is related to the risk in the underlying spot price, while risk in the term structure is statistically significant but constant for most maturities. The variation in the AFD, an ex-ante measure of expected returns, is pro cyclically related to macroeconomic variables. Furthermore, macroeconomic variables contain information about expected returns that are not captured by the AFD and therefore, predict spot price growth. The results survive a variety of robustness test.

The results have several important implications. First, I find that the characteristics of commodity futures portfolio returns can be embedded within the more recent literature on discount rate variation as surveyed by Cochrane (2011). In particular I show that a valuation ration, the AFD, predicts expected returns and not spot price growth similar to the findings in other asset classes. Second, from an investment perspective the results can be utilized to implement commodity portfolio investment strategies that take advantage of the predictive relations. In this chapter I scratched the surface of this task by implementing very simple trading strategies to exploit the predictability. Third, understanding expected returns of commodity futures is important for policy makers and commodity producers to make informed business and policy decisions. Fourth, I show that the variation of the risk premia is pro-cyclical, contradicting the results of bond and stock markets. This might be due to the fact that commodities are viewed as inflation hedges. Future endowers to explain this observation can have important implications for economic theory. Fourth, a large part of the variation of the AFD remains unexplained suggesting that more commodity specific sources of risk play an important role in explaining the fluctuations.

CHAPTER 3
PREDICTING COMMODITY RETURNS AND VOLATILITY: A BAYESIAN APPROACH

3.1 Introduction

The effects of changes in economic state variables for commodity returns and volatility have been studied extensively in international finance and within production and storage frameworks. Commodity return predictability from a financial and asset allocation perspective is of a more recent date and has received only little recognition in the academic literature (see Rouwenhorst (2013) for a review). This lack of attention is surprising given the dramatic increase in financial investors' interest in commodity index investments over the last 10 years[1]. In this chapter, I provide a comprehensive analysis of the relative statistical contribution of macroeconomic, financial, and commodity-specific variables in predicting commodity sector and characteristics-based portfolio returns and volatility in a data-rich bayesian model averaging framework. Second, I examine the out-of-sample economic value of the predictive relations within a conditional mean variance portfolio allocation.

Various authors have shown that commodity returns are predictable (e.g. Szymanowska et al. (2013) and Hong and Yogo (2013)). Economic theory is, however, ambiguous regarding the sources of commodity return variation. In particular, it remains difficult to determine whether commodities are integrated within traditional asset classes, influenced by commodity market-specific sources of risk, or a combination of both. From an equilibrium model perspective, systematic sources of risk, common across asset classes, should lead to variation in expected commodity returns with respect to economic state variables, e.g., Richard and Sundaresan (1981). A separate, yet equally important stream of literature argues that commodity market-specific sources of risk, such as the hedging demand and inventory decisions of commodity producers (De Roon et al. (2000), Gorton et al. (2007)), the limited risk bearing capacity of speculators (Acharya et al. (2013)), and the trading behaviour of commodity index investors (Hamilton and Wu (2013) and Basak and Pavlova (2013)) lead to predictable variation in expected returns and segment the return premia in commodity markets from other asset classes. The various theoretical explanations of commodity return

1. According to the CFTC (2008) staff report the total value of commodity index investments increased from 15 billion dollar to more than 200 billion by the end of 2008 (Rouwenhorst and Tang 2012).

variation, the multiplicity of resulting explanatory variables, and the changes in the market participants behaviour over time makes the empirical evidence of predictability difficult to interpret and introduces significant uncertainty regarding the correct model specifications.

I propose to formally assess the value a particular predictor adds relative to other predictors by taking the view of an economic agent forming beliefs about commodity activity given macroeconomic, financial, and commodity-specific information available at a given point in time using a Bayesian model averaging (BMA) framework (Fernandez et al. (2001), Wright (2008), and Binsbergen et al.(2012)). The economic agent forms beliefs about a set of candidate forecasting models and assigns posterior probabilities to various sets of competing predictive models. This framework explicitly incorporates model uncertainty and is therefore robust to model misspecifications at least within linear models. To understand predictability in a BMA framework, I focus my study on three key questions. First, are commodity portfolio returns predictable at a monthly frequency, and if so, what is the relative contribution of commodity market-specific, macroeconomic, and financial variables? Second, is commodity volatility predictable (relative to an autoregressive benchmark model) and do the same variables that predict returns also predict volatility? Third, do the return and volatility predictions add economic value to a real-time investor in a mean/variance asset allocation framework?

In answering these questions, I aim to shed light on the sources of price variation in commodity markets and the relative contribution of macroeconomic, financial and commodity-specific variables. For the choice of predictors, I do not propose any new set of variables but rely on instruments that have previously been shown to contain information about future commodity returns[2]. In addition, economic theory suggests that variables capturing time-varying risk premia in returns are also candidates for understanding predictability of volatility (Mele 2007). However, predictability of returns does not immediately imply volatility predictability. I therefore make volatility predictability a complimentary task in this study. Finally, statistical in-sample predictability does not imply

2. The variables I study are: aggregate basis, open interest growth (Hong and Yogo (2013)), hedging pressure (De Roon et al. (2000), Gorton and Rouwenhorst (2004)), as well as a measure for speculative capital flow (Acharya et al. (2013), Etula (2014)), dividend yield, default spread, term spread, short rate, TED spread (Bessembinder and Chan (1992), Bailey and Chan (1993)), global and US specific industrial production growth (Bakshi et al. (2011)), and inflation changes Gargano and Timmermann (2013). The first four instruments are local to commodity markets and are constructed for each commodity sector separatly, the next four are financial variables and the last three are macroeconomic variables.

out-of-sample predictability and economic value of the predictive relations, e.g., Goyal and Welch (2008) and Thornton and Valente (2012). I use economic criteria to quantify how much a risk-averse utility maximizing investor is willing to pay to switch from a dynamic portfolio strategy, based on a autoregressive model[3], to BMA combination forecast models that use commodity market-specific, macroeconomic, and financial variables (Fleming et al. (2001), Della Corte, Sarno, and Thornton (2008) and Della Corte, Sarno, and Tsiakas (2008)). I use mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility, which allows me to quantify how risk aversion affects the economic value of return and volatility predictability, building on empirical studies of volatility timing in stock returns by Fleming, Kirby, and Ostdiek (2001) and Marquering and Verbeek (2004).

In addressing these questions, I rank model specifications by calculating the posterior probabilities of each model and calculate weighted mean and t-statistics for each coefficient estimate. To preview my results I first find, that returns and volatility of a global commodity portfolio as well as the corresponding commodity sectors; agriculturals, energy, livestock, metals and precious metals, are highly predictable within a BMA combination forecasting framework. The in-sample results suggest that macroeconomic variables related to global and US industrial production growth have strong forecasting power for future portfolio returns. Furthermore, the futures portfolio basis adds value in terms of cumulative posterior inclusion probability suggesting that commodity-specific sources of risk influence expected return variation. On the other hand, I find that financial variables, in particular the dividend yield, default spread, and inflation growth are mostly responsible for the predictability of commodity portfolio volatility. These results are consistent across commodity sectors. For the characteristics based portfolios, I do not find significant in-sample predictability. Second, I assess the statistical and economic out-of-sample value of the BMA combination forecasts against the null of a pure autoregressive forecasting models with only lagged values of portfolio returns and volatility. I find evidence for both statistical and economic value of the combination forecasts for a real-time investor. These results are robust to reasonably high transaction costs. I, however, do not find any economic value or statistical significance for the out-of-sample volatility

3. In this article, I study commodity predictability relative to a model that includes lagged values of returns and volatility (autoregressive model). I make this choice to focus on information that is not already incorporated in commodity prices but adds additional statistical and economic value.

predictions.

The recent paper by Gargano and Timmermann (2013) comes closest to this chapter in that it analyses predictability of commodity returns and volatility using a large set of financial and macroeconomic variables. Here, I go beyond their work in several regards. First, I study a more diverse set of predictive variables by adding commodity market specific variables in addition to macro-finance predictors. Second, I study predictability of sector and characteristics based portfolios constructed from futures contracts. Gargano and Timmermann study predictability of commodity spot price indices only. Spot indices in commodity markets are, however, not traded liquidly which would make our economic assessment not feasible. Third, I study the economic value of mean and variance predictability in a comprehensive mean-variance portfolio allocation framework. Finally and most importantly, I consider the effect of model uncertainty using the BMA approach when studying and comparing the relative statistical and economic value added from a large set of variables. Avramov (2002), Wright (2008) and Cremers (2002) introduced BMA to study predictability in equity and exchange rate markets. In a related vain Daskalaki and Skiadopoulos (2011) show that including commodities in a portfolio that consists of traditional asset classes, without conditional variables, is not beneficial in an out-of-sample setting.

This chapter is more broadly but closely related to a large stream of literature that explores predictable time variation in expected returns across asset classes. Predictable time series variation of aggregated expected returns using price based valuation measures is by now understood to be a common feature across equity, bond, currency and housing markets (see Cochrane (2011) for an excellent review). Cooper and Priestley (2009), Bessembinder and Chan (1992), Ludvigson and Ng (2009) and Lustig et al. (2013) show that measures of macroeconomic activity have incremental power to predict expected returns in bond, equity and currency portfolios. On the other hand, Goyal and Welch (2008) and Thornton and Valente (2012) show that some of the in sample predictability does not extend to out-of-sample tests and does not generate economically significant value from an asset allocation perspective questioning the usefulness of the predictive variables. Recently, Christiansen at al (2013) explore the predictability of volatility across various asset classes.

The remainder of this chapter is structured as follows. In section 3.2 I describe the commodity data I use, the construction of the commodity portfolios and predictive variables. Section 3.3 lays out the Bayesian estimation procedure and combination forecasts based on the BMA framework.

Furthermore, I describe the framework to assess the economic value of commodity return and volatility predictability for a risk averse investor with a dynamic portfolio allocation strategy. Section 3.4 presents the main empirical results in a statistical in-sample and economic out-of-sample assessment. Section 3.5 concludes.

3.2 Data and Descriptive Statistics

3.2.1 Commodity Data

Commodity Futures Data The data sample consists of daily prices of commodity futures contracts that cover a cross section of 22 futures markets. These data are collected from various futures exchanges via Bloomberg over the sample period from January 1973 to September 2012. I collect price data of futures contracts across the term structure of up to 12 months. My total futures sample consists of *energy* contracts, including Crude Oil, Heating Oil, RBOB Gas, and Natural Gas, *agricultural* contracts, including Corn, Soybeanm Chicago Wheat, Soybean Oil, Kansas Wheat, Coffee, Cotton, Sugar, Cocoa, *livestock* contracts, including Feeder Cattle, Live Cattle, Lean Hogs, *industrial metal* contracts, including Copper, Zinc, Aluminum and Nickel, and *precious* metals, including Gold and Silver. These contracts and maturities were chosen because they are traded in sufficient volumes and open interest to implement dynamic trading strategies.

Commodity Returns and Volatility The main variables of interest in my analysis are returns and volatilities of commodity futures portfolios. I construct returns by assuming that an investor opens a futures position in the most liquid futures contract, which is typically the nearest or second closest to maturity contract, and is not within the delivery month. Furthermore, I assume that the futures investor is fully collateralised in that he always posts capital in a margin account that is as high as the futures price at initiation. Given these assumptions the excess futures return for contract j at time $\tau + 1$ during month t is defined as,

$$r_{j;t,\tau+1} = \frac{F^T_{j;t,\tau+1} - F^T_{j;t,\tau}}{F^T_{j;t,\tau}} \tag{3.1}$$

where $F^T_{j;t,\tau}$ is the price of the futures contract j at time τ, during month t, maturing at time T

$(t < T)$.

Having constructed daily return series of individual futures contracts, I construct equally weighted commodity sector portfolios, a global portfolio, as well as two dynamic long/short portfolios. To construct sector portfolios, I take the equally weighted average of individual futures returns within the commodity classes agriculturals, energy, livestock, industrial metals and precious metals. The global portfolio is an equally weighted portfolio of the sector portfolios. This approach of constructing commodity portfolios is similar to Hong and Yogo (2013) and Gorton et al. (2007). Furthermore, I construct basis and momentum sorted commodity portfolios as in Szymanowska et al. (2013). Each month, I sort the 22 commodities according to the strength of the basis and momentum signal in descending order[4]. The basis signal is the log difference in prices between front and third months to maturity futures contract, while the momentum signal is the cumulative past 12 months performance of each commodity contract. Having sorted the individual commodities, I take a long position in the top half of commodities and a short position in the bottom half. Then I construct the equally weighted average across all commodities. Finally, I compound the daily portfolio return series to monthly frequencies.

Following French et al. (1987) and Schwert (1989), I compute the realized variance of each commodity portfolio i in month t as the sum of squared daily returns: $\sum_{\tau=1}^{M_t} r_{i;t,\tau}^2$, where $r_{i;t,\tau}$ is the τth daily return during month t, as defined above for commodity portfolio i and M_t are the trading days during month t. I then follow Christiansen et al. (2012) and calculate the log of the square root of the realized variance as

$$RV_{i,t} = \ln \sqrt{\sum_{\tau=1}^{M_t} r_{i;t,\tau}^2} \tag{3.2}$$

for each month t. Andersen et al. (2003) show that the realized volatility is an accurate proxy for the true integrated volatility as the number of intra-period observations becomes large.

Panel A of Table 3.1 reports the summary statistics for the returns and (log) realised volatility of sector, global, and characteristics-based portfolios of commodities. Within the sector and global portfolios, all annualised mean returns are positive with a maximum of 3 percent for energy com-

4. The commodity with the highest momentum and basis signal is first followed by the commodity with the second strongest signal and so on.

modities. The annualised standard deviations of the returns of the sector portfolios are comparable to equity portfolios, ranging from 16 percent to 26 percent for the livestock and energy portfolio, respectively. The global portfolio has a standard deviation of 14 percent suggesting a diversification benefit when aggregating across sectors. I find annualised mean returns of 8 and 6 percent with standard deviations of only 17 and 14 percent for momentum and basis portfolios, respectively. The high returns suggest that commodity returns are indeed predictable by the momentum and basis characteristic similar to the findings by Szymanowska et al. (2013). All portfolios exhibit negative skewness, but the agriculturals, and a kurtosis that indicates non-normal tail distribution. These findings are consistent with Szymanowska et al. (2013). The autocorrelation coefficients of the return series are below 20 percent for all commodity portfolio suggesting low predictive power of past returns for future returns.

The means of the (log) commodity portfolio volatility display, trivially, the same characteristics as the volatility of the returns. The standard deviation of the (log) commodity portfolio volatility is relatively low and the kurtosis and skewness display close to normal characteristics. As is well-known, realized volatility is highly persistent and I find this behavior for all commodity portfolios under investigation as indicated by the high positive autocorrelation coefficients.

Panel B of Table 3.1 presents contemporaneous correlations among portfolio returns and volatilities, which are all positive. This outcome indicates commonalities across the returns and volatilities of different commodity sectors. The only portfolio that deviates from this observation are basis sorted portfolios that are close to uncorrelated to the other strategies.

Table 3.1: **Summary statistics for Commodity Portfolio Returns and Portfolio Volatility:** Panel A of this table presents the mean, standard deviation, skewness, kurtosis and autocorrelation coefficients for the returns and (log) realized volatility of sector, global, and characteristics-based commodity portfolios over the sample period from 1973 to 2012. The sector portfolios are constructed as equally weighted averages of individual commodity excess returns within the commodity sectors energy, agriculturals, livestocks, precious metals and metals. The global portfolio is an equally-weighted average of the sector portfolios. The characteristics based portfolios are long/short portfolios based on momentum and basis signals. At the end of each month, individual commodities are ranked in ascending order according to the strength of their basis (difference between log spot and futures price) and momentum (past 12 months cumulative returns) signal, respectively. A long position is assigned to the top half of the individual commodities, while the bottom half are assigned a short position. Then an equally-weighted average across long and short position is built. Panel B reports the correlations between the different return and volatility series. The total number of individual commodities is 22.

Panel A: Summary Statistics

	Portfolio Returns							**Portfolio (Log) Volatility**						
	Mean	Std.Dev	Skew.	Kurt.	AR(1)	AR(2)	AR(3)	Mean	Std.Dev	Skew.	Kurt.	AR(1)	AR(2)	AR(3)
Global														
Global	0.02	0.14	-0.06	6.02	0.09	0.09	0.03	-2.15	0.26	0.50	3.94	0.51	0.45	0.37
Sector Portfolios														
Agriculturals	0.01	0.17	0.18	4.90	0.04	0.11	0.01	-2.13	0.27	0.54	3.95	0.52	0.44	0.30
Energy	0.03	0.26	-0.22	5.03	0.19	-0.00	0.04	-2.15	0.53	-0.58	4.69	0.73	0.67	0.59
Livestock	0.03	0.16	-0.75	5.96	-0.03	-0.04	0.07	-2.51	0.40	0.32	3.56	0.27	0.38	0.18
Metals	0.02	0.22	-0.14	8.23	0.10	-0.02	0.11	-2.17	0.41	0.43	2.62	0.75	0.69	0.67
Prec.Metals	0.00	0.24	-0.23	8.16	0.05	-0.08	0.04	-2.47	0.46	0.61	4.00	0.66	0.58	0.57
Characteristics Portfolios														
Momentum	0.06	0.13	-0.20	3.94	0.01	-0.07	-0.10							
Basis	0.08	0.15	-0.11	4.80	-0.04	-0.12	-0.03							

Panel B: Correlation

	Portfolio Returns								**Portfolio (Log) Volatility**					
	Global	Agri.	Ener.	Live.	Metal	Prec.	Mom.	Basis	Global	Agri.	Ener.	Live.	Metal	Prec.
Global														
Global	1.00								1.00					
Sector Portfolios														
Agriculturals	0.91	1.00							0.85	1.00				
Energy	0.48	0.14	1.00						0.55	0.10	1.00			
Livestock	0.34	0.12	0.09	1.00					0.28	0.07	0.04	1.00		
Metals	0.52	0.41	0.38	0.11	1.00				0.47	0.34	0.30	-0.02	1.00	
Prec.Metals	0.56	0.30	0.24	0.07	0.41	1.00			0.50	0.11	0.13	0.01	0.62	1.00
Characteristics Portfolios														
Momentum	0.10	0.01	0.03	0.02	0.14	0.29	1.00							
Basis	-0.01	-0.05	0.03	0.08	-0.04	0.05	0.19	1.00						

3.2.2 Predictors Data

I collect data on predictive variables that have been previously shown to forecast commodity returns, albeit mostly in-sample. The set of predictive variables consists of commodity-market specific variables, macroeconomic variables, and financial variables. In the following, I provide a brief overview and some further details regarding the theoretical motivation behind some specific variables. In total, I consider a set of 29 state variables.

Commodity market specific variables The set of commodity specific variables contains a global as well as sector specific predictive variables for each predictor that I consider. I start by constructing the equally weighted average *basis* across commodity contracts and maturities. The basis for commodity i at time t is defined as

$$basis_{i,t,T_j} = log(F_{i,t}^{T_1}/F_{i,t}^{T_j}) \tag{3.3}$$

where $F_{i,t}^{T_1}$ is the futures price of the contract closest to maturity and $F_{i,t}^{T_j}$ is a contract with maturity in T_j periods where $T_1 < T_j$. In the aggregate version of this measure, I take the average across all commodity contracts i within a sector and available maturities T_j. The basis as a predictor for commodity returns is an ex-ante measure of risk premia and has been previously used by Fama and French (1987), Szymanowska et al. (2013) and in its aggregate version by Hong and Yogo (2013). Based on the Theory of Storage, the basis has also a direct relation to the volatility of futures prices. Routledge et al. (2002) show that low inventory levels are related to a low basis and high volatility. Whether this relation holds in a predictive way is not clear ex ante.

The next predictive variable I use is the growth rate in commodity market *open interest*. Hong and Yogo (2013) show theoretically and empirically that open interest contains information about future returns in the presence of downward sloping demand curves. To construct the variable, I compute, for each commodity, the logarithmic growth in dollar open interest by multiplying the spot price with the quantity of futures contracts outstanding and taking the first difference. The total futures contracts outstanding (open interest) stem from the Commitments of Traders Report published by the Commodity Futures Trading Commission (CFTC). Having constructed the individual growth in open interest series, I take the equally weighted average of the time series

Table 3.2: **Summary Statistics for Predictive Variables:** This table presents the summary statistics of the predictive variables that are available at a business cycle (monthly) frequency. The reported statistics include the variable name, its abbreviation, mean, standard deviation, skewness, kurtosis and autocorrelation coefficient for commodity market specific variables, macroeconomic variables, and financial variables. Data is available from 1973 to 2012

Panel A: Business Cycle (Full Sample)

No.	Variable	Abbrev.	Mean	Std.Dev.	Skewness	Kurtosis	AR(1)
Commodity Market Specific Variables							
1	Basis Agriculturals	BaAg	-0.01	0.03	0.76	4.00	0.83
2	Basis Energy	BaEn	0.01	0.05	1.28	8.05	0.67
3	Basis Livestock	BaLi	0.00	0.05	-0.05	3.50	0.63
4	Basis Metal	BaMe	-0.01	0.03	0.76	4.42	0.82
5	Basis Precious	BaPr	-0.02	0.02	-0.84	4.45	0.85
6	Hedging Pressure Agri.	HPAg	0.06	0.08	0.80	3.27	0.79
7	Hedging Pressure Ener.	HPEn	0.03	0.06	0.57	3.70	0.56
8	Hedging Pressure Live..	HPLi	0.05	0.14	1.01	4.89	0.76
9	Hedging Pressure Metal	HPMe	0.07	0.08	0.53	2.10	0.86
10	Hedging Pressure Prec.	HPPr	0.19	0.21	0.68	2.02	0.85
11	Open Interest Agri.	OIAg	0.00	0.03	0.06	4.94	0.32
12	Open Interest Ener.	OIEn	0.00	0.05	0.26	5.32	0.35
13	Open Interest Live.	OILi	0.00	0.04	-0.28	6.45	0.33
14	Open Interest Metal	OIMe	0.00	0.02	-0.10	5.17	0.33
15	Open Interest Prec.	OIPr	0.00	0.03	0.12	5.04	0.22
16	Speculators Capital Agri.	SCAg	0.01	0.10	-0.17	5.63	0.21
17	Speculators Capital Ener.	SCEn	0.02	0.23	0.53	6.88	0.12
18	Speculators Capital Live.	SCLi	0.01	0.12	-0.17	6.58	0.16
19	Speculators Capital Metal	SCMe	0.01	0.08	-0.23	5.33	0.23
20	Speculators Capital Prec.	SCPr	0.01	0.12	0.65	6.52	0.17
Macroeconomic Variables							
21	Kilian Global Ind. Prod.	KGIP	-0.00	18.28	0.79	6.14	0.93
22	Baltic Dry Index	BDI	-0.00	0.00	-4.47	56.48	0.40
23	Inflation Rate	CPI	0.00	0.00	1.13	5.79	0.80
24	US Industrial Production	USIP	0.00	0.02	-0.04	15.25	0.51
Financial Variables							
25	Dividend Price Ratio	DivY	1.18	1.68	1.15	3.04	0.99
26	Term Spread	TSpr	0.58	1.08	1.59	4.06	0.96
27	Default Spread	DefS	0.45	0.63	1.33	4.42	0.96
28	Short Rate	ShRt	2.10	3.34	1.59	5.00	0.98
29	TED Spred	TED	0.18	0.37	2.82	13.36	0.89

across commodities within each sector.

I use *hedging pressure* as a measure of imbalances between hedgers and speculators in commodity futures markets. I follow De Roon et al. (2000), Bessembinder (1992) and Gorton et al. (2007) and construct hedging pressure as the difference between the outstanding number of short positions by hedgers and the number of long positions divided by the total number of hedging positions. Long and short hedging positions for each commodity are identified from the Commitments of Traders Report. I construct an aggregate imbalance measure by taking the equally weighted average of individual hedging pressure variables. Keynes (1923) and Hicks (1939) argue that futures prices should embed a risk premium to entice speculators to offset the spot price risk that commodity producers (hedgers) are facing. This "Theory of Normal Backwardation" postulates that aggregate futures positions of hedgers are net short and the risk premium increases or decreases with the amount of net short positions.

The final commodity specific variable is a proxy for *speculative capital flow*. Similar to the hedging pressure variable, I take the difference between the outstanding number of long positions by speculators and the number of short positions divided by the total number of speculative positions. To aggregate this measure, I take the spot price weighted average of the individual speculative capital flows. The speculative capital flow measure has a correlation of more than 0.7 with other measures of speculative activity, such as hedge fund capital flows used by Etula (2014) and Acharya et al. (2013). In contrast to hedge fund capital measures, my measure is observable in real time.

Macroeconomic Variables The second set of predictive variables are macroeconomic state variables. While there are potentially many variables that are indicative of the state of the economy, I will focus on four main variables to omit the peculiar danger of data-mining. Furthermore, it is questionable if adding more macroeconomic variables does really span a larger information space to understand economic activity[5]. I use three measures of *industrial production growths*: two global and one US specific measure. For the first global version of industrial production, I follow Bakshi et al. (2011) and construct the logarithmic growth rate of the Baltic Dry Index (BDI). The BDI

5. Ludvigson and Ng (2009) show that using principal components of a set of 132 measures of economic activity has predictive power for bond risk premia. At the same time however, they also show that the first principal component has a correlation of more than 91% with a simple industrial production growth variable

measures the average shipping costs for raw materials that are largely variable due to changes in the worldwide demand for raw materials. This link to global demand suggests the use of the BDI as an indicator for global economic activity. Bakshi et al. (2011) show that the BDI growth rate has predictive power for stocks, bonds, and commodities, however only for spot indices. The second global variable is the global oil demand variable (KGIP) by Kilian (2009), which has been used to determine the role of demand and supply shocks on commodity price dynamics. The US specific measure of economic activity is the logarithmic growth rate of industrial production (USIP) as reported by the Federal Reserve Bank of St. Louis. Gargano and Timmermann (2013) show that this variable has predictive power for commodity spot indices. Furthermore, the US industrial production growth has been used by Cooper and Priestley (2009) to forecast stock and bond returns. Lastly, I use the inflation growth rate (CPI) that has been linked to futures discount movements by Gospodinov and Ng (2011).

Financial Variables Third, I discuss common systematic price-based measures of risk premia from bond and equity markets. These variables include the three months T-bill or *short rate* that is negatively related to future stock market returns and serves as a proxy for expectations of future economic activity. The *dividend yield* on the market, defined as the total dividend payments accruing to the CRSP value-weighted index over the previous 12 months divided by the current level of the index. The *term spread*, measured as the difference between the average yields of 10-year Treasury bonds and 3 months T-bills, based on the findings of Fama and French (1988), who provide evidence that this measure forecasts short-term business cycles. And the *default spread*, defined as the difference in yield between the highest and lowest rated investment grade bonds. Specifically, I calculate the default spread as the average yield of bonds rated BAA by Moodys minus bonds rated AAA by Moodys. Fama and French (1988) demonstrate that the spread tracks long term business cycles. Furthermore, I include the TED spread as a measure of aggregate market liquidity. Bessembinder and Chan (1992) show that these equity and bond return predictors do also have forecast power for commodity futures returns. Bailey and Chan (1993) show that these variables are correlated with movements in the commodity basis.

Because most macroeconomic and trading activity variables are noisy, I smooth them by taking the rolling 3-months average in time series. I use 3-months consistently across variables to avoid

data-mining concerns. Table 3.2 reports the summary statistics of the predictive variables. The means and standard deviation of the variables are comparable to previous studies. Importantly, I find that most predictive variables do not display high negative or positive correlations and are therefore good candidates for combination forecasts.

3.3 Bayesian and Economic Framework

3.3.1 Bayesian Setup

The main variables of interest are volatilities and returns of the commodity portfolios, as defined in the previous section, which serve as dependent variables in my predictive regressions. I formally assess the value a particular predictor may add relative to other predictors by using a Bayesian econometrics approach following the recent work by Fernandez et al. (2001), Wright (2008), and Binsbergen et al.(2012). In the financial economics literature the BMA approach was initially developed to explicitly incorporate model uncertainty, which is the uncertainty about the true set of predictive variables, and is therefore robust to model misspecification see e.g Avramov (2002) and Cremers (2002).

Predictive regression specification I follow Fernandez et al. (2001), Wright (2008), Avramov (2002) and Cremers (2002) and consider a set of k linear models $M_1, ...M_k$ with a total of K predictive variables. Each of these linear models obey the form

$$y_{t+1} = \alpha + X_{i,t}\beta_i + \epsilon_{t+1} \tag{3.4}$$

where X_i is a $T \times K_i$ matrix containing all regressors specified in model M_i. I specify $y_{t+1} \epsilon \{r_{t+1}, RV_{t+1}\}$ depending on whether I study return or volatility forecasts. The T-vector of errors, ϵ, is assumed to be normally distributed as $N(0, \sigma^2 I_T)$ where I_T is the identity matrix. The regression coefficients β_i and α_i are a $K_i \times 1$ vector and a constant coefficient, respectively.

Within the Bayesian econometrics framework, I am interested in learning about β_i and α_i given the data y_{t+1} and $X_{i,t}$ for model M_i. The logic of Bayesian econometrics suggests that I use Bayes'

rule to derive a probability statement about what I do not know (the correct specification of the model), conditional on what I know (the data). Formally, I derive the *posterior density* of model i, M_i, being the true model as

$$p(M_i|y) = \frac{p(y|M_i)p(M_i)}{p(y)} = \frac{p(y|M_i)p(M_i)}{\sum_{j=1}^{k} p(y|M_j)p(M_j)} \tag{3.5}$$

where y is the return or volatility data the econometrician observes. $p(M_i)$ is the *prior model density*, and, since it does not involve any observed data, the choice is subject to the econometricians beliefs of M_i being the correctly specified model. $p(y|M_i)$ is the *marginal likelihood* of M_i and can be calculated using some straightforward algebraic manipulations. Using Bayes' rule, I can write the *posterior density* of the model parameters α_i,β_i and σ_i as

$$p(\alpha_i, \beta_i, \sigma_i|y, M_i) = \frac{p(y|\alpha_i, \beta_i, \sigma_i, M_i)p(\alpha_i, \beta_i, \sigma_i|M_i)}{p(y|M_i)} \tag{3.6}$$

where $p(y|\alpha_i, \beta_i, \sigma_i, M_i)$ is the *likelihood function* and $p(\alpha_i, \beta_i, \sigma|M_i)$ the *prior density* of the parameters of model M_i. Integrating both sides of Equation (3.6) with respect to α_i,β_i and σ_i and using the fact that the posterior density integrates to one, I can rearrange the equation and obtain the following expression for the *marginal likelihood*,

$$p(y|M_i) = \int\int\int p(y|\alpha_i, \beta_i, \sigma_i, M_i)p(\alpha_i, \beta_i, \sigma_i|M_i)d\alpha_i d\beta_i d\sigma_i \tag{3.7}$$

which consists only upon the prior and the likelihood of the parameters. With particular choices of the priors the *marginal likelihood* can be calculated in closed form.

Prior selection In the context of model uncertainty, the choice of the prior distribution can have substantial impact on posterior model probabilities (see, e.g., George (1999) and Cremers (2002)). For my study, I choose priors that are easy to interpret and provide analytical solutions, which are denoted as *natural conjugate priors*. Furthermore, it is desirable for data information to be predominant over prior information as I do not have any intuition or theoretical justification on how to choose the parameters ex-ante. In choosing the prior distribution for α_i,β_i and σ_i, I follow Fernandez et al. (2001), who propose a benchmark prior distribution that has little influence on

the posterior distribution. In particular, they propose to use *improper noninformative priors* for the parameters that are common to all models, namely α_i and σ_i. This corresponds to choosing[6]

$$p(\alpha, \sigma) \propto \frac{1}{\sigma}. \tag{3.8}$$

For the β_i parameters, I choose the *g-prior* that was first introduced in Zellner (1986). In particular, I am setting

$$p(\beta_i|\alpha, \sigma, M_i) = f_N^{K_i}(\beta_i|0, \sigma^2(gX_i'X_i)^{-1}), \tag{3.9}$$

where $f_N^q(w|m, V)$ denotes the density function of a q-dimensional normal distribution with mean m and variance covariance matrix V. In summary, the g-prior says that the prior covariance of β is proportional to the comparable data-based quantity. For more details, I refer to Zellner (1986). The final specification to identify the prior distribution of the parameters is the choice of the g-parameter in equation (3.9). The value of $g = 0$ corresponds to a perfectly non-informative prior, while a value of $g = 1$ implies that prior and data information are weighted equally in the posterior covariance matrix. While both previous cases are extreme cases, a natural choice would be to choose a value that lies somewhere in between zero and one according to an information criterion. In this study, I follow Fernandez, Ley and Steel (2001) and choose $g = 1/\max\{n, k^2\}$, a value that has been proposed after extensive simulation experiments. In summary, as Fernandez, Ley and Steel (2001) show, the priors in equation (3.8) and (3.9) allow me to compute the marginal likelihoods analytically while leading to satisfactory results in simulation test from a posterior and predictive point of view.

So far I have focused on the specifications for the parameter priors given model M_i. Equation (3.5) shows that, besides the marginal likelihood, I need to specify the prior model distribution over the space of all models. I assume, as in Cremers (2002) that each variable is assigned an equal and independent prior probability of inclusion. This choice is reasonable because I do not have any a priori belief about which model is more likely to be the right one. More formally, if model M_i

6. This prior choice is called improper, because it is not, in fact, a valid density, in that it does integrate to one. This is why I use $\propto$ instead of equality. Nevertheless, it has many desirable properties and provides a close bridge between Bayesian and frequentist econometrics. I will come back to this point at a later stage.

consists of K_i explanatory variables out of K total variables, I assign the prior probability

$$p(M_i) = \rho^{K_i}(1-\rho)^{K-K_i} \tag{3.10}$$

where ρ is the probability of inclusion of a particular variable. For example, a choice of $\rho = 0.5$ assigns equal probability to each model. In the case of, e.g., $\rho = 0.25$ a model including less variables is a priori more likely to be the right model. Therefore, the higher the choice of ρ the more confident I am that a model with a lot of variables is the correct one. I will focus on $\rho = 0.5$ given that I do not want to impose to much structure on the model.

Marginal Likelihood Having specified the structure of the model and parameter priors, I can derive the marginal likelihood in equation (3.7) and consequently the posterior density in equation (3.5). While many Bayesian models require numerical computations to derive the values of the marginal likelihood in equation (3.7), my choice of priors allows for a closed form solution. The marginal likelihood in equation (3.7) can be written as

$$p(y|M_i) = \int\int\int p(y|\alpha, \beta_i, \sigma, M_i)p(\alpha, \sigma|M_i)p(\beta_i|\alpha, \sigma, M_i)d\alpha d\beta_i d\sigma \tag{3.11}$$

where $p(y|\alpha, \beta_i, \sigma, M_i)$ is the normally distributed density of the sampling model in equation (3.4), $p(\alpha, \sigma|M_i)$ is the prior density in equation (3.8), and $p(\beta_i|\alpha, \sigma, M_i)$ is the g-prior in equation (3.9). Fernandez et al. (2001) show that for these prior specifications I can write

$$p(y|M_i) \propto (\frac{g}{g+1})^{k_i/2}[\frac{1}{g+1}y'P_{X_i}y + \frac{g}{g+1}(y - \overline{y}\mathbf{1}_T)'(y - \overline{y}\mathbf{1}_T)]^{-\frac{N-1}{2}}, \tag{3.12}$$

where $P_{X_i} := I_N - X_i(X_i'X_i)^{-1}X_i'$ and $\overline{y} := \mathbf{1}_T y/n$. Since the marginal likelihoods can be computed analytically, the same holds for the posterior model probabilities given in equation (3.5).

Mean and Variance of Parameter Space and Forecasts The appeal of the Bayesian framework is that the posterior distribution of any quantity of interest, denoted as ϕ[7], is an average of the posterior distributions of that quantity under each of the models M_i with weights given by the

7. ϕ could be for example all the parameters of the model in equation (3.4), or just one of the parameters.

posterior model probabilities. In particular, I can write

$$p(\phi|y) = \sum_{i=1}^{K} p(\phi|y, M_i)p(M_i|y) \tag{3.13}$$

where $p(\phi|y, M_i)$ can be typically calculated in standard form and $p(M_i|y)$ is the posterior density of model i. Following from the rules of probability, I can use the result to calculate the mean and variance of the quantity of interest as

$$\mathbb{E}(\phi|y) = \sum_{i=1}^{K} \mathbb{E}(\phi|y, M_i)p(M_i|y) \tag{3.14}$$

and

$$Var(\phi|y) = \sum_{i=1}^{K} Var(\phi|y, M_i)p(M_i|y), \tag{3.15}$$

respectively. An important example for my study is the choice of $\phi = \beta_i$. With this choice, I can show that the posterior distribution of $p(\beta_i|y, M_i)$ follows a multivariate t distribution with mean $\mathbb{E}(\beta_i|y, M_i) = \overline{V}_i X_i' y$, and variance $Var(\phi|y, M_i) = \frac{\overline{\nu}\overline{s}_i^2}{\overline{\nu}-2}\overline{V}_i$ where

$$\overline{V}_i = [(1+g)X_i'X_i]^{-1} \tag{3.16}$$

and

$$\overline{s}_i^2 = \frac{\frac{1}{g+1}y'P_{X_i}y + \frac{g}{g+1}(y - \overline{y}\mathbf{1}_T)'(y - \overline{y}\mathbf{1}_T)}{\overline{\nu}} \tag{3.17}$$

where P_{X_i} is defined as above. Another example are mean and variance forecasting exercises. Let $\phi = r_{t+1}$ be a $K \times 1$ vector of commodity risky commodity returns and choose $y_t = r_t$. I can characterise the predictive distribution of r_{t+1} as

$$\begin{aligned} p(r_{t+1}|r_t) = \sum_{i=1}^{2^k} & f_S(r_{t+1}|n-1, \overline{r} + \frac{1}{g+1}x_{t,i}'\beta_i^*, \\ & \frac{n-1}{d_j^*}\{1 + \frac{1}{n} + \frac{1}{g+1}x_{t,i}'(X_i'X_i)^{-1}x_{t,i}\}^{-1})p(M_i|r), \end{aligned} \tag{3.18}$$

where $f_S(x|\nu, b, a)$ denotes the probability density of a univariate Student-t distribution with ν degrees of freedom, location b and precision a evaluated at x. $x_{t,i}$ is the observation used to predict

r_{t+1} for model i, $\beta_i^* = (X_i'X_i)^{-1}X_i'r$ and

$$d_i^* = \frac{1}{g+1}y'M_{Z_i}y + \frac{g}{g+1}(r - \overline{r}\mathbf{1}_T)'(r - \overline{r}\mathbf{1}_T). \tag{3.19}$$

Given that the posterior of the predictive model has a univariate Student-t distribution, I can immediately write the conditional mean as

$$\mu_{j,t+1|t} := E_t[r_{t+1}] = \sum_{i=1}^{2^k}(\overline{r} + \frac{1}{g+1}x_{t,i}'\beta_i^*)p(M_i|r) \tag{3.20}$$

for portfolio j. This equation shows that the mean is a weighted average and scaled version of the ordinary least square regression framework. This example can be trivially extended to the volatility forecast by choosing $\phi = RV_{t+1}$ and $y_t = RV_t$, which results in the equation

$$\ln \sigma_{j,t+1|t} := E_t[RV_{t+1}] = \sum_{i=1}^{2^k}(\overline{RV} + \frac{1}{g+1}x_{t,i}'\beta_i^*)p(M_i|RV). \tag{3.21}$$

Markov Chain Monte Carlo Model Composition (MC^3) Given the specific setup of my model (linear normal and g-priors), I can calculate all quantities of interest in closed form. Unfortunately, this procedure is computationally demanding in rolling or expanding regressions (used for out-of-sample forecasts) that re-estimate the model at each time step. To handle the large number of models, I follow Fernandez et al. (2001) and Cremers (2002) and use the Markov Chain Monte Carlo Model Composition (MC^3) sampling approach, which has been shown to be well suited for the problem at hand. In short, for each run the algorithm draws a candidate model M^* from the model space $M_1, ...M_k$ where $k = 2^{39}$, accepts it if the model improves the marginal likelihood from the previous model selection and rejects it otherwise. If the model is rejected it remains at its previous model draw chain. If the number of Monte Carlo draws is sufficiently large the fraction of draws for the different models converges to the posterior model probability. To ensure a close approximation, I sample 500000 models and use the first 50000 draws as a burn-in period to avoid sensibility of the results with respect to the starting model. More information on the algorithm

can be found in Fernandez et al. (2001) and references therein.

Statistical Inference and Model Comparison Having derived all the posterior probabilities, the weighted mean and variance equations, and the mean and variance forecasting equations, I proposes various statistics to investigate the robustness of explanatory variables in predictive regressions and their relative statistical importance in- and out-of-sample. The first statistic is the *posterior t-ratio*, which is defined as the posterior mean of each of the slope coefficients in the weighted model given by equation (3.14) divided by its corresponding standard error obtained from equation (3.15). The interpretation of the posterior t-ratio is analog to the classical frequentist t-ratio but explicitly incorporates model uncertainty. Second, I propose the *posterior odd ratio* that is a key tool for model comparison in Bayesian econometrics. The posterior odd ratio is defined as $\frac{p(M_i|r)}{p(M_i|r)}$ that compares the posterior density of model i to the posterior density of model j. I will use various specifications of this measure. The third statistic is the cumulative posterior probability of the predictive variables. I add the model posterior probabilities of each variable across models in order to measure the probabilities that each of the predictive variables appears in the weighted forecasting model. For more details on the statistics, I refer to Avramov (2002). The next statistics are the in- and out-of-sample Bayesian R^2, the adjusted mean-squared prediction model (MSPE), and the root mean square error (RMSE). The definition of the statistics is analogue to the regular in- and out-of-sample statistics of Campbell and Thompson (2008), Clark and West (2007) and Goyal and Welch (2008), respectively, where the conditional mean forecasts and consequently the mean square errors are calculated with the Bayesian forecasting methodology.

3.3.2 Portfolio Choice and Economic Value

In the previous section, I developed the Bayesian estimation framework for an investor who is uncertain regarding the right choice of the predictive model. In this section, I will apply the Bayesian forecasting within a dynamic asset allocation framework. Specifically, I consider an investor with a one-month horizon, who wants to maximize conditional expected returns of a commodity portfolio by dynamically rebalanceing the portfolio weights to achieve a particular conditional volatility. I therefore use a classic dynamic mean-variance portfolio choice problem similar to Della Corte et al. (2008), Thornton and Valente (2012) and Fleming et al. (2001).

Portfolio Choice Problem The classical mean variance portfolio choice problem states that for each month t an investor solves the problem

$$\begin{aligned} \max_{w_t} & [w_t' \mu_{t+1|t} + (1 - w_t' \mathbf{1}) r_f] \\ s.t. & (\sigma_p^*)^2 = w_t' \Sigma_{t+1|t} w_t \end{aligned} \tag{3.22}$$

where w_t is an $K \times 1$ vector of portfolio weights on the risky assets, r_f is the return on a risk free asset, $\mathbf{1}$ is a $K \times 1$ vector of ones. σ_p^* is the target volatility of the portfolio that will be specified ex-ante. It is well known that the optimisation problem has the solution for the weights, w_t, given by

$$w_t = \frac{\sigma_p^* \Sigma_{t+1|t}^{-1} (\mu_{t+1|t} - \mathbf{1} r_f)}{(\mu_{t+1|t} - \mathbf{1} r_f)' \Sigma_{t+1|t}^{-1} (\mu_{t+1|t} - \mathbf{1} r_f)}. \tag{3.23}$$

In this article, I am working with futures contracts only, which are zero cost securities. As a consequence, the return on a futures position is the excess return, which simplifies the formula for the weights further. I can rewrite equation (3.23) as

$$w_t = \frac{\sigma_p^* \Sigma_{t+1|t}^{-1} \mu_{t+1|t}}{\mu_{t+1|t}' \Sigma_{t+1|t}^{-1} \mu_{t+1|t}}. \tag{3.24}$$

It should be noted that I do not constrain the weights in any way and therefore allow for unlimited short and long positions. Given that I trade in futures markets only, which do not present any natural short selling constraints, the assumption is not restrictive.

The task at hand is to determine the optimal weights of the portfolios, given condition mean return and volatility predictions. I consider the value of mean and volatility predictions *separately* to understand which predictive model adds economic value. Again, this is not restrictive, because using both predictive models at the same time will only increase the value of the portfolio allocation. When I consider the value of conditional mean predictions, I use equation (3.20) to determine $\mu_{t+1|t}$ and estimate $\Sigma_{t+1|t}$ using a simple rolling scheme, where at each time t, I estimate the rolling variance/covariance matrix using only returns up to time t. When I consider the economic value of conditional variance predictions I set the conditional mean equal to

the sample mean $\mu_{t+1|t} = \frac{1}{t}\sum_{i=1}^{t} r_{t+1}$. For the estimation of $\Sigma_{t+1|t}$, I follow Bollerslev (1990) and specify the conditional variance/covariance matrix for each element, at location i and j, as $\Sigma_{i,j,t+1|t} = \sigma_{j,t+1|t}\sigma_{i,t+1|t}\rho_t$, where $\sigma_{j,t+1|t}$ is estimated as specified in equation (3.21). For ρ_t I use an expanding window of sample correlations of portfolio i with portfolio j. A similar approach has been used by Della Corte et al. (2008) and Fleming et al. (2001).

Quadratic Utility and Economic Value I use mean-variance analysis with quadratic utility to rank the performance of competing commodity models, as in West et al. (1993), Fleming et al. (2001), Della Corte et al. (2008) and Thornton and Valents (2012). The resulting realised utility of the investor in period $t+1$ can be written as

$$U(W_{t+1}) = W_t R_{p,t+1} - \frac{\lambda W_t^2}{2} R_{p,t+1}^2 \tag{3.25}$$

where W_{t+1} is the investor's wealth at $t+1$, λ is his risk aversion, and

$$R_{p,t+1} = (1 - w_t'\mathbf{1})r_f + w_t' r_{t+1} \tag{3.26}$$

is the period $t+1$ return on this portfolio. Under some regularity assumptions the expected utility can be shown to be linear homogenous in wealth. When I set the initial level of wealth $W_0 = 1$, I can write the average realised utility as

$$\overline{U}(.) = \sum_{t=0}^{T-1} \{R_{p,t+1} - \frac{\lambda}{2(1+\lambda)} R_{p,t+1}^2\}. \tag{3.27}$$

The appeal of using average realised utility is that I can equate two competing commodity trading strategies to estimate the economic value as in Fleming et al. (2001). In particular, I assess the value added from conditional mean prediction model and the conditional volatility model against weights that are based on the benchmark model. Assume, for example, that a benchmark model yields the same average utility as holding a conditional model that is subject to additional monthly expenses of ϕ, expressed as a fraction of wealth invested. Since the investor would be indifferent between these two strategies, I interpret ϕ as the performance fee the investor will pay to switch from the benchmark model to the conditional model. I estimate this fee by finding the value ϕ that

satisfies the equation

$$\sum_{t=0}^{T-1}\{(R_{p,t+1}^{C}-\phi)-\frac{\lambda}{2(1+\lambda)}(R_{p,t+1}^{C}-\phi)^2\}=\sum_{t=0}^{T-1}\{R_{p,t+1}^{RW}-\frac{\lambda}{2(1+\lambda)}(R_{p,t+1}^{RW})^2\} \tag{3.28}$$

where R_p^C and R_p^B denote the portfolio returns using a dynamic strategy based on conditional means or volatility and the benchmark, respectively. If $\phi > 0$ then the conditional mean or volatility prediction add economic value to the model, while $\phi \leq 0$ is an indication of no predictability. Another commonly used measure of economic performance in the context of mean-variance analysis, is the Sharpe ratio. The Sharpe ratio is calculated as the average portfolio excess return to its standard deviation. Given that the Sharpe ratio can underestimate the performance of a dynamic strategy (Marquering and Verbeek (2004) and Han (2006)), I will only report the Sharpe ratios for comparison and not as a main decision tool.

Transaction Costs Given the nature of conditional trading strategies, I have to control for the possible effect of transaction costs. Since, static strategies typically generate a lower turnover than dynamic strategies, the gain from the dynamic strategy could be partly offset by transaction costs. It should be noted, however, that I implement the trading strategies in futures markets that require a rollover of a futures position even for the static portfolio. Nevertheless, it is instructive to consider transaction costs. I follow Han (2006) and calculate break even transaction costs. Using break even transaction costs avoids the difficult inference of real transaction costs that are usually dependent on the type of investor, the value of the transaction, and the nature of the broker. To accomplish this, I assume that transaction costs equal a fixed proportion, τ, of the value traded in different commodity futures. The average transaction cost is calculated as $\tau \times V$ where

$$V=\frac{1}{T}\sum_{t=0}^{T}\sum_{i=1}^{k}|w_{i,t}-w_{i,t-1}\frac{(1+r_t)}{1+r_{p,t}}|. \tag{3.29}$$

The break even transaction cost can then be calculated as (e.g., Jondeau and Rockinger (2008))

$$\tau^{BE}=\frac{(\frac{1}{T}\sum_{t=0}^{T}R_{p,t+1}^{CMP})-(\frac{1}{T}\sum_{t=0}^{T}R_{p,t+1}^{RW})}{V^{CMP}-V^{RW}}. \tag{3.30}$$

I will report the transaction costs if, and only if, the performance fees ϕ are positive. If the performance fees are smaller than zero then the dynamic strategy would not be optimal, independent of the effect of transaction costs.

3.4 Empirical Results

In this section, I present the main empirical results of the study. All results are based on the Bayesian Model Averaging (BMA) approach, which selects the best predictive variables in terms of posterior model probabilities. I study the relative importance of predictive variables for return and volatility commodity portfolios separately. Furthermore, I distinguish between results for the global portfolio, sector portfolios and characteristics based portfolios.

3.4.1 Bayesian Model Averaging: In-Sample Analysis

Predicting the Global Portfolio In this section, I present the in-sample baseline results, which entail the predictability results for global commodity returns and volatilities using the BMA methodology. For the Bayesian estimation of global portfolio return and volatility predictions, I restrict the analysis to global commodity-specific, macroeconomic and financial variables as described in the data section. I refrain from including sector-specific variables. Furthermore, it is well known that volatility is fairly persistent. Therefore, in addition to the predictive variables described above, I include autoregressive terms of volatility (of order one) and for completeness also lagged return values (of order one) in the model selection setup to focus on information other than that already incorporated in prices. For the global model setup, I use in total $K = 15$ variables that result in $k = 2^{15}$ potential predictive models. This number of models is computationally tractable. I therefore calculate the posterior probabilities for *each* model in closed form.

The main results for predicting global commodity portfolio returns and volatility are reported in Table 3.3, which presents the cumulative posterior probability of each variable being included in the model. From an economic agent point of view, the cumulative posterior probabilities of inclusion reflect the belief of how likely it is he should include a variable in the model after seeing the data. The agent specifies that each model is equally likely to be the right one before seeing

the data, which implies an ex-ante inclusion probability of each predictive variable of 0.5. If the cumulative posterior probability for a particular variable is higher than 0.5 the prior belief of the economic agent is revised upwards and he should include or assign this variable a higher weight in his model setup. I mark these instances of probabilities that are higher than 0.5 in bold. To get a better overview of the results, I further sort the variables by their posterior inclusion probability in descending order. I also display the posterior means and posterior t-ratios of each variable to make a comparison to classical frequentist econometric frameworks easier. It should be noted, however, that the posterior t-ratio includes an adjustment for model uncertainty. The frequentist t-ratio and the Bayesian t-ratio will consequently differ in value.

I find that 4 out of 15 predictive variables have cumulative posterior probabilities that are higher than 0.5 and therefore add predictive value to the model, while the other variables are discarded as worthless. The low number of significant predictors is surprising given that all variables have been previously shown to be significant predictors for commodity portfolio returns by other authors. The finding is, however, consistent with the view that the BMA methodology includes uncertainty in the model setup (Avramov (2002)) and therefore, can reduce the statistical value a particular predictor would have in a classical regression framework. I find that the BMA methodology emphasises the importance of macroeconomic variables in predicting commodity returns. US industrial production growth (USIP) and the growth in Baltic Dry Index (BDI) are included in almost all models, having cumulative posterior inclusion probabilities of 1 and 0.9, respectively. The positive Bayesian coefficient estimates imply that growth in USIP and BDI is positively related to future commodity returns, which is in line with the findings of Bakshi et al. (2011). Also the posterior t-ratios are significant with values of 3.1 and 2.2 for USIP and BDI, respectively. The only financial ratio that adds predictive value to the model according to BMA is the short rate (ShRt), with a cumulative posterior probability of 0.97. The short rate is negatively related to global commodity returns with a posterior t-ratio of -2.55 similar to the findings of Hong and Yogo (2013). The positive relation between variables of macroeconomic activity and the negative relation to the short rate implies that global commodity portfolio returns are pro-cyclical. The pro-cyclical relation is in sharp contrast to the findings in stock and bond markets that find a counter-cyclical relation (Cooper and Priestley (2009) and Ludvigson and Ng (2009)). Since expected commodity returns are high when macroeconomic activity is high, commodities can be viewed as a good hedge

Table 3.3: **Predictive Regressions for Global Commodity Portfolio Returns and Portfolio Volatility using BMA estimation:** This table reports in-sample predictability results for the global commodity portfolio returns and portfolio volatility, estimated using the Bayesian Model Averaging (BMA) approach in closed form. The global commodity portfolio is an equally weighted average across the commodity sectors agricultural, energy, livestock, metals, and precious metals. The set of predictive variables contains commodity specific variables (global variables only), macroeconomic variables and financial variables as well as lagged dependent returns and volatility (RV_{t-1}). In total the number of predictors is 15. The table presents the results separately for portfolio returns on the 5 left hand columns and for the portfolio volatility on the 5 right hand columns. The table displays the results for the 15 predictors ranked according to the cumulative posterior probability of inclusion. From left to right the table reports rank, variable name, cumulative posterior probability of inclusion (Prob), posterior means and posterior t-ratios for each variable. The posterior means of slope coefficients are obtained by averaging slope estimates by the posterior model probabilities across all possible model combinations (BMA approach). The posterior t-ratios are adjusted to account for model uncertainty. In particular, the t-ratios are obtained by dividing the posterior model averaged means across models by its posterior standard error averaged over all models, including model uncertainty that summarizes the dispersion in slopes across models. Variables with cumulative posterior probabilities of more than 50 percent are marked bold. The sample period spans from 1973 to 2012.

Portfolio Returns					**Portfolio (Log) Volatility**				
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	USIP	**1.00**	1.43	3.11	1	Vol	**1.00**	0.29	4.27
2	ShRt	**0.97**	-0.00	-2.55	2	CPI	**0.99**	21.32	3.31
3	BDI	**0.90**	0.82	2.18	3	DefS	**0.99**	0.14	3.21
4	BaGl	**0.62**	0.40	1.97	4	DivY	**0.98**	-0.08	-3.56
5	KGIP	0.21	0.03	0.30	5	BaGl	**0.93**	1.56	2.23
6	TSpr	0.14	-0.00	-0.16	6	HPGl	**0.53**	-0.62	-1.59
7	DivY	0.09	0.00	0.07	7	USIP	0.28	-1.26	-0.43
8	TED	0.09	-0.00	-0.08	8	SCGl	0.18	-0.05	-0.23
9	CPI	0.09	0.06	0.07	9	TSpr	0.14	-0.00	-0.16
10	HPGl	0.08	-0.01	-0.07	10	ShRt	0.12	-0.00	-0.15
11	OIGl	0.07	-0.00	-0.05	11	BDI	0.08	-0.17	-0.07
12	Ret	0.06	-0.00	-0.03	12	OIGl	0.08	-0.02	-0.07
13	DefS	0.06	-0.00	-0.02	13	Ret	0.06	0.01	0.04
14	Vol	0.05	-0.00	-0.02	14	TED	0.06	0.00	0.03
15	SCGl	0.05	-0.00	-0.01	15	KGIP	0.05	-0.00	-0.02

for time-varying investment opportunities in stock and bond markets. An alternative explanation for the positive relation is that macroeconomic information are not immediately incorporated into commodity prices due to informational frictions and underreacting agents lead to delayed price reactions. This explanation corroborates with behavioural theories of Hong and Stein (1999) and Hong et al. (2000). The final variable of use for the predictive regression is the commodity-specific futures discount (BaGl). The BaGl enters the regression with a cumulative posterior probability of 0.62 and has a positive posterior mean coefficient and significant posterior t-ratio of 1.97. The positive relation is consistent with the view of the futures discount or basis being an ex-ante measure of expected returns (Szymanowska et al. (2013) and Fama and French (1987)) and contains information about commodity risk that is not contained in macroeconomic variables. Surprisingly, most commodity market specific variables such as open interest growth (OIGl) (Hong and Yogo (2013)), hedging pressure (De Roon et al. (2000) and speculative capital (Acharya et al. (2013) do not add any value to the predictive regressions. Furthermore, most financial variable do not enter the predictive regressions suggesting that commodity markets are partially segmented from other asset classes contrasting the findings of Bailey and Chan (1993) when including model uncertainty.

Next, I study the predictive BMA results using the global (log) commodity volatility as the endogenous variable presented in the right hand columns of Table 3.3. I find that several variables add significant value to the predictive model. First, the lagged value of commodity volatility (vol) enters all predictive regressions consistent with earlier findings that volatility is highly persistent. The posterior mean is positive and the t-ratio highly significant with a value of 4.27. More interesting, however, is the finding that variables other than pure price-based measure enter the predictive model with high posterior inclusion probabilities. Growth in inflation (CPI) has a posterior inclusion probability of 0.99, a positive posterior mean coefficient and a highly significant posterior t-ratio of 3.31, suggesting a very strong positive relation between inflation and volatility of commodity returns at short horizons. The positive effect can be related to a higher trading activity during times of inflation since commodities are commonly known to be good inflation hedge. Mele (2007), on the other hand, argues that inflation has an effect on volatility due to its relation to risk premia. The latter explanations finds further support in the finding that the financial variables default spread (DefS) and the dividend yield (DivY) enter the predictive regressions with posterior

inclusion probabilities of 0.99 and 0.98, respectively. These two variables have been shown to be ex-ante measures of stock and bond risk premia. The positive (negative) posterior coefficient of the default spread (dividend yield) suggests that commodity volatility is counter-cyclical; a finding similar to observations in stock and bond markets (Mele (2007)). The counter-cyclical behaviour is further confirmed by negative, although insignificant, relations between USIP and BDI with commodity volatility. The final two variables entering the predictive regressions are the commodity market-specific variables global futures discount and global hedging pressure. The basis enters the regression with an inclusion probability of 0.93, positive posterior mean coefficient and a significant posterior t-ratio of 2.23. The positive relation is consistent with the predictions of the Theory of Storage (Ng and Pirrong (1994) and Routledge et al. (2002)). On the other hand, this relation also corroborates with the view that the futures discount is an ex-ante measure of risk premia. Hedging pressure is only weakly related to future volatility with inclusion probability of only 0.53. Furthermore, the negative sign of the posterior coefficient seems troubling as higher hedging pressure is associated with lower inventory levels and higher volatility (Gorton et al. (2007) and De Roon et al. (2000)). All other variables are discarded as worthless from a posterior perspective.

In Table 3.4, I extend the results from the previous section and report the variables that are included within the top 5 models according to the posterior model probabilities. Inclusion of a variable in the return and volatility predicting model is indicated with a bold one and zero otherwise. Consistent with the previous results, I find in Panel A that the macroeconomic variables, BDI and USIP, and the financial variable; short rate, are included in all top 5 models. The futures discount is selected in only 2 out of the top 5 models. I also report the frequentist R^2 and adjusted R^2 values. All top 5 models have R^2 values that are higher than 8 percent and adjusted R^2 values that depart at most 1 percent from their unadjusted counterpart. Eight percent R^2 values are high on a monthly frequency and indeed suggest significant predictive power of the selected models. Finally, I report the posterior odds ratios. The posterior odd ratios compare posterior model probabilities of models including only lagged volatility and returns (price based models) with the posterior model ratios of the top 5 models. All posterior odd ratios indicate significant predictive ability of the selected models in contrast to the pure price based models. Table 3.4 panel B shows that 5 variables are included in all top 5 models for the volatility predictability. These variables are, consistent with the previous results, lagged volatility, global futures discount, inflation growth,

dividend yield, and default spread. The other significant variables enter in at most 2 out of the 5 top models. Each top 5 model explains more that 25 percent of the variation in volatility and display significant posterior odd ratios. The R^2 values are lower than results from stock and bond markets (e.g. Christiansen et al (2012)), however, are still highly significant.

By comparing the results from the return and the volatility predictions several findings are worth mentioning. First, both global portfolio returns and portfolio volatility are predictable. The results of predictability are even stronger for the volatility predictions than for the portfolio returns, which is in part attributable to the higher idiosyncratic volatility of returns and the persistence of volatility. Second, a relatively large portion of the total monthly variation of returns and volatility can be explained by the top predictive variables in terms of posterior probabilities. I show that almost 10 percent of return and more than 25 percent of variation in volatility can be explained by the top 5 models. Third, from an economic perspective, I show that most predictable variation of portfolio returns can be attributed to macroeconomic variables. In particular business cycles proxies are pro-cyclically related to future returns. On the other hand, commodity portfolio volatility is mostly related to financial risk premia proxies that also have predictive power for stocks and bonds. In particular, the dividend yield and default spread have high explanatory power for future volatility. The only variable that is significant for both returns and volatility is the global futures discount. The relation of futures discount with returns is consistent with the view that futures entail a time-varying risk premium, however also indicates that I have omitted important sources of expected commodity return variation. For the volatility, the findings of a significant basis relation are consistent with the Theory of Storage, however might also be attributable to risk premia variation. In summary, the results indicate that returns and volatility of global commodity futures portfolios are driven by different economic sources. In this article, I only document and compare the stylised facts of commodity return and volatility variation but a lot more theoretical work is necessary to explain these observations.

Sector Portfolios In this section, I repeat the previous analysis by focussing on predictive relations in commodity sector portfolios. In particular, I decompose the global commodity portfolio, into agricultural, energy, livestock, metal and precious metal portfolios. I include commodity sector

and cross-sector specific, macroeconomic, and financial variables in the predictive model setup. The main difference to the global analysis is that I allow all commodity-specific variables to enter the predictive models across commodity sectors. This decision is motivated by the findings of De Roon et al. (2000) and Basak and Pavlova (2013), who show that spill-over effects in commodity markets exist. For example, De Roon et al. (2000) show that hedging pressure effects of particular commodities have an effect on other commodity returns. This specification increases the total number of predictive variables from $K = 15$ to $K = 41$ (including lagged values of returns and volatility). Given the high number of variables, I do not resort to closed form solutions, which are computationally intractable, but rely on MC^3 methods as described previously. The Monte Carlo experiment is based on 500000 repetitions for each commodity sector model and I allow for a burn-in period of 50000 steps to avoid sensibilities to the starting point. Furthermore, I use a shorter sample period from 1983-2012 given that some sectors, in particular energy commodities, have a later starting date.

Table 3.4: **Top 5 Models, Variable Inclusions and Statistical Inference of Bayesian Estimation:** This table reports the top 5 predictive models for the global commodity portfolio returns and volatilities, estimated using the Bayesian Model Averaging (BMA) approach in closed form. Panel A reports results for the portfolio returns and Panel B for the portfolio volatility as endogenous variables. The set of predictive variables contains commodity specific variables (global variables only), macroeconomic variables and financial variables as well as lagged dependent return and volatility (RV_{t-1}). Inclusion of a particular variable in the top 5 models is indicated by a bold 1 and selected according to the top posterior model probability. The last three columns of the table report the R^2, adjusted R^2 and the posterior odd ratios (PO) of the top 5 models. The PO ratio is calculated by summing up the posterior probabilities of models, containing only lagged endogenous variables or the random walk benchmark and dividing it by the posterior model probability of the particular top 5 model. The sample period ranges from 1973 to 2012.

Panel A: Portfolio Returns

Top	Ret(t-1)	Vol(t-1)	BaGl	HPGl	OIGl	SCGl	KGIP	BDI	CPI	USIP	DivY	TSpr	DefS	ShRt	TED	R^2	Adj R^2	PO
1	0	0	**1**	0	0	0	0	**1**	0	**1**	0	0	0	**1**	0	0.09	0.08	0.00
2	0	0	0	0	0	0	0	**1**	0	**1**	0	0	0	**1**	0	0.08	0.08	0.00
3	0	0	**1**	0	0	0	**1**	**1**	0	**1**	0	0	0	**1**	0	0.09	0.08	0.00
4	0	0	0	0	0	0	**1**	**1**	0	**1**	0	0	0	**1**	0	0.09	0.08	0.00
5	0	0	0	0	0	0	0	**1**	0	**1**	0	**1**	0	**1**	0	0.09	0.08	0.00

Panel B: Portfolio (Log) Volatility

Top	Ret(t-1)	Vol(t-1)	BaGl	HPGl	OIGl	SCGl	KGIP	BDI	CPI	USIP	DivY	TSpr	DefS	ShRt	TED	R^2	Adj R^2	PO
1	0	**1**	**1**	**1**	0	0	0	0	**1**	0	**1**	0	**1**	0	0	0.27	0.26	0.00
2	0	**1**	**1**	0	0	0	0	0	**1**	0	**1**	0	**1**	0	0	0.25	0.25	0.00
3	0	**1**	**1**	**1**	0	0	0	0	**1**	**1**	**1**	0	**1**	0	0	0.27	0.26	0.00
4	0	**1**	**1**	0	0	0	0	0	**1**	**1**	**1**	0	**1**	0	0	0.26	0.25	0.00
5	0	**1**	**1**	0	0	**1**	0	0	**1**	0	**1**	0	**1**	0	0	0.26	0.25	0.00

Table 3.5 presents the sector-specific results. As a first observation, I find that almost all sector portfolio returns and volatilities have some predictive variables with cumulative inclusion probabilities above 0.5. The only exception are livestock portfolio returns. The general observations of the global portfolio analysis also hold on a sector portfolio level. I find that the baltic dry and/or US industrial production growth variables are among the best predictors in most sectors according to the cumulative posterior inclusion probabilities. Furthermore, I find some evidence that the sector-specific futures discount and, in some cases, cross-sector futures discounts predict sector level returns. For the commodity portfolio volatility the default spread, the dividend yield and the inflation growth emerge as key predictive variables. The only sector portfolio that deviates from this general conclusion are precious metals. I find that the precious metals basis and the dividend yield are key predictors for portfolio returns with inclusion probabilities of 1 and 0.84 and posterior t-ratios of 3.97 and 1.82. Macroeconomic variables do not seem to be relevant to predict precious metal returns. This finding is consistent with the view that gold and silver are mostly pure financial investments. Proxies for systematic risk premia from stock and bond markets, such as the dividend yield, are therefore important to describe return variation for precious metals.

Besides these commonalities across commodity sector portfolios, I also find cross-sectional differences. First, I find that one of the key predictors of portfolio volatility of agricultural commodities is the aggregate open interest growth (OIAg). In particular, I find that the variable has a negative posterior mean coefficient. This finding is interesting as it contributes to the literature on the financialisation of commodity markets. An increase in capital committed to agricultural markets (as measured by open interest) leads to a decrease in volatility. This finding supports the view that an increase in capital invested in agricultural commodity markets actually decreases price volatility. Second, I find that the TED spread, which is a proxy for funding liquidity, has high explanatory power for future energy portfolio volatility. In particular, the posterior coefficient estimate predicts that a decrease in funding liquidity (increase in TED spread) leads to higher energy portfolio volatility. This finding is consistent with the view that capital constraint speculators cannot fully arbitrage inefficiencies in energy markets, which leads to an increase in volatility (Acharya et al. (2013) and Etula (2014)).

Table 3.5: **Predictive Regressions for Sector Commodity Portfolio Returns and Portfolio Volatility using BMA estimation:** This table reports in-sample predictability results for the sector commodity portfolio returns and portfolio volatility, estimated using the Bayesian Model Averaging (BMA) approach with an MC^3 algorithm. The sector portfolios are constructed as equally weighted averages within the sectors agricultural, energy, livestock, metals, and precious metals. The set of predictive variables contains commodity specific variables (cross commodity predictive variable), macroeconomic variables and financial variables as well as lagged dependent return and volatility (RV_{t-1}). In total the number of predictors is 31. The table presents the results separately for portfolio returns on the 5 left hand columns and for portfolio volatility on the 5 right hand columns. The table displays the results for the top 6 predictors ranked according to the cumulative posterior probability of inclusion. From left to right the table reports rank, variable name, cumulative posterior probability of inclusion (Prob), posterior means and posterior t-ratios for each variable. The posterior means of slope coefficients are obtained by averaging slope estimates by the posterior model probabilities across all possible model combinations (BMA approach). The posterior t-ratios are adjusted to account for model uncertainty. In particular, the t-ratios are obtained by dividing the posterior model averaged means across models by its posterior standard error averaged over all models, including model uncertainty that summarizes the dispersion in slopes across models. Variables with cumulative posterior probabilities of more than 50 percent are marked bold. The sample period for all sectors is from 1983-2012.

Portfolio Returns					**Portfolio (Log) Volatility**				
Agriculturals									
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	USIP	**0.99**	1.76	3.84	1	Vol	**1.00**	0.37	8.37
2	KGIP	**0.96**	0.00	2.83	2	OIAg	**1.00**	-1.08	-4.58
3	ShRt	**0.52**	-0.00	-0.89	3	DefS	**0.99**	0.14	4.10
4	BDI	0.29	0.23	0.57	4	CPI	**0.99**	18.12	3.81
5	BaGl	0.20	0.04	0.43	5	DivY	**0.97**	-0.05	-3.42
6	BaPr	0.11	0.02	0.15	6	BaAg	**0.83**	1.58	1.92
Energy									
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	BDI	**0.99**	2.43	3.61	1	Vol	**1.00**	0.43	9.67
2	USIP	0.50	1.38	0.87	2	HPEn	**1.00**	0.67	4.33
3	OIEn	0.38	0.11	0.70	3	KGIP	**1.00**	-0.00	-4.09
4	Ret	0.35	0.05	0.65	4	TED	**0.99**	0.19	3.91
5	HPEn	0.25	0.02	0.51	5	DivY	**0.99**	-0.09	-3.44
6	HPAg	0.11	0.01	0.27	6	BaEn	**0.95**	1.09	1.13
Livestock									
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	USIP	0.50	0.50	0.88	1	CPI	**0.99**	32.22	3.73
2	BaLi	0.20	0.02	0.43	2	DefS	**0.94**	0.14	2.28
3	BaEn	0.11	0.01	0.29	3	DivY	**0.90**	-0.06	-2.16
4	KGIP	0.09	-0.00	-0.26	4	ShRt	**0.83**	-0.03	-1.78
5	BDI	0.09	0.03	0.23	5	BDI	0.45	2.99	0.80
6	SCLi	0.08	-0.00	-0.24	6	Ret	0.32	0.08	0.60

Metals									
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	BDI	**1.00**	2.29	4.72	1	Vol	**1.00**	0.45	6.73
2	BaPr	0.48	0.31	0.82	2	DivY	**0.97**	-0.08	-3.34
3	USIP	0.25	0.37	0.51	3	BDI	**0.92**	-6.55	-2.35
4	HPGl	0.13	0.01	0.31	4	BaEn	**0.69**	-0.81	-1.29
5	Vol	0.12	0.00	0.30	5	DefS	0.32	0.04	0.60
6	ShRt	0.11	-0.00	-0.26	6	BaGl	0.30	-1.09	-0.56
Precious Metals									
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	BaPr	**1.00**	1.61	3.97	1	Vol	**1.00**	0.50	11.06
2	DivY	**0.84**	0.01	1.82	2	Ret	**1.00**	1.02	4.17
3	TSpr	0.13	-0.00	-0.32	3	DefS	**0.99**	0.17	3.55
4	SCPr	0.12	-0.00	-0.30	4	KGIP	**0.81**	0.00	1.68
5	KGIP	0.11	0.00	0.29	5	BDI	**0.70**	-4.82	-1.28
6	BDI	0.08	0.05	0.23	6	BaPr	**0.66**	-0.61	-1.20

Characteristics Based Strategies Finally, I investigate the value of commodity market specific variables, macroeconomic variables, and financial variables in predicting characteristics-based portfolio returns. I focus on the momentum and basis characteristics that has been shown to generate significant positive returns in cross-sectional sorts (Szymanowska et al. (2013) and Gorton et al. (2007)). I focus on returns in the analysis and leave the setup of the estimation methodology unchanged from the sector portfolio analysis.

Two results emerge: First, I find that basis portfolio returns are not related to any of the aggregate predictive variables I consider in this study. This finding corroborates with the finding of Gorton et al. (2007), who show that idiosyncratic sources of risk such as individual inventory levels drive individual basis movements and therefore, cross-sectional return premia. Another explanation is that I have omitted important variables in my study or the predictability emerges at another frequency than monthly. Second, I find that momentum returns are predictable by the short rate (ShRt) with a posterior inclusion probability of only 0.55. The short rate has a positive posterior coefficient estimate, which is opposite to the findings for the global portfolio, suggesting counter-cyclical movements. Consistent with this I find that US industrial production growth is also negatively related to momentum returns as well, however insignificantly.

Table 3.6: **Predictive Regressions for Characteristics Based Commodity Portfolio Returns and Portfolio Volatility using BMA estimation:** This table reports in-sample predictability results for characteristics based commodity portfolio returns, estimated using the Bayesian Model Averaging (BMA) approach with an MC^3 algorithm. The characteristics based portfolios are constructed by ranking commodities according to their individual basis and momentum signal and taking long positions in the top half and a short position in the bottom half. Then an equally weighted average across all commodities is taken. The set of predictive variables contains commodity specific variables (cross commodity predictive variable too), macroeconomic variables and financial variables as well as lagged dependent return variables. In total the number of predictors is 30. The table displays the results for the top 10 predictors ranked according to the cumulative posterior probability of inclusion. From left to right the table reports rank, variable name, cumulative posterior probability of inclusion (Prob), posterior means and posterior t-ratios for each variable. The posterior means of slope coefficients are obtained by averaging slope estimates by the posterior model probabilities across all possible model combinations (BMA approach). The posterior t-ratios are adjusted to account for model uncertainty. In particular, the t-ratios are obtained by dividing the posterior model averaged means across models by its posterior standard error averaged over all models, including model uncertainty that summarizes the dispersion in slopes across models. Variables with cumulative posterior probabilities of more than 50 percent are marked bold. The sample period for all sectors ranges from 1983 to 2012.

Basis Returns					**Momentum Returns**				
No.	Variable	Prob	Post.Mean	Post t-ratio	No.	Variable	Prob	Post.Mean	Post t-ratio
1	BaPr	0.26	0.08	0.47	1	ShRt	**0.55**	0.01	1.82
2	OIEn	0.19	0.01	0.28	2	BaPr	0.35	0.98	1.29
3	KGIP	0.15	-0.00	-0.15	3	USIP	0.19	-0.12	-0.39
4	USIP	0.15	-0.10	-0.36	4	Ret	0.17	-0.01	-0.35
5	DivY	0.14	-0.00	-0.29	5	CPI	0.15	-0.00	-0.27
6	SCPr	0.12	0.00	0.31	6	BaGl	0.13	0.01	0.30
7	OIEn	0.12	-0.00	-0.12	7	SCPr	0.12	0.00	0.27
8	BaEn	0.12	-0.00	-0.10	8	SCGl	0.09	-0.00	-0.07
9	SCEn	0.11	-0.00	-0.09	9	SCAg	0.07	-0.00	-0.07
10	HPGl	0.10	-0.01	-0.30	10	DivY	0.07	0.00	0.11

3.4.2 Bayesian Model Averaging: Out-of-sample Analysis

Thus far my analysis reveals significant in-sample predictability of portfolio returns and portfolio volatility for global, sector and characteristics-based commodity portfolios and examines the relative importance of the predictive variables. However, several recent studies find that significant in-sample predictability does not necessarily translate into any predictive or economic value for a real time investor or forecaster (Goyal and Welsh (2007) and Thornton and Valente (2012)). In this section, I comprehensively test if the predictive results of the previous section generate significant out-of-sample (OOS) result and if they add economic value for a real time mean variance investor as described in section 3.3.2. Importantly, I evaluate all forecasts and portfolio allocation strategies against the null hypothesis of a forecast model that is based on an autoregressive linear specification including lagged values of portfolio returns and volatility. The choice of the autoregressive benchmark model is more robust than a pure random walk benchmark (Della Corte et al (2006)) and corroborates with the intention of the paper to evaluate the value added from non-pure price based measures to forecast commodities. In the statistical OOS tests and the economic value assessment, I rely on recursive estimation techniques with a learning period. In particular, I first estimate the models with data from 1 to P and conduct a mean or volatility forecast for time $P+1$, next I estimate the models with data from 1 to $P+1$ and obtain forecasts for time $P+2$. The final estimation period is from 1 to $T-1$. In the model estimation I rely on the MC^3 algorithm for all commodity portfolios given the computational intensity. I use 100000 Monte Carlo runs and a burn-in period of 10000 runs.

Statistical Measures: I start examining the statistical out-of-sample value of commodity-specific variables, macroeconomic variables, and financial variables in predicting global, sector, and characteristic-based commodity portfolio returns and portfolio volatility. Table 3.7 and 3.8 reports the results from several out-of-sample statistics including the out-of-sample R^2, by Campbell and Thompson (2008), the adjusted mean-squared prediction model, by Clark and West (2007), and the root mean square error. I evaluate the out-of-sample forecasts based on the Bayesian Model Averaging (BMA) approach and the Best Bayesian Model (BBM), which is chosen according to the highest posterior probability across model combinations. I use the autoregressive baseline model as the

Table 3.7: **Out-of-sample Statistical Criteria for BMA Combination Forecasts of Global Portfolio Returns and Portfolio Volatility:** This table presents the out-of-sample R^2 statistic (OOS R^2), the MSPE-adjusted t-statistic (MSPE), and root mean square error (RMSE) for combination forecasts of global commodity returns and volatility using the Bayesian Model Averaging approach (BMA) and the Best Bayesian Model (BBM). The BBM is chosen according to the highest posterior probability across all model combinations. The set of predictive variables contains commodity specific variables (cross commodity predictive variable), macroeconomic variables and financial variables as well as lagged dependent return and volatility. All forecasts are evaluated against the null hypothesis that the BMA model does not significantly improve a forecast based solely on a constant and lagged values of global portfolio returns and portfolio volatility. Results are based on an expanding estimation window that starts at $T_0 = 1990, 1995, 2000$ and 2005.

	Global Portfolios			
	Returns		**(Log) Volatility**	
	BMA	BBM	BMA	BBM
$T_0 = 1990$				
OOS R^2	0.071	0.055	0.055	0.049
MSPE	3.126	3.522		
RMSE	0.001	0.001	0.005	0.005
$T_0 = 1995$				
OOS R^2	0.080	0.064	-0.015	-0.029
MSPE	2.954	3.392		
RMSE	0.002	0.001	-0.001	-0.002
$T_0 = 2000$				
OOS R^2	0.077	0.057	-0.049	-0.072
MSPE	2.673	2.989		
RMSE	0.002	0.001	-0.004	-0.005
$T_0 = 2005$				
OOS R^2	0.097	0.069	-0.093	-0.128
MSPE	2.163	2.383		
RMSE	0.003	0.002	-0.007	-0.010

null hypothesis of no predictive value added.

Table 3.7 reports the out-of-sample forecast evaluation results for global portfolio returns on the left hand columns and global portfolio volatility on the right hand columns. The results are reported for different starting values for the first forecast ($T_0 = 1990, 1995, 2000$ and 2005). I find that predictive models augmented with commodity specific variables, macroeconomic variables, and financial variables generally outperform the autoregressive bench mark model for portfolio returns as indicated by the out-of-sample test statistics. The out-of-sample R^2 are consistently positive and the adjusted mean-squared prediction models consistently statistically significant across the various forecasting periods. This behaviour is true for the BMA combination forecast model as well as the BBM, however the BBM values are generally less significant than the BMA results. Another observation is that the forecasting performance increases with the starting date. The results for the global portfolio volatility are less encouraging. Even though I find that the out-of-sample statistics are significant for the first starting date $T_0 = 1990$, all other forecasting horizons show that the autoregressive model outperforms the BMA and BBM model. Table 3.8 reports the out-of-sample forecast evaluation results for sector and characteristics based portfolio returns in Panel A and sector portfolio volatility in Panel B. The results are reported for different starting values for the first forecast ($T_0 = 1990, 1995$ and 2000). Similarly to the global results, I find that BMA and BBM forecasts generally outperform the autoregressive model for sector portfolio returns. Again the out-of-sample R^2 are consistently positive and the adjusted mean-squared prediction models consistently statistically significant across the various forecasting periods. The best out-of-sample statistics can be observed for energy commodity portfolios, while livestock portfolios are only marginally predictable out-of-sample.

Table 3.8: **Out-of-sample Statistical Criteria for BMA Combination Forecasts of Sector and Characteristics Based Portfolio Returns and Portfolio Volatility:** This table presents the out-of-sample R^2 statistic (OOS R^2), the MSPE-adjusted t-statistic (MSPE),, and root mean square error (RMSE) for predicting sector and characteristics based commodity portfolio returns (Panel A) and portfolio volatility (Panel B) using the Bayesian Model Averaging approach (BMA) and the Best Bayesian Model (BBM). The BBM is chosen according to the highest posterior probability across all model combinations. The set of predictive variables contains commodity specific variables (cross commodity predictive variable), macroeconomic variables and financial variables as well as lagged dependent return and volatility variables. All forecasts are evaluated against the null hypothesis that the BMA model does not significantly improve a forecast based solely on a constant and lagged values of global portfolio returns and volatility. Results are based on an expanding estimation window that starts at $T_0 = 1990, 1995$ and 2000.

Panel A: Portfolio Returns														
	Sector Portfolio										Characteristics Based			
	Agricult.		Energy		Livestock		Metal		Precious		Basis		Momentum	
	BMA	BBM	BMA	BBM	BMA	BBM	BMA	BBM	BMA	BBM	BMA	BBM	BMA	BBM
$T_0 = 1990$														
OOS R^2	0.042	0.034	0.074	0.049	0.004	0.001	0.068	0.045	0.050	0.025	0.007	-0.002	0.025	0.020
MSPE	3.110	2.639	2.777	2.397	1.961	1.737	2.176	2.012	3.611	3.113	1.375	0.397	3.581	2.826
RMSE	0.001	0.001	0.004	0.003	0.002	0.002	0.012	0.001	0.002	0.001	0.000	-0.000	0.000	0.000
$T_0 = 1995$														
OOS R^2	0.046	0.041	0.080	0.052	0.020	0.016	0.085	0.060	0.059	0.031	0.000	-0.010	0.018	0.011
MSPE	2.953	2.554	2.543	2.192	2.010	1.859	2.010	1.826	3.245	2.732	0.562	-0.410	2.602	1.699
RMSE	0.001	0.001	0.004	0.003	0.000	0.000	0.002	0.001	0.002	0.002	0.000	-0.000	0.000	0.000
$T_0 = 2000$														
OOS R^2	0.034	0.029	0.106	0.085	0.029	0.029	0.111	0.085	0.061	0.046	-0.000	-0.013	0.009	0.001
MSPE	2.512	2.189	2.255	2.192	2.556	2.917	2.046	1.928	2.845	2.503	0.383	-0.530	1.431	0.361
RMSE	0.001	0.001	0.005	0.004	0.001	0.001	0.004	0.003	0.002	0.002	-0.000	-0.000	0.000	0.000
Panel B: Portfolio Volatility														
	Agricult.		Energy		Livestock		Metal		Precious					
	BMA	BBM	BMA	BBM	BMA	BBM	BMA	BBM	BMA	BBM				
$T_0 = 1990$														
OOS R^2	0.013	-0.009	0.108	0.061	0.056	0.048	0.041	0.009	0.041	0.024				
RMSE	0.001	-0.001	0.018	0.010	0.011	0.009	0.007	0.001	0.007	0.004				
$T_0 = 1995$														
OOS R^2	0.020	0.004	0.013	-0.021	0.055	0.043	0.018	-0.002	0.011	-0.016				
RMSE	0.002	0.000	0.002	-0.003	0.010	0.008	0.003	-0.000	0.002	-0.003				
$T_0 = 2000$														
OOS R^2	0.015	-0.002	-0.074	-0.103	0.061	0.043	-0.019	-0.037	-0.050	-0.102				
RMSE	0.001	-0.000	-0.011	-0.015	0.010	0.007	-0.003	-0.006	-0.007	-0.015				

I cannot observe the pattern of increasing statistical value with starting date, as in the global portfolios, for the sector portfolios. Within the characteristics-based portfolios, I do not find any evidence of out-of-sample predictive power for the basis sorted portfolio, consistent with the findings in the in-sample section. On the other hand, I find that the BMA and BBM models have some incremental power in predicting momentum portfolios out-of-sample. In particular the forecasting results starting at $T_0 = 1990$ show significant test statistics but deteriorate with starting date. The results for out-of-sample combination forecasts of BMA and BBM for sector portfolio volatility are generally less significant than the return forecasts or not significant at all. Again, I find that the forecasting power deteriorates with the starting date. Nevertheless, some sector portfolios such as agricultural and livestock commodities show significant predictive power across sub samples.

In summary, I find significant results that portfolio returns are predictable out-of-sample for global, sector, and momentum portfolios. On the other hand, predicting portfolio volatility does not add significant value in comparison to an autoregressive model.

Economic Value: Next, I study the economic value of the forecasting models as outlined in section 3.3.2. In particular, I analyse the performance of a dynamically rebalanced mean variance portfolio, constructed using the Bayesian forecasts for conditional means and conditional volatility. The comparability of this analysis is based on the performance fees, ϕ, a US investor is willing to pay for switching from a autoregressive forecasting model that includes only autoregressive terms of portfolio returns and portfolio volatility to a Bayesian model that weights commodity specific, macroeconomic variables and financial variables according to their posterior probabilities in the forecasting regression. I present the results separately for conditional mean and volatility forecasts. For the economic evaluation of the conditional mean forecasts I estimate the predictive model of the mean using the BMA and BBM approach, compare it to the value of the pure autoregressive model and estimate in both cases the conditional variance-covariance matrix using an expanding window of the sample variance-covariance estimator. For the economic evaluation of the conditional volatility forecasts I use the predictive models of conditional volatility using the BMA and BBM approach and estimate the correlation using an expanding window of sample correlations as described in section 3.3.2 in more detail. In this setting I set the conditional mean equal to its sample mean. I present the results for different degrees of relative risk aversion equal to either $\delta = 3$ or $\delta = 8$.

Table 3.9: **Economic Value of Sector Portfolio Return and Portfolio Volatility Predictability using BMA and BBM Combination Forecasts:** This table shows the results from a conditional multivariate mean variance portfolio strategy, based on conditional mean and volatility forecasts, for all sector commodity portfolios (Agriculturals, Energy, Livestock, Metals and Precious Metals) together. Reported are the out-of-sample performance fees (ϕ) and break even transaction costs (τ_{BE}) for three ex-post target portfolio volatilities (8%, 10%, and 12%). The fees represent the maximum percentage an investor with quadratic utility and a degree of relative risk aversion equal to either $\delta = 3$ or $\delta = 8$ is willing to pay to switch from a conditional mean and volatility forecast model with lagged returns and volatility to a model that utilizes all commodity market specific, macroeconomic and financial variables in a BMA or BBM combination forecast framework. BMA denotes the combination forecasts based on the Bayesian model averaging framework and BBM is the best Bayesian model that considers only the forecast variables that were included in the model with the highest posterior probabilities. Panel A reports the results for conditional mean forecasts only (the variance-covariance is estimated using an expanding window for the sample variance-covariance estimator) and Panel B reports the results for fixed conditional mean (conditional means are estimated as sample averages) but utilizing the conditional volatility forecasts. τ_{BE} is defined as the minimum monthly proportional costs which would cancel out the utility advantage of the given strategy. τ_{BE} is only reported whenever ϕ is positive. The performance fees are in annual basis points and the transaction costs in monthly basis points. All out-of-sample results are based on the period from January 1990 to September 2012.

Panel A: Sector Portfolio Returns

	$\delta = 3$						$\delta = 8$					
	BMA			BBM			BMA			BBM		
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
8%	434.84	270.92	0.64	539.11	320.61	0.70	388.21	229.62	0.57	529.03	320.61	0.62
10%	551.65	280.02	0.60	670.86	334.13	0.66	490.77	234.15	0.53	653.08	334.13	0.59
12%	671.70	292.14	0.56	801.39	351.50	0.63	595.52	240.00	0.50	773.81	351.50	0.58

Panel B: Global Portfolio Volatility

	$\delta = 3$						$\delta = 8$					
	BMA			BBM			BMA			BBM		
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
8%	-47.42	-	0.35	-47.73	-	0.34	-57.83	-	0.37	-54.34	-	0.36
10%	-53.11	-	0.31	-53.70	-	0.30	-40.89	-	0.33	-66.71	-	0.32
12%	-58.94	-	0.20	-59.08	-	0.21	-56.39	-	0.20	-71.61	-	0.20

Panel A of Table 3.9 presents the out-of-sample performance fees for a dynamic strategy with BMA and BBM forecasted mean returns for commodity sector portfolios in a multivariate portfolio allocation approach. The results provide robust evidence against the null hypothesis of an autoregressive benchmark model as all performance fees based on the out-of-sample Bayesian forecasts, BMA and BBM, are positive and high. For example, when selecting $\delta = 3$ and $\sigma_p^* = 10\%$ the annual performance fees from switching away from the benchmark autoregressive model is 551bps for BMA and 670bps for BBM forecasts. In all possible combinations of risk aversions and volatilities the BBM model generally outperforms the BMA model. The highest annual performance fee overall is 801bps for the BBM forecast, $\delta = 3$ and $\sigma_p^* = 12\%$. Furthermore, I find that the performance fees for BMA and BBM forecasts are decreasing with the level of risk aversion. For comparison, I also display the annualised Sharpe ratios of the dynamic portfolio models and find that they are as high as 0.7 and 0.64 for the BBM and BMA mean forecast, respectively. The values of the Sharpe ratios decrease with increasing volatility. In summary, there is clear economic evidence on the improvement in performance using commodity specific variables, macroeconomic variables and financial variables relative to the autoregressive benchmark. Figure 3.1 presents the cumulative returns of the dynamic strategies generated with BMA and BBM forecasts. For comparison, I also plot the cumulative returns of the autoregressive benchmark and the returns to basis and momentum trading strategies. Miffre and Rallis (2007) and Fuertes et al. (2011) show that basis and momentum strategies generate large economic profits in comparison to a long only equally weighted commodity strategy. A visual inspection of figure 3.1 reveals that the dynamic strategies based on BMA and BBM are as profitable or even more profitable than basis and momentum strategies. All strategies generate high economic profits in comparison to the autoregressive benchmark.

To better understand the differences between dynamic strategies based on BMA and BBM forecasts and the autoregressive benchmark model, I plot in Figure 3.2 the estimated portfolio weights for each dynamic strategy. In both cases I observe that the volatility of the BMA and BBM portfolio weights is much higher than for the autoregressive model. In particular, in Figure 3.2, it can be seen that for all sectors the weights switch frequently from long to short position and vice versa. In the autoregressive model the portfolio weights remain constantly short or long along the time period with the agricultural commodities being an exception.

The volatile portfolio weights can be dangerous for the dynamic BMA and BBM strategy. The

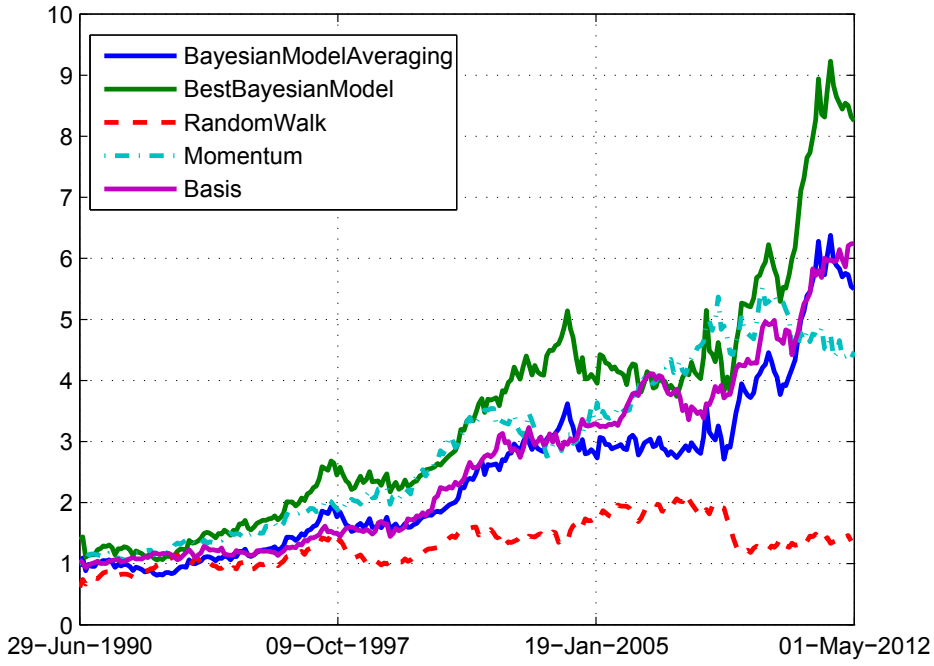


Figure 3.1: **Cumulative Returns of Global Portfolio Allocation Using BMA and BBM Forecasts and Comparisons:** This figure displays the out-of-sample cumulative returns of a mean/variance optimal portfolio with conditional expected BMA and BBM forecasts, the random walk benchmark, the basis and the momentum portfolio. The maximum return strategy builds an efficient portfolio by investing in monthly returns of agricultural, energy, livestock, precious and metal portfolios considering all predictive variables. The sample runs from 1990 to 2012.

frequent changes in portfolio weights can render the strategy to be to costly to implement relative to the autoregressive benchmark. In particular, if transaction costs are sufficiently high they can wipe out the performance gains from a dynamic strategy. I address this issue similar to Della Corte et al. (2006) by computing the break-even transaction cost, τ_{BE}, as the minimum monthly proportional cost, which cancels out the utility advantage of a Bayesian strategy relative to the autoregressive strategy. An investor, who pays lower transaction cost than, τ_{BE}, will prefer the Bayesian strategy. The, τ_{BE}, are expressed as monthly basis points. Panel A of Table 3.9 presents the break even transaction costs for each risk aversion and volatility combination. The τ_{BE} are all positive, large, and are around 280bps and 330bps per month for BMA and BBM forecasts, respectively. In summary, I conclude that the economic value of the dynamic commodity strategies I reported is robust to the inclusion of transaction costs.

Panel B of Table 3.9 presents the out-of-sample performance fees for a dynamic portfolio allo-

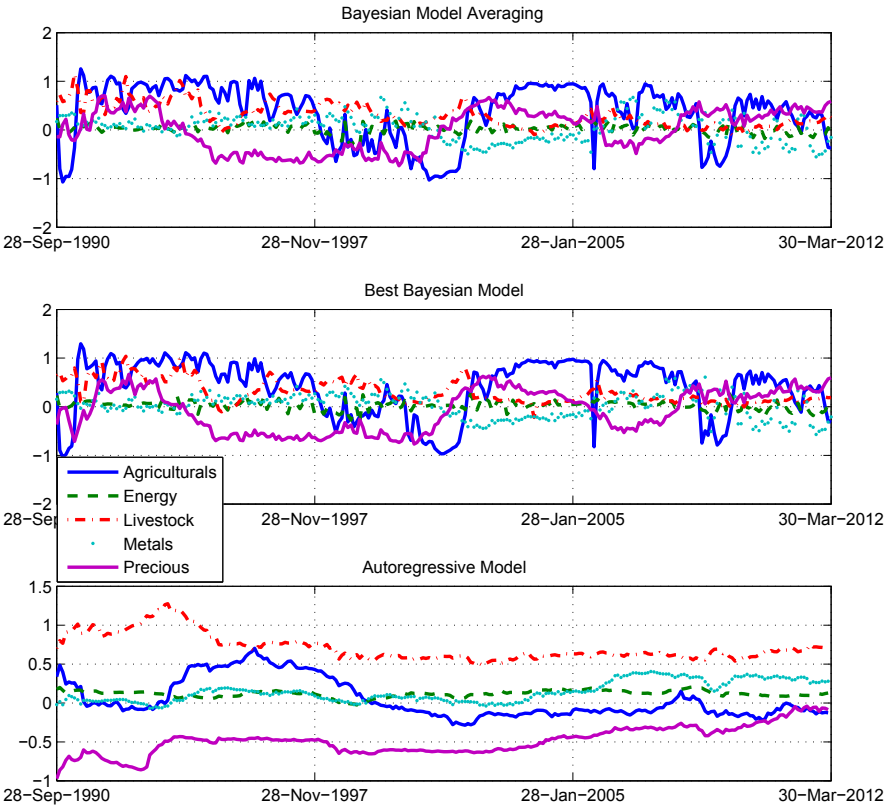


Figure 3.2: **Dynamics of Portfolio Weights of Global Portfolio Allocation Using BMA and BBM Forecasts and Weights of Autoregressive Benchmark:** The figure displays the time variation of the optimal weights of a mean/variance optimal portfolio with conditional expected BMA and BBM forecasts and for the autoregressive benchmark. The optimal weights are displayed for agricultural, energy, livestock, precious and metals portfolio for the time period of 1990 to 2012.

cation strategy with forecasted volatilities. The finding is consistent with the statistical evidence in that I do not find any positive performance fee and therefore, no economic gain of including commodity specific variables, macroeconomic variables, and financial variables relative to the autoregressive benchmark in the predictive regressions. In fact, the performance fees are even negative, however, very close to zero.

Next, I examine the relative economic value that each of the predictive variable classes add to the overall performance of a dynamically rebalanced mean variance portfolio. To do so, I include only either commodity market-specific variables, macroeconomic variables, or financial variables to generate BMA and BBM forecasts and estimate the portfolio weights. Table 3.10 reports the results

Table 3.10: **Economic Value of Global Portfolio Return and Portfolio Volatility Predictability using BMA Combination Forecasts for Specific Predictor Classes:** This table shows the results from a conditional multivariate mean variance portfolio strategy, based on conditional mean forecasts for specific forecasting variables, for all sector commodity portfolios (Agriculturals, Energy, Livestock, Metals and Precious Metals) together. Reported are the out-of-sample performance fees (ϕ) and break even transaction costs (τ_{BE}) for three ex-post target portfolio volatilities (8%, 10%, and 12%). The fees represent the maximum percentage an investor with quadratic utility and a degree of relative risk aversion equal to either $\delta = 3$ or $\delta = 8$ is willing to pay to switch from a conditional mean forecast model with lagged returns and volatility to a model that utilizes only macroeconomic (Panel A), commodity market specific (Panel B), and financial variables (Panel C) in a BMA or BBM combination forecast framework. BMA denotes the combination forecasts based on the Bayesian model averaging framework and BBM is the best Bayesian model that considers only the forecast variables that were included in the model with the highest posterior probabilities. τ_{BE} is defined as the minimum monthly proportional costs which would cancel out the utility advantage of the given strategy. τ_{BE} is only reported whenever ϕ is positive. The performance fees are in annual basis points and the transaction costs in monthly basis points. All out-of-sample results are based on the period from January 1990 to September 2012.

Panel A: Macroeconomic Predictors

	$\delta = 3$						$\delta = 8$			
	BMA			BBM			BMA		BBM	
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	ϕ	τ_{BE}
8%	384.22	194.37	0.58	394.13	187.85	0.59	341.49	194.26	370.83	187.85
10%	466.82	194.59	0.53	484.95	188.10	0.54	399.46	194.49	447.87	188.10
12%	544.00	194.79	0.51	572.67	188.32	0.52	446.26	194.69	518.56	188.32

Panel B: Commodity Market Specific Predictors

	$\delta = 3$						$\delta = 8$			
	BMA			BBM			BMA		BBM	
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	ϕ	τ_{BE}
8%	101.55	24.41	0.34	54.57	11.79	0.29	79.52	24.37	48.71	11.79
10%	120.53	24.45	0.29	66.34	11.80	0.25	86.47	24.41	57.09	11.80
12%	136.93	24.47	0.27	77.36	11.80	0.22	88.20	24.44	63.95	11.80

Panel B: Financial Predictors

	$\delta = 3$						$\delta = 8$			
	BMA			BBM			BMA		BBM	
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	ϕ	τ_{BE}
8%	19.99	3.71	0.25	64.92	51.19	0.30	36.65	5.14	47.65	51.19
10%	27.45	4.03	0.21	74.22	56.40	0.26	50.95	5.59	45.51	56.40
12%	35.90	4.48	0.18	80.74	63.63	0.23	67.31	6.20	37.67	63.63

from the restricted Bayesian forecasting and economic value assessment. As a first observation I find that for each predictor class the performance remains positive and independent from the risk aversion and volatility level. These findings indicate that each variable class has some economic value to generate the overall performance. However, I find large cross-sectional differences. Table 3.10 Panel A to Panel C reports the performance fees and transaction costs for macroeconomic, commodity-market specific and financial variables, respectively. The macroeconomic variables contribute most of the overall performance fees with values of up to 572bps per year, followed by the commodity market specific variables that generate a performance fee of up to 136bps and the financial predictors having a performance fee of only up to 80bps. These findings resemble the results from the statistical analysis that suggests macroeconomic variable being the main predictors of commodity returns. When observing the break even transaction costs, the only economically viable strategy is the macroeconomic strategy with τ_{BE} of up to 200bps, while τ_{BE} is only up to 90bps for commodity specific and 80bps for financial variables. In this final section, I present two robustness test. First, Table 3.11 repeats the previously performed analyses for four sub-samples (1990-1995, 1995-2000, 2000-2005, and 2005-2012). While the general conclusions of the previous sections remain the same, some addition features of the dynamic portfolio allocations appear. I find that the predictive relations do not add economic value in the early sample from 1990-1995 for all predictors and predictor classes. The performance of the predictive relations increases, however, in the periods from 1995 to 2012 and is in most cases positive. In particular the late period from 2005 to 2012 is characterised by high economic profits using macroeconomic predictors. The economic evaluation of the BMA and BBM forecasts is, furthermore, independent of the Sharpe ratio the strategy generated in the particular time period. Second, Table repeats the dynamic portfolio allocation using the global commodity portfolio as the only risky asset. I, therefore, first equally weight the sector portfolios, then estimate the BMA and BBM return and volatility forecasts and finally calculate the optimal mean/variance weights to generate the dynamic strategy. I find that the performance fees to switch from the autoregressive benchmark model to the dynamic BMA and BBM strategy are higher than for the multivariate sector portfolio strategy. In particular the performance fees accumulate up to 1119bps per year for $\sigma_p^* = 12\%$ and $\delta = 3$. At the same time the break even transaction costs are as high as 520bps per month and the Sharpe ratio is almost 0.9. The high performance is related to the higher statistical significance of the predictive relations

Table 3.11: **Economic Value of Sector Portfolio Return and Portfolio Volatility Predictability in Subsample Analysis:** This table shows the results from a conditional multivariate mean variance portfolio strategy, based on conditional mean and volatility forecasts, for all sector commodity portfolios (Agriculturals, Energy, Livestock, Metals and Precious Metals) for 5 year sub-samples. Reported are the out-of-sample performance fees (ϕ), break even transaction costs (τ_{BE}) and Sharpe ratio for ex-post target portfolio volatility of 10%. The fees represent the maximum percentage an investor with quadratic utility and a degree of relative risk aversion equal $\delta = 3$ is willing to pay to switch from a conditional mean and volatility forecast model with lagged returns and volatility to a model that utilizes all commodity market specific, macroeconomic and financial variables in a BMA or BBM combination forecast framework. Panel A reports the results when all predictors are included and Panel B-D when only, macroeconomic, financial and commodity market specific variables are included. τ_{BE} is defined as the minimum monthly proportional costs which would cancel out the utility advantage of the given strategy. The performance fees are in annual basis points and the transaction costs in monthly basis points. The results are based on the period from January 1990 to September 2012.

Panel A: All Predictors

	BMA			BBM		
	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
1990-1995	-117.76	-	0.22	-2.67	-	0.26
1995-2000	440.94	113.34	0.59	433.82	138.19	0.63
2000-2005	488.86	35.20	0.80	509.97	56.40	0.88
2005-2012	1165.43	452.75	0.43	1132.61	531.35	0.40

Panel B: Macroeconomic Predictors

	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
1990-1995	-140.12	-	0.37	-222.12	-	0.33
1995-2000	335.57	151.58	1.01	628.72	130.61	0.99
2000-2005	234.39	85.46	0.34	232.57	79.38	0.35
2005-2012	1267.52	417.39	0.91	1491.78	400.35	0.93

Panel C: Commodity Market Specific Predictors

	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
1990-1995	42.41	27.13	0.67	-652.51	-	0.21
1995-2000	793.99	134.24	0.94	1004.36	133.06	0.95
2000-2005	-136.12	-	0.40	-332.32	-	0.30
2005-2012	485.41	93.76	0.04	574.37	77.95	-0.03

Panel D: Financial Predictors

	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
1990-1995	-635.93	-	-0.00	-686.14	-	-0.04
1995-2000	-442.73	-	-0.05	-56.38	-	0.32
2000-2005	-80.74	-	0.51	21.61	-	0.56
2005-2012	198.74	135.01	-0.25	190.89	123.48	-0.25

when examining the global portfolio in the previous section. I do not find any economic significance for volatility predictions.

Table 3.12: **Economic Value of Global Portfolio Return and Portfolio Volatility Predictability using BMA Combination Forecasts:** This table shows the results from a conditional mean variance portfolio strategy, based on conditional mean and volatility forecasts, with the global commodity portfolio as the only risky asset. Reported are the out-of-sample performance fees (ϕ) and break even transaction costs (τ_{BE}) for three ex-post target portfolio volatilities (8%, 10%, and 12%). The fees represent the maximum percentage an investor with quadratic utility and a degree of relative risk aversion equal to either $\delta = 3$ or $\delta = 8$ is willing to pay to switch from a conditional mean and volatility forecast model with lagged returns and volatility to a model that utilizes all commodity market specific, macroeconomic and financial variables in a BMA or BBM combination forecast framework. BMA denotes the combination forecasts based on the Bayesian model averaging framework and BBM is the best Bayesian model that considers only the forecast variables that were included in the model with the highest posterior probabilities. Panel A reports the results for conditional mean forecasts only (the variance-covariance is estimated using an expanding window for the sample variance-covariance estimator) and Panel B reports the results for fixed conditional mean (conditional means are estimated as sample averages) but utilizing the conditional volatility forecasts. τ_{BE} is defined as the minimum monthly proportional costs which would cancel out the utility advantage of the given strategy. τ_{BE} is only reported whenever ϕ is positive. The performance fees are in annual basis points and the transaction costs in monthly basis points. All out-of-sample results are based on the period from January 1990 to September 2012.

Panel A: Global Portfolio Returns

	$\delta = 3$						$\delta = 8$					
	BMA			BBM			BMA			BBM		
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
8%	642.62	517.22	0.86	587.33	562.24	0.80	696.12	437.69	0.89	600.35	562.24	0.80
10%	880.43	516.69	0.81	811.00	561.71	0.76	951.49	437.81	0.85	830.26	561.71	0.76
12%	1119.09	520.79	0.77	1035.39	567.20	0.72	1209.34	441.55	0.81	1062.08	567.20	0.72

Panel B: Global Portfolio Volatility

	$\delta = 2$						$\delta = 6$					
	BMA			BBM			BMA			BBM		
σ_p^*	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe	ϕ	τ_{BE}	Sharpe
8%	-11.45	-	0.62	-7.73	-	0.62	-27.83	-	0.61	-24.34	-	0.62
10%	-15.31	-	0.43	-10.70	-	0.43	-40.89	-	0.43	-36.71	-	0.43
12%	-19.58	-	0.30	-14.08	-	0.31	-56.39	-	0.30	-51.61	-	0.31

3.5 Conclusion

This chapter provides a comprehensive analysis of return and volatility predictability in commodity markets by commodity market specific, macroeconomic and financial variables. My main goal is to shed light on the relative statistical value of each predictive variable and to conduct robust out-of-sample inference using statistical and economic criteria. Compared to previous literature I extend the setup of the analysis in three key dimensions: First, I allow for a comprehensive set of predictive variables, that have previously been shown to possess forecast power for commodity returns, to enter the predictive regressions. Second, I use Bayesian model averaging to conduct a comprehensive analysis of selection and combination forecasts to infer the value added from each predictive variable. Third, I use dynamic portfolio strategies to assess the economic value added from commodity return and volatility forecasts.

I find that portfolio returns and volatility are predictable using economically motivated variables. I draw this conclusion from significant posterior inclusion probabilities of predictive variables within the Bayesian model averaging methodology. In particular, for commodity returns, I find that macroeconomic variables of industrial production growth play a dominant role in the predictive regression. For the predictability commodity volatility I find information content in inflation growth, dividend yields, and default spreads. The finding suggests that commodity returns and volatility are driven by distinct information sources. The results for commodity return predictability are also supported in realistic out-of-sample settings. I find that commodity return predictions add robust statistical and economic value in a real time setting.

CHAPTER 4

OPTIMAL FUTURES PORTFOLIOS AND HEDGE FUND CAPITAL

4.1 Introduction

Futures contracts have become an integral part in many diversified portfolios. In particular dynamic trading strategies that exploit the predictable variation of futures returns such as carry and momentum strategies have attracted the interest of a variety of institutional investors. Carry strategies relate the directly observable current spread between futures and spot price to the expected return on an asset (Koijen et al. (20012)), while momentum strategies resort to a purely statistical measure that relates an asset's expected return to its own past return (Moskowitz et al. (2011)). Moskowitz et al. (2011) and Koijen et al. (20012) show that carry and momentum signals predict expected returns in futures markets consistently across asset classes leading to significant trading profits. The high profitability of carry and momentum strategies has generated a wide debate as to the underlying explanations for the significant return premia, and presents a strong challenge to standard finance theory[1].

In this article we contribute to the literature by studying the economic anatomy of optimal portfolios based on carry and momentum signals in a unifying framework. We take the viewpoint of a risk averse investor who maximizes his expected utility of a portfolio of futures contracts across equities, bonds, currencies, commodities and metals. We model the portfolio weights of each asset as a function of the asset's momentum and carry signal using the parametric portfolio approach by Brandt et al. (2009). Our data covers the period from January 1980 to January 2012, and we study the cross section of up to 54 futures contracts spanning across the asset classes of currencies, commodities, equity indices and bonds.

Besides using the parametric portfolio choice framework, we go beyond previous research in carry and momentum strategies by (a) providing an in-depth analysis of the performance of the optimal

1. Most research on the profitability of carry and momentum strategies has focused on momentum in equity markets, Fama and French (1992), Jegadeesh and Titman (1992) and Grinblatt and Moskowitz (2004), and carry in currency markets, Burnside et al. (2008), Lustig and Verdelhan (2008) and Menkhoff et al. (2011). More recently the concepts have been generalized for other asset classes or across asset classes, e.g., Menkhoff et al. (2012) and Gorton et al. (2007) for momentum in currency and commodity markets, respectively, or Asness et al. (2013),Moskowitz et al. (2011) and Koijen et al. (20012) for momentum and carry strategies across asset classes. Theoretical explanations for the return premia, generally fall into the categories, rational risk-based (e.g. Fama and French (1996), Jegadeesh and Titman (2002), Sarno et al. (2012)), or characteristics-based models (e.g. Moskowitz and Grinblatt (1999)), and explanations based on limits to arbitrage (e.g. Jylhä and Suominen (2011), Lesmond et al. (2004))

portfolios, including a detailed description of the characteristics of the strategy returns and portfolio weights, the sensitivity of the returns to changes in risk aversion levels, and the return profile over time, (b) studying the relative contribution of carry and momentum signals for the return attribution in a statistical framework, (c) studying the diversification benefit of including optimal futures portfolios in traditional and non-traditional equity and bond portfolios,(d) quantifying the importance of transaction costs, (e) study the relation to standard and non standard risk factors and investigate the performance over the business cycle, and (f) study the influence of institutional investors, in particular Hedge Funds, on the profitability of the trading strategies.

We find large and significant return premia to optimal futures portfolios based on carry and momentum, when estimating the parametric portfolio policy in- and out-of-sample. The certainty-equivalent (CE) gain from investing in optimal carry and momentum strategies relative to the risk-free rate is 15 percent per annum (p.a.) and has a Sharpe ratio of 1.2. The CE and Sharpe ratio are 8 percent and 0.5 higher, than the performance of a volatility weighted long-only portfolio, suggesting that a large part of the additional return premia can be attributed to the predictable variation of excess futures returns due to the carry and momentum signal. We find that the optimal out-of-sample portfolio weight of the momentum signal is 1.5 times the carry signal, nevertheless, both signals contribute statistically significantly to the performance of optimal carry and momentum strategies. Increasing the level of risk aversion reduces the allocation to carry and momentum strategies, however the relative importance of carry and momentum signal in the portfolio remains constant. Furthermore, we find that the diversification benefit of including optimal carry and momentum strategies in portfolios of stock market factors and bonds, as in Barroso and Santa-Clara (2013) and Burnside et al. (2006), is highly desirable from an investors perspective. We find that including optimal carry and momentum strategies in a traditional portfolio increases the Sharpe ratio on average by 0.7 out-of-sample. Even more important, the inclusion reduces the skewness and kurtosis of the original portfolios and therefore minimizes the overall risk.

In order to understand the high returns to optimal futures portfolios, we investigate if they can be explained by (i) the inclusion of transaction costs, (ii) traditional risk factors, business cycle and liquidity risk, and (iii) institutional investors influences related to segmented markets. First, we find that transaction costs do not play a significant role in explaining the significant return premia. Even though the inclusion of transaction costs marginally reduces the Sharpe ratio of the optimal

futures portfolios they cannot account for the full return premia. The result is not surprising given that transaction costs in futures markets are much smaller than in equity and currency markets.

Second, we test if time variation in optimal futures portfolios can be explained by traditional risk factors, business cycles (e.g., Cooper and Priestley (2009), Ludvigson and Ng (2009) and Chordia and Shivakumar (2002)) and liquidity risk (e.g., Brunnermeier et al. (2008)). While tradition risk factors such as Fama and French (1992) three factor model, the value and momentum "everywhere" factors of Asness et al. (2013) and liquidity risk (Pastor and Stambaugh (2002)) do not explain the full extend of the return premia we find that macroeconomic and common systematic variables predict the returns to optimal carry and momentum portfolios. We find that the relation between macroeconomic variables and optimal futures returns is pro-cyclical consistent with the findings of Chordia and Shivakumar (2002) but inconsistent with the tradition macro finance view that asset risk premia are countercyclical (Cochrane (2011)).

Third, we test if institutional investors and in particular hedge fund activity has an effect on the return premia to optimal futures portfolios. Recently, the literature in financial economics has started to investigate the influence of segmented markets on the return generating process, e.g., Brunnermeier (2008), Gabaix et al. (2007) and Daniel (2011). Segmented markets have an effect on expected returns because financial agents, who are active in a particular market, have only limited risk sharing abilities. Most theoretical models predict that expected returns are related to capital and capital flows deployed in the asset classes. The limited risk bearing capacity of speculators generate downward sloping demand curves that result in a positive relation between capital flows and contemporaneous returns. Consequently, expected returns are negatively correlated with the total capital committed within an asset class. We view the returns of carry and momentum strategies through the lens of segmented markets, exploring the relation between speculative capital and the profitability of the trading strategies.

Consistent with the theoretical predictions of segmented market models, we find that speculative capital of hedge funds is a strong determinant of the time-varying profitability of optimal carry and momentum strategies. We first show that the returns to carry and momentum strategies account for a large fraction of aggregated hedge fund returns, with 40 percent of hedge fund return variance across styles explained by market and optimal carry and momentum factors. Although it is not the main focus of the paper, we note that the optimal carry and momentum strategies explain on

average a larger fraction of hedge funds returns, and further reduce the statistical significance of the residual alpha, than the factors of Fung and Hsieh (2004). We next regress optimal carry and momentum returns on contemporaneous capital flows and lagged total assets under management (AUM) in the hedge fund industry. We find that hedge fund capital flows are positively related to the strategy returns, while total AUM negatively predicts expected returns, consistent with the theoretical predictions of segmented market theories.

To understand the relative contribution of business cycle risk and hedge fund capital flows, we run regressions of strategy returns on hedge fund variables together with macroeconomic business cycle indicators. We find that the hedge fund variables completely drive out the predictability of the macroeconomic variables. This suggests that the strategy return and business cycle relation is due to pro-cyclical capital flows into the hedge fund industry. For robustness we re-run all regressions on a long-only portfolio. Here, we find that expected returns are counter-cyclically related to business cycle conditions, consistent with macroeconomic theory, and are not related to speculative capital of hedge funds.

This chapter differs from the recent literature in five main aspects. First, we study the returns of optimal dynamic futures portfolios conditioned on carry and momentum signals from the perspective of a utility maximizing risk averse investor. Most studies on trading strategies focus on simple, equally or cross-sectionally weighted, portfolios. This choice is guided by evidence that simple allocation rules tend to outperform more complex optimized portfolios (e.g., DeMiguel et al. (2009)). However, this simplified framework ignores uncertainties that investors face. Namely, investors have to deal with the signal choice, the weighting of each signal, their risk preferences and be aware of changing cross and inter asset class correlation changes. In the setting of Brandt et al. (2009) we can accommodate preference structures other than simple mean-variance utility functions. The main advantage of using a preference structure such as utility with constant relative risk aversion (CRRA) is that it incorporates preferences to higher-order moments in a parsimonious manner. This feature is particularly desirable when studying momentum and carry strategies that are exposed to crash risks leading to negatively skewed return distributions, e.g., Brunnermeier et al. (2008). Second, the parametric portfolio approach avoids the need to model all the moments of the conditional return distribution. The methodology makes the portfolio estimation computationally tractable and avoids the estimation of a high dimensional parameter space. Third, our framework

allows us to study the joint contribution of carry and momentum signals to the return process together in a unifying framework. Most recent studies of carry and momentum have focused on the strategies individually, however there are indications that the correlation between carry strategies and momentum strategies is very low, e.g., Menkhoff et al. (2012) and Koijen et al. (20012). Analogous to Asness et al. (2013), who combine value and momentum strategies, combining carry and momentum strategies can not only add to the diversification benefit from a portfolio allocation perspective, but also helps us to understand the common drivers of strategies from an asset pricing perspective. Fourth, besides the in-sample results we also report extensive out-of-sample statistics. This is important give that for example Daskalaki and Skiadopoulos (2011) show that the benefit of including commodities in a diversified portfolio of a utility maximizing investor is significant in-sample, but commodities do not add significant value out-of-sample. Fifth, we study the predictable variation of futures returns across a wide range of asset classes; commodities, currencies, bonds, and equities. Most previous studies on momentum and carry strategies resort to individual asset classes. Studying return variation across asset classes reduces the influence of idiosyncratic asset class specific sources of risk, and helps us to identify systematic sources of variation. We therefore contribute to the recent asset pricing literature that attempts to develop financial theories that are applicable to a variety of asset classes, markets and strategies together in a unifying framework, see Asness et al. (2013), Koijen et al. (20012) and Moskowitz et al. (2011).

The remainder of this chapter proceeds as follows. In Section 4.2 we describe the construction of optimal parametric portfolios by Brandt et al. (2009) that we adapt to accommodate asset classes with various levels of volatility, the data we use, how we construct futures returns and our construction of the momentum and carry signal. In section 4.3 we describe the data used. In Section 4.4 we characterize optimal carry and momentum portfolios in terms of profitability, changing risk aversion, risk exposure, transaction costs and diversification benefit. Section 4.5 presents the empirical results of the theoretical explanations for time variation in returns and Section 4.6 concludes.

4.2 Portfolio Construction and Data

In this section we describe the general framework of parametric portfolio construction introduced by Brandt et al. (2009) and adapt it to futures markets that feature different level of contract specific volatilities. Furthermore, we provide details on the data we use and the construction of tradable futures return series, as well as the construction of carry and momentum characteristics.

4.2.1 Optimal Futures Portfolios with Predictable Returns

We optimize futures portfolios from the perspective of a risk averse utility maximizing investor, who has access to futures markets across a range of asset classes. The investor's problem is to maximize conditional expected utility of the portfolio return, $r_{p,t+1}$, by choosing portfolio weights $w_{i,t}$ for N_t futures contracts, where each futures contract, i, has a return of $r_{i,t+1}$ from date t to $t+1$,

$$\max E_t\left[u(r_{p,t+1})\right] = \max_{\{w_{i,t}\}_{i=1}^{N_t}} E_t\left[u\left(r_t^f + \sum_{i=1}^{N_t} w_{i,t} r_{t+1}^i\right)\right]. \tag{4.1}$$

We assume the investor has constant relative risk aversion (CRRA) utility, with risk aversion factor γ. The main advantage of CRRA utility is that it incorporates preferences to higher-order moments. In contrast to standard mean variance utility, which focuses purely on preferences in the first two moments of the distribution of portfolio returns, CRRA utility penalizes kurtosis and skewness (Brandt et al. (2009)).

Optimizing the portfolio using a traditional optimization approach requires the investor to model the joint distribution of returns and characteristics. The investor is required to estimate the conditional means, variances, and covariances of all asset returns as a function of their characteristics. If the investor uses a utility function other than the quadratic one the parameter estimation for the skewness, kurtosis and even higher moments, by portfolio optimization becomes extremely challenging and robust estimates are difficult to obtain.

To address this problem we adapt the recently developed portfolio optimization methodology by Brandt et al. (2009), who parameterize the optimal portfolio weights as a linear function of

asset's characteristics. In this setting the portfolio weight on asset i at time t is defined by,

$$w_{i,t} = \theta^\top \frac{x_{i,t}}{\sigma_{i,t} N_t}, \tag{4.2}$$

where $x_{i,t}$ is a vector of characteristics of futures contract i at time t. The portfolio policy, θ, is a vector of parameters to be estimated. In addition to scaling the weights by N_t as in Brandt et al. (2009), which allows us to optimize the portfolio with an arbitrary and time varying number of assets without changing the characteristics of the allocation, we scale the weights and therefore the returns by their respective volatilities. The volatility scaling parameter $\sigma_{i,t}$, which is defined as the past 3 months absolute value of daily returns[2], ensures that the performance of the optimal portfolio is not dominated by the returns of higher volatility assets (Moskowitz et al. (2011)). It is important to note that we do not scale the individual returns to optimize the portfolio returns, but rather to take into account the very different levels of volatility across diverse asset classes such as bonds, equities, currencies and commodities. To obtain a meaningful comparison between, and portfolio allocation across, the different asset classes the dynamic volatility scaled weights considered in this article reduce the leverage of assets with a high volatility and increase the leverage of assets with low volatility to reach an equal level of volatility for all asset classes. Even though we impose a linear structure of portfolio weights, this parametrization is not restrictive in the sense that we allow for positive or negative weights and therefore, long and short futures positions.

For the choice of characteristics, $x_{i,t}$, we restrict ourself to the use of carry and momentum that have previously been shown to be related to expected futures returns (e.g. Koijen et al. (20012) and Moskowitz et al. (2011)). Carry for asset i at time t is typically defined as the relative difference between the front month futures contract and second to maturity contract, $\frac{F^{T+1}_{i,t-j}}{F^{T}_{i,t-j}} - 1$, where $F^{T}_{i,t}$ and $F^{T+1}_{i,t}$ are the prices of the time t front and second to maturity futures contracts, respectively. We adapt the definition of the carry characteristic to make it suitable for the presented portfolio optimization framework. We require that the measure is (1) proportional to the strength of the futures basis, (2) comparable across the cross-section of futures contracts for assets across investment classes, and (3) considered in isolation for each asset rather ranked over the cross-

2. Using the absolute value, instead of the squared returns reduces the impact of extreme realization on the volatility estimator.

section. To generate the measure we take the monthly futures basis of each asset divided by mean absolute deviation of the futures basis in the previous 12 months.

$$x_{i,t}^{carry} = \frac{F_{i,t}^{T} - F_{i,t}^{T+1}}{\sigma_t^{255}(F_{i,t}^{T} - F_{i,t}^{T+1})} \tag{4.3}$$

where $\sigma^{255}(F_{i,t}^{T} - F_{i,t}^{T+1})$ is defined as the annualized daily mean absolute deviation of the futures basis for the previous 255 days. This normalization is an important feature that ensures that for a given carry portfolio policy, our returns are not dominated by assets with high carry volatility.

The momentum characteristic is constructed using a similar method to Brandt et al. (2009) and Moskowitz et al. (2011), using the 12 month average of the excess return defined in Equation 4.9. However, for our framework, due to the range in volatility of excess returns between futures contracts, we normalise the average excess return by the annualized mean absolute deviation of the past 255 days excess returns. This ensures that higher volatility futures do not have disproportionately large characteristics that would dominate the portfolio weighting. The momentum characteristic for the contract i at time t is therefore

$$x_{i,t}^{\text{mom}} = \frac{1}{12}\sum_{s=1}^{12} \frac{r_{i,t-s}}{\sigma^{255}(r_{i,t})}, \tag{4.4}$$

where $\sigma^{255}(r_{i,t})$ is the annualised daily standard deviation of excess returns of the ith futures contract for the past 255 days. The momentum characteristic, analogously to carry, deviates from the literature since it is not ranked as in cross-sectional studies (Carhart (1997)), and takes the magnitude of excess returns, rather than only the sign, of excess returns (Moskowitz et al. (2011))[3].

3. By scaling the futures basis and the 12 month excess return for individual carry and momentum characteristics we ensure that no one asset has a disproportionably large characteristics, and therefore excessive extreme portfolio weight. However, although this re scaling ensures characteristics are consistent in the cross-section of assets for each portfolio parameter, the necessary and fundamentally different scaling methods between carry and momentum characteristics lead to dissimilar variation over all assets between characteristics. To ensure that portfolio coefficients are comparable, we follow Brandt et al. (2009) and re normalize over all assets for each portfolio coefficient, dividing each characteristic by its respective standard deviation measured across all assets. It is important to realize that this has no effect on optimal portfolio weights, or performance metrics, but simplifies the interpretation of portfolio policy coefficients to allow the comparison of their relative magnitudes.

4.2.2 Portfolio Estimation

The parametric structure of portfolio weights introduced in the previous section implies that we can rewrite the conditional portfolio optimization with respect to the portfolio weights $w_{i,t}$ as the unconditional optimization with respect to the coefficient θ, as

$$\max_{\theta} E\left[u(r_{p,t+1})\right] = E\left[u\left(r_t^f + \sum_{i=1}^{N_t} \theta^\top \frac{x_{i,t}}{\sigma_{i,t} N_t} r_{t+1}^i\right)\right]. \tag{4.5}$$

which can be estimated by maximizing the corresponding sample analog with respect to θ, given by

$$\max_{\theta} \frac{1}{T} \sum_{t=0}^{T-1} u(r_{p,t+1}) = \frac{1}{T} \sum_{t=0}^{T-1} u\left(r_t^f + \sum_{i=1}^{N_t} \theta^\top \frac{x_{i,t}}{\sigma_{i,t} N_t} r_{t+1}^i\right). \tag{4.6}$$

To obtain our estimate of the optimal parametric portfolio we therefore reach the optimization problem,

$$\hat{\theta} = \arg\max_{\theta} \sum_{t=0}^{T-1} u\left(r_t^f + \sum_{i=1}^{N_t} \theta^\top \frac{x_{i,t}}{\sigma_{i,t} N_t} r_{t+1}^i\right). \tag{4.7}$$

The chosen parametrization of the optimal portfolio problem has several main advantages to the traditional Markowitz approach, particularly when applied to optimal futures portfolios. First, the parameter space to be estimated is reduced to the dimension of θ, which is only as high as the number of characteristics included in the problem. Second, the constant coefficient θ implies that the portfolio weight in each asset depends only on the futures's characteristics and not its historic return. This is important as we are not interested in which asset class or individual asset performed best but only on the relative contribution of momentum and carry signal. Third, the parametric portfolio policy optimizes a utility function and not a measure of the distance between forecasted and realized return. Therefore, θ conveys information about higher order moments, even if it does not convey any information for the first moment. Fourth, we can easily embed the volatility scaling as in Moskowitz et al. (2011) and Baltas and Kosowski (2013) that makes the returns in each asset class comparable. Finally, by posing the portfolio problem as an estimation problem, we can conduct statistical inference for the estimated parameters $\hat{\theta}$ using simple generalized methods of moments (GMM) procedure. More details on the statistical inference are in Appendix A and in Brandt et al. (2009).

4.2.3 Construction of Tradable Futures Return Series

The construction of individual excess futures returns, r^i_{t+1}, for the estimation of the portfolio policy in Equation (4.7), is non-trivial. In contrast to equities, futures contracts have a finite life defined by the contractual delivery date and are 'zero cost' securities that do not require any initial cash payments, besides margins, which constitute a fixed percentage.

We define excess futures returns as the change in prices from date t to $t+1$ for a futures contract with the same delivery date, T. We assume an investor allocates $M_{i,t}$ dollars of capital today, at time t, to finance a futures contract on asset i[4]. For the future i, the contract at time t that expires in period T is denoted as $F^T_{i,t}$ where $t < T$. Suppose the investor opens a futures position in the most liquid futures contract, which is typically the nearest or next nearest to delivery contract, and is not within the delivery month[5]. At time $t+1$ the value of the initial dollar allocation $M_{i,t}$, assuming it can be invested in a risk-free money account r^f, and the price change in the futures contract results in a total value of $M_{i,t}(1+r^f_t) + F^T_{i,t+1} - F^T_{i,t}$, leads to a futures return of

$$R_{i,t+1} = \frac{M_{i,t}(1+r^f_t) + F^T_{i,t+1} - F^T_{i,t} - M_{i,t}}{M_{i,t}} = \frac{F^T_{i,t+1} - F^T_{i,t}}{M_{i,t}} - r^f_t. \tag{4.8}$$

Consequently, the return in excess of the risk free rate is given by

$$r_{i,t+1} = \frac{F^T_{i,t+1} - F^T_{i,t}}{M_{i,t}}. \tag{4.9}$$

In defining the return on futures positions we do not use the spot price to close the position, but take the view of an investor, who trades in futures contracts only. If we assume a fully collateralized futures position, we set $M_{i,t}$ equals to $F^T_{i,t}$ and the excess futures return is defined as the scaled price change.

4.2.4 Transaction Costs

For each underlying asset contract positions must be 'rolled' before expiry, if an equal exposure needs to be maintained using futures contracts. Rolling futures contracts requires closing the full

4. $M_{i,t}$ must be at least as big as the margin requirement set by the respective futures exchanges

5. We appropriately adjust for rollovers in order to reflect a tradable return series

futures position in advance of each maturity, whilst simultaneously opening a position in the next nearest maturity contract. For the majority of futures contracts this requires monthly trading of the whole position, suggesting a monthly cost equal to that of the bid-ask spread. However, Garleanu and Pedersen (2009) argue that trading costs from rolling are negligible, due to a separate roll market which entails smaller costs than independently selling the old contract and buying the new one. Consequently, although in implementation the entire portfolio is traded before each maturity, we only consider the *difference* in monthly allocation when calculating trading costs. Analogous to Brandt et al. (2009) we therefore calculate costs based on monthly turnover defined as

$$T_t = \sum_{i=1}^{N_t} |w_{i,t} - w_{i,t-1}|, \tag{4.10}$$

where $w_{i,t}$ refers to the monthly weighting of assets defined above. The optimal portfolio equation is now,

$$\hat{\theta} = \arg\max_{\theta} \sum_{t=0}^{T-1} u \left(r_t^f + \sum_{i=1}^{N_t} \theta^\top \frac{x_{i,t}}{\sigma_{i,t} N_t} r_{t+1}^i - \sum_{i=1}^{N_t} c_{i,t} \theta^\top \left| \frac{x_{i,t}}{\sigma_{i,t} N_t} - \frac{x_{i,t-1}}{\sigma_{i,t-1} N_{t-1}} \right| \right) \tag{4.11}$$

where $c_{i,t}$ is the fractional cost of buying (or selling) asset i at time t, defined as

$$c_{i,t} = \frac{F_{i,t,T}^{\text{Bid}} - F_{i,t,T}^{\text{Ask}}}{2(F_{i,t,T}^{\text{Bid}} + F_{i,t,T}^{\text{Ask}})} \tag{4.12}$$

for futures contracts on the i-th asset at time t with maturity T. The factor of two in Equation 4.12 corresponds to the one-way transaction cost incurred when the weighting is changed for buying and selling, as defined in Equation 4.11. For a given time-series and cross section of transaction costs we can therefore find optimal portfolio characteristics by optimizing the average utility of returns net of transaction costs.

4.3 Data

In this section we describe the price data we use and provide preliminary descriptive statistics. Furthermore, we describe the intuition and construction of the macroeconomic and hedge fund data.

4.3.1 Price Data

The data we use covers a cross section of 55 global futures contracts spanning the asset classes commodities, currencies, bonds, equities and metals during the sample period from January 1980 to January 2012. Summary statistics are reported in Table (4.1). The analysis is carried out on a monthly frequency, although we collect daily data to construct proxies for futures market volatility. We obtain daily settlement prices for the first and second nearest to maturity futures contract from different futures exchanges via Bloomberg. It is important to note that some asset classes have futures contracts expiring every month, whilst others have contract months that are regularly spaced (every 3 months, for example) or irregularly spaced (7 months out of the year). We acknowledge this fact by conducting robustness tests with different specifications of the carry measure. Furthermore, while spot data is available for some futures contracts, we follow the literature, e.g., Moskowitz et al. (2011), and use near-month futures prices instead of spot prices. This is particulary important for commodity futures contracts. Using the near-month futures price as a measure of spot price allows us to focus on the change in futures price over time, unclouded by other differences between futures and cash prices, such as differences in transaction or transportation costs, delivery cost, or delivery location. The near-month futures contract also mitigates problems of low liquidity and irregular price behavior. One weakness in using near-month prices is that due to the irregular spacing of futures expiry, for a given futures contract the near month price may reflect a contract that expires after a few months, rather than within a few days. To remain consistent, we apply this methodology to all asset classes.

We note that our effective sample size varies over time as not all futures contracts are available from the beginning of the sample in January 1980. To illustrate this point we plot the number of futures contracts available through time in Figure 4.1. We distinguish the number of available assets by asset class. As can be inferred from the data in Figure 4.1, our sample starts with 12 contracts including mostly commodities. Currencies are available from April 1986 onwards. The number of contracts increases monotonically and from the early 1990s there are more than 40 contracts available. This number only increases again in 2000 when additional contracts across asset classes are introduced. The time-varying nature of the contract availability will not alter the results significantly as more than 4/5th of the contracts are available for more than 2/3rds of the

Table 4.1: **Summary Statistics for Individual Futures Contracts:** The table presents summary statistics for 55 individual futures contracts across the asset classes bonds, commodities, currencies, equities and metals. From left to right the summary statistics include the instruments' name, the start date of the price series, the annualized mean of the fully collateralized excess return, and its annualized standard deviation. All series end in September 2012, but have different starting dates.

Instrument	Start	Returns		Instrument	Start	Returns	
		Mean	Std			Mean	Std
Bond				**Currency**			
Canadian Bill	Aug-90	0.8	1.3	AUD/USD	Jan-87	4.5	12.1
Canadian 10Yr	Sep-89	4.1	6.4	CAD/USD	Apr-86	2.0	7.5
Euro Bobl	Oct-91	3.1	3.5	CHF/USD	Apr-86	1.4	11.9
Euro Bund	Nov-90	4.1	5.5	EUR/USD	May-98	0.6	10.3
Euro Buxl	Oct-98	4.2	11.1	GBP/USD	May-86	2.2	10.0
Euro Schatz	Mar-97	1.1	1.4	JPY/USD	May-86	0.4	11.2
Japan 10Yr	Oct-85	3.4	5.2	NOK/USD	May-02	4.7	13.7
Long Gilt	Nov-82	2.8	7.6	NZD/USD	May-97	3.6	13.9
US Govt Long	Jan-80	4.8	11.6	SEK/USD	May-02	4.4	13.5
US Govt 10Yr	May-82	5.3	7.0	**Equity**			
US Govt 2Yr	Jun-90	1.8	1.7	ASX SPI 200	May-00	1.8	16.8
US Govt 5Yr	May-88	3.3	4.2	AEX	Jan-89	0.8	25.1
Commodity				CAC 40	Dec-88	2.4	22.9
Cattle, Live	Jan-80	4.2	14.8	DAX Index	Nov-90	4.2	23.6
Cocoa	Jan-80	-5.6	30.7	FTSE 100	Mar-88	2.5	18.7
Coffee, 'C'	Jan-80	-2.6	36.0	FTSE MIB	Mar-04	-3.5	24.8
Corn	Jan-80	-5.9	23.6	IBEX 35	Jul-92	5.2	24.7
Cotton, No. 2	Jan-80	-3.4	24.6	S&P 500	Apr-82	6.0	19.9
Crude, Brent	Jun-88	12.2	33.8	Toronto 60	Sep-99	3.4	21.0
Crude, WTI	Apr-83	6.9	34.4	TOPIX	May-90	-5.8	23.8
Gasoil	Jul-89	12.2	32.3	**Metals**			
Gasoline	Oct-05	5.1	37.8	Copper	Jul-97	8.3	28.0
Heating Oil	Jul-86	8.3	33.8	Gold, 100 oz	Jan-80	-0.7	19.4
Lean Hogs	Apr-86	-0.4	22.8	Nickel	Jul-97	8.0	38.5
Natural Gas	Apr-90	-24.2	48.4	Platinum	Jan-84	3.2	23.3
Soybean	Jan-80	0.7	22.7	Aluminum	Jul-97	-3.1	21.8
Soybean Meal	Jan-80	5.4	24.5	Silver	Jan-80	-4.6	31.6
Soybean Oil	Jan-80	-2.5	23.8	Zinc	Jul-97	-3.4	31.1
Sugar, #11	Jan-80	-5.5	38.6				
Wheat	Jan-80	-7.2	26.4				

sample.

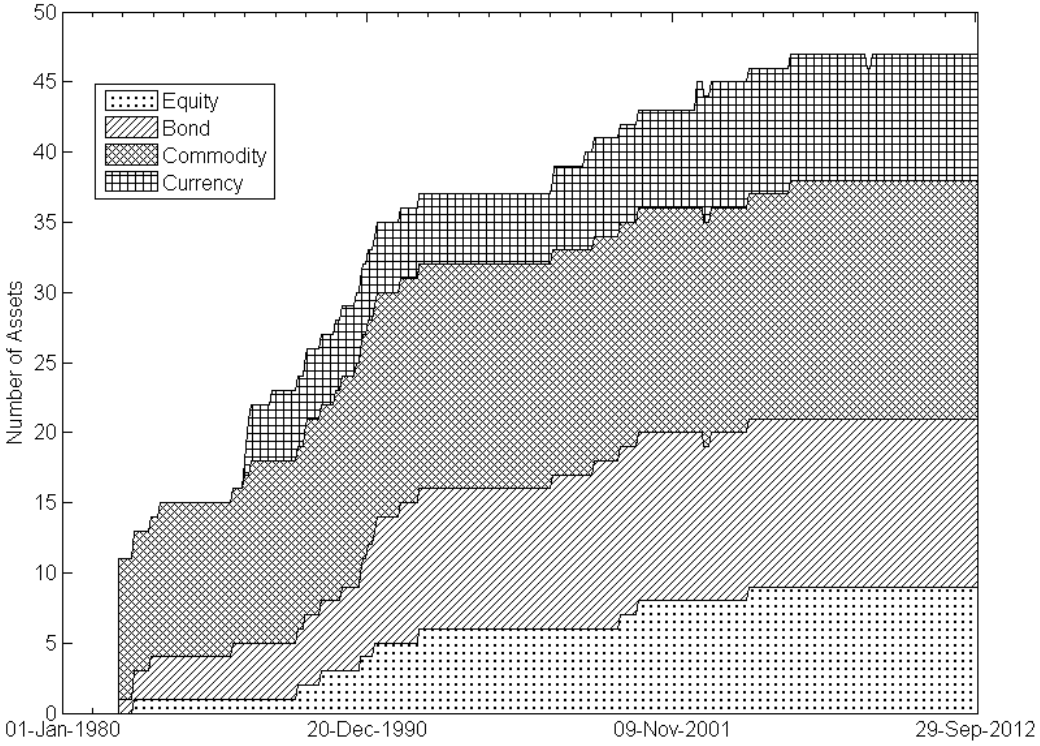


Figure 4.1: **Number of Available Futures Contracts over Time:** This figure shows the cumulative number of futures contracts that are available for the optimal portfolio allocation over the sample period January 1980 to September 2012. The figure also shows the share of available contracts distinguished by the asset classes currency, commodity, bond and equity index futures.

Figure (4.2) reports the cross-sectional means and standard deviations of the carry and momentum signal across time. The carry and momentum signal are time varying and display episodes of high and low volatilities. The cross sectional standard deviation of the carry signal is higher than for the momentum signal. Figure (4.2) also displays the cross sectional correlation of the signals. We can observe that the correlation is mostly positive but highly time varying. The correlation varies mostly between 0 and 0.5. Interestingly the correlation decreased since the early 2000's with a short interruption during the Lehman Brothers crisis. This indicates that including carry and momentum signals in an optimal portfolio portfolio framework might add a substantial diversification benefit.

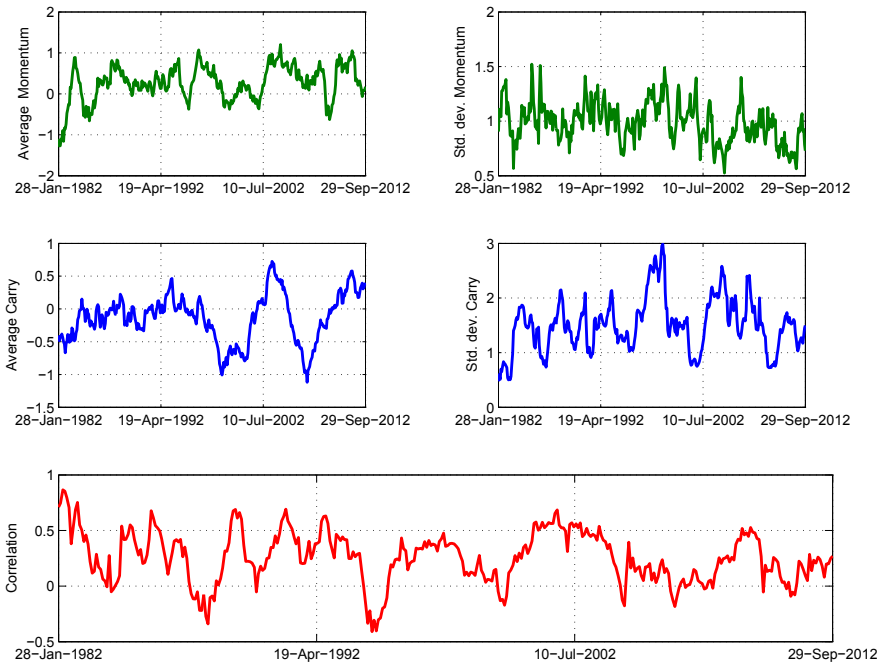


Figure 4.2: **Cross Sectional Summary Statistics of Carry and Momentum Signal:** The figure displays cross-sectional means and standard deviations of the carry and momentum signal every month from January 1982 to September 2012. We also report the cross sectional correlation of carry and momentum signal.

4.3.2 Macroeconomic and Liquidity Data

Macroeconomic Proxies As proxies for the state of the economy we choose three sets of macroeconomic variables. First, we follow Cooper and Priestley (2009) and choose the output gap as an empirical proxy for business cycles. The output gap (IP Gap) is measured as the deviation of the logarithm of total industrial production from a trend that includes both a linear component and a quadratic component. Cooper and Priestley (2009) show that the output gap is a prime business cycle indicator and predicts stock and bond risk premia counter-cyclically, consistent with theoretical predictions.

Second, we use a set of common systematic price-based variables that are related to the business cycle. These variables include the three month T-bill (Yield), or short rate, that is negatively related to future stock market returns and serves as a proxy for expectations of future economic activity. The dividend yield (Div Yield) on the market, defined as the total dividend payments

accruing to the CRSP value-weighted index over the previous 12 months divided by the current level of the index, which has been shown to be associated with slow mean reversion in stock returns across several economic cycles. The term spread (TSpread), measured as the difference between the average yields of 10 year Treasury bonds and 3 month T-bills, is based on the findings of Fama and French (1988), who provide evidence that this measure forecasts short-term business cycles. The final business cycle indicator is the default spread (Def), defined as the difference in yield between the highest and lowest rated investment grade bonds. Specifically, we calculate the default spread as the average yield of bonds rated BAA by Moodys minus bonds rated AAA by Moodys. Fama and French (1988) demonstrate that the spread tracks long term business cycles. The authors find that it is higher in recessions and lower in expansions, and therefore proxies for the unobservable default risk premium embedded in corporate debt.

Finally, we use business cycle indicators from the National Bureau of Economic Research (NBER), who publish a range of ex-post information on business cycles. Notably, they report the dates that are related to an economic expansion, defined from trough to peak, and economic recession, defined from peak to trough. We further subdivide the expansion and recession periods into equal halves and denote them early and late expansions/recessions. These indicators are not used for forecasting regressions as they are only reported ex-post by NBER.

Liquidity Proxies As a state variable for changing liquidity environments we use the TED spread. The TED spread is defined as the difference between the 3 months LIBOR Eurodollar rate and the 3 months T-Bill rate. The LIBOR rate reflects uncollateralized lending in the interbank market, which is subject to default risk, while the T-Bill rate is assumed risk-less since it is guaranteed by the U.S. government. When banks face liquidity problems the TED spread typically increases, as bank funding cost increase whilst the T-Bill yield falls due to a "flight-to-liquidity" or "flight-to-quality".

4.3.3 Hedge Fund Data

We construct proxies for the total capital available and capital flow to the hedge fund industry as measures for arbitrage capital similar to Jylha et al. (2010), Jylhä and Suominen (2011), and Naik et al. (2007). For both measures we use monthly return and total asset under management (AUM) data from a large cross-section of hedge funds and funds-of-funds from January 1994 to December

2012 obtained from Barclay-Hedge databases. We select the subset of funds from this data, that are self reporting global macro funds and commodity trading advisors (CTAs)to construct the aggregate measures.

Total Arbitrage Capital We define total arbitrage capital for each time t as the sum of AUM across N global macro funds and commodity trading advisors. Furthermore we scale the aggregate number by M2 money supply of the US to get a relative measure of hedge fund/arbitrage capital,

$$HFAUM_t = \frac{\sum_{i=1}^{N} AUM_{i,t}}{M2}. \tag{4.13}$$

Total Arbitrage Flows We measure the total arbitrage flow as the AUM-weighted net flow of CTA and global macro funds scaled by lagged AUM. For each fund i, we construct time t net flows as,

$$Flow_{i,t} = AUM_{i,t} - AUM_{i,t-1}(1 + r_{i,t}). \tag{4.14}$$

Then we define the AUM-weighted net flow scaled by lagged AUM as

$$Flow_t = \sum_{i=1}^{N} w_{i,t-1} \frac{Flow_{i,t}}{AUM_{i,t-1}}. \tag{4.15}$$

where $w_{i,t-1} = \frac{AUM_{i,t-1}}{\sum AUM_{i,t-1}}$. Having described the portfolio allocation and the data we use we will continue with the main empirical results.

4.4 Characterising Optimal Futures Portfolios

In this section we present the performance characteristics of optimal futures portfolios based on the carry and momentum signal across all asset classes. Furthermore, we study the role of changing risk aversion on the profitability of the optimal portfolios and study the diversification benefit of optimal futures for classical stock and bond portfolios.

4.4.1 Carry and Momentum Portfolio Returns Across Asset Classes

Table 4.2 presents summary statistics for the parameter estimates, portfolio weights, and returns of four base-case portfolio parametrizations, which are estimated in- and out-of-sample. The in-

sample exercise uses all available data to estimate the parameters of the model. The out-of-sample methodology uses a recursive procedure by first estimating the parameters in an initial configuration period of 120 months and then re-estimating the model every month, using an expanding window of data until the end of the sample is reached. The out-of-sample estimation is less prone to look ahead biases and resembles the parameter estimation of a real-time investor. The first three portfolio parametrizations are formulated by including either the futures contract's carry signal and momentum signal in isolation or a combination of carry and momentum signals, using the portfolio policy function in Equation 4.2. These parametrizations are compared to a long only benchmark portfolio[6]. The table is split into three sections: The first set of rows report parameter estimates and associated student t-statistics, the following set describe the distribution of portfolio weights and the remaining rows report summary statistics of the optimized portfolio returns. The sample period is from January 1982 to September 2012.

The first 4 rows of the first column present the estimates of the in-sample policy parameters, θ, and the corresponding t-statistics for a combined parametric portfolio including both the carry and momentum characteristic. Without imposing any ex-ante relation between signals and future returns, we find that the θ estimates for both characteristics are positive and significant, with a GMM t-stat of 2.65 and 4.29 for the carry and momentum signal, respectively. The positive relation confirms the general findings of the literature: a positive carry signal, that is when the near contract trades above the far contract, predicts positive future returns and contracts that had a positive performance over the previous 12 months (momentum) continue to perform well. It should be stressed again that the relation between futures returns and trading signals is a result of the estimation procedure and was not assumed ex-ante, as most cross-sectional sorts resort to. For comparison, we also report the results of the portfolio parametrization using only individual signals in the second and third column. For each signal we find positive coefficient estimates with highly significant t-statistics (4.65 and 5.39 for the carry and momentum signal, respectively). The fourth column reports the summary statistics of the long only portfolio. This portfolio policy is nested in the parametric function of Equation 4.2 by removing the momentum and carry signal and

6. For the long only portfolio we still estimate the θ parameter but set the characteristics to $x_i = 1$. We therefore estimate the optimal weights for a inverse volatility weighted portfolio of futures contracts, that is a popular investment strategy denoted as risk parity portfolios.

Table 4.2: **Optimal Futures Strategy Portfolio Returns** This table shows estimates of the portfolio policy for individual carry, momentum, and joint carry and momentum characteristics across all available assets. The portfolio policy is defined in Equation 4.7. The reported values are estimated by optimizing a power utility function with a relative risk aversion of five. We use futures prices from Bloomberg from January 1980 to September 2012. The first four columns report in-sample results where optimal characteristics are calibrated once for the entire sample period. The remaining columns report out-of-sample results. Here we use data from January 1982 to January 1992 (120 months) to estimate the initial coefficients and corresponding portfolio weights to generate next month returns. The coefficients are then re-estimated each subsequent month for the remainder of the sample, yielding an out-of-sample return. The first 6 rows refer to the optimal portfolio coefficients and their associated student t-statistic for the carry characteristic, momentum characteristic and long-only strategy, in descending order. The following 4 rows report weighting parameters, the average absolute portfolio weight, the maximum weight, the minimum weight, and the turnover of the portfolio. The remaining columns report the portfolio return statistics, annualized mean return, annualized standard deviation of returns, skewness of returns, excess kurtosis of returns, Sharpe ratio, and certainty equivalent return excess of the risk-free rate.

	In sample				Out of sample			
	C&M	Car.	Mom.	Long	C&M	Car.	Mom.	Long
θ_{carry}	2.036	3.195			1.714	3.073		
(t-stat)	(2.65)	(4.65)			(1.70)	(3.51)		
θ_{mom}	3.157		3.853		3.514		4.197	
(t-stat)	(4.29)		(5.39)		(3.59)		(4.61)	
θ_{long}				2.862				2.572
(t-stat)				(4.13)				(3.00)
$\lvert w_i \rvert$	0.072	0.055	0.062	0.082	0.050	0.038	0.046	0.052
$\max(w_i)$	0.308	0.305	0.239	0.082	0.174	0.171	0.141	0.052
$\min(w_i)$	-0.301	-0.381	-0.197	0.082	-0.158	-0.230	-0.100	0.052
$\sum w_i I(w_i < 0)$	-1.658	-1.686	-1.173	0.000	-1.002	-1.133	-0.752	0.000
$\sum I(w_i < 0)/N_t$	0.427	0.534	0.375	0.000	0.271	0.351	0.241	0.000
$\sum \lvert w_{i,t} - w_{i,t-1} \rvert$	1.119	0.743	1.163	0.000	0.825	0.511	0.865	0.001
$\bar{r}$	0.368	0.202	0.307	0.179	0.364	0.190	0.307	0.171
$\sigma(r)$	0.232	0.166	0.210	0.150	0.239	0.160	0.224	0.131
Skewness	0.105	0.256	0.084	-0.218	0.061	0.364	-0.119	-0.421
Excess kurtosis	1.097	2.097	1.237	0.627	1.059	2.674	0.671	1.114
Sharpe	1.162	0.821	1.056	0.779	1.110	0.791	0.993	0.838
CE	0.151	0.074	0.123	0.064	0.151	0.085	0.119	0.087

introducing a signal of unity across all assets and months. The reported θ_{long} of 2.86 demonstrates that ex-post allocation to an equally weighted portfolio of futures over the sample period is 2.86 times the initial wealth for an investor with power utility of risk aversion five. The t-statistics are based on errors calculated using the covariance estimator of Equation 5.2.

As a consequence of standardizing the portfolio characteristics cross-sectionally, the magnitudes of the coefficients can be compared to each other. In-sample we find that for the optimal futures portfolios that the momentum characteristic is assigned a higher weight than the carry, 3.157 compared to 2.036. Nevertheless, both coefficients are highly significant, and again the comparable magnitudes support the inclusion of both signals in an optimal futures portfolio.

Compared to the momentum-only and carry-only parametrizations, the sum of coefficients for the optimal carry and momentum portfolio is bigger, indicating that a diversification benefit leads the optimal portfolio to have larger net exposure. This feature is reflected in the next few rows of Table 4.2, which characterize the distribution of the portfolio weighting.

For portfolios based on carry and momentum signals, we find that the time series average absolute volatility adjusted weight[7] of futures contracts is 0.072, slightly larger than that for the individual carry and momentum signals of 0.055 and 0.062. The next two rows report the time series average of the maximum and minimum volatility-adjusted contract weighting. The range of weights describes the average maximum exposure the strategy has to any one contract. For the in-sample carry and momentum portfolio we find the maximum and minimum allocation of 0.308 and -0.301, respectively. On first inspection, these values seem high compared to Brandt et al. (2009). However, we have far fewer assets, 55 futures contracts compared to over 5500 stocks, and in the beginning of our sample we have only 13 available assets as can be seen in Figure 4.1. As such, the limits of the optimal portfolio weight is only three times that of the equally weighted portfolio, and therefore we can conclude that the portfolio remains well diversified given the number of available assets.

The next two rows report the time series average of the sum of negative volatility-adjusted weights, and time series average fraction of short positions, respectively. We find that the combined

7. For each weighting statistic we report the product of the actual futures weight and the annualized 3 months absolute mean deviation, $\sigma_i w_i$, to compare the economic significance of the futures weights within the optimal portfolio for contracts over a wide range of volatilities.

signals portfolio has a smaller fraction of short positions compared to that of the individual signals portfolios. However, the magnitude of the difference in negative weights is negligible.

The last weighting statistic is the average monthly turnover for each contract. We use this measure to calculate the transaction costs from modifying contract weighting. We find a turnover of 1.119 for the in-sample optimal carry and momentum portfolio[8]. Although such a turnover would be high for a stock portfolio, this value is less than the inherent turnover of a futures trading strategy, or that required for the long-only futures strategy.

The last 6 rows of Table 4.2 present summary statistics for portfolio returns of the optimal strategies. The four left hand columns demonstrate the in-sample portfolio to have very impressive performance. The leveraged optimal futures portfolio based on the carry and momentum characteristic has an annualized return of 36.8 percent with a Sharpe ratio of 1.162. The last row of Table 4.2 reports the certainty equivalent return (CE), at which an investor would value the returns series, above that of the risk free rate. The CE measure provides a parsimonious perspective, from which to understand the return distribution, since it takes into account the entire distribution of returns and is monotonic with our objective function. As a result the CE penalizes large values in the higher moments of the distribution of returns. The certainty equivalent of the optimal carry and momentum portfolio is 15.1 percent above the risk-free rate. Here, we can also see the benefit of combining both carry and momentum signals to form the optimal portfolio. The carry-only portfolio has a high excess kurtosis of 2.097, which is reflected in a lower CE return of 7.5%. In the combined portfolio this is reduced to 1.097 leading to a 6.7 percent increase in CE. The four right hand side columns of Table 4.2 repeat the summary statistics of the in-sample estimation for the out-of-sample exercise. The first rows of Table 4.2 report the average out-of-sample coefficient estimates. We use the first ten years of data, from January 1982 to September 1992, to estimate the coefficients of the initial portfolio policy. Each month, we then re-estimate the parameters using an expanding window of data until the end of the sample is reached. Figure 4.3 displays the θ estimates and the corresponding confidence bands for the carry and momentum coefficients over time. The coefficients of the characteristics are roughly similar to the in-sample estimates and do not display much time variation. Furthermore, the confidence bands show that the coefficient

8. In results not reported in the table this is found as 29.0% of the mean absolute weight.

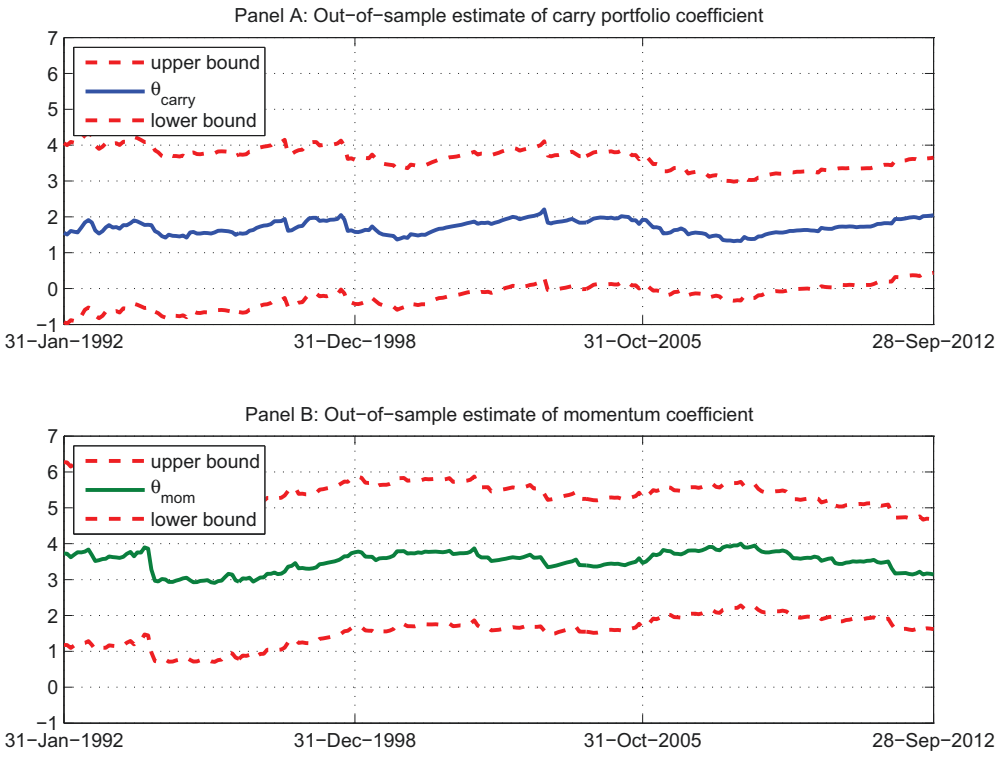


Figure 4.3: **Out of Sample Parameter Estimates over Time (θ estimates):** These figures display the θ estimates for the out of sample parametric portfolio choice problem. The displayed parameters are estimated for all futures assets and both signals in the portfolio optimization of Equation 4.7. Panel A reports the estimates for the carry signal (θ_{carry}), Panel B reports for the momentum signal (θ_{mom}). We also report the upper and lower 95 percentile confidence bands for each parameter. The standard errors are estimated form the asymptotic covariance matrix within the framework of general methods of moment (GMM) by Hansen (1982). The model is calibrated for the in sample period from January 1980 to December 1990 and then the θ's re-estimated every month, using an expanding window of data until September 2012.

estimates are statistically significant, most of the time, for both characteristics. In the contrast to the in-sample results, we find a slight widening of the difference between the carry and momentum coefficients, with the ratio of momentum-to-carry coefficient increasing from 1.6 to 2.1, between in- and out-of-sample estimates, respectively.

The strong performance of the in-sample optimized portfolio is to some extent unsurprising, since the coefficients are chosen ex-post to maximize the CE return. What is surprising is the similarity to the out-of-sample results. For all portfolio configurations, we report a smaller than 1% difference in CE, when returns are calculated based on out-of-sample portfolio characteristics. Figure 4.4 Panel A displays the cumulative returns of the optimal out-of-sample carry and

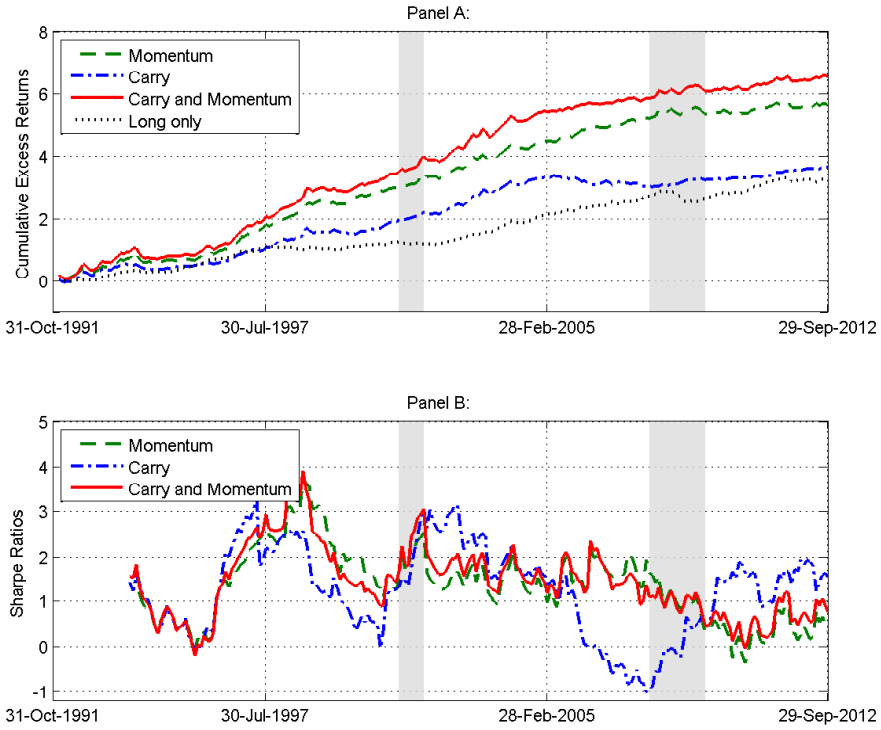


Figure 4.4: **Historical Performance of Carry, Momentum and Carry and Momentum:** These plots present the cumulative logarithmic returns and rolling Sharpe ratios for the out of sample estimation from January 1990 to September 2012. The portfolios are optimized for a power utility function with a relative risk aversion of $\gamma = 5$. Panel A presents the cumulative returns for the carry, momentum, and the combined carry and momentum portfolio. For comparison Panel A includes the cumulative return series of an optimized long only portfolio. Panel B presents the 24 month rolling Sharpe ratios of the return series for carry, momentum, and the combined carry and momentum portfolio. The grey bands in both panels indicate the NBER recession periods (business cycle peak to trough).

momentum strategies over time and compares it to the optimal carry only, momentum only and long only portfolios. We find that the optimal carry and momentum portfolio has a significantly higher performance than the individual carry and momentum portfolios. This again indicates a diversification benefit of including both characteristics. Furthermore, all strategies have a better performance than the long only strategy which is consistent with their statistically significant θ estimates.

Comparing the distribution of adjusted portfolio weights between in- and out-of-sample results, we find out-of-sample metrics to be generally more conservative. For instance, for combined signals

the monthly mean absolute adjusted weight is 0.05 out of sample compared to 0.07 in sample. The difference is stark when we consider the extremes of the distribution. Going from in- to out-of-sample, we find the maximum and minimum weights halved from 0.308 to 0.174, and from -0.301 to -0.158, respectively. These findings give us confidence that the strong performance of the out-of-sample portfolio is not dependent on a few lucky trades. Although the time series average of portfolio turnover is smaller, we find that the turnover as percentage of total absolute portfolio weight remains nearly unchanged at 32%. We report analogous results when the in- and out-of-sample carry-only and momentum-only portfolios are compared.

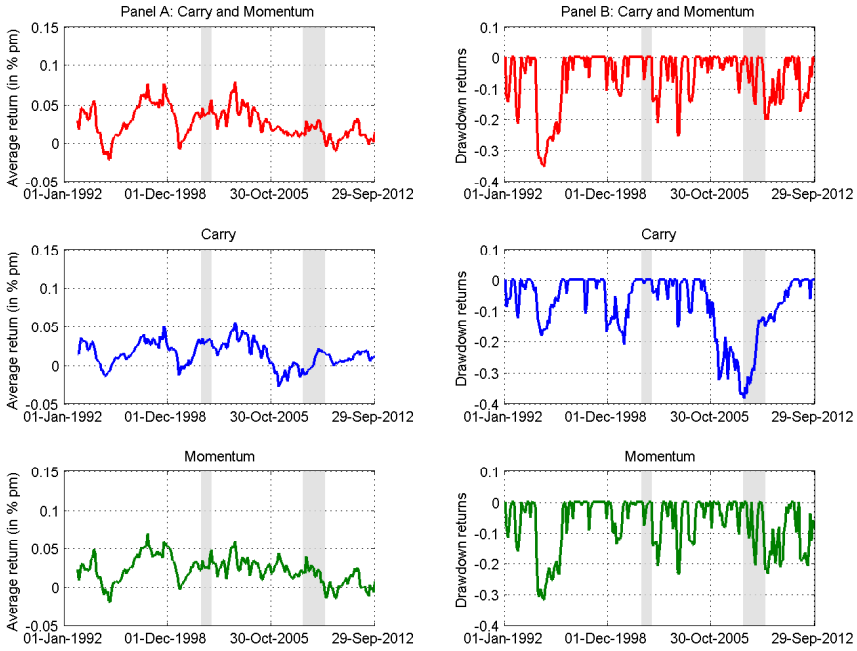


Figure 4.5: **Rolling Average Returns and Draw-down Dynamics:** These figures show average monthly excess returns over rolling windows of 12 months and the maximum draw-down dynamics for the optimal portfolios. The portfolios are optimized out of sample for the estimation from January 1990 to September 2012 and a power utility function with a relative risk aversion of $\gamma = 5$. Panel A reports the rolling averages for momentum and carry as well as for carry and momentum separately. Panel B reports the draw-downs for the same set of trading strategies and draw-down is defined as: $D_t := \sum_{s=1}^{t} r_s - \max_{u \epsilon \{1,\ldots,t\}} \sum_{s=1}^{u} r_s$ and r_t are the returns to the trading strategies.

Finally, we study the dynamic behaviour of the out-of-sample carry and momentum returns from 1992 to 2012. Figure 4.4 Panel B presents the 24 months rolling Sharpe ratio for carry and momentum strategies. The Sharpe ratio is highly time-varying, even though it is positive for the

entire period, a negative trend is observable. While the Sharpe ratio varied around 2 and 3 in the period from 1997 to 2005 it deteriorated to a level of around 1 or lower since then. The same can be observed for the rolling 12 months average return of the carry and momentum strategy in Figure 4.5. While the returns are highly time varying in the early sample, they deteriorate towards the end. Lastly, Figure 4.5 reports the draw-down dynamics of the strategy. While a clear trend is not observable, we consider two sample events, and their effect on the optimal portfolio. Looking at the optimal momentum portfolio, we find the largest draw-down from December 1993 to August 1994, losing 31.3 percent. In Figure 4.3, we see the policy response to this event: we observe a small but sharp decrease of the momentum coefficient around the same period. The second event we consider is the dramatic draw-down of the carry strategy in the run-up to the sub-prime crisis of 2008, where it suffered a loss of 38.1 percent between April 2005 and December 2007. However, for the carry event we observe a different portfolio response. The carry event estimation sample was nearly double the size of the momentum event. As a consequence of this and its occurrence towards the end of the sample, we observe a smaller magnitude, and more gradual, reduction in theta estimate compared to the earlier momentum event.

4.4.2 Effects of Changing Risk Aversion

The optimal portfolio policy depends critically on the preferences of the investor. Throughout this paper we optimize a portfolio for an investor with a CRRA utility function with relative risk aversion of $\gamma = 5$. In this section, to give a better sense of the role of the utility function in the portfolio allocation, we present the in- and out-of-sample results for varying degrees of risk aversion. In addition to the base-case coefficient $\gamma = 5$, we report the estimation results for a CRRA risk aversion coefficient of $\gamma = 2$, corresponding to an investor with a low risk aversion, and a risk aversion of $\gamma = 100$, corresponding to an investor who is very sensitive to even small loses.

The results for varying relative risk aversion are reported in Table 4.3. Panel A of Table 4.3 reports statistics for the optimal carry and momentum portfolio, Panel B reports analogous results for an optimal allocation to an equal volatility-weighted long-only portfolio, which serves as a benchmark for comparison.

Looking first at the optimal portfolio of momentum and carry characteristic, for small values of γ we find that the coefficients of both signals are large. As we increase risk aversion a higher γ

Table 4.3: **Varying Risk Aversion** Reported are estimates of futures portfolio policy with the characteristic carry, momentum, carry and momentum combined, and long-only optimized for different power utility functions with a relative risk aversion of one, five (used in previous tables), and 100. Policy estimates and associated return statistics are reported for in-sample and out-of-sample portfolios. For the in-sample section, we estimate an optimal portfolio policy based on data from January 1982 to September 2012. Out-of-sample optimisation is calibrated over the period from January 1982 to January 1992. We then form out-of-sample monthly portfolios by enlarging the sample period and re-estimating the portfolio policy for each subsequent month until the end of the sample. In the upper panel, the first 4 rows refer to the optimal portfolio characteristics and their associated bootstrapped student t-statistic for the carry, and momentum, in descending order. The following 4 rows report weighting parameters, the average absolute portfolio weight, the maximum weight, the minimum weight, and the turnover of the portfolio. In these measures, each weighting parameter is multiplied by its respective futures volatility to ensure comparability. The remaining columns report portfolio the return statistics, annualized mean return, annualized standard deviation of returns, skewness of returns, excess kurtosis of returns, Sharpe ratio, and certainty equivalent return excess of the risk-free rate. The lower panel reports analogous statistics for the long-only portfolio.

	In sample			Out of sample		
Variable	$\gamma = 2$	$\gamma = 5$	$\gamma = 100$	$\gamma = 2$	$\gamma = 5$	$\gamma = 100$
Panel A: Carry and Momentum						
θ_{carry}	3.542	2.043	0.106	4.855	1.713	0.083
(t-stat)	(3.63)	(2.66)	(0.52)	(2.25)	(1.70)	(1.73)
θ_{mom}	4.528	3.147	0.154	9.793	3.515	0.172
(t-stat)	(4.27)	(4.28)	(0.80)	(4.77)	(3.59)	(3.58)
$\lvert w_i \rvert$	0.111	0.071	0.004	0.139	0.050	0.002
$\max(w_i)$	0.491	0.308	0.016	0.488	0.174	0.009
$\min(w_i)$	-0.493	-0.301	-0.015	-0.445	-0.158	-0.008
$\sum w_i I(w_i < 0)$	-2.649	-1.658	-0.084	-2.814	-1.002	-0.049
$\sum I(w_i < 0)/N_t$	0.436	0.427	0.430	0.271	0.271	0.270
$\sum \lvert w_{i,t} - w_{i,t-1} \rvert$	1.689	1.117	0.055	2.312	0.825	0.041
$\bar{r}$	0.572	0.368	0.064	1.194	0.364	0.048
$\sigma(r)$	0.356	0.231	0.015	0.670	0.240	0.013
Skewness	0.120	0.105	0.264	0.052	0.065	-0.139
Excess kurtosis	1.084	1.097	1.413	1.024	1.068	0.632
Sharpe	1.161	1.162	0.892	1.141	1.110	-0.118
CE	0.371	0.151	0.003	0.481	0.151	0.005
Panel B: Long only						
θ_{long}	8.098	2.862	0.162	7.351	2.572	0.156
(t-stat)	(5.37)	(4.13)	(0.16)	(3.92)	(3.00)	(3.63)
$\lvert w_i \rvert$	0.233	0.082	0.005	0.149	0.052	0.003
$\max(w_i)$	0.233	0.082	0.005	0.149	0.052	0.003
$\min(w_i)$	0.233	0.082	0.005	0.149	0.052	0.003
$\sum w_i I(w_i < 0)$	0.000	0.000	0.000	0.000	0.000	0.000
$\sum I(w_i < 0)/N_t$	0.000	0.000	0.000	0.000	0.000	0.000
$\sum \lvert w_{i,t} - w_{i,t-1} \rvert$	0.001	0.000	0.000	0.002	0.001	0.000
$\bar{r}$	0.453	0.179	0.057	0.469	0.171	0.042
$\sigma(r)$	0.426	0.150	0.011	0.374	0.131	0.009
Skewness	-0.229	-0.218	-0.056	-0.380	-0.421	-0.598
Excess kurtosis	0.573	0.627	0.889	0.992	1.114	1.223
Sharpe	0.777	0.779	0.568	0.916	0.838	-0.834
CE	0.195	0.064	0.000	0.269	0.087	0.003

gives lower portfolio coefficients and therefore allocations. These observations demonstrate that the allocation is related to both risk and return, since for larger levels of risk aversion the utility cost of risk subsumes benefits associated returns until the allocation, and thus the portfolio coefficients, are reduced.

Note that the tilt towards the momentum signal is maintained with the ratio of momentum-to-carry coefficients of around 1.5 and 2.0 for in-sample and out-of-sample, respectively, across the three levels of risk aversion. The portfolio weighting statistics paint the same picture: The weights are scaled down equally with increasing levels of risk aversion. This effect is seen for both time series averages of the mean absolute weight, minimum weight, maximum weight and the portfolio turnover metric. Accordingly, the time series average faction of short positions is nearly unchanged over the different γ levels maintaining approximately 0.43 and 0.270 for in- and out-of-sample, respectively. We therefore conclude that the relative importance of the carry and momentum characteristics in the investors optimal portfolio remains independent of varying degrees of risk aversion.

The last 6 rows of Table 4.3 Panel A report performance statistics. As we would expect from its higher characteristic coefficients the $\gamma = 2$ case generates very high returns. With an annualized return of 57 percent in sample and 119 percent out of sample. These returns are accompanied by similarly high annualized standard deviations of 36 percent and 67 percent. With increasing γ, the coefficients, the portfolio weights, and therefore portfolio return and volatility, are reduced towards zero. Indeed, with $\gamma = 100$ annualized return is reduced to 6.4 percent in sample and 4.8 percent out of sample. The associated annualized standard deviations of returns are less than 1.5 percent for both sampling methodologies.

The last row of Table 4.3 Panel A reports the CE return in excess of the risk-free rate. The CE cannot be compared across levels of risk aversion. However, note that for each level of risk aversion the CE is improved for the out-of-sample methodology.

4.4.3 Diversified Investor with Optimal Futures Portfolio

Given the recent trend of including dynamic futures portfolios in diversified portfolios of institutional investors, we next investigate the diversification benefit of including optimal futures portfolios based on the carry and momentum characteristics in such allocations. Explicitly, we test the hypothesis that a diversified investor, who already has an optimal allocation to popular tradable risk

factors or asset classes, would benefit out of sample from the inclusion of an optimal carry and momentum strategy in his portfolio. This relatively unexplored topic of including dynamic trading strategies in traditional asset allocation mixes, notable exceptions are Barroso and Santa-Clara (2013) and Burnside et al. (2006), is not only important from an investors perspective, but will give us an indication of the economic value of higher moments of the return distribution with respect to the standard risk factors that cannot be observed in the linear regression framework.

To do so, we modify our return on wealth, $r_{p,t+1}$, in Equation 4.6 to incorporate classes of risk factor returns containing M tradable assets. The return on the wealth of the investor is now given by

$$r_{p,t+1} = rf_t^{US} + \sum_{j=1}^{M} w_j F_{j,t+1} + \sum_{i=1}^{N_t} w_{i,t} r_{i,t+1}, \tag{4.16}$$

where w_j is the portfolio weight on the tradable factor $F_{j,t+1}$, expressed as excess returns, and $w_{i,t}$ depends on the carry and momentum signals through θ characteristics, as previously described. We estimate the optimal portfolio by maximizing the CRRA utility of returns from Equation 4.16 for θ and w_j, with a $\gamma = 5$.

Table 4.4 reports comparative return statistics for out-of-sample optimal allocations to four popular risk factor classes: a passive index investment in the S&P 500, Fama and French factors taken from Fama and French (1992), equity value and momentum factors of Asness et al. (2013), and a portfolio containing the S&P 500 and the 10 year US government treasury bond. To provide a benchmark case, for each risk factor class, we first consider performance of an out-of-sample allocation to just the tradable risk factor(s) in that class, effectively forcing θ coefficients to be zero.

The first and second columns of Table 4.4 report return statistics for our optimal carry and momentum strategy base case, and the optimal carry and momentum with long-only, respectively. The remaining columns are in pairs, and report performance statistics of each class of risk factors for the benchmark case, on the left, and diversified portfolio case, on the right.

The opportunity to invest in futures is clearly of significant value to investors. Figure 4.6 displays performance metrics for each risk factor class, without and then with the addition of the carry and momentum optimal futures strategy. The inclusion of futures in the optimal portfolio enhances performance by nearly all metrics. Amongst the benchmark portfolios the highest out-

Table 4.4: **Diversification Benefit of the Optimal Futures Portfolio** Reported are out-of-sample performance statistics for optimal portfolios containing popular classes of tradable risk factors. For each class of risk factors we compare out-of-sample portfolio performance with and without the inclusion of the carry and momentum futures portfolio. Optimal portfolios are calculated by incorporating traditional strategy performance into the portfolio return as described in Equation 4.16. The first column reports performance of our base case optimal carry and momentum strategy. The second column reports the same strategy with the inclusion of an equally weighted long only futures return. The remaining columns display statistics for the four strategy classes we consider; a passive index strategy represented by the excess return of the S&P 500 index, the Fama-French factors (Market Excess Return, High minus Low, Big minus Small) taken from Eugene Famas website, Equity value and momentum factors taken from Tobias Moskowitzs website, and a portfolio consisting of S&P 500 excess return and the excess return on the 10 year treasury bill. For each class we report a 'benchmark' case, where we force all futures weights to zero, so that the reported performance is the optimal allocation for that risk factor class in isolation. This is reported in the left hand column of each class. On the right hand column of each class we include the optimal carry and momentum strategy. For each strategy class out-of-sample optimisation is calibrated over the period from January 1982 to January 1992. The reported performance is then calculated using optimised weights and characteristics for an expanding window for each subsequent month. Performance statistics are in rows, reporting in descending order: maximum monthly return, minimum monthly return, annualised mean return, annualised standard deviation, monthly excess skewness, monthly kurtosis, annualised Sharpe ratio and the certainty equivalent return in excess of the risk free rate.

	Optimal Futures		Risk Factor Classes							
	Long only		Stocks		Fama & French		Value & Mom.		Stocks & Bonds	
	w/o	w	w/o	w	w/o	w	w/o	w	w/o	w
$\max(r)$	0.276	0.285	0.075	0.267	0.206	0.295	0.573	0.314	0.165	0.259
$\min(r)$	-0.188	-0.234	-0.129	-0.194	-0.162	-0.210	-0.306	-0.193	-0.117	-0.205
$\bar{r}$	0.364	0.529	0.086	0.406	0.162	0.462	0.258	0.420	0.094	0.403
$\sigma(r)$	0.239	0.301	0.106	0.246	0.204	0.265	0.253	0.265	0.132	0.241
Skewness	1.059	0.928	1.323	1.242	0.800	1.402	16.538	1.037	1.707	1.437
Ex. Kurtosis	0.060	-0.026	-0.568	0.086	0.154	0.106	1.829	0.092	-0.025	0.121
Sharpe	1.110	1.276	0.318	1.209	0.502	1.275	0.724	1.163	0.310	1.224
CE	0.151	0.186	0.021	0.179	0.016	0.198	0.049	0.164	0.013	0.184

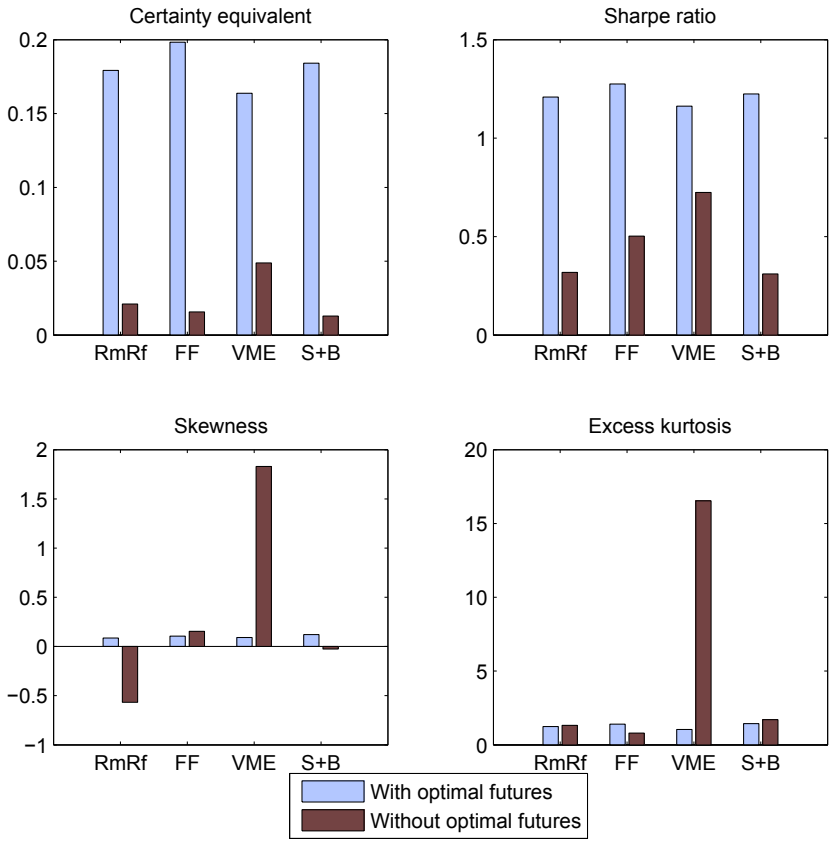


Figure 4.6: **Diversified Investor** Reported are out-of-sample performance statistics for optimal portfolios containing popular classes of tradable risk factors. For each class of risk factors we compare out-of-sample portfolio performance with and without the inclusion of the carry and momentum futures portfolio. Optimal portfolios are calculated by incorporating traditional risk factor performance into the portfolio return as described in Equation 4.16. The figure contains four plots of performance measures, clockwise from the upper left depicting, certainty equivalent, Sharpe ratio, skewness and excess kurtosis. Each plot displays statistics for the four strategy classes we consider; a passive index strategy represented by the excess return of the S&P 500 index, denoted "RmRf", the Fama-French factors (Market Excess Return, High minus Low, Big minus Small),"FF", taken from Eugene Famas website, Equity value and momentum factors,"VME", taken from Tobias Moskowitzs website, and a portfolio consisting of S&P 500 excess return and the excess return on the 10 year treasury bill, "S+B". For each class we report a benchmark' case, where we force all futures weights to zero, so that the reported performance is the optimal allocation for that risk factor class in isolation. This is displayed by the red bars in the plots, blue bars show performance statistics when we include the optimal carry and momentum strategy.

of-sample certainty equivalent was the equity value and momentum factors of Asness et al. (2013) with 4.9 percent above the risk free rate. In contrast, when futures are included in the diversified portfolios even the lowest CE, the diversified value and momentum portfolio, is 16.4 percent.

Looking across Table 4.4 we can see the CE gains of including futures, diminishes those of

popular risk factors over the stock market significantly. The highest certainty equivalent gain amongst our benchmark factor portfolios is the one based on the value and momentum factors of Asness et al. (2013), corresponding to 2.7 percent gain over the stock market. Interestingly, we obverse that, although, the Fama and French factors have a higher-than-market Sharpe ratio, the inclusion of the HmL and SmB factors actually decreases CE, due to the higher kurtosis of these factors. We also find that the addition of bonds decreases the CE compared to only the stock market.

In summary we, find that including optimal carry and momentum portfolios into the asset allocation of a diversified investor is optimal, both from the perspective of increasing Sharpe ratios and decreasing higher moments, such as kurtosis and skewness.

4.5 Business Cycles, Limits to Arbitrage and Hedge Fund Capital

The impressive performance of optimal futures portfolios across asset classes and the significant diversification benefit of including optimal futures portfolios in traditional portfolio allocations presents yet another challenge to existing financial economic theories. The performance of the optimal portfolios is driven to a large extent by the significant predictable variation in excess returns of futures contracts conditioned on the price-based asset characteristics carry and momentum. Optimal futures portfolios improve the Sharpe ratio of the long-only benchmark, which is by itself a dynamic strategy due to the volatility adjustment, by almost 0.5. Furthermore, the diversification benefit of including optimal carry and momentum portfolios within a stock and bond portfolio are highly desirable from an investors perspective. They increase the Sharpe ratios of the total portfolios on average by 0.6 and reduce skewness and kurtosis. In this section we contribute to the understanding of dynamic trading strategies by exploring the role of transaction costs, traditional risk factors, business cycle, liquidity risk and the role of institutional investment, in particular hedge fund activity, on the profitability of optimal futures portfolios based on carry and momentum characteristics.

4.5.1 Optimal Futures Returns with Transaction Costs

In this first section we will explore the role of transaction costs on the profitability of optimal futures portfolios. Transaction costs have been shown to significantly reduce (Korajczyk and Sadka (2004)) or even remove (Lesmond et al. (2004)) abnormal profits. For carry strategies, Burnside et al. (2006) find the inclusion of costs substantially reduces excess returns, although they remain significant. Therefore, it is a natural starting point to introduce transaction cost adjusted portfolio returns. Futures markets have low transaction costs compared to equity markets and currency forwards, ex-ante we therefore expect their effect to be small. Furthermore, due to a separate roll market (Garleanu and Pedersen (2009)) transaction costs are only apposite when the portfolio is modified. Thus, with a typical one-way trading cost of 0.05 percent (Garleanu and Pedersen (2009)) combined with monthly turnover of approximately 25 percent (see Table 4.2), annual costs would be expected of the order 0.1 percent, small compared to strategy returns in excess of 35 percent.

Table 4.5.1 reports statistics for optimal portfolios where the portfolio policy is determined by optimizing the utility of net returns, given in equation 4.11, for two different cost specifications, alongside the zero cost case for comparison. The first specification considers a conservative estimate for one-way transaction costs of 0.05 percent, constant across time series and cross-section. The second specification measures transaction costs directly from the futures bid-ask spread of monthly close prices.

The difference in optimal portfolio policy with transaction costs is negligible, even with the introduction of conservative costs. For the in-sample methodology, the addition of a flat cost across all contracts in the utility maximisation leads to a minor increase in the magnitude of the carry characteristic, from 2.036 to 2.090, and a slight decrease in the momentum characteristic, from 3.137 to 3.135. Due to the similarity in portfolio characteristics the futures weighting dynamics, reported in rows five to ten of Table 4.5.1, are virtually identical between the three cost cases.

The last seven rows of the table report return statistics. As conjectured, the return of the flat transaction cost optimal portfolio is reduced the most, albeit still small compared to overall return, with a reduction of 1.6 percent from 36.8 to 35.2 percent. The different portfolio policy leads the bid-ask cost specification to have a slight increase in returns of 0.4 percent over the zero cost base case. The consistency of the last three rows demonstrates changes in certainty equivalent due to

Table 4.5: **Optimal Futures Returns with Transaction Costs** This table shows estimates of portfolio policy for joint carry and momentum characteristics across all available assets for the utility function specified in Equation 4.7. Average utility of return is maximized for a gamma five power utility function for two cost specifications. In the first specification transaction costs are a constant 0.1 percent across assets and time. In the second, costs are based on the monthly close bid-ask spread for the corresponding asset. For comparison, statistics for zero transaction costs are also reported. The first three columns report "in-sample results where optimal characteristics are calibrated once for the entire sample period. The remaining columns report "out-of-sample results, here we use data from January 1982 to January 1992 to estimate the initial coefficients and corresponding portfolio weights to generate next month return. The coefficients are then re-estimated each subsequent month for the reminder of the sample, yielding an out of sample return. The first 4 rows refer to the optimal portfolio coefficients and their associated student t-statistic for the carry characteristic and momentum characteristic, in descending order. The following 6 rows report weighting parameters, the average absolute portfolio weight, the maximum weight, the minimum weight, and the turnover of the portfolio. Each weighting parameter is multiplied by volatility to ensure comparability. The remaining rows report the portfolio return statistics, annualized mean return, annualized standard deviation of returns, skewness of returns, excess kurtosis of returns, Sharpe ratio, and certainty equivalent return excess of the risk free rate. The certainty equivalent return is reported for returns net of costs for each of the three cost specifications.

	In sample			Out of sample		
	$c_{i,t} =$			$c_{i,t} =$		
	0.000	0.001	$f(me_{i,t}, t)$	0.000	0.001	$f(me_{i,t}, t)$
θ_{carry}	2.036	2.090	2.074	1.714	1.714	1.714
(t-stat)	(2.65)	(2.75)	(2.75)	(1.70)	(1.68)	(1.68)
θ_{mom}	3.157	3.135	3.297	3.514	3.514	3.514
(t-stat)	(4.29)	(4.21)	(4.46)	(3.59)	(3.54)	(3.54)
$\lvert w_i \rvert$	0.072	0.072	0.074	0.050	0.050	0.050
$\max(w_i)$	0.308	0.311	0.317	0.174	0.174	0.174
$\min(w_i)$	-0.301	-0.305	-0.309	-0.158	-0.158	-0.158
$\sum w_i I(w_i < 0)$	-1.658	-1.675	-1.709	-1.002	-1.002	-1.002
$\sum I(w_i < 0)/N_t$	0.427	0.428	0.426	0.271	0.271	0.271
$\sum \lvert w_{i,t} - w_{i,t-1} \rvert$	1.119	1.120	1.161	0.825	0.825	0.825
$\bar{r}$	0.368	0.352	0.372	0.364	0.346	0.358
$\sigma(r)$	0.232	0.231	0.238	0.239	0.238	0.238
Skewness	0.105	0.103	0.099	0.061	0.033	0.037
Excess kurtosis	1.097	1.098	1.093	1.059	1.001	1.008
Sharpe	1.162	1.115	1.145	1.110	1.058	1.101
CE$\lvert c_{i,t} = 0.000$	0.152	0.152	0.152	0.084	0.084	0.084
CE$\lvert c_{i,t} = 0.001$	0.138	0.138	0.138	0.072	0.072	0.072
CE$\lvert c_{i,t} = f(me_{i,t}, t)$	0.146	0.146	0.146	0.080	0.080	0.080

transaction costs are only related to the reduction in returns, ceteris paribus, and not the impact of costs over the optimal portfolio policy. Taking costs into account, we find a 1.4 and 0.6 percent reduction in certainty equivalent net return for the flat cost and bid-ask specification, respectively, compared to the no cost case. Columns four to seven of Table 4.5.1 show differences in portfolios and returns are further reduced for the out-of-sample methodology.

In summary transaction cost play only a minor role in explaining the performance of optimal futures portfolios.

4.5.2 Risk Factor Exposure

To better understand the sources of the risk-return profile, we next evaluate the risk-adjusted performance of optimal carry and momentum portfolios within the framework of standard, and non-standard, linear risk-factor models, typically used in the asset pricing literature. Table 4.6 reports results from regressing individual optimal carry and momentum returns, as well as a combination of the two, on various model specifications.

The first three columns of Table 4.6 report the results from regressing the optimal strategy returns on factors from the Value and Momentum Everywhere model by Asness et al. (2013). The value and momentum factors are diversified, volatility weighted portfolios of cross-sectional value and momentum strategies in global equities, equity indices, bonds, commodities, and currencies. We add the optimal long-only portfolio as a benchmark market factor to the model. In all three cases, the strategies deliver a large and highly significant alpha, or intercept, with respect to the value and momentum factors of about 19, 14 and 21 percent for optimal carry, momentum and combined strategies, respectively. Both, the value and momentum factors are statistically significant and positively related to all optimal strategies. The factor exposure of the long-only benchmark shows some cross-sectional differences. While it is negatively and insignificantly related to carry strategies it is highly significantly and positively related to the optimal momentum strategy. This high positive significance carries over to the combination of optimal carry and momentum strategies. These findings confirm the results by Moskowitz et al. (2011), who show that timeseries momentum strategies are related to cross-sectional strategies, although, it is in contrast to Koijen et al. (20012), who do not find significant exposure of carry strategies with respect to value and momentum factors.

The next three columns in Table 4.6 present results from regressing the optimal strategy re-

turns on Fama and French (1992)'s market, high-minus-low and small-minus-big factors. With the exception of the momentum strategy, we find that all factors have a significant negative relation with optimal carry, momentum, and carry and momentum returns. The negative factor exposure explains the highly significant alphas of 21, 21 and 37 percent.

Table 4.6: **Risk Factors and Optimal Futures Portfolios:** This table reports the regression coefficients (alpha is % annualized) and the respective Newey and West (1987) adjusted t-statistics from regressing the returns to optimal out-of-sample momentum, carry and combined carry and momentum (Car&Mom) strategies on four factor model specifications: (a) A version of Asness et al. (2013) "Value and Momentum everywhere" model that is a volatility weighted cross sectional momentum portfolio across asset classes (Mom. Every.) and a volatility weighted cross sectional value portfolio across asset classes (Value Every.) which we augment with the volatility weighted long only portfolio as a market proxy, (b) a version of the Fama and French (1992) model that uses the high minus low (FF HML) and the small minus big (FF SMB) portfolios as factors and includes the MSCI index as a market proxy, (c) a liquidity risk model that includes the Pastor and Stambaugh (2002) market liquidity proxy (Market Liq.) and (d) a combination of (a) and (c). The last row reports the adjusted R^2's of the regressions. The sample is from January 1992 to September 2012.

	Carry	Mom	Car& Mom	Carry	Mom	Car&Mom	Carry	Mom	Car&Mom	Carry	Mom	Car&Mom
Alpha	19.272	14.412	21.698	21.353	31.472	37.012	20.016	29.965	34.946	19.273	14.412	21.699
tstat	(4.89)	(3.14)	(4.12)	(6.09)	(6.64)	(7.17)	(5.40)	(6.24)	(6.66)	(4.89)	(3.14)	(4.12)
Long	-0.169	1.001	0.729							-0.151	1.013	0.748
tstat	(-1.34)	(6.80)	(4.32)							(-1.19)	(6.82)	(4.40)
Value Every.	0.336	0.269	0.410							0.335	0.268	0.409
tstat	(3.96)	(2.73)	(3.62)							(3.94)	(2.71)	(3.60)
Mom. Every.	0.259	0.468	0.535							0.252	0.463	0.528
tstat	(3.67)	(5.71)	(5.69)							(3.56)	(5.62)	(5.59)
MSCI				-0.281	-0.006	-0.167						
tstat				(-4.51)	(-0.07)	(-1.82)						
FF HML				0.002	-0.003	-0.002						
tstat				(2.07)	(-3.04)	(-1.57)						
FFSMB				-0.002	-0.003	-0.004						
tstat				(-2.42)	(-2.21)	(-2.68)						
Market Liq							-0.071	-0.023	-0.057	-0.151	1.013	0.748
tstat							(-1.46)	(-0.37)	(-0.83)	(-1.19)	(6.82)	(4.40)
R^2	0.061	0.238	0.160	0.118	0.035	0.160	0.004	-0.003	-0.001	0.062	0.236	0.160

The fourth set of regressions in Table 4.6 reports the exposure of optimal carry, momentum and carry and momentum returns to the market liquidity risk factor by Pastor and Stambaugh (2002). In all three strategy cases the beta coefficient estimates remain statistically insignificant. In the last three columns of Table 4.6 we repeat the the exercise but include the Value and Momentum Everywhere factors. While the coefficient estimates of the Asness et al. (2013) factors remain almost unchanged from the first regression exercise, we find that the market liquidity risk factor is now highly positive and significantly related to optimal momentum and optimal carry and momentum returns. The coefficient for optimal carry strategies remains unchanged.

Even though we find some exposure of optimal carry and momentum returns to standard and non-standard risk factors, a large proportion of the returns remains unexplained, as measured by alpha. In particular, the negative correlation of optimal carry and momentum returns with market factors is difficult to reconcile with standard risk-based explanations. However, these findings may be masked by significant dynamic exposure of these factors driven by changing economic environments that are related to time-varying beta estimates or market frictions. We explore this hypothesis in the next two sections.

4.5.3 Macroeconomic Influences and Funding Liquidity

Even though, the literature in financial economics has provided evidence of time variation in risk premia due to changing business cycle and liquidity environments in a variety of asset classes, it is not clear ex ante if these sources of time-varying risk also explain, or are related to, the returns of optimal carry and momentum strategies in futures markets. Equilibrium models based on the assumption that all claims, including futures contracts, are costlessly marketable, e.g., Richard and Sundaresan (1981), indicate that risk premia in both futures and other asset markets are sensitive to the same sources of economic uncertainty. This similarity implies that state variables that possess forecast power in stock and bond markets should also forecast expected futures returns for other assets. Imposing macroeconomic theories of discount rate variation would therefore imply that the returns to optimal carry and momentum strategies are counter-cyclically related to macroeconomic state variables. In this section we test the theoretical conjectures on the effect of macroeconomic risk and liquidity risk on optimal carry and momentum portfolios are reflected in the data.

Macroeconomic and Liquidity Risk in Optimal Carry and Momentum Portfolios Table 4.7 Panel A reports the performance of the optimal carry and momentum strategy across different economic environments and serves as a preliminary test for our economic analysis. The first three columns report the Sharpe ratios of the trading strategy in three different time periods. The first time period ranges from 1982 to 1992, the second from 1992 to 2002, and the final and most recent time period ranges from from 2002 to 2012. We find that the Sharpe ratios in the first two time periods remain almost unchanged at 1.51 and 1.53, respectively. However, in the third period the performance deteriorates significantly to 1.07. While this is still an impressive Sharpe ratio, it is almost one third smaller than in the previous twenty years. This suggests that significant time variation in the performance of the trading strategies exists. We compare these results to the long-only strategy and find that the effect is exactly the opposite. The Sharpe ratios are lower in the early sample but increase steadily.

Next, we relate the performance of the optimal carry and momentum strategy to variations in business cycle conditions. In particular, we study the Sharpe ratios during the early and late phases of recessions and expansions. We first observe that Sharpe ratios are substantially higher during recessions, averaging 2.2, than in expansions, averaging 1.3. Second, the difference between the Sharpe ratios during the early and late stage of recession and expansion periods, respectively, is only significant during expansions. While early expansions have on average a Sharpe ratio of 1.07 this value increases to 1.53 in late expansions. Again comparing the results to the long-only strategy, we find exactly the opposite trend. During expansions the Sharpe ratio is on average 1.3, while during recessions the returns are even negative.

Table 4.7: **Macroeconomic and Funding Liquidity Influences on Optimal Carry and Momentum Portfolios:** This table reports the performance of optimal carry and momentum portfolios across different economic environments (Panel A) and the results from regressing optimal carry and momentum portfolios on economic predictor variables (Panel B). The first three columns of Panel A show the Sharpe ratios of the trading strategies across the total sample period divided in three subperiods (1982-1992, 1992-202, and 2002-2012). The next four columns report the Sharpe ratios across different business cycle states. Early expansions and late expansions are defined as the first half and second half, respectively, of an expansion period as defined by NBER (from trough to peak). Early recessions and late recessions are defined as the first half and second half, respectively, of an recession period as defined by NBER (from peak to trough). The final two columns report the Sharpe ratios of improving and worsening liquidity environment, defined as positive and negative changes in the TED spread, respectively. Panel B presents the coefficients and Newey and West autocorrelation consistent standard errors when optimal carry and momentum returns are regressed against lagged values of macroeconomic and funding liquidity variables. First, optimal carry and momentum returns are regressed on lagged values of the industrial production gap (IP Gap) as a measure of business cycle conditions. IP Gap is the deviation of the logarithm of total industrial production from a trend that includes both a linear component and a quadratic component.Second, trading returns are regressed on lagged values of common systematic variables including the dividend Yield (Div Yield), the short yield (yield), the default spread (Def) and the term spread (TSpread). Third, trading returns are regressed on a measure of funding liquidity, the TED spread (TED). In all specifications adjusted R^2's are reported. In all panels the table also reports the results for a long only, carry only and momentum only optimized portfolio.

Panel A: Sharpe Ratios across Different Economic Environments

	Timeperiod			Business Cycle				Liq. Environment	
				Expansion		Recession			
	82-92	92-02	02-12	Early	Late	Early	Late	Worse.	Improv.
Carry & Mom	1.51	1.53	1.07	1.07	1.53	2.12	2.37	1.59	1.18
Long	0.95	1.02	1.35	1.43	1.26	-0.14	-0.43	1.47	1.13
Mom	1.46	1.50	0.92	0.96	1.60	1.58	2.10	1.56	1.09
Carry	1.15	1.32	0.87	1.08	0.88	2.71	2.84	1.24	1.01

Panel B: Statistical Significance

	Macro			Common Factors						Liquidity		
	α	IP Gap	$\bar{R}^2$	α	Div Yield	Yield	Def	TSpread	$\bar{R}^2$	α	TED	$\bar{R}^2$
Carry & Mom	0.022	0.149	0.021	0.012	-0.025	0.009	0.013	0.012	0.013	0.025	-0.013	-0.002
(t-stat)	(4.59)	(2.63)		(0.41)	(-2.09)	(1.96)	(0.82)	(1.82)		(5.23)	(-0.75)	
Long	0.014	-0.022	-0.003	0.031	0.003	-0.003	-0.014	-0.001	-0.004	0.022	-0.016	0.023
(t-stat)	(4.69)	(-0.62)		(1.52)	(0.47)	(-0.89)	(-1.34)	(-0.19)		(3.94)	(-1.48)	
Carry	0.014	0.033	-0.001	-0.023	-0.019	0.009	0.015	0.017	0.042	0.015	-0.016	0.003
(t-stat)	(4.73)	(0.91)		(-1.30)	(-2.76)	(3.43)	(2.14)	(3.71)		(4.73)	(-1.26)	
Momentum	0.022	0.119	0.014	0.012	-0.012	0.007	0.003	0.007	0.003	0.023	-0.005	-0.004
(t-stat)	(4.99)	(2.22)		(0.40)	(-1.24)	(1.53)	(0.21)	(1.12)		(5.36)	(-0.31)	

Finally, we test how changing liquidity environments affect the performance of the optimal carry and momentum strategy. We find that during worsening liquidity environments the Sharpe ratios are higher than during improving liquidity conditions[9]. In particular, we find that the Sharpe ratio is 1.59 on average during worsening conditions and 1.18 during improving liquidity conditions.

All these preliminary results indicate that expected returns are time-varying and the time variation is related to macroeconomic state variables. In this section we investigate if the findings are consistent with theoretical predictions. Panel B of Table 4.7 examines more formally how optimal carry and momentum returns are related to macroeconomic and liquidity risk. We use predictive regressions within the context of a multi-beta framework with linear time-varying risk premia. We follow the approach by Chordia and Shivakumar (2002) and Bessembinder and Chan (1992), who show that, assuming a linear multivariate proxy for the pricing kernel, the expected return is the one-period-ahead forecast from the regression,

$$R_t = \alpha + \beta Z_{t-1} + \epsilon. \tag{4.17}$$

where Z_{t-1} is a vector of economic state variables and R_t are the portfolio returns. By following this approach, we do not test whether the payoff of the momentum strategy is related to its covariation with risk factors. Nevertheless, the regression framework allows us to test whether momentum payoffs are captured by a common set of standard macroeconomic variables. As mentioned earlier we use two sets of macroeconomic variables. First, we use the US industrial production gap (IP gap) for which Cooper and Priestley (2009) show a strong relation to business cycles and to have strong predictive power for expected bond and stock returns. Second, we use the common systematic price-base factors, dividend yield, term spread, yield and default spread, as measures for macroeconomic conditions.

Panel B of Table 4.7 presents the regression coefficients and the Newey and West adjusted t-statistics from regressing the optimal carry and momentum returns on the lagged sets of macroeconomic variables. In the macroeconomic model, the estimated coefficient for the IP gap is positively and statistically significant related to the strategy returns. The adjusted R^2 value of the predictive

9. Improving and worsening liquidity environments are defined as positive and negative changes in the TED spread, respectively.

regression is 2 percent. From the results for the common systematic risk factors, we find that coefficients for the dividend yield are statistically negative related to the future strategy returns. The default spread is positively related to expected returns, however, not statistically significant. The short-rate is again statistically significant and positively related to the strategy returns. Finally, the term spread is positively and marginally significantly related to the optimal carry and momentum strategy. The adjusted R^2 value is only 1.3 percent.

When interpreting these results, we note that the two ex-ante stock and bond variables, default spread and dividend yield, are counter-cyclical: when real activity is expected to be high, these variables take low values. Conversely, the term spread and yield are pro-cyclical: high values coincide with high expected real activity. Examining the data, we find that the IP gap is positively related to expected optimal carry and momentum returns. This relationship suggests that the variation of the expected returns is pro-cyclical. The same result is applicable, when observing the common systematic factors. While the dividend yield is negatively related to expected returns, the yield and the term spread are positively related to future returns. Again this suggests that the return premia are pro-cyclical. These findings, however, are inconsistent with macroeconomic theories of discount rate variation, which predict counter-cyclical variation of risk premia. These results may be interpreted as investors using futures trading strategies as a source of hedging, or as a misspecification due to the macroeconomic variable being related to variables that are pro-cyclically correlated.

4.5.4 Hedge Fund Activity and Capital Flows

Recently, a large stream of literature in financial economics emphasizes segmented markets and frictions as a source of predictable variation in asset prices. Segmented markets are important for expected returns because financial agents, who are active in a particular market, have only limited risk-sharing abilities. Theoretical examples of segmentation in diverse asset classes are given by Brunnermeier (2008), Gabaix et al. (2007) and Daniel (2011). Most of the theoretical models predict that expected returns are related to capital and capital flows deployed in the asset classes. The limited risk-bearing capacity of speculators generates downward sloping demand curves that result in a positive relation between capital flows and contemporaneous returns. Consequently, expected returns are negatively correlated with the total capital committed within an asset class.

Jylhä and Suominen (2011) show theoretically and empirically, in a similar vain to our study, that returns to carry strategies in currency markets are predictable by the speculative capital of hedge funds active in these markets.

In this section we test if financial market frictions, in particular speculative capital flows by hedge funds, are related to the returns of optimal carry and momentum portfolios. We proceed as follows. First, we test if hedge funds indeed trade in carry and momentum strategies within the framework of a linear factor model. Second, we test if speculative capital by hedge funds is related to optimal carry and momentum returns.

Hedge Fund Returns and Optimal Carry and Momentum Strategies Table 4.8 presents the results from regressing the net-of-fee monthly returns of three asset-under-management weighted hedge fund indices on the returns of the optimal futures portfolios, and various combinations of factor benchmark models. We include a hedge fund return index that covers all hedge fund styles and two sub indices: one macro hedge fund index and a managed futures hedge fund index. We choose the Fung and Hsieh (2004) (FH7, hereafter) 7-factor model, and a simple CAPM model based on the returns of the MSCI, as benchmark models. The FH7 incorporates five primitive trend-following factors for bonds, commodities, currencies, interest rates and equity classes. Furthermore, we include the term spread and the default spread as control variables for changing macroeconomic environments. For each hedge fund index or subindex, columns (1)-(4) report the alpha, which can be interpreted as the proportion of returns that cannot be explained by the factors, and the beta coefficients denoting the factor loadings. In brackets we report the statistical significance in terms of Newey-West adjusted t-stats for each coefficient. In the last row we report the adjusted R^2 values.

Column (1) and (2) of Table 4.8 report the results of regressing the hedge fund index returns, first on the MSCI returns, and then on a combination of MSCI returns and the optimal carry and momentum returns. In agreement with the literature, we observe that all factor coefficient estimates are positive and highly significantly related to hedge fund index returns. However, a surprising feature is the explanatory power of the factors measured by adjusted R^2 values. We find that, while the MSCI explains around 28 percent of the variation in all hedge fund returns, this value increases by almost 10 percent to a total of 40 percent when including the optimal carry

Table 4.8: **Return Decomposition of Hedge Fund Indices:** This table reports the results of regressing the returns of hedge fund indices on various return factors. The hedge fund indices include an index across all strategies (HF Index), a macro strategy sub index (HF Macro), and a managed futures sub index (HF CTA). For each hedge fund index return regression we use three factor specifications: (1) includes the returns to the MSCI, (2) the returns to the MSCI and the optimal carry and momentum strategy (Car & Mom), (3) corresponds to an extended Fung and Hsieh (2004) model using all primitive trend-following Fung and Hsieh (2001) factors, the S&P500, the term spread (TSpread), which is the difference between a 10 year and a 3 month bond yield, and the default spread (DefSpread), and (4) which is a combination combination of (3) including the optimal carry and momentum strategy (Car & Mom). The data period used for the regressions is restricted by the data availability of the Fung and Hsieh (2004) factors, from January 1994 to December 2011.

	HF Index				HF Macro				HF CTA			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Alpha	0.01	0.00	0.00	-0.00	0.01	0.01	0.01	0.01	0.01	-0.00	0.02	0.01
tstat	(4.87)	(3.13)	(0.35)	(-0.10)	(4.89)	(2.61)	(1.25)	(0.64)	(2.46)	(-0.73)	(1.10)	(0.56)
MSCI	0.25	0.26	0.23	0.24	0.11	0.14	0.10	0.12	-0.03	0.01	0.05	0.08
tstat	(6.39)	(7.06)	(5.61)	(6.36)	(2.05)	(2.67)	(1.82)	(2.38)	(-0.59)	(0.32)	(0.95)	(2.11)
Car&Mom		0.10		0.09		0.17		0.16		0.28		0.26
tstat		(4.03)		(4.52)		(5.04)		(5.14)		(7.48)		(7.93)
PTFSbd			-0.03	-0.03			-0.02	-0.02			0.04	0.04
tstat			(-2.36)	(-2.22)			(-1.21)	(-1.13)			(2.39)	(2.75)
PTFScom			0.02	0.01			0.02	0.00			0.05	0.02
tstat			(1.73)	(0.98)			(1.20)	(0.33)			(3.03)	(1.51)
PTFSfx			0.01	0.01			0.01	0.01			0.04	0.04
tstat			(1.03)	(1.08)			(0.78)	(0.81)			(2.89)	(3.03)
PTFSir			-0.02	-0.01			-0.02	-0.01			-0.02	-0.00
tstat			(-3.58)	(-3.18)			(-1.85)	(-1.12)			(-2.14)	(-0.72)
PTFSstk			0.02	0.01			0.02	0.01			0.05	0.02
tstat			(1.66)	(0.92)			(1.35)	(0.42)			(2.99)	(1.75)
TSpread			-0.00	-0.00			-0.00	-0.00			0.00	0.00
tstat			(-1.90)	(-1.44)			(-1.47)	(-1.00)			(0.01)	(0.81)
DefSpread			0.00	0.00			0.00	0.00			-0.00	-0.00
tstat			(0.84)	(0.95)			(0.07)	(0.11)			(-0.59)	(-0.68)
R^2	0.29	0.38	0.36	0.43	0.03	0.19	0.05	0.19	-0.00	0.30	0.18	0.42

and momentum factor. For robustness, we report in column (3) and (4) the benefit of including the optimal carry and momentum factor into the FH7 model. The FH7 model already explains 35 percent of variation in hedge fund returns. When including the optimal carry and momentum factor this value increases again by almost 8 percent to a total of 43 percent.

In the next four columns we repeat the exercise, but use global macro hedge fund returns as the dependent variable. The MSCI only explains 3 percent and the FH7 factors account for only 5 percent variability of global macro returns. When including the optimal carry and momentum returns into the regression the R^2 values increase in both cases to more than 19 percent. A similar finding emerges for the last four columns. Here, we present the results for the managed futures index. The MSCI is almost unrelated, or even negatively related, to the managed futures returns. Including the optimal carry and momentum factor raises the explanatory power to 30 percent. Even more unexpected are the results for the FH7 factors. The FH7 factors are constructed to resemble the risk in trend-following strategies, which is exactly what managed futures invest in. While these factors indeed explain 18 percent of the variance, the inclusion of the optimal carry and momentum factor raises this by 25 percent to a total of 42 percent.

In summary, the presented results suggest that hedge funds do indeed follow optimal carry and momentum strategies and their trading behaviour, or capital flow in particular, can have an effect on the profitability of optimal carry and momentum strategies as conjectured earlier. We explore this feature further in the following section.

Hedge Fund Capital and Optimal Carry and Momentum Strategies Table 4.9 tests if optimal carry and momentum returns are related to financial frictions, and, in particular, to speculative capital of hedge funds. Panel A reports the performance of the optimal carry and momentum strategy conditioning on different hedge fund related states and serves as a preliminary investigation. First, the Sharpe ratios are reported during times of capital flows, into and out of the hedge fund sector. We find that the Sharpe ratios of optimal carry and momentum strategies are higher during times of capital outflows. The Sharpe ratios are 1.32 and 1.49 during times of inflows and outflows, respectively. These observations are consistent with the implications of markets with frictions. During times of capital outflows, less capital is available to arbitrage inefficiencies, which in turn implies higher returns. For comparison, we also report the performance of a long-only strategy

Table 4.9: **Hedge Funds, Speculative Capital and the Optimal Carry and Momentum Strategy:** This table reports the performance of optimal carry and momentum portfolios across different hedge fund related economic environments (Panel A), the results from regressing optimal carry and momentum portfolios on proxies for speculative capital of hedge funds (Panel B), the interaction of macroeconomic environments and speculative capital in explaining optimal carry and momentum returns (Panel C). The first two columns of Panel A present the performance, measured as Sharpe ratio, of optimal carry and momentum portfolios during increasing and decreasing hedge fund flow (HF Flow) environments. HF Flow is measured as the sum of the past 12 months capital flows into the hedge fund industry. The next two columns report the Sharpe ratios before and after the LTCM crisis and the last four columns interact improving and worsening liquidity environments before and after the LTCM crisis. Panel B provides statistical evidence of optimal carry and momentum returns regressed on total hedge fund asset under management scaled by M2 (AUM), HF Flow and a time trend variable. The table only reports the coefficient estimates as well as Newey-West adjusted test statistics for the total AUM and HF Flow variables. The table also presents regression results when including combinations of TED, IP Gap, Yield and Tspread. Panel C presents the coefficient estimates and test statistics of optimal Carry and Momentum returns regressed on HF AUM and HF Flow interacted with growing and decreasing business conditions measured by IP Gap growth and NBER recession periods. In all panels the results for long only, carry only and momentum only are presented for comparison.

Panel A: Sharpe Ratios in Different Hedge Fund States

	HF Flow		LTCM		LTCM and Liq.			
	Incr.	Decr.	pre 08/98	post 08/98	pre 08/98		post 08/98	
					Wors.	Impr.	Wors.	Impr.
Carry & Mom	1.32	1.49	1.67	1.15	1.83	1.49	1.33	0.95
Long	1.13	1.35	1.41	1.10	1.73	1.27	1.23	0.96
Mom	1.16	1.49	1.60	1.04	1.76	1.49	1.34	0.80
Carry	1.12	1.14	1.51	0.92	1.58	1.18	0.92	0.92

Panel B: Speculative Capital and Optimal Carry and Momentum Returns

	α	HF AUM_{t-1}	HF $Flow_t$	TED	IP Gap	Yield	TSpread	$\bar{R}^2$
(1)	0.043	-0.178	0.402					0.026
(t-stat)	(5.86)	(-3.97)	(3.24)					
(2)	0.038	-0.194	0.431	0.010				0.026
(t-stat)	(4.75)	(-4.16)	(3.83)	(1.17)				
(3)	0.043	-0.182	0.417		-0.006			0.022
(t-stat)	(5.45)	(-3.03)	(2.05)		(-0.09)			
(4)	0.021	-0.178	0.462			0.002	0.005	0.015
(t-stat)	(0.50)	(-2.50)	(2.13)			(0.35)	(0.61)	
Long	0.027	-0.068	0.31	-0.024	-0.104			0.036
(t-stat)	(4.21)	(-1.13)	(1.52)	(-1.81)	(-2.38)			
Carry	0.031	-0.189	0.40	0.003	-0.0852			0.029
(t-stat)	(4.50)	(-2.3)	(1.86)	(0.33)	(-1.23)			
Mom	0.029	-0.146	0.43	0.012	0.006			0.061
(t-stat)	(-1.75)	(-1.98)	(1.86)	(1.55)	(0.07)			

Panel C: Macro Environments and Speculative Capital

		HF AUM×		HF Flow×		
	α	$I_{(IPGap>0)}$	$I_{(IPGap<0)}$	$I_{(IPGap>0)}$	$I_{(IPGap<0)}$	$\bar{R}^2$
Carry and Mom	0.043	-0.159	-0.224	0.134	0.776	0.025
(t-stat)	(5.08)	(-2.58)	(-2.48)	(0.62)	(1.92)	
Long	0.016	0.007	-0.046	0.086	0.132	-0.010
(t-stat)	(3.63)	(0.11)	(-0.85)	(0.43)	(0.61)	
Carry and Mom	0.045	-0.188	-0.188	0.255	0.898	0.027
(t-stat)	(5.27)	(-3.06)	(-1.98)	(1.19)	(1.88)	
Long	0.014	0.037	-0.093	0.021	0.075	0.020
(t-stat)	(3.46)	(0.93)	(-1.18)	(0.14)	(0.19)	

during these states of the world. We find exactly the opposite relation: Sharpe ratios are higher for increasing states, than for decreasing states.

In the next columns of Table 4.9 we split the data sample into two periods, prior to and after August 1998, which roughly corresponds to the peak of the funding crisis following the collapse of Long Term Capital Management (LTCM). A variety of academic studies find a structural break in the return generating process of hedge funds during this time, e.g., Meligkotsidou and Vrontos (2008) or Sadka (2010). We find that the performance of optimal carry and momentum strategies before 1998 is significantly higher than afterwards, with Sharpe ratios of 1.67 and 1.15, respectively.

Finally, in Table 4.9 Panel A, we investigate the effects of changing liquidity conditions on the performance of optimal carry and momentum before and after August 1998 (LTCM crash). Similarly to Panel A Table 4.7 we find that, during times of worsening funding liquidity conditions, the Sharpe ratios are higher than during improving funding liquidity conditions. This effect is symmetric for the time before and after August 1998. Changing liquidity conditions have the same effect on long-only strategies. These two findings suggest that, while funding liquidity is an important source of expected returns in futures markets, it is not specific to optimal carry and momentum strategies.

Table 4.9 Panel B examines more formally the relation between optimal carry and momentum strategies and speculative capital, which we define as the capital available to hedge funds. We regress the trading strategies return series' on the contemporaneous month's hedge fund flow and the previous month's hedge fund total AUM (HF AUM). This regression follows the methodology of Jylhä and Suominen (2011) and Jylha et al. (2010). We also include a time trend variable to account for the positive trend in total AUM and the negative trend in the optimal portfolio returns. Panel B of Table 4.9 presents the coefficient estimates, the Newey-West adjusted test statistics and the adjusted R^2 values for the regression. We find that the coefficients to both HF AUM and HF Flow are statistically highly significant. Furthermore, the signs of the coefficients confirm the conjectures. We find that the HF Flow variable is contemporaneously positive related to optimal carry and momentum returns, consistent with a downward sloping demand curve. However, this effect results in the lower than expected returns with respect to total capital committed in the next period, which is confirmed by the data; the lagged HF AUM variable is negatively related to optimal carry and momentum returns.

To disentangle the role of liquidity risk, macroeconomic risk and speculative capital, we run several control regressions. First, we add the lagged TED spread into the regression framework. As expected, and consistent with the results in Table 4.7 Panel A, the coefficient estimate for the TED spread is statistically insignificant, while the coefficients of HF AUM and HF Flow remain significant. Then we test the explanatory power of lagged macroeconomic variables used in the previous section together with the hedge fund flow variables. Panel B of Table 4.9 shows that the coefficient estimate of the IP Gap becomes insignificant once we use this variable together with the HF AUM and HF Flow variables. This influence indicates that the pro-cyclical behaviour of expected returns of optimal carry and momentum strategies documented in the previous section is not driven by changing risk preferences of economic agents, but rather due to the correlation of the HF Flow and IP Gap variable. In undocumented results we find that the correlation between HF Flow and IP Gap is 55 percent. Finally, we regress optimal carry and momentum returns on lagged yield and term spread. As for the IP Gap we find that the coefficients are not significant once we include the hedge fund capital variables. For comparison, we also test the regressions for the long-only strategy. We find that the HF AUM and HF Flow variables are not significantly related to the long strategy. On the other hand, we find that IP Gap variable predicts significantly counter-cyclical risk premia for the long-only strategy, consistent with macroeconomic risk theories. Furthermore, we find that the TED spread is weakly but significant related to the long-only strategy, consistent with liquidity risk theories.

Table 4.9 Panel C reports the results from regressing optimal carry and momentum strategies on the HF AUM and HF Flow variable conditioned on improving and weakening macroeconomic environments. First, we find that HF AUM is always negatively and statistically significantly related to optimal carry and momentum strategies regardless of the macroeconomic states. For the HF Flow variable, on the other hand, we find that a significant flow return relation is only given in weakening macroeconomic environments, while this relation is insignificant during economic upturns. For long-only strategies the relation between HF AUM and HF Flow variables and returns remains insignificant.

In summary, the results in this section suggest that a large proportion of the returns to optimal carry and momentum strategies are due to the segmentation of futures markets from other asset classes. In particular, the strong involvement of hedge funds limits the risk sharing ability of market

participants with other asset classes. Consequently, only increasing capital available to hedge funds reduces the returns to the trading strategies and increases the efficiency of futures markets.

4.6 Conclusion

We have empirically investigated optimal futures portfolios based on the momentum and carry characteristic across currencies, commodities, equity indices and bond futures. We adapt the optimal parametric policy methodology by Brandt et al. (2009) to include assets with varying levels of volatility. We find that the optimal carry and momentum strategy yields high Sharpe ratios and certainty equivalent returns. The returns are hard to understand in a framework that relies on covariance risk with standard risk factors and cannot be explained by time varying macroeconomic and funding liquidity risk. In contrast, we find that the returns to optimal carry and momentum returns are highly related to speculative capital of hedge funds. Contemporaneously larger capital flows lead to higher optimal carry and momentum returns, implying that expected returns decrease with the total amount of assets under management of hedge funds.

These findings have important implications for investors and policy makers. First, investors, who invest in carry and momentum strategies, should be aware that a large proportion of the performance relies on the size of the hedge funds industry and, in particular, the speculative capital available to arbitrage inefficiencies. The high returns of optimal carry and momentum strategies might deteriorate in the future with increasing hedge fund industry size. From a policy perspective it seems disadvantageous to restrict hedge funds in their activity in futures markets. In contrast, the more capital hedge funds allocate to futures markets the more efficient they become.

CHAPTER 5
CONCLUSIONS, POLICY IMPLICATIONS AND FUTURE RESEARCH

This thesis builds on the literature of predictability in financial markets and, in particular, time variation in expected commodity returns and other global futures markets. I identify several new stylized facts of aggregate commodity returns that are consistent with the literature of financial economics of equity, bond and currency markets. At the same time, however, I highlight distinguishing features that are local to commodities. I use novel econometric and portfolio allocation frameworks to uncover and extend the understanding of predictability and its drivers in global futures markets. In this final section I provide a brief overview of my contributions, outline similarities and differences to the findings in other asset classes, highlight the policy implications and propose future research paths.

The **first topic** of this thesis focuses on predictability of aggregate commodity futures returns using price-based valuation ratios. I use the average futures discount (AFD) as an ex-ante measure of expected returns and run predictive time series regressions of future returns and spot changes on the AFD. The coefficient estimates for the predictive regressions at various horizons range from 0.5 to 1.1 and are all statistically significant and robust. The R^2 values of the predictive regression increase from 1 percent on a monthly horizon to almost 10 percent at the 9 months horizon. I decompose the total expected returns into spot and term risk premia and find that the time variation of total expected returns is mostly attributable to time-varying spot risk premia, while term premia satisfy a version of the expectation hypothesis with constant risk premia. My findings extend the literature on time variation of expected returns, as surveyed by Cochrane (2011), to aggregate commodity returns. I show that the main characteristics of discount rate variation observable in a variety of asset classes extend to aggregate commodity markets regarding (a) robust predictability of expected returns by yield-based valuation ratios, (b) increasing R^2 values with longer predictive horizons, and (c) no or weak predictability of spot price growth, extend to aggregate commodity markets.

Several findings in this chapter, however, deserve special attention. First, I find that the coefficient estimates of the predictive regression are below 1 at short horizons. This finding indicates mean reversion in spot rates and therefore, diminishes the value of the predictive regressions. An

explanation for the mean reversion is provided by The Theory of Storage. This theory predicts that low inventory levels lead to high contemporaneous futures discounts and high spot prices for individual commodities and subsequent mean reversion. I show that this behaviour is not only observable on an individual commodity level, but also for commodity portfolios. Second, and most importantly, I find that the AFD and consequently expected returns vary pro-cyclical with macroeconomic conditions. Asset pricing theory suggests, however, that risk premia should move counter-cyclically. Two potential explanations for my findings are feasible. On the one hand, economic agents view commodities as hedges against stock and bond return fluctuations. This reasoning would result in increasing risk premia in commodities when bond and stock risk premia are low and vice versa. On the other hand, pro-cyclical expected commodity returns could be a consequence of under-reacting economic agents to new macroeconomic information and therefore, fall into the behavioural category (Hong and Stein (1999)). These arguments are encouraged by the finding that IP growth predicts spot price growth in addition to changes in the AFD. In summary, more theoretical models of commodity return variation are needed to understand these complex dynamics.

The **second topic** offers statistical and economic evidence of commodity return and volatility predictability using a multitude of predictive variables including commodity market-specific, macroeconomic, and financial variables. To accommodate the many potential model outcomes, I use a Bayesian Model Averaging (BMA) methodology for the combination forecasts. The BMA methodology explicitly incorporates model uncertainty in the predictive regressions, a desirable feature given the ambiguous empirical and theoretical evidence on commodity return prediction models. I find, that at most 4 out of 41 variables that have previously been shown to predict commodity returns by other authors, provide statistical evidence on return and volatility predictability according to posterior inclusion probabilities. The in-sample BMA tests reveal a prominent role for macroeconomic variables, such as global and US specific IP growth, for predicting portfolio commodity returns across commodity sectors. At the same time, the short rate and futures discount are significant commodity return predictors. The results for predicting commodity portfolio volatility are somewhat different. I find that lagged values of volatility predict future volatility a result that is not surprising given its high persistence. Next, I find that the dividend yield and default spread, which are measures of ex-ante stock and bond risk premia, have posterior inclusion probabilities

that indicate predictive value. Again the futures discount also offers predictive power for future commodity volatility. While the in-sample results show clear pattern of predictability, they do not guarantee that a real-time investor or forecaster can benefit from the predictive relations. I assess the out-of-sample performance of commodity return and volatility predictability using statistical and economic criteria. I find that the return predictions generate significant statistical and economic value for a real-time investor or forecasts. This result is robust for various sub-samples and commodity sector portfolios.

Findings within this topic contribute to the literature on robust inference of predictive relations and sheds light on the underlying sources of return and volatility variation in commodity markets. In assessing the predictive performance of the commodity returns and volatility models I implement a BMA methodology that is robust to parameter and model uncertainty. I conduct robust out-of-sample forecasting exercises by providing a comprehensive economic evaluation of the forecasting models in the context of dynamic asset allocation strategies. In contrast to studies in currency, equity and bond markets, which find that most information content for future returns is in price-based predictive measures such as the dividend yield, default spread or futures discount (Avramov (2002) and Della Corte et al. (2009)), I find that future commodity returns are significantly related to macroeconomic information. These findings are similar to the recent studies by Cooper and Priestley (2009), Ludvigson and Ng (2009) and Lustig et al. (2013). Furthermore, Christiansen et al. (2012) show that predictors of return premia of financial assets also predict their volatility. I find a somewhat different result for commodity markets. In particular, I find that returns and volatility of commodities are related to different sources of information.

The **third topic**, which was written in collaboration with James B. Grant, evaluates the performance of combined carry and momentum strategies in a parametric portfolio choice problems across the asset classes commodities, currencies, bonds, and equities. We find that the portfolio returns are high, with Sharpe ratios exceeding 1.2 and are robust to transaction costs. The diversification benefit of including the carry and momentum strategy in a traditional stock and bond portfolio or other factor portfolios reduces skewness and kurtosis of the traditional portfolios, while increasing the combined Sharpe ratio. The high returns to carry and momentum strategies pose yet another challenge to existing financial theories. To shed light on the sources of the high returns, we further investigate risk based explanations, macroeconomic theories and friction based theories. We

find that traditional and non-traditional risk factors are not able to explain carry and momentum returns. However, we also find that hedge fund capital flow is contemporaneously related to the returns to carry and momentum strategies and that the amount of total asset under management in the hedge fund industry negatively predicts future returns. These findings are consistent with predictions of friction based theories with segmented markets and limits to arbitrage. Our results are robust to the inclusion of macroeconomic and liquidity variables.

Findings in this topic contribute to the literature on explanations for positive return premia across asset classes (Asness et al. (2013), Koijen et al. (20012), Moskowitz et al. (2011) and Asness et al. (2012)). By assessing the performance of carry and momentum strategies across asset classes and relating it to the activity of the hedge fund industry, we provide empirical evidence of frictions in futures markets. In contrast to previous studies relating hedge fund activity to return premia, such as Baltas and Kosowski (2013), we do not rely on equally- or volatility-weighted strategies with only one particular signal but use a more comprehensive and realistic portfolio construction approach that optimally trades of the momentum and carry signal in a unifying framework. The success of this investigation is related to the inclusion of important limitations in the portfolio allocation process such as transaction cost, optimal weighting, and utility functions, other than the mean/variance approach.

The findings in the three main chapters have important policy implications. First, despite the general view that commodity market prices are related to inflation, US Real GDP, macroeconomic activity and stock market movements, the formal link is not well understood[1]. I provide novel evidence for such links in this thesis (topic one and two). I show that macroeconomic variables predict commodity returns, but commodity returns do not lead macroeconomic activity. Second, the aggregate commodity portfolios that I study contain a large fraction of agricultural and energy commodities. Understanding agricultural commodity price movements has recently become highly relevant for developing countries to alleviate poverty. Chen et al (2010) argue that many developing economies rely heavily on commodity production for growth and export and governments often distribute food grains at subsidized prices to help combat poverty. A better understanding of

1. Recently, Gospodinov and Ng (2011) show that the first three principal components of the cross section of futures discounts have predictive power for future inflation, Kilian and Vigfusson (2013) find that the oil price has predictive power for US GDP and Boons et al (2012) find that commodity risk is priced in the cross section of stock returns.

energy prices, on the other hand, is important for policy makers, who introduced emission trading schemes to combat climate change such as the European Union Emission Trading Schemes (EU ETS). The EU ETS is built on the premise that CO_2 emission can be priced (emission permits) and traded freely, just as other commodities. Alberola et al. (2008) and reference therein show that the main driver of emission permit prices are energy commodity fluctuations. Third, the recent surge in commodity prices and subsequent sharp decline during 2008 has triggered an extensive policy debate about the role of speculative trading activity in commodity markets. The large capital inflow into commodity markets from non-user speculators and passive investors concurrently with the large price fluctuations led to the hypothesis that speculators changed the behaviour of commodity markets (Masters (2009) and Singleton (2013)). On the other hand, various authors attribute the large price movements to fundamentals of demand and supply and actually argue that the increasing influence of speculators has moderated the volatility of commodities (see Irwin et al. (2009) for a review). In my thesis I present evidence for the latter argument. In Chapter 3, I show that predictive relations are to the most part attributed to macroeconomic variables and measure of capital flows of speculators do not have a distorting effect on commodity prices. In Chapter 4, I find that the increasing presence of speculators (hedge funds) leads to a decrease in deviations from efficient markets as measured by carry and momentum strategies.

Findings presented in this thesis give rise to various future research directions to improve current methodologies and provide more detailed and robust explanations for the empirical results. First and foremost, commodity and futures market research embedded within financial economic theory of discount rate variation and asset allocation decisions is at its infancy (Rouwenhorst and Tang (2012)). As a first step, I provide new stylised facts regarding the characteristics of aggregate return variation of commodity and global futures markets. In the future these stylised facts of return variation in aggregate commodity markets need to be accommodated in sound theoretical models, an endower that has already started to explain oil market risks (Baker and Routledge (2012)). One particularly interesting avenue of future research is to better understand the link between trading activity of commodity market participants, macroeconomic fundamentals, and commodity returns. So far the data sample on commodity trading activity is too short to provide meaningful results but will become interesting in the future. Finally, more research is warranted to understand common characteristics across asset classes, as reviewed by Cochrane (2011). The

third topic of my research provides first evidence in this direction.

5.1 Appendix A

5.1.1 Statistical Inference of optimal future portfolio coefficients

By posing the portfolio problem as an estimation problem, we can conduct statistical inference for the estimated parameters $\hat{\theta}$. In particular, the optimization problem (4.7) with the parametric and linear portfolio policy (4.2) satisfies the first-order conditions,

$$\frac{1}{T}\sum_{t=0}^{T-1} h(r_{t+1}, x_t; \theta) \equiv \frac{1}{T}\sum_{t=0}^{T-1} u'(r_{p,t+1}) \left(\frac{1}{N_t \sigma_{i,t}} \hat{x}_t^\top r_{t+1}\right) = 0. \tag{5.1}$$

Interpreting this relation in terms of the general method of moments estimator (GMM) by Hansen (1982), we have the asymptotic covariance matrix of $\hat{\theta}$ as,

$$\Sigma_\theta = \text{AsyVar}[\hat{\theta}] = \frac{1}{T}\left[G^\top V^{-1} G\right]^{-1} \tag{5.2}$$

where

$$G \equiv \frac{1}{T}\sum_{t=0}^{T-1} \frac{\partial(r_{t+1}, x_t; \theta)}{\partial \theta} = \frac{1}{T}\sum_{t=0}^{T-1} u''(r_{p,t+1}) \left(\frac{1}{N_t \sigma_{i,t}} \hat{x}_t^\top r_{t+1}\right) \left(\frac{1}{N_t \sigma_{i,t}} \hat{x}_t^\top r_{t+1}\right)^\top \tag{5.3}$$

and V is a consistent estimator of the covariance matrix of $h(t, x; \theta)$. If we assume that the portfolio policy is correctly specified and implements an unconstrained optimization then, by construction, the marginal utilities are uncorrelated and we can consistently estimate V by,

$$\frac{1}{T}\sum_{t=0}^{T-1} h(r_{t+1}, x_t; \hat{\theta}) h(r_{t+1}, x_t; \hat{\theta})^\top. \tag{5.4}$$

To accommodate for a misspecified portfolio policy or weighting constraints, this estimator can be substituted for an autocorrelation adjusted estimator, such as that of Newey and West (1987).

As an alternative to asymptotic estimators altogether, the covariance matrix can be estimated using a bootstrap method. This approach consists of taking a large number of randomized samples (with replacement) of price and signal data. For each sample, the θ coefficients are re-estimated, collectively yielding a distribution of coefficients to compute the covariance matrix. This method has the advantage of not being subject to difficulties associated with asymptotic results, and accom-

modates any non-normal features of the data. We use bootstrapped covariance matrix estimates in the analysis below.

Using the covariance matrix, we test individual and joint hypotheses on the portfolio policy estimates to determine whether an investor finds it optimal to assign portfolio weights according to historical characteristics of each futures contract.

It is important to recognise that the significance of the portfolio coefficients does not correspond to whether a characteristic is cross-sectionally related to the conditional moments of excess futures returns. Instead, significance implies that optimized portfolio characteristic generates a statistically significant increase in the investors utility of wealth. Therefore, if the contribution of two moments are correlated such that their effects offset each other in the investor utility function, the optimal portfolio will be found to be independent of that characteristic.

REFERENCES

Viral V Acharya, Lars A Lochstoer, and Tarun Ramadorai. Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, 2013.

Tobias Adrian, Erkko Etula, and Tyler Muir. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, 2011.

Emilie Alberola, Julien Chevallier, and Benoıt Chèze. Price drivers and structural breaks in european carbon prices 2005–2007. *Energy Policy*, 36(2):787–797, 2008.

Andrew Ang and Geert Bekaert. Stock return predictability: Is it there? *Review of Financial Studies*, 20(3):651–707, 2007.

Cliff Asness, Andrea Frazzini, and Lasse Heje Pedersen. Leverage aversion and risk parity. *Financial Analysts Journal*, 68(1):47–59, 2012.

Clifford S Asness, Tobias J Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013.

Doron Avramov. Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64(3):423–458, 2002.

Doron Avramov, Robert Kosowski, Narayan Y Naik, and Melvyn Teo. Hedge funds, managerial skill, and macroeconomic variables. *Journal of Financial Economics*, 99(3):672–692, 2011.

Doron Avramov, Laurent Barras, and Robert Kosowski. Hedge fund return predictability under the magnifying glass. *Journal of Financial and Quantitative Analysis*, 2012.

Warren Bailey and K.C. Chan. Macroeconomic influences and the variability of the commodity futures basis. *The Journal of Finance*, pages 555–573, 1993.

Malcolm Baker and Jeffrey Wurgler. Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, 61(4):1645–1680, 2006.

Steven D Baker and Bryan R Routledge. The price of oil risk. *Working paper*, 2012.

Gurdip Bakshi, George Panayotov, and Georgios Skoulakis. The baltic dry index as a predictor of global stock returns, commodity returns, and global economic activity. *Working paper*, 2011.

Akindynos-Nikolaos Baltas and Robert Kosowski. Momentum strategies in futures markets and trend-following funds. *Working paper*, 2013.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

Ravi Bansal, Dana Kiku, and Amir Yaron. Long run risks, the macroeconomy, and asset prices. *The American Economic Review*, 100(2):542–546, 2010.

Nicholas Barberis and Richard Thaler. A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128, 2003.

Pedro Barroso and Pedro Santa-Clara. Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, 2013.

H. Bessembinder. Systematic risk, hedging pressure, and risk premiums in futures markets. *Review of Financial Studies*, 5(4):637–667, 1992.

H. Bessembinder and K. Chan. Time-varying risk premia and forecastable returns in futures markets. *Journal of Financial Economics*, 32(2):169–193, 1992.

Van Binsbergen and Ralph SJ Koijen. Predictive regressions: A present-value approach. *The Journal of Finance*, 65(4):1439–1471, 2010.

Peter Bossaerts and Pierre Hillion. Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies*, 12(2):405–428, 1999.

Michael W Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22(9): 3411–3447, 2009.

M.J. Brennan. The supply of storage. *The American Economic Review*, 48(1):50–72, 1958.

Markus K Brunnermeier. Deciphering the liquidity and credit crunch 2007-08. *National Bureau of Economic Research*, 2008.

Markus K Brunnermeier and Lasse H Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.

Markus K Brunnermeier, Stefan Nagel, and Lasse H Pedersen. Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23(1):313–348, 2008.

Andrea Buraschi, Robert Kosowski, and Worrawat Sritrakul. Incentives and endogenous risk taking: A structural view of hedge funds alphas. *The Journal of Finance*, 2013.

Craig Burnside, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. The returns to currency speculation in emerging markets. *The American Economic Review*, 2006.

Craig Burnside, Martin Eichenbaum, and Sergio Rebelo. Carry trade: The gains of diversification. *Journal of the European Economic Association*, 6(2-3):581–588, 2008.

John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behaviour. *The Journal of Political Economy*, 107(2):205–251, 1999.

John Y Campbell and Robert J Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228, 1988.

John Y Campbell and Samuel B Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4):1509–1531, 2008.

John Y Campbell and Motohiro Yogo. Efficient tests of stock return predictability. *Journal of Financial Economics*, 81(1):27–60, 2006.

Mark M Carhart. On persistence in mutual fund performance. *The Journal of Finance*, 52(1): 57–82, 1997.

Jaime Casassus, Pierre Collin-Dufresne, and Bryan R Routledge. Equilibrium commodity prices with irreversible investment and non-linear technology. *Working Paper*, 2005.

Tarun Chordia and Lakshmanan Shivakumar. Momentum, business cycle, and time-varying expected returns. *The Journal of Finance*, 57(2):985–1019, 2002.

Todd E Clark and Kenneth D West. Approximately normal tests for equal predictive accuracy in nested models. *Journal of econometrics*, 138(1):291–311, 2007.

John H Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008.

John H. Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4):1047–1108, 2011.

John H. Cochrane and Monica Piazzesi. Bond risk premia. *American Economic Review*, 2002.

I. Cooper and R. Priestley. Time-varying risk premiums and the output gap. *Review of Financial Studies*, 22(7):2801–2833, 2009.

K.J.M. Cremers. Stock return predictability: A bayesian model selection perspective. *Review of Financial Studies*, 15(4):1223–1249, 2002.

Kent Daniel. Momentum crashes. *Working Paper*, 2011.

Charoula Daskalaki and George Skiadopoulos. Should investors include commodities in their portfolios after all? new evidence. *Journal of Banking & Finance*, 35(10):2606–2626, 2011.

Charoula Daskalaki, Alexandros Kostakis, and George Skiadopoulos. Are there common factors in commodity futures returns? *Journal of Banking & Finance*, 40(2):346–363, 2014.

Frans A. De Roon, Theo E. Nijman, and Chris Veld. Pricing term structure risk in futures markets. *Journal of Financial and Quantitative Analysis*, 33(1), 1998.

Frans A. De Roon, Theo E. Nijman, and Chris Veld. Hedging pressure effects in futures markets. *The Journal of Finance*, 55(3):1437–1456, 2000.

A. Deaton and G. Laroque. Competitive storage and commodity price dynamics. *Journal of Political Economy*, pages 896–923, 1996.

Pasquale Della Corte, Lucio Sarno, and Ilias Tsiakas. An economic evaluation of empirical exchange rate models. *Review of Financial Studies*, 22(9):3491–3530, 2009.

Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953, 2009.

Darrell Duffie and Bruno Strulovici. Capital mobility and asset pricing. *Econometrica*, 80(6): 2469–2509, 2012.

K. Dusak. Futures trading and investor returns: An investigation of commodity market risk premiums. *The Journal of Political Economy*, pages 1387–1406, 1973.

Claude B Erb and Campbell R Harvey. The strategic and tactical value of commodity futures. *Financial Analysts Journal*, pages 69–97, 2006.

Erkko Etula. Broker-dealer risk appetite and commodity returns. *Journal of Econometrics*, 2014.

E.F. Fama and K.R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993.

Eugene Fama. Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3):319–338, 1984.

Eugene Fama and Kenneth French. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, pages 55–73, 1987.

Eugene Fama and Kenneth French. Business cycles and the behavior of metals prices. *The Journal of Finance*, 43(5):1075–1093, 1988.

Eugene Fama and Kenneth French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1), 1996.

Eugene F Fama. The behavior of stock-market prices. *The Journal of Business*, 38(1):34–105, 1965.

Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.

Eugene F Fama. Term premiums and default premiums in money markets. *Journal of Financial Economics*, 17(1):175–196, 1986.

Eugene F Fama. Efficient capital markets: 2. *The Journal of Finance*, 46(5):1575–1617, 1991.

Eugene F Fama and Robert R Bliss. The information in long-maturity forward rates. *The American Economic Review*, pages 680–692, 1987.

Eugene F Fama and Kenneth R French. The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465, 1992.

C. Fernandez, E. Ley, and M.F.J. Steel. Model uncertainty in cross-country growth regressions. *Journal of Applied Econometrics*, 16(5):563–576, 2001.

Jeff Fleming, Chris Kirby, and Barbara Ostdiek. The economic value of volatility timing. *The Journal of Finance*, 56(1):329–352, 2001.

Ana-Maria Fuertes, Joëlle Miffre, and Georgios Rallis. Tactical allocation in commodity futures markets: Combining momentum and term structure signals. *Journal of Banking & Finance*, 34 (10):2530–2548, 2010.

W. Fung and D.A. Hsieh. Hedge fund benchmarks: A risk-based approach. *Financial Analysts Journal*, pages 65–80, 2004.

Xavier Gabaix. Power laws in economics and finance. *Working paper*, 2008.

Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700, 2012.

Xavier Gabaix, Arvind Krishnamurthy, and Olivier Vigneron. Limits of arbitrage: Theory and evidence from the mortgage-backed securities market. *The Journal of Finance*, 62(2):557–595, 2007.

Antonio Gargano and Allen Timmermann. Predictive dynamics in commodity prices. *Working paper*, 2013.

Nicolae B Garleanu and Lasse H Pedersen. Dynamic trading with predictable returns and transaction costs. *Journal of Finance*, 2009.

Gary B Gorton and K Geert Rouwenhorst. Facts and fantasies about commodity futures. Technical report, 2004.

Gary B Gorton, Fumio Hayashi, and K Geert Rouwenhorst. The fundamentals of commodity futures returns. Technical report, 2007.

N. Gospodinov and S. Ng. Commodity prices, convenience yields and inflation. *The Review of Economics and Statistics*, (0), 2011.

Amit Goyal and Ivo Welch. Predicting the equity premium with dividend ratios. *Management Science*, 49(5):639–654, 2003.

Mark Grinblatt and Tobias J Moskowitz. Predicting stock price movements from past returns: The role of consistency and tax-loss selling. *Journal of Financial Economics*, 71(3):541–579, 2004.

Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, pages 1029–1054, 1982.

Lars Peter Hansen and Robert J Hodrick. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *The Journal of Political Economy*, pages 829–853, 1980.

Zhigu He and Arvind Krishnamurthy. A model of capital and crises. *The Review of Economic Studies*, 79(2):735–777, 2012.

John R Hicks. The foundations of welfare economics. *The Economic Journal*, 49(196):696–712, 1939.

David Hirshleifer. Residual risk, trading costs, and commodity futures risk premia. *Review of Financial Studies*, 1(2):173–193, 1988.

H. Hong and J.C. Stein. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 54(6):2143–2184, 1999.

H. Hong and M. Yogo. What does futures market interest tell us about the macroeconomy and asset prices? *Journal of Financial Economics*, 2013.

Harrison Hong, Terence Lim, and Jeremy C. Stein. Bad news travels slowly: Size, analyst coverage and the profitability of momentum strategies. *The Journal of Finance*, 55, 2000.

Harrison Hong, Walter Torous, and Rossen Valkanov. Do industries lead stock markets? *Journal of Financial Economics*, 83(2):367–396, 2007.

Atsushi Inoue and Lutz Kilian. On the selection of forecasting models. *Journal of Econometrics*, 130(2):273–306, 2006.

Scott H Irwin, Dwight R Sanders, Robert P Merrin, et al. Devil or angel? the role of speculation in the recent commodity price boom (and bust). *Journal of Agricultural and Applied Economics*, 41(2):377–391, 2009.

N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1992.

Narasimhan Jegadeesh and Sheridan Titman. Cross-sectional and time-series determinants of momentum returns. *Review of Financial Studies*, 15(1):143–157, 2002.

Michael C Jensen. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6(2):95–101, 1978.

J Dave Jobson and Bob M Korkie. Performance hypothesis testing with the sharpe and treynor measures. *The Journal of Finance*, 36(4):889–908, 1981.

P. Jylha, K. Rinne, and M. Suominen. Do hedge funds supply or demand immediacy? 2010.

Petri Jylhä and Matti Suominen. Speculative capital and currency carry trades. *Journal of Financial Economics*, 99(1):60–75, 2011.

N. Kaldor. Speculation and economic stability. *The Review of Economic Studies*, 7(1):1–27, 1939.

J.M. Keynes. Treatise on money: Pure theory of money vol. i. *Macmillan, London*, 1930.

Lutz Kilian. The economic effects of energy price shocks. *Working Paper*, 2007.

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *The American Economic Review*, 99(3):1053–1069, 2009.

Lutz Kilian and Logan T Lewis. Does the fed respond to oil price shocks? *The Economic Journal*, 121(555):1047–1072, 2011.

R. Koijen, T. Moskowitz, L. Pedersen, and E. Vrugt. Carry. 20012.

Robert A Korajczyk and Ronnie Sadka. Are momentum profits robust to trading costs? *The Journal of Finance*, 59(3):1039–1082, 2004.

Robert Kosowski, Narayan Y Naik, and Melvyn Teo. Do hedge funds deliver alpha? a bayesian and bootstrap analysis. *Journal of Financial Economics*, 84(1):229–264, 2007.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.

David A Lesmond, Michael J Schill, and Chunsheng Zhou. The illusory nature of momentum profits. *Journal of Financial Economics*, 71(2):349–380, 2004.

S.C. Ludvigson and S. Ng. Macro factors in bond risk premia. *Review of Financial Studies*, 22(12): 5027–5067, 2009.

H. Lustig and A. Verdelhan. The cross-section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 2008.

H. Lustig, N. Roussanov, and A. Verdelhan. Countercyclical currency risk premia. *Journal of Financial Economics*, 2013.

Hanno N Lustig and Stijn G Van Nieuwerburgh. Housing collateral, consumption insurance, and risk premia: An empirical perspective. *The Journal of Finance*, 60(3):1167–1219, 2005.

Loriano Mancini, Angelo Ranaldo, and Jan Wrampelmeyer. Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *The Journal of Finance*, 2013.

Wessel Marquering and Marno Verbeek. The economic value of predicting stock index returns and volatility. *Journal of Financial and Quantitative Analysis*, 39(02):407–429, 2004.

Michael McCracken and Giorgio Valente. Testing the economic value of asset return predictability. *Working paper*, 2012.

Loukia Meligkotsidou and Ioannis D Vrontos. Detecting structural breaks and identifying risk factors in hedge fund returns: A bayesian approach. *Journal of Banking & Finance*, 32(11): 2471–2481, 2008.

Christoph Memmel. Performance hypothesis testing with the sharpe ratio. *Finance Letters*, 1(1), 2003.

L. Menkhoff, L. Sarno, M. Schmeling, and A. Schrimpf. Carry trades and global foreign exchange volatility. 2011.

L. Menkhoff, L. Sarno, M. Schmeling, and A. Schrimpf. Currency momentum strategies. 2012.

Joëlle Miffre and Georgios Rallis. Momentum strategies in commodity futures markets. *Journal of Banking & Finance*, 31(6):1863–1886, 2007.

T. Moskowitz, Y.H. Ooi, and L.H. Pedersen. Time series momentum. *Journal of Financial Economics*, 2011.

Tobias J Moskowitz and Mark Grinblatt. Do industries explain momentum? *The Journal of Finance*, 54(4):1249–1290, 1999.

Narayan Y Naik, Tarun Ramadorai, and Maria Stromqvist. Capacity constraints and hedge fund strategy returns. *European Financial Management*, 13(2):239–256, 2007.

Charles R Nelson and Myung J Kim. Predictable stock returns: The role of small sample bias. *The Journal of Finance*, 48(2):641–661, 1993.

Whitney K Newey and Kenneth D West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, pages 703–708, 1987.

Lubos Pastor and Robert F Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 2002.

S.F. Richard and M. Sundaresan. A continuous time equilibrium model of forward prices and futures prices in a multigood economy. *Journal of Financial Economics*, 9(4):347–371, 1981.

B.R. Routledge, D.J. Seppi, and C.S. Spatt. Equilibrium forward curves for commodities. *The Journal of Finance*, 55(3):1297–1338, 2002.

K Geert Rouwenhorst and Ke Tang. Commodity investing. *Working paper*, 2012.

Ronnie Sadka. Liquidity risk and the cross-section of hedge-fund returns. *Journal of Financial Economics*, 98(1):54–71, 2010.

Lucio Sarno, Paul Schneider, and Christian Wagner. Properties of foreign exchange risk premiums. *Journal of Financial Economics*, 2012.

Robert J Shiller. Rational expectations and the dynamic structure of macroeconomic models: A critical review. *Journal of Monetary Economics*, 4(1):1–44, 1978.

K.J. Singleton. Investor flows and the 2008 boom/bust in oil prices. *Management Science*, 2013.

R.F. Stambaugh, J. Yu, and Y. Yuan. The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, 2011.

Robert F Stambaugh. Predictive regressions. *Journal of Financial Economics*, 54(3):375–421, 1999.

Lawrence H Summers. Does the stock market rationally reflect fundamental values? *The Journal of Finance*, 41(3):591–601, 1986.

M. Szymanowska, F. De Roon, T. Nijman, and R. Van den Goorbergh. An anatomy of commodity futures risk premia. *Journal of Finance*, 2013.

L.G. Telser. Futures trading and the storage of cotton and wheat. *The Journal of Political Economy*, 66(3):233–255, 1958.

Daniel L Thornton and Giorgio Valente. Out-of-sample predictions of bond excess returns and forward rates: An asset allocation perspective. *Review of Financial Studies*, 25(10):3141–3168, 2012.

Allan G Timmermann. How learning in financial markets generates excess volatility and predictability in stock prices. *The Quarterly Journal of Economics*, 108(4):1135–1145, 1993.

Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. 2009.

Ivo Welch and Amit Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008.

H. Working. The theory of price of storage. *The American Economic Review*, 39(6):1254–1262, 1949.

J.H. Wright. Bayesian model averaging and exchange rate forecasts. *Journal of Econometrics*, 146 (2):329–341, 2008.