



City Research Online

City St George's, University of London

Citation: Piana, J. (2017). Expectations, fundamentals, and asset returns: evidence from the commodity markets. (Unpublished Doctoral thesis, Cass Business School)

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/19656/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Expectations, Fundamentals, and Asset Returns

Evidence from the Commodity Markets



Cass Business School
CITY UNIVERSITY LONDON

Jacopo Piana

Cass Business School, Department of Finance
City University London

Thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

November 2017

To my loving parents

Table of contents

List of figures	ix
List of tables	xi
Introduction	1
1 Returns Expectations: Evidence from Commodity Survey Price Forecasts	11
1.1 Introduction	11
1.2 Data and Preliminaries	15
1.2.1 Survey Data	16
1.2.2 Futures Prices, Fundamentals and Trading Flows	18
1.3 Results	21
1.3.1 Survey Expectations and the Futures Curve	21
1.3.2 Predictive Regressions	23
1.3.3 Survey Forecasts, Time-Series Momentum and Value	26
1.3.4 Survey Forecasts and COT Positions	29
1.3.5 Forecasts Dispersion and Options Implied Volatility	30
1.4 Interpretation	33
1.4.1 Predictability	35
1.4.2 Determinants of Expectations of Returns: Momentum and Value	36
1.5 Concluding Remarks	38
Appendix A Appendix and Tables	41
A.1 Model Description	41
A.1.1 Rational Expectations Traders	42
A.1.2 Feedback Traders	42
A.1.3 Market Clearing	43
A.2 Tables	45

2	Expected Spot Prices and the Dynamics of Commodity Risk Premia	77
2.1	Introduction	77
2.2	Motivation	80
2.3	Learning and Expectations	83
2.3.1	Comparison with Survey Expectations	86
2.4	Empirical Analysis	90
2.4.1	Data	91
2.4.2	Dissecting Ex-Ante Risk Premia	93
2.5	Concluding Remarks	104
	Appendix B Appendix and Tables	105
B.1	A Simple Model of Learning	105
B.1.1	Introducing a Futures Market	106
B.1.2	Learning Dynamics	108
B.2	Econometric Design	109
B.2.1	Step 1. Sampling the Conditional Factor Sensitivities $\theta_{1:T}$	111
B.2.2	Step 2. Sampling the State Variance Parameters W	112
B.2.3	Step 3. Sampling the Idiosyncratic Risk H	112
B.3	Forecasts and Learning	113
3	From Prices to Fundamentals: A Round-Trip	129
3.1	Introduction	129
3.2	Data	132
3.3	From Prices to Fundamentals	133
3.3.1	Identification	133
3.3.2	Filtering	134
3.3.3	Predicting Fundamental News	135
3.4	Returns Predictability	137
3.4.1	Simulated Trading Strategy	137
3.4.2	Returns Response to News	139
3.5	Simulation	141
3.5.1	Returns Response to News: Interpretation	141
3.5.2	Look-Back Period	142
3.5.3	Portfolio Size	143
3.6	Concluding Remarks	144

Appendix C	Appendix and Tables	145
C.1	Model	145
C.1.1	Spot Market	146
C.1.2	Futures Market	147
C.2	Tables and Figures	150
C.3	Spot Prices Data	162
C.4	Look-Back Period and Signal Extraction	164
4	Concluding Remarks	165
	References	169

List of figures

A.1	Agreement	49
A.2	Bloomberg Survey Forecasts: Crude Oil (WTI)	50
A.3	Consensus Survey Forecasts: Crude Oil (WTI)	51
A.4	Bloomberg Survey Forecasts: Copper	52
A.5	Consensus Survey Forecasts: Copper	53
A.6	Bloomberg Forecasts Dispersion: WTI	64
A.7	Consensus Forecasts Dispersion: WTI	65
A.8	Implied Volatility and Bloomberg Forecasts Dispersion: WTI	66
A.9	Implied Volatility and Consensus Forecasts Dispersion: WTI	67
B.1	Expectations Errors for Future Spot Prices	114
B.2	Ex-Ante vs Realised Risk Premia	117
B.3	Spot Prices and Aggregate Demand	118
B.4	Model-Implied vs Survey Expectations ($h = 2$)	119
B.5	Model-Implied vs Survey Expectations ($h = 4$)	120
B.6	R^2 Decomposition	121
B.7	Time-Varying Betas for WTI Crude Oil	122
B.8	Time-Varying Betas for Copper	123
B.9	Time-Varying Betas for Corn	124
B.10	Time-Varying Betas for Silver	125
B.11	Ex-Ante vs Realised Risk Premia (Horizon $h = 2$)	126
B.12	Ex-Ante vs Realised Risk Premia (Horizon $h = 4$)	127
C.1	Filtered Fundamentals vs. Actuals	150
C.2	Filtered - Actual Fundamentals Correlation	154
C.3	1st Principal Component Weights	154
C.4	Trading Strategy Performance	155
C.5	Trading Strategy Statistical Significance	156

C.6	Trading Strategy and Lagged Prices	157
C.7	Returns Response to News about Fundamentals	158
C.8	Delayed Information and Trading Performance	158
C.9	Simulated Returns Response to News	159
C.10	Signal Extraction and Look-Back Period	160
C.11	Signal Extraction and Portfolio Size	161

List of tables

A.1	Summary Statistics	45
A.2	Bloomberg's Survey Participants	46
A.3	Multi-Horizon Expectations of Returns and Futures Curve	47
A.4	Correlations Between Surveys	48
A.5	Predictive Regressions	54
A.6	Inventories Predictive Regressions	55
A.7	Determinants of Analysts' Expectations	56
A.8	Survey Forecasts and Net Long Positions	60
A.9	Crude Oil WTI Implied Volatility and Forecasts Dispersion	68
A.10	Implied and Realized Volatility Spread and Forecasts Dispersion	69
A.11	Simulation: Predictive Regressions - Rational Traders	70
A.12	Simulation: Predictive Regressions - Feedback Traders	71
A.13	Simulation: Expectations of Returns and Past Cumulated Returns	72
A.14	Determinants of Analysts' Expectations: Look-Back	73
A.15	Determinants of Analysts' Expectations: Perfect Foresight	74
A.16	Determinants of Analysts' Expectations: Skipping Most Recent Returns	75
B.1	Summary Statistics	115
B.2	Static Regression Analysis	116
B.3	Future Spot Prices Out-of-Sample Forecasting Comparison	125
B.4	Forecasts and Learning	128
C.1	Summary Statistics	151
C.2	1st Principal Component Loadings on Commodity Returns	152
C.3	Out-of-Sample Fundamentals Forecasts	153

Acknowledgements

I am hugely indebted to my supervisor Alessandro Beber and Daniele Bianchi for their constant advice, help and mentorship. Furthermore, I am grateful to my parents and friends. This long journey would have been impossible without their continuous support and patience. A special thank goes to my PhD classmates for their help, cooperation, insightful comments and for the wonderful time spent together. I would also like to thank Jules van Binsbergen, Michael Brandt, Max Bruche, Giovanni Cespa, Carlo A. Favero, Itay Goldstein, Massimo Guidolin, Giorgio S. Questa, Richard Payne, Tarun Ramadorai, Lucio Sarno, Maik Schmeling and Paolo Volpin for their helpful comments and suggestions. I finally thank Cass Business School for the generous financial support, the Wharton School and Bocconi University for hosting me and their faculty members and PhD students for their numerous helpful comments.

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. I further declare that whenever the results of this thesis are used for co-authored research papers this is clearly stated in the document. Powers of discretion are hereby granted to the University Librarian to allow this thesis to be copied in whole or in part without further reference to the author. This permission covers only simple copies made for study purpose, subject to the normal conditions of acknowledgement.

Jacopo Piana
November 2017

Abstract

This Thesis contributes to the study of the links between expectations, fundamentals, and asset returns using the rich empirical setup offered by commodity markets. The three Chapters, constituting this work, analyse empirically how expectations are formed and what are the implications of departures from perfect rational expectations on returns predictability. Monte-Carlo experiments are also used to rationalise and to give an economic interpretation to the empirical findings. In the First chapter, we show that survey-based expectations of returns are strongly correlated with past price variations, but not with fundamentals and can be largely explained by time-series momentum and value factors. Furthermore, we find that expectations have positive, but not significant correlation with future realised returns, which implies little predictive power. Using a Monte-Carlo experiment, we show that both results can also be generated by rational individuals provided the existence of extrapolative momentum traders and little predictability of fundamentals. Finally, our analysis also suggests that survey-based expectations can have a crucial role to understand better the dynamics of trading flows and the drivers of the option implied volatility risk premium. In the second Chapter, we investigate the dynamics of the ex-ante risk premia for different commodities and maturities through the lens of a model of rational learning in which expected future spot prices are revised in line with past prediction errors and changes in aggregate economic growth. The main results show that time-varying risk premia are predominantly driven by market activity and financial risks. More generally, we provide evidence of heterogeneity in the dynamics of factor loadings, both across commodities and time horizons. Finally, we show that the model-implied expectations are consistent with the cross-sectional average of Bloomberg professional analysts' forecasts. In the third Chapter, we exploit the peculiarities of commodity markets to show that fundamental news about global growth is reflected into prices, but not instantaneously. News on economic activity can be filtered in real-time from commodity prices, but such news takes several months before being fully incorporated into prices, leading to returns predictability. Coherently with the theories of overreaction and underreaction to news, we show using simulated data that the results obtained can be explained by the existence of latecomers, who process information with a delay, and momentum traders.

Introduction

This Thesis contributes to the study of the links between expectations, fundamentals, and asset returns using the rich empirical setup offered by commodity markets. The three Chapters, constituting this work, analyse empirically how expectations are formed and what are the implications of departures from perfect rational expectations on returns predictability.

Expectations about future returns are the key building block of any asset pricing model, and not surprisingly an important part of the asset pricing literature focuses on their empirical analysis. Unfortunately, expectations are not directly observable and must be approximated somehow. In fact, when no data measuring directly investors' expectations are available, the only viable option is estimating them by making assumptions on how expectations are formed. However, this approach has the not negligible shortcoming that the obtained approximations are model dependent, i.e. the estimated expectations of returns depend on modelling assumptions.¹ Conversely, when available, survey data on returns expectations represent an appealing way to circumvent such problem as they provide a direct approximation without relying on any specific assumption on how individuals shape their views about future prices and therefore returns. Leveraging this advantage, in the first Chapter, we contribute to the empirical study of returns expectations by analysing an unexplored dataset on commodity price predictions reported by professional forecasters to Bloomberg and Consensus Economics. More specifically, we study the predictive power of survey-based price forecasts, their determinants and relationship with options implied volatility and trading flows for the

¹In fact, this method consists, first in deriving a theoretical proxy for expectations by making assumptions on how they are formed, and then in making further assumptions to estimate them empirically.

two most traded industrial commodities, namely: crude oil and copper.

Survey data on commodities are not only interesting because they have not been analysed in this context so far, but also because the peculiarities of the data allow us to expand the existing empirical analysis of survey-based expectations of returns over two new dimensions. First, the existence of survey forecasts data at multiple forecasting horizons enables us to study how the properties of expectations and their predictive power change as the forecasting horizon changes. Second, the fact that the target of the forecast is the price of the underlying of derivatives futures contracts, makes possible the study of the relationship between expectations of returns and the futures curve, which is key in the pricing of commodity derivatives.² In addition, an appealing feature of the data is that the surveyed population includes only professional forecasters specialised in the commodity market. This peculiarity, combined with the clarity and quantitative nature of the questions posed, reduces decisively the room for the measurement error problem, which is the main critique raised against survey data.

As a first result of our empirical analysis, we find by estimating predictive regressions that survey price forecasts contain little predictive power across forecasting horizons for both crude oil and copper. Conversely, we find that expectations of returns, consistently with the theory's predictions stating that an expected increase in the spot price of a commodity should induce an increase in inventory levels (Pindyck [83]), tend to predict significantly future inventory levels.

From the correlation analysis it emerges that that price forecasts are highly correlated across different forecasting horizons and data providers, which suggests the existence of common drivers that can explain how forecasts are formed. In this respect, in order to understand how analysts' price predictions are generated, we follow Greenwood and Shleifer [49] by regressing survey expectations of returns on fundamentals factors; time-series momentum and a proxy for value. The inclusion of time series momentum and the proxy for value

²The risk premium hypothesis states that the price of a futures contract is equal to the expected future spot price plus a risk premium. See, for example, Cootner [26], Fama and French [41], Dusak [35].

is motivated by the possibility that analysts might form their expectations by looking at historical price trends and levels. In this regard, in the last part of the first Chapter we show that in a model with heterogeneous beliefs, where we allow for the presence of extrapolative traders (feedback / momentum traders) and mean reversion in fundamentals, time-series momentum and value can also affect the expectations of rational investors.³

In fact, the empirical analysis of the determinants of analysts' expectations reveals three main results. First, time-series momentum and value factors explain a large part of forecasts variation (around 70% of the time-series variation in the case of crude oil and around 60% in the case of copper). Second, when we include the fundamentals factors in the regressions, they add little explanatory power and they do not eliminate the significance of the momentum and value factors. Third, the momentum and value factors are negatively correlated with survey-based expectations of returns across forecasting horizons for both crude oil and copper, meaning that analysts tend on average to expect negative (positive) price changes when prices are relatively high (low) and have been increasing (decreasing) in the recent past. On this point, using simulated data from the model with heterogeneous beliefs, we show that the correlations between past returns and survey-based expectations of returns do not necessarily imply that on average professional forecasters are not rational. Conversely, we show that such correlations can be obtained also in rational expectations when we assume both mean reversion in fundamentals and the presence of momentum traders that push prices away from fundamentals by buying in periods of past positive trends and selling when prices have fallen. Furthermore, the experiment shows that the predictive power of rational forecasts purely depends on the predictability of the fundamentals and the asset demand of other types of agents.

Moving to the study of the relationship between expectations of returns and the futures curve, we find that analysts' expectations are highly correlated with the ones extrapolated from the slope of the futures curve and that such correlation tend to increase with the forecast-

³In the first Chapter, we use the terms extrapolative, momentum and feedback traders interchangeably.

ing horizon. This empirical finding has two important implications. First, that futures prices on crude oil and copper can be considered a proxy for analysts' average price expectations. Second, this result provides evidence consistent with the hypothesis that commodity futures prices should equal expectations of future spot prices plus a risk premium (e.g. Cootner [26], Dusak [35] and Fama and French [41])

Using trading flows data from the Commodity Futures Trading Commission (CFTC), we analyse the relationship between futures net positions of hedgers and speculators and returns expectations. By regressing changes in net long positions on expectations of returns, we find that on average hedgers' variation in long position tends to correlate positively with expectations of returns. Conversely, large speculator trading activity is negatively correlated with analysts' expectations of returns. These empirical findings, in light of the the negative relationship between past returns (momentum) and expectations of returns, are therefore consistent with the results presented by Kang et al. [64] and Moskowitz et al. [77] showing that speculators follow momentum strategies, whereas market makers and hedgers accommodate their short term liquidity needs by trading as contrarians.

As a last step, we exploit the availability in the dataset of the cross-section of analysts' forecasts to study the relationship between analysts' forecasts dispersion and options implied volatility. Buraschi and Jiltsov [18] and Beber et al. [11] point to the importance of considering heterogeneity in investors returns expectations to learn more about asset prices by showing that differences in analysts' beliefs about future returns have an impact on option implied volatility. We show the importance of differences in returns expectations by showing that these findings do not break down also in a setting where price variations can be explained more easily by fundamentals shocks and where investors might rely less on analysts' opinions as a source of information *per se*.

In terms of contribution, the first Chapter enriches the strand of literature on asset pricing and survey forecasts by extending the studies of Koiijen et al. [70]; Greenwood and Shleifer

[49] and Beber et al. [11] to the commodity markets and by adding the analysis on multiple forecasting horizons. Most importantly, our findings shed some light on the puzzling evidence of contrarian predictive power of survey forecasts and positive correlation with past returns found by Greenwood and Shleifer [50], by showing that when a population of professional forecasters is interviewed, the contrarian predictive power disappears and the sign of the correlation with past returns changes. Our simulation experiment also shows that both results can be obtained depending on which of the two populations, rational or extrapolative, is surveyed. Therefore, we hypothesize that our and the previous empirical evidence is driven by the population being interviewed, with professional forecasters being close to rational expectations agents and the non-professional forecasters close to the extrapolative type.

If in the first Chapter, we have shown that the dependence of expectations on past prices, can also arise under the limiting case of perfect rational expectations when extrapolative agents are present in the market; in the second Chapter we model price expectations to study the dynamics of commodity risk premia, by still taking into account past price dynamics but inside a more realistic framework of rational learning. More specifically, we assume that investors do form their expectations rationally according to a rational Muth's model, but they learn adaptively about the coefficients of the prices perceived law of motion from past prediction errors and as new fundamental data and prices become available.

Such model of learning, allows us to approximate the time-varying *ex-ante* risk premia – calculated as the spread between the futures price as of date t with maturity $t + h$ and expectations at time t on future spot prices over the same time-horizon – for a reasonably long sample period. Once the dynamics of ex-ante risk premia is extracted from futures prices, we investigate their determinants across investment horizons and commodities by using a dynamic linear regression framework, which features random-walk betas on a set of widely discussed economic risk factors. This approach allows us to shed some light on the

dynamics and determinants of risk premia, that have been controversial in the literature.⁴

Our main results show that risk premia are time-varying, both across commodities and time-horizons, and their dynamics is predominantly driven by risks sharing mechanisms and the changing nature of market activity, as proxied by Open Interest (OI henceforth), Hedging Pressure (HP henceforth) and time-series Momentum (TSMOM henceforth). These results hold after controlling for a variety of other commonly used proxies for risk factors, e.g. changes in inventories and realised volatility. Yet, we show that emerging markets, as proxied by the MSCI Emerging Market Index (MXEF), plays a sensible role for both WTI Oil and Copper, which is coherent with the increasing weight of emerging economies in the global economic growth and the presence of potential spillover effects to be associated with concerns about a worldwide economic slowdown.⁵ More generally, we provide evidence of heterogeneity in the dynamics of factor loadings in the time-series of commodity risk premia across both products and maturities.

Also, we compare the expected future spot prices obtained from our model with the cross-sectional average of survey forecasts provided by Bloomberg. In this respect, we show that, although with differences across commodities, our model generates conditional expectations which are broadly consistent with the average survey forecast from two to four quarters ahead.

Finally, we show that our model, in which individuals learn over time about the prices law of motion, compares favourably against alternative specifications for forecasting future spot prices. More precisely, an out-of-sample comparison of mean squared prediction errors against models in which expectations are based on either futures or current spot prices, or a spread of the two, demonstrates that the forecasts generated by our model reach a statistically

⁴See, e.g. Keynes [66], Hicks [55], Kaldor [61], Working [98], Brennan [17], Hsieh and Kulatilaka [60], Fama and French [40], Fama and French [42], Gorton et al. [47], Singleton [89], Szymanowska et al. [93] and Bakshi et al. [8] just to cite a few. The controversy stems from the fact that investors' expectations are not directly observable.

⁵China itself is the second largest economy and the second largest importer of both goods and commercial services.

significance 1% higher predictive R^2 on average across commodities and maturities. This result, possibly, rules out the concern that the model-implied ex-ante risk premia merely represent forecast errors which have nothing to do with investors' preferences or the actual expectations formation process. As a matter of fact, a further analysis clarifies that the expected payoffs extracted in this Chapter are highly correlated to the actual, realised, excess rolling returns in the same-maturity generic futures contract.

In terms of contribution, the second Chapter builds on a number of existing works such as Nerlove [80], Evans and Honkapohja [39], Sargent [87], Sargent and Williams [88], and Malmendier and Nagel [72], who consider a model of adaptive learning to explain the dynamics of expectations on inflation and more general macroeconomic outcomes. Also, our study is related to recent research that posits trading activity is the result of a learning process in which hedgers and speculators update their views as news about economic fundamentals, and prices become available (see, e.g. Singleton [89]). Finally, the second Chapter contributes also to the recent literature that aims at understanding the origins of unconditional realised commodity risk premia such as Carter et al. [22], Bessembinder [15], De Roon et al. [29], Acharya et al. [1], Hong and Yogo [59], Asness et al. [5], Basu and Miffre [10], Hamilton and Wu [53], Szymanowska et al. [93], and Bakshi et al. [8].

In the third Chapter we investigate the existence of delays in the processing of fundamental news as a possible cause of the persistence in asset returns and the time-series momentum anomaly (see Moskowitz et al. [78]).

Over the last decades, different theories have been proposed to explain the puzzling empirical evidence of the persistence in asset returns over short horizons (up to twelve months) and reversal at longer horizons. The seminal works by Cutler et al. [27] and Hong and Stein [58] propose the existence of delays in the processing of fundamental news by rational traders and the activity of momentum traders as the main explanation of this phenomenon.

In their framework, the persistence in returns is caused by an initial underreaction of prices to fundamental news and an overreaction induced by momentum traders. Unfortunately, it is not easy to test directly these predictions empirically by studying the correlation between news and returns.⁶ This complexity stems from the fact that it is difficult, not only to agree on what can be considered a stable fundamental pricing factor for a given asset class, but also to agree on what can be considered a news.⁷ Even if these obstacles can be overcome, it remains complicated to find a time-series of fundamental news sufficiently long to perform a meaningful analysis.

To understand better the causes of this phenomenon, in this Chapter, we exploit the peculiarities of commodity markets to test this mechanism more directly and in an innovative way. In fact, our approach does not consist in collecting news on supply or demand to test if this is reflected into prices, but we do the opposite. We devise an identification strategy, which extracts news from prices of 48 different commodities, and then we test if such information can predict future official statistical publications on industrial production, which is a highly regarded indicator of global demand growth. In this way, we first test if prices carry information about fundamentals, and then we use the filtered news to study the speed at which such information is reflected into different speculative assets.

The news identification strategy is based on the very simple assumptions that commodity prices increase (decrease) across the board after a positive (negative) news about global growth, but supply shocks are uncorrelated and affect a single commodity or substitutes.⁸ For instance, an unexpected acceleration in China's economic activity will increase commodities

⁶When it is not possible to study directly the correlation between news and returns, the empirical analysis must rely on indirect effects. For instance studying the autocorrelation of returns itself is an indirect way to test for the possibility that news are not incorporated instantaneously into prices. Unfortunately, such indirect test cannot rule out the possibility of alternative explanations to the phenomenon.

⁷The stability in terms of the importance of the fundamental pricing factor is crucial for time-series analysis. If the pricing factors follow switching regimes, the analysis of the relationship between news and returns becomes much more difficult in small samples.

⁸The role of news about global growth on commodity demand is undisputed. Deaton and Laroque [34] explicitly model the effect of global income dynamics on commodity demand and show how it generates co-movements in commodity prices. The role of fundamental factors in commodity markets has been analysed by: Alquist and Kilian [4]; Bakshi et al. [9]; Bhardwaj et al. [16]; Brennan [17]; Deaton and Laroque [31];

demand, which in turn will lead to an increase in commodity prices (with price variations that depend on each commodity sensitivity to economic growth). To the contrary, a drought in South America will affect agricultural commodity prices, but not metals and energy commodities. Given these assumptions, we show that price co-movements are caused by news about global growth, and that news can be filtered using Principal Component Analysis (PCA).⁹

Our empirical analysis shows that prices do reflect fundamental information about global growth. The news extracted from prices correlates well with the OECD publications on industrial production and the level of correlation changes with the basket of commodities used as predicted by the theoretical model we use to guide us in the empirical analysis. Furthermore, we show that prices incorporate also non-publicly available information by showing that the news extracted from prices helps in predicting the official statistical publications that are usually released with three to six months of delay. In fact, in a out-of-sample exercise, the inclusion of the filtered news in the forecasting model reduces the RMSE by 18% compared to a benchmark forecasting model.

In Section 3.4 we investigate the existence of delays in the information processing as the possible cause of the persistence and predictability of returns by studying if the news extracted from commodity prices are incorporated instantaneously into different speculative assets. To this end, we study the delay in the reaction of returns to fundamental news.

Our findings show that such news are reflected into speculative assets but not instantaneously and with a delay of several months. More specifically, we demonstrate the existence of a delay in information processing, by showing that the correlation between news and asset returns is positive for several periods and goes gradually to zero to become negative only after fifteen months. Furthermore, to quantify the economic significance of such delay in

Deaton and Laroque [32]; Deaton and Laroque [34]; Fama and French [40]; Gorton et al. [47]; Hamilton [52]; Kilian and Hicks [67]; Szymanowska et al. [93]; Turnovsky [95]; Working [98]

⁹See Section 3.3.1 for a detailed discussion of the assumptions.

the information processing, we simulate an out-of-sample trading strategy in which we use news with different delays as a trading signal. The results show that also in economic terms the absorption of news about global growth into prices is rather slow, leading to positive excess returns of 12% after one month, and yet 5% six months after the news has occurred.¹⁰ Furthermore, similar, but negative performances obtained with a delay between fifteen and eighteen months support the existence of positive feedback traders causing overreaction to news as proposed by Cutler et al. [27] and by Hong and Stein [58].

In the last part of the Chapter, we contribute to the literature by rationalising our empirical findings in light of the theories of prices under and overreaction to news by analysing the data generated from a modified version of the Cutler et al. [27] model. The model, which features fully rational traders, latecomers (who are rational but observe fundamental information with a delay), and positive feedback traders (who engage in momentum strategies) allows us to study quantitatively how the interaction among the different types of traders can alter the speed of information processing leading to returns persistence and the time-series momentum anomaly. The results of the Monte-Carlo simulation exercise show that the predictability found in the data is consistent with a delay in the information processing generated by the presence of rationally-bounded traders. More specifically we find that our empirical results can be obtained when around 70% of the traders process news with a delay of several periods. Finally, we show that also momentum traders, but to a lesser extent (20%), are necessary to generate the over-reaction mechanism to news we have found in the data.

¹⁰This refers to the case of commodity futures.

Chapter 1

Returns Expectations: Evidence from Commodity Survey Price Forecasts¹

1.1 Introduction

Expectations about future returns are the key building block of any asset pricing model, and not surprisingly an important part of the asset pricing literature focuses on their empirical analysis. Unfortunately, expectations are not directly observable and must be approximated somehow. In fact, when no data measuring directly investors' expectations are available, the only viable option is estimating them by making assumptions on how expectations are formed. However, this approach has the not negligible shortcoming that the obtained approximations are model dependent, i.e. the estimated expectations of returns depend on modelling assumptions.² Conversely, when available, survey data on returns expectations represent an appealing way to circumvent such problem as they provide a direct approximation without relying on any specific assumption on how individuals shape their views about future prices and therefore returns. Leveraging this advantage, in this Chapter, we contribute to the empirical study of returns expectations by analysing an unexplored dataset on commodity price predic-

¹A research paper joint with my PhD supervisor, Prof. A. Beber, entitled "Expectations, Fundamentals, and Asset Returns: Evidence from the Commodity Markets" is based on this Chapter. The paper has been presented at the 2016 Cass Research Days and the World Finance Conference at Cagliari 2017.

²In fact, this method consists, first in deriving a theoretical proxy for expectations by making assumptions on how they are formed, and then in making further assumptions to estimate them empirically.

tions reported by professional forecasters to Bloomberg and Consensus Economics. More specifically, we study the predictive power of survey-based price forecasts, their determinants and relationship with options implied volatility and trading flows for the two most traded industrial commodities, namely: crude oil and copper.

Survey data on commodities are not only interesting because they have not been analysed in this context so far, but also because the peculiarities of the data allow us to expand the existing empirical analysis of survey-based expectations of returns over two new dimensions. First, the existence of survey forecasts data at multiple forecasting horizons enables us to study how the properties of expectations and their predictive power change as the forecasting horizon changes. Second, the fact that the target of the forecast is the price of the underlying of derivatives futures contracts, makes possible the study of the relationship between expectations of returns and the futures curve, which is key in the pricing of commodity derivatives.³ In addition, an appealing feature of the data is that the surveyed population includes only professional forecasters specialised in the commodity market. This peculiarity, combined with the clarity and quantitative nature of the questions posed, reduces decisively the room for the measurement error problem, which is the main critique raised against survey data.

As a first result of our empirical analysis, we find by estimating predictive regressions that survey price forecasts contain little predictive power across forecasting horizons for both crude oil and copper. Conversely, we find that expectations of returns, consistently with the theory's predictions stating that an expected increase in the spot price of a commodity should induce an increase in inventory levels (Pindyck [83]), tend to predict significantly future inventory levels.

From the correlation analysis it emerges that that price forecasts are highly correlated across different forecasting horizons and data providers, which suggests the existence of common drivers that can explain how forecasts are formed. In this respect, in order to under-

³The risk premium hypothesis states that the price of a futures contract is equal to the expected future spot price plus a risk premium. See, for example, Cootner [26], Fama and French [41], Dusak [35].

stand how analysts' price predictions are generated, we follow Greenwood and Shleifer [49] by regressing survey expectations of returns on fundamentals factors; time-series momentum and a proxy for value. The inclusion of time series momentum and the proxy for value is motivated by the possibility that analysts might form their expectations by looking at historical price trends and levels. In this regard, in the last part of this Chapter we show that in a model with heterogeneous beliefs, where we allow for the presence of extrapolative traders (feedback / momentum traders) and mean reversion in fundamentals, time-series momentum and value can also affect the expectations of rational investors.⁴

In fact, the empirical analysis of the determinants of analysts' expectations reveals three main results. First, time-series momentum and value factors explain a large part of forecasts variation (around 70% of the time-series variation in the case of crude oil and around 60% in the case of copper). Second, when we include the fundamentals factors in the regressions, they add little explanatory power and they do not eliminate the significance of the momentum and value factors. Third, the momentum and value factors are negatively correlated with survey-based expectations of returns across forecasting horizons for both crude oil and copper, meaning that analysts tend on average to expect negative (positive) price changes when prices are relatively high (low) and have been increasing (decreasing) in the recent past. On this point, using simulated data from the model with heterogeneous beliefs, we show that the correlations between past returns and survey-based expectations of returns do not necessarily imply that on average professional forecasters are not rational. Conversely, we show that such correlations can be obtained also in rational expectations when we assume both mean reversion in fundamentals and the presence of momentum traders that push prices away from fundamentals by buying in periods of past positive trends and selling when prices have fallen. Furthermore, the experiment shows that the predictive power of rational forecasts purely depends on the predictability of the fundamentals and the asset demand of other types of agents.

⁴In this Chapter, we use the terms extrapolative, momentum and feedback traders interchangeably.

Moving to the study of the relationship between expectations of returns and the futures curve, we find that analysts' expectations are highly correlated with the ones extrapolated from the slope of the futures curve and that such correlation tends to increase with the forecasting horizon. This empirical finding has two important implications. First, that futures prices on crude oil and copper can be considered a proxy for analysts' average price expectations. Second, this result provides evidence consistent with the hypothesis that commodity futures prices should equal expectations of future spot prices plus a risk premium (e.g. Cootner [26], Dusak [35] and Fama and French [41])

Using trading flows data from the Commodity Futures Trading Commission (CFTC), we analyse the relationship between futures net positions of hedgers and speculators and returns expectations. By regressing changes in net long positions on expectations of returns, we find that on average hedgers' variation in long position tends to correlate positively with expectations of returns. Conversely, large speculator trading activity is negatively correlated with analysts' expectations of returns. These empirical findings, in light of the negative relationship between past returns (momentum) and expectations of returns, are therefore consistent with the results presented by Kang et al. [64] and Moskowitz et al. [77] showing that speculators follow momentum strategies, whereas market makers and hedgers accommodate their short term liquidity needs by trading as contrarians.

As a last step, we exploit the availability in the dataset of the cross-section of analysts' forecasts to study the relationship between analysts' forecasts dispersion and options implied volatility. Buraschi and Jiltsov [18] and Beber et al. [11] point to the importance of considering heterogeneity in investors returns expectations to learn more about asset prices by showing that differences in analysts' beliefs about future returns have an impact on option implied volatility. We show the importance of differences in returns expectations by showing that these findings do not break down also in a setting where price variations can be explained more easily by fundamentals shocks and where investors might rely less on analysts' opinions

as a source of information *per se*.

In terms of contribution, this Chapter enriches the strand of literature on asset pricing and survey forecasts by extending the studies of Koijen et al. [70]; Greenwood and Shleifer [49] and Beber et al. [11] to the commodity markets and by adding the analysis on multiple forecasting horizons. Most importantly, our findings shed some light on the puzzling evidence of contrarian predictive power of survey forecasts and positive correlation with past returns found by Greenwood and Shleifer [50], by showing that when a population of professional forecasters is interviewed, the contrarian predictive power disappears and the sign of the correlation with past returns changes. Our simulation experiment also shows that both results can be obtained depending on which of the two populations, rational or extrapolative, is surveyed. Therefore, we hypothesize that our and the previous empirical evidence is driven by the population being interviewed, with professional forecasters being close to rational expectations agents and the non-professional forecasters close to the extrapolative type.

The rest of the Chapter is organised as follows. Section 1.2 describes the data on survey forecasts, futures prices, fundamentals factor and trading activity. Section 1.3 describes the empirical analysis and presents the results. In Section 1.4 we analyse simulated data from the heterogeneous beliefs model to give an economic interpretation to the empirical results obtained analysing actual data. Section 1.5 concludes. We leave the model derivation and further results to the Appendix.

1.2 Data and Preliminaries

In this Section we describe briefly the data sources we use in our analysis.

[Insert Table A.1 about here]

1.2.1 Survey Data

We obtain commodities survey price forecasts from the Bloomberg's commodity price forecasts database and from Consensus Economics. The Bloomberg's commodity price forecasts database contains analysts' price expectations at multiple quarterly forecasting horizons and across diverse commodities from 2006 to 2015. In this Chapter, we focus on crude oil and copper because they are the most traded consumption commodities with the most complete sample of survey data.

[Insert Table A.2 about here]

The survey respondents are commodity market analysts, mainly from banks and commodity consultancy companies. The survey is quantitative in nature as participants are asked to provide a point forecasts on the average quarterly commodity price for a specified futures contract. The database allows to retrieve for each analyst the historical price forecasts and the related publication date. Since the analysts' forecasts submissions are recorded daily and not evenly spaced in time, we sample the data at monthly frequency to reduce the difference in the market information-set available between early and late submitters. Then, to prepare the data for time-series analysis, we transform the discontinuous fixed-horizon quarterly price forecasts into continuous quasi-constant forecasting horizon price forecasts. More specifically, we first calculate the forecasting horizon in months for each forecast at each point in time. Then, we generate continuous quasi-constant horizon series by stacking the forecasts that belong to the following forecasting horizons (in months): $+2Q \in [4; 6]$; $+3Q \in [7; 9]$; $+4Q \in [10; 12]$.⁵ As a last step, we compute the cross-sectional average forecast as the simple mean forecast at each point in time by analyst and by forecasting horizon. Formally, we

⁵The forecasting horizon is calculated with respect to the end of the month of the last month in the quarter which is object of the forecast. For instance, the forecasting horizon of the forecasts for Q4 2014 recorded in June 2014 is 6 months. We discard the survey forecasts with forecasting horizon between 1 to 3 months as prices that contribute to the quarterly average are partially already observed.

define the average price survey forecasts as

$$\hat{S}_t^{(m)} = \sum_{i=1}^N \hat{S}_{i,t}^{(m)} / N \text{ with } m \in \{+2Q, +3Q, +4Q\} \quad (1.1)$$

where m is the forecasting horizon and N is the number of analysts submitting the forecast at time t given the forecasting horizon m .

The Consensus Economics survey price forecasts database covers energy and metals commodities since 1995.⁶ Similarly to Bloomberg, Consensus Economics collects commodity price forecasts from financial institutions and consultancy companies at the quarterly forecasting horizons. Differently from the Bloomberg's survey, in which forecasters can submit their forecasts at any time, the Consensus Economics survey procedure asks the participants to submit their forecasts within a surveying period of a few days. This aims to minimise the difference in the amount of information available between early and late submitters when they are asked to make their price predictions.

The sample period is longer than the Bloomberg's commodity price forecasts as the publication of the Consensus' forecasts started in 1995. However, the publications occurred at irregular frequency and no forecasts were published between August 2002 and March 2004, and during the third quarter of 2007. Before April 2012 the forecasts were published at quarterly frequency, after that date the publication frequency has become bi-monthly.

[Insert Table A.3 about here]

We follow the same procedure described above to generate quasi-constant forecasting horizons data series at monthly frequency. For the months where no publication was available, we do not make any interpolation, and treat those points as missing values.

[Insert Table A.4 about here]

⁶Also in this case we focus on WTI and Copper as they have the most complete sample of survey data.

As shown in Table A.3, survey-based expectations of returns (as defined in equation 1.4) are highly correlated across forecasting horizons. Furthermore, Table A.4 shows that the two measures of expectations of returns based on Bloomberg and Consensus Economics' survey data are highly correlated. In fact, the pairwise correlation (given the same forecasting horizon) is always higher than 90% and 75% for crude oil and copper respectively.

Analysing the forecasts at the analyst level, we observe that there is on average an high level of agreement among analysts about the sign of the expected returns. Figure A.1 shows that the time-series average of the agreement statistic is close to 80% for both crude oil and copper, meaning that on average 80% of the analysts agree that prices will move in one specific direction (up or down compared to current spot prices). This evidence, corroborates the hypothesis that there is a common set of factors driving analysts' expectations of returns.⁷

[Insert Figure A.1 about here]

1.2.2 Futures Prices, Fundamentals and Trading Flows

Futures Prices and Implied Volatility

We obtain futures prices data on Crude Oil (WTI) from the New York Mercantile Exchange (NYMEX). Copper data are from the Commodity Exchange (COMEX). Crude oil futures are in U.S. Dollars, whereas copper prices are transformed from USD Cents/Pound to USD/Tonne to match the measurement unit used in the survey forecasts. Since futures prices do not have constant maturity, we calculate quasi-constant maturity monthly futures returns by cumulating daily log-returns with average time to maturity at 2-3-4 quarters ahead. On the settlement date, all positions are closed and rolled over the following day using the next reported contract. This procedure ensures that returns are always computed using the same contract and not across two different contract maturities on the settlement date. Formally, we

⁷If analysts generated forecasts randomly, or according a different prediction rules, we should have observed instead a very low level of agreement.

define log-returns as

$$\begin{cases} r_t^{(m)} = 0 & \text{if } t - 1 \text{ is expiration date} \\ r_t^{(m)} = \ln(F_t^{(m)}) - \ln(F_{t-1}^{(m)}) & \text{otherwise} \end{cases} \quad (1.2)$$

where $F_t^{(m)}$ is the futures price at time t for the contract with average time to maturity m .

The spot price is approximated using the nearest to maturity contract, as in de Roon et al. [30]. Formally, we define

$$r_t = \ln(S_t) - \ln(S_{t-1}) \quad (1.3)$$

where $S_t = F_t^{(+1Q)}$.

Our measure of crude oil implied volatility is constructed as the average of the daily at-the-money (ATM) put and call implied volatility on the futures options closest to maturity on the NYMEX WTI contracts from 2000M2 to 2015M2.

Fundamentals

To control for fundamentals factors that can affect both commodity prices and analysts' expectations of returns, we include in our analysis data on industrial production and inventory levels.

Deaton and Laroque [33] show in their model of commodity prices that the demand for commodities is increasing in global economic activity. As industrial production is an highly regarded indicator of global commodity demand by market analysts, we proxy for global economic activity using data on world industrial production from the Netherlands Bureau for Economic and Policy Analysis. More specifically, we use the import weighted, seasonally adjusted, world industrial production index. The sample is monthly and starts in January 1991.

As predicted by the traditional Theory of Storage, originally put forth by Kaldor [62] and then corroborated by empirical evidence in Gorton et al. [48], inventory levels are a key fundamental variable in the determination of commodity futures risk premiums. Furthermore, the level of inventories is directly linked to futures prices and expected spot prices via the “cost of carry” relationship (e.g. Gorton et al. [48], Pindyck [83]). To control for this fundamental factor, as in Gorton et al. [48], we collect data on crude oil and copper inventories from the Energy Information Administration (EIA) and London Metal Exchange (LME) respectively. Crude oil inventory data publication is monthly and starts on January 1945. Copper inventory levels are published daily by the LME and the publication starts in June 1974. In this case we sample the data at monthly frequency using the inventory level reported on the last business day of the month.⁸

Trading Flows

We use data on futures contract positions of hedgers, large and small speculators collected by the Commodity Futures Trading Commission (CFTC). Since each group can hold long and short positions, we analyse net long positions, i.e. the number of long futures positions in excess of the short ones. Given the strong non-stationarity of the net long position data, we analyse the monthly change in net long positions divided by the previous period open interest as in Kang et al. [64]. The open interest adjustment effectively neutralizes the strong trending effect of commodity derivatives since the early 2000’s. As in Moskowitz et al. [77], we identify hedgers with what the CFTC defines commercial traders. Large and small speculators refer to non-commercial and not-reportable traders CFTC’s definition.

Summary statistics are reported in Table A.1.

⁸Crude oil inventories are the U.S. total inventory data (excluding strategic petroleum reserves) measured in thousands of barrels. Copper inventories data are the London Metal Exchange world inventory levels.

1.3 Results

In this Section we present the results of our empirical analysis. We start by presenting the correlation analysis between survey-based expectations of returns and the futures curve. In Subsection 1.3.2 we illustrate the results of the predictive regressions on spot price returns and inventory levels. Subsections 1.3.3 and 1.3.4 describe the analysis of the determinants of survey-based expectations of returns and their relationship with the trading flows, respectively.

1.3.1 Survey Expectations and the Futures Curve

One of the peculiar characteristics of the Bloomberg and Consensus Economics survey forecasts on commodity prices is the availability of forecasts at multiple horizons. This distinctive characteristic of our dataset allows us to analyse the correlation structure of the survey expectations of returns across different forecasting horizons. Furthermore, since the objective of the price forecast are well defined futures contract, we can document the correlation between the expectation of returns derived from the survey forecasts and the one extrapolated from the slope of the future curve.

The correlation analysis is important for several reasons. First, if expectations of returns for a specific commodity tend to be highly correlated across forecasting horizons and data providers, we can hypothesize that it exists a set of common drivers that can explain how analysts' forecasts are formed. Second, if survey-based expectations of returns are correlated positively with the ones extrapolated over the same horizon from futures curve, we provide evidence consistent with the hypothesis that commodity futures prices should equal expectations of future spot prices plus a risk premium (e.g. Cootner [26], Dusak [35] and Fama and French [41]). Furthermore, if this correlation is very high, we can approximate the average analysts' forecasts with observations from the futures curve.

Table A.3 reports the correlation of expectations of returns across forecasting horizons and between survey-based expectations and the slope of the futures curve for both the

Bloomberg survey (Panel A) and the Consensus Economics survey (Panel B).⁹ We define survey-based expectations of returns and the slope of the future curves with the two following equations:

$$\hat{r}_t^{(m)} = \ln(\hat{S}_t^{(m)}) - \ln(S_t) \text{ with } m \in \{+2Q, +3Q, +4Q\} \quad (1.4)$$

$$slope_t^{(m)} = \ln(F_t^{(m)}) - \ln(S_t) \text{ with } m \in \{+2Q, +3Q, +4Q\} \quad (1.5)$$

where m defines the forecasting horizon.

[Insert Figures A.2 A.3 A.4 A.5 about here]

The top left corner of the correlation matrix confirms the visual impression from Figure A.2 to A.5 that the expectations of returns tend to be highly correlated across forecasting horizons. When survey forecasts point to an increase in the spot price over the next two quarters, then it is very likely that a spot price increase is expected also at three and four quarters horizons. In fact, the pairwise correlation between expectations of returns at different horizons is almost in all cases higher than 90% and never lower than 80% for both crude oil and copper.

The top right corner of the matrix in Table A.3 reports the correlation between the survey-based expectations of returns and the expectations extrapolated from the slope of the future curve for oil and copper. As mentioned earlier, if futures prices are equal to the expected future spot price plus a risk premium, as postulated by Cootner [26], we expect survey expectations of returns and the ones derived from the slope of the futures curve to be highly and positively correlated.¹⁰ Consistent with Cootner's hypothesis, we find that the correlation between the expectations of returns and the slope of the futures curve is always positive and tend to increase with the forecasting horizon, reaching 80% for both crude oil and copper in the case of the Consensus Economics forecasts. This result suggests that the

⁹We estimate the correlation using forecasted returns rather than price levels to avoid spurious correlation, as price predictions display a unit root. When running Augmented Dickey-Fuller tests on returns forecasts instead, we strongly reject the null hypothesis of a unit root across forecasting horizons for both crude oil and copper.

¹⁰Under the assumption that the risk premium is either constant or its time-series variance is much smaller than the variance of the expectations term.

slope of the futures curve can be used to proxy for analysts' average expectations of returns.

1.3.2 Predictive Regressions

Future Returns

In this Section, we test the forecasting power of survey forecasts by estimating predictive regressions at different forecasting horizons. We specify the predictive regressions as:

$$r_{t+m:t} = c + \beta_1 \hat{r}_t^{(m)} + \varepsilon_{t+m:t} \quad (1.6)$$

$$r_{t+m:t} = \ln(S_{t+m}) - \ln(S_t) \text{ with } m \in \{+2Q, +3Q, +4Q\} \quad (1.7)$$

where $r_{t+m:t}$ is the realised future spot return between time t and $t + m$.¹¹

[Insert Table A.5 about here]

Table A.5 reports the results of the predictive regressions for crude oil and copper across forecasting horizons using both Bloomberg and Consensus Economics mean survey forecasts as the forecasting variables. Survey-based expectations of returns do not show significant predictive power for future realised spot returns over different horizons. In the case of crude oil, expectations of returns and future returns are positively correlated at 3 and 4 quarters forecasting horizons, but their correlation is low and not statistically significant. Conversely, for copper, the correlation is positive only in the case of the Bloomberg's survey (which spans the sample from 2006 to 2014). In all cases, the explanatory power is low and increases only slightly with the forecasting horizon. As we will show in Section 1.4, this result is not necessarily at odds with the assumption that forecasters hold rational expectations. In fact, their forecasting power depends on both the predictability of the fundamentals and on the importance of the trading activity of other investors with predictable investing strategies. In

¹¹In order to match the average forecasting horizon of the quasi-constant horizon survey forecasts series we use the central month of each quarterly horizon, namely: 5th; 8th and 11th month ahead.

other words, the low forecasting power of survey forecasts does not necessarily imply that such forecasts do not reflect efficiently the information, but it can be explained simply by the fact that fundamentals are very hard to predict and that only a small share of the trading activity by speculators is predictable.

These results are interesting because, differently from Greenwood and Shleifer [49] and Kojen et al. [70] that find contrarian forecasting power when analysing survey forecasts in the context of stocks, bonds and currencies, we find instead positive but not significant correlation between expectations of returns and future returns, which is easier to reconcile with rational expectations asset pricing models. We hypothesize that such difference is not due to the peculiar characteristics of the commodity markets, but rather to the different survey methodology and to the fact that in the case of commodities only professional forecasters are being surveyed. In fact, this latter characteristic can increase the number of “rational” forecasters in the survey population compared to the case of other markets where not necessarily the population surveyed is composed by professional forecasters. If non-professional forecasters formed their expectations in a more simple way, by for instance extrapolating the future from the past, then their contrarian predictive power could also stem by the little negative autocorrelation that appears when analysing lagged returns over the past 13 to 24 months (see Cutler et al. [28]).

Future Inventory Levels

An expected increase in the spot price of a storable commodity should induce an increase in inventory levels as it decreases the opportunity cost of holding inventories (see Pindyck [83]). Furthermore, if the expected increase in the spot price is large enough to compensate for the cost of storage and the risk of future negative price variation, it can induce a speculative inventories build-up. An increase in inventories can also stem from traditional speculative carry strategies that entail selling futures contracts and storing the commodity to lock-in

a risk-less profit.¹² For the reasons mentioned above, and given the positive correlation between expected returns and the slope of the futures curve, we expect that a predicted increase in spot prices should correlate with a future build up in inventories and vice versa. To test this hypothesis, we analyse the predictive power of expectations of returns on future inventory levels by estimating the following regression

$$inventories_{t+m} = c + \beta_1 \hat{r}_t^{(m)} + \beta_2 t + \beta_3 momentum_{t:t-12} + \varepsilon_{t+m} \quad (1.8)$$

where $inventories_{t+m}$ is the logarithm of the inventory levels at time $t + m$, $momentum_{t:t-12}$ is the spot log-return between time t and $t - 12$, which controls for past spot price variation, and t is a time trend. Given that both the dependent and independent variables are stationary but persistent, inducing positive autocorrelation in the error term ε , we compute t-statistics using Newey and West [81] standard errors.

[Insert Table A.6 about here]

Table A.6 reports the results of the predictive regressions for both surveys data sources and commodities. We find that expectations of returns tend to anticipate inventory levels across forecasting horizons and commodities. As expected, the regression coefficients on the expectations term are always positive, meaning that an expected increase in the spot price leads to future inventories build up and vice versa. The results are statistically significant for both commodities and both samples, also when we control for a time trend and past price variations. The only exception is the case of copper, where the coefficients on the expectations term are positive but not statistically significant when the Bloomberg's sample period is analysed.

¹²Obviously, this trading strategy can be profitable only if the future curve is in contango and the slope is steep enough to compensate for the cost of storage.

1.3.3 Survey Forecasts, Time-Series Momentum and Value

We follow the approach of Greenwood and Shleifer [49] to understand how expectations of returns are formed. More specifically, we regress survey expectations of returns on a time-series momentum factor, a proxy for value, and a set of fundamentals variables. The inclusion of the momentum factor is motivated by previous empirical research (e.g. Greenwood and Shleifer [49], Kojien et al. [70], Cutler et al. [28], Moskowitz et al. [77]) and aims at capturing the existence of investors holding extrapolative expectations. Formally, we define the time-series momentum factor as:

$$momentum_{t:t-b} = \ln(S_t) - \ln(S_{t-b}) \text{ with } b > 0 \quad (1.9)$$

We define our proxy for value following the approach of Asness et al. [6] and define value as the cumulated log-returns over the past 5 years:

$$value_{t:t-B} = \ln(S_t) - \ln(S_{t-B}) \quad (1.10)$$

where $B = 60$, *Months*. We use the proxy for value to capture long term trends and deviations from average historical prices. In the case of commodity markets, such deviations from the historical mean price level can be a relevant factor as commodity prices can be mean reverting around the marginal costs of production, as in the Deaton and Laroque [33] model. Intuitively, the importance of the value factor in explaining analysts' expectations is due to the possibility that forecasters believe in prices mean reversion. In such case, our definition of the value factor should correlate negatively with the expectations of returns, as analysts should expect prices above (or below) the historical average to revert to the marginal costs of production.

[Insert Table A.7 about here]

Table A.7 shows the results of the following time-series regression:

$$\begin{aligned} \hat{r}_t^{(m)} = & c + \beta_1 momentum_{t:t-b} + \beta_2 value_{t:t-B} + \dots \\ & \dots + \beta_3 slope_t^{(m)} + \beta_4 inv_t + \beta_5 \Delta ind.prod_{t:t-b} + \epsilon_t \end{aligned} \quad (1.11)$$

where the momentum and value factors are defined by equations (1.9) and (1.10) with $b = 3$ and $B = 60$ months, respectively.¹³ Inv_t is the commodity specific logarithm of inventory levels and $\Delta ind.prod_{t:t-b}$ are the log-differences in world industrial production calculated over the previous b months.

Table A.7 shows that, both the momentum and value factors are negatively correlated with expectations of returns. This result is the opposite of what found by Greenwood and Shleifer [49], where they find that expectation of returns on the stock market are positively correlated with past cumulated returns and a proxy for value.

In Section 1.4, using simulated data from a model with heterogeneous beliefs, we show that these two opposite results can arise when two different samples of the population are surveyed. More specifically, our results are consistent with the predictions of rational expectations analysts when we allow for mean reversion in fundamentals and for the presence of positive feedback traders. Conversely, the results obtained by Greenwood and Shleifer [49], are consistent with the ones obtained when a sample of positive feedback traders is surveyed.

The results in Table A.7 show that large part of the time-series variation of expectation of returns across forecasting horizons is explained by the momentum factor itself, which explains around 50% and 40% of the variation for crude oil and copper respectively. Conversely, the proxy for value, despite explaining further variability in the expectations of returns, plays a smaller role. Moreover, the explanatory power of the value factor tends to increase with the forecasting horizon, which is consistent with a model in which forecasters make their predictions according to a mean reverting pricing model. The statistical significance of the momentum and value factors holds also when we include the slope of the future curve

¹³Asness et al. [6] define value as minus the cumulated log-return over the past 5 years. We reports results using time-series momentum computed over the previous three months $b = 3$ as modelled in Cutler et al. [27]. Table A.14 reports results on the relationship between survey-based expectations of returns and the time-series momentum factor computed using different look-back periods.

and fundamentals as control variables. Interestingly, inventory levels and world industrial production growth are not statistically significant and do not explain an additional important variation in the survey-based expectations of returns. The slope of the futures curve is statistically significant only at longer forecasting horizon, but it is able to explain only a marginal amount of the time-series variation in the expectations of returns. As a robustness check, we perform the same regressions but assuming that analysts have perfect foresight on futures inventory and industrial production growth.¹⁴ Table A.15 shows that also when we control for perfect expectations about future fundamentals, the momentum and the value factors remain significant. Only future inventory levels become statistically significant, but they do not improve the explanatory power in a significant way.

To understand the importance of the look-back period used to compute the momentum factor, we regress survey-based expectations on a time-series momentum factor computed, not only using the last three months observations as modelled by Cutler et al. [27], but also using returns over the last six and twelve months. Table A.14 shows that, despite the fact that momentum remains an important and significant factor in explaining survey expectations of returns, its explanatory power decreases with the length of the window over which it is computed, with more recent past price variations that seem to matter more. As a last robustness check Table A.16 shows that the value and momentum factors remain significant also when computed skipping the most recent return.

Overall, the time-series momentum and value factors explain around 70% of the expected returns variation in the case of crude oil and around 60% in the case of copper, at different forecasting horizons.

¹⁴This is to make sure that our results are not due to a misspecified model that does not take into account the fact that in reality analysts form their returns expectations using forecasted fundamentals and not historical values.

1.3.4 Survey Forecasts and COT Positions

In this Section, we extend our empirical analysis to include commodity derivatives trading flows and study the relationship between survey-based expectations of returns and trading flows. To this end, we use data on futures contracts positions of hedgers, large and small speculators collected by the Commodity Futures Trading Commission (CFTC). Since each group can hold long and short positions, we analyse net long positions, that define the number of long futures positions in excess of the short ones. Given the high trending behaviour of the time-series data on net long positions, we analyse the monthly change in net long positions divided by the previous period open interest as in Kang et al. [64]. The open interest adjustment serves the purpose of neutralizing the trending effect of the increasing open interest that has been experienced by the commodity derivatives since the early 2000's.

To study the contemporaneous relationship between expectations of returns and trading activity we run the following time-series regressions

$$\Delta Net Long_{i,t}^* = c + \beta_1 \hat{r}_t^{(m)} + \beta_2 dummy_t \times \hat{r}_t^{(m)} + \varepsilon_{i,t} \quad (1.12)$$

where $\Delta Net Long_{i,t}^*$ is the “net trading activity” as in Kang et al. [64], i.e. the change in the net long position divided by the previous period open interest

$$\Delta Net Long_{i,t}^* = \frac{Net Long_{i,t} - Net Long_{i,t-1}}{Open Interest_{i,t-1}} \quad (1.13)$$

where i indicates hedgers, speculators, and small traders. We control for the 2008 financial crises by adding a dummy variable, which interacts with the expectations of returns. The dummy variable is equal to one from January 2008 to December 2009.

[Insert Table A.8 about here]

Table A.8 shows that hedgers and market makers “net trading activity” is contemporaneously positively correlated with survey-based expectations of returns. Conversely, large

speculator trading activity is negatively correlated with expectations of returns. Therefore, in light of the negative relationship between past returns (momentum) and expectations of returns presented above, the indirect positive relationship between momentum and large speculators trading activity is consistent with the results presented by Kang et al. [64] and Moskowitz et al. [77] showing that speculators follow momentum strategies, whereas market makers and hedgers accommodate their short term liquidity needs by trading as contrarians.

1.3.5 Forecasts Dispersion and Options Implied Volatility

In this Section we study the relationship between analysts' forecasts dispersion and options implied volatility. Buraschi and Jiltsov [18] and Beber et al. [11] point to the importance of considering heterogeneity in investors returns expectations to learn more about asset prices by showing that differences in analysts' beliefs about future returns have an impact on option implied volatility. More specifically, they show that survey forecasts dispersion on future foreign exchange rates has a significant impact on both implied volatility levels and the implied volatility risk premium. These results hold also when taking into account the volatility of fundamentals macroeconomic variables. In their empirical setting, the absence of an undisputed pricing model and the limited importance of fundamentals factors in explaining exchange rates variations (see e.g. Meese and Rogoff [74]), creates a favourable setting in which analysts' opinions might draw more attention than in other asset classes. In fact, one might argue that differences in beliefs are less relevant in other markets where the role of fundamentals on asset prices is more clear and pronounced. Here, in order to evaluate such hypothesis, we exploit the commodity markets settings to test if the relationship between uncertainty, proxied by survey forecasts dispersion and implied volatility, breaks down in a market where price variations can be explained more easily by fundamentals shocks.

We restrict the analysis to the Crude Oil (WTI) as its futures options market liquidity allows us to construct a long enough measure of ATM implied volatility. We measure survey forecasts dispersion as the cross-sectional forecasts standard deviation divided by the spot

price

$$\hat{\sigma}_t^{(m)} = st.dev. \left(\hat{S}_{i|t}^{(m)} \right) / S_t \quad (1.14)$$

where m is the forecasting horizon and the subscript i identifies the $i - th$ forecaster.

[Insert Figures A.6 and A.7 about here]

Figures A.6 and A.7 show the evolution of survey forecasts dispersion, together with the price of oil, for both the Bloomberg and Consensus survey forecasts data respectively. From both Figures emerges that higher levels of forecasts dispersion follow large price swings. For instance, survey forecasts dispersion jumped to more than twenty and twelve percent following the oil price collapses of 2008 and 2014 respectively.

[Insert Figures A.8 and A.9 about here]

Figures A.8 and A.9 show clearly that futures options implied volatility jumped simultaneously with the forecasts dispersion during the same periods of time. In order to test formally the importance of analysts' differences of beliefs we estimate the following time-series regression

$$IV_t = c + \beta_1 \hat{\sigma}_t^{(+2Q)} + \beta_2 \sigma_t + \beta_3 slope_t^{(+2Q)} + \beta_4 IV_{t-1} + \beta_5 inventories_t + \varepsilon_t \quad (1.15)$$

where IV_t represents the crude oil implied volatility index at time t of the WTI futures options and $\hat{\sigma}_t$ represents the analysts forecasts dispersion defined in equation 1.14.¹⁵ We additionally control for other factors that can explain the implied volatility levels and that might also be correlated with the forecasts dispersion. More specifically, σ_t is the realised volatility measured as the monthly average of daily squared log-returns on the front month WTI contract, $slope_t$ is the slope of the futures curve as defined in equation 1.5. We also include past levels of implied volatility to take into account the relatively high persistence of the the series. Finally, in order to control also for fundamentals, we include the logarithm

¹⁵We use the forecasts dispersion on the $+2Q$ horizon because is the one with the horizon closer to the maturity of the options used to compute the implied volatility.

of crude oil inventories levels, which variations are a reflection fundamentals supply and demand shocks that are difficult to measure at global level.

[Insert Table A.9 about here]

Table A.9 reports the results of estimating Equation 1.15. The results show that the forecasts dispersion coefficient is always positive and strongly statistically significant also when the control variables are included in the regressions. In fact, $\hat{\sigma}_t$ alone explains 66% and 47% of the implied volatility variance respectively, when the Bloomberg and Consensus forecasts sample are considered. As expected, the sign of the coefficient is positive, meaning that an increase in analysts' opinions dispersion corresponds to higher implied volatility levels. Not surprisingly, the coefficients of realised volatility and past levels of implied volatility are also positive and statistically significant. Conversely, the coefficient of inventories levels is negative and significant, reflecting the fact that lower levels of inventories are on average related to higher levels of price volatility. The slope of the futures curve instead is never significant in any of the regression specifications, showing that implied volatility does not correlate with expected returns.¹⁶

In this Section, we also analyse the relationship between survey forecasts dispersion and the volatility risk premium, measured as the spread between the options implied volatility and realised volatility. The volatility risk premium reflects the difference between the risk-neutral and the physical probability measure, which depends on the risk aversion of the representative agent. In a neoclassical framework in which individuals have heterogeneous beliefs, the representative agent discount factor depends on the relative weight of the individuals, which in turn depends on the level of dispersion of beliefs.¹⁷ In order to study the impact of forecasts dispersion on the implied volatility risk premium we estimate the following time-series

¹⁶In Section 1.3.1 we showed that survey-based expectations of returns and the slope of the future curve are positively correlated.

¹⁷See e.g. Beber et al. [11] and Buraschi and Jiltsov [18].

regression

$$spread_t = c + \beta_1 \hat{\sigma}_t^{(+2Q)} + \beta_2 slope_t^{(+2Q)} + \beta_3 spread_{t-1} + \beta_4 inventories_t + \varepsilon_t$$

where $spread_t$ represents the difference between crude oil implied volatility index at time t of the WTI futures options and the realised volatility on the front month contract.¹⁸

[Insert Table A.10 about here]

As before, the results in Table A.10 show that the dispersion in analysts' forecasts have a significant and positive impact. More specifically, the spread between implied volatility and realised volatility tends to increase with the level of analysts' disagreement on futures oil price variations. The significance level remains strong across the two different survey data samples and does not vanish when the additional control variables are included in the regression specification.

In summary both findings, are similar to the results obtained by Beber et al. [11], confirming the importance of differences in beliefs (proxies by analysts' forecasts dispersion) in explaining both the level of options implied volatility and the volatility risk premium also in a market where the link between prices and fundamentals is much more clear than in the foreign exchange market.

1.4 Interpretation

In this Section, we show using simulated data from a modified version of the Cutler et al. [28] Model of speculative prices, that the predictive power of rational traders' returns expectations can depend on the level of prices mean reversion and on the existence of feedback traders. In fact, we show that their predictive power increases not only with the speed of mean reversion

¹⁸The realised volatility is measured as the monthly average of daily squared log-returns on the front month WTI contract.

in prices but also with the share of trading activity of feedback (momentum) traders.¹⁹ This latter result stems from the fact that rational traders can anticipate asset demand and take advantage of momentum traders by looking at past price variations. Conversely, in our model, feedback traders' expectations correlate negatively with future realised returns, as on average rational traders make profits at their expenses.

Also, we show that the correlation between expectations of returns and past returns found analysing survey data, which seems at odds with standard rational expectations frameworks, can easily arise in rational traders' expectations with very little departures from standard assumptions. In fact, in our model, it is sufficient a very slow mean reversion in fundamental prices and the existence of little trading activity by feedback traders to generate a rational dependence of expectations on past price variations. These results suggest that the sign of such correlation found analysing survey data can depend on the share of "rational" and extrapolative individuals in the population being surveyed. This would imply that the different correlation sign obtained in this analysis and the one conducted by Greenwood and Shleifer [50] could be due to the fact that we use forecasts by professional forecasters, that should be close to the rational expectation type, whereas Greenwood and Shleifer [50] use data of non-professional forecasters, that could use simple extrapolative expectations to form their forecasts.

The model we use to generate data under different assumptions about prices mean reversion and the amount of trading activity of feedback traders builds on Cutler et al. [28]. This setting allows us to generate both futures prices and expectations about future prices. The model, described in detail in the Appendix, assumes the existence of two types of agents holding rational and extrapolative expectations.²⁰ In our model, rational expectation

¹⁹When the fundamental price mean reversion and the feedback traders activity are both zero, our model boils down to the standard random walk model, where rational agents have constant expectations and zero predictive power.

²⁰The inclusion of extrapolative expectation agents (positive feedback traders) is motivated by previous research. Cutler et al. [28] show that the alternation of positive and negative autocorrelation in returns can arise in presence of feedback traders and rational traders receiving information about the fundamentals with delay.

agents observe the fundamental value underlying the futures price. Furthermore, to be able to generate data under different assumptions about the predictability of fundamentals, we assume that the fundamental process is governed by an autoregressive process of order one. By changing the value of the autoregressive coefficient α , we are able to generate data, either from a model with no predictability in fundamentals ($\alpha = 1$), or with different degrees of predictability ($0 < \alpha < 1$). This choice allows us to switch easily from a model in which fundamentals follow a random walk to one in which they are mean reverting (as for instance in the Deaton and Laroque [33] model of commodity prices).

The parameter γ is also crucial in the simulation exercise as it governs the percentage of trading activity undertaken by the two types of agents. When $\gamma = 1$, only rational expectation agents operate in the market. Therefore, the expectations of rational agents can be affected only by different levels of mean reversion in the fundamentals and no room is left for any momentum effect generated by feedback traders. Conversely, when γ is smaller than one, the share $1 - \gamma$ of the trading activity is undertaken by positive feedback traders that form their expectations about future prices by looking at historical trends. In this case, rational expectation agents form their expectations taking into account the fact that a share of the trading activity is generated by time-series momentum strategies.

1.4.1 Predictability

The predictive regressions results reported in Table A.5 show that, in the actual data sample analysed, expectations of returns in most of the cases correlate positively but not significantly with future returns, i.e. they do not carry any real predictive power. Thus, the first question we want to answer by analysing the simulated data is under what theoretical assumptions we should observe such results.

[Insert Tables A.11 and A.12 about here]

More recently, Greenwood and Shleifer [49] show that positive correlation between survey expectations of returns and past returns in the stock market can stem from the existence of positive feedback traders.

Tables A.11 and A.12 report the result of predictive regressions of expectations of returns generated by rational and feedback (momentum) traders respectively under different calibration of α and γ . As we can see from the Tables, the predictive power is very limited under all the calibrations analysed and tend to increase, as expected, with the speed of the mean reversion in fundamentals (parameter α) and with the share of trading activity undertaken by the feedback traders holding extrapolative expectations. The latter result is due to the fact that, given the predictable trading activity of the feedback traders and given that rational expectations traders take this into account, an increase in the share of feedback traders activity makes future price movements more predictable.

In terms of statistical significance, we see that when the observations number is limited and close to the one analysed using real survey data, the statistical significance tend to be very weak. Only if the fundamentals are enough predictable and the feedback trading activity is large enough we can observe t-statistics level slightly higher than 2. As expected, the sign of the expectation coefficient is positive in the case of rational agents and negative for positive feedback traders, meaning that even if the predictive power is very poor, rational agents tend to predict market movements in the right direction, whereas feedback traders tend to have contrarian predictive power.

In conclusion, from this experiment, we can affirm that the result obtained in Section 1.3.2 do not necessarily imply neither that survey forecasts are pure noise nor that the analysts taking part in the survey are not rational. Thus our interpretation is that the lack of predictive power is not due to non-rationality of the survey participants, but simply to the high level of unpredictability of the underlying commodity fundamentals process.

1.4.2 Determinants of Expectations of Returns: Momentum and Value

In Section 1.3.3 we show that survey expectations of returns are well explained by a time-series momentum and a value factor. Both of them correlate significantly and negatively with expectations of returns for both crude oil and copper also when controlling for the slope of

the futures curve and other fundamentals variables. In fact, our results show that when prices have been growing strongly there is a high chance that on average analysts will predict a future price drop and vice versa. At first sight, it seems that the analysts taking part to the survey bet on prices mean reversion and against market over-reactions.

To understand better some of the possible theoretical assumptions under which we can observe such empirical evidence, we run the same analysis presented in Table A.7 with simulated data from the theoretical model. More specifically, we regress simulated rational expectations on the time-series momentum and value factors under different assumptions about the speed of mean reversion in fundamentals and the share of trading activity undertaken by feedback traders.²¹

[Insert Table A.13 about here]

Table A.13 (calibration II and III) shows that slow mean reversion in fundamentals or the presence of a limited trading activity by positive feedback traders is a sufficient condition to have the momentum and value factors displaying a negative correlation with expectations of returns. However, in small samples and with reasonable levels of mean reversion and feedback trader activity, the two assumptions taken independently are not sufficient to generate the strong combined statistical significance of both the momentum and value factors reported in Table A.7. Furthermore, by allowing for mean reversion only and not for the presence of positive feedback traders activity, we are not able to generate a larger coefficient on the momentum factor as observed when analysing the actual data.

In conclusion, calibrations IV and V show that in small samples the coexistence of both mean reversion in fundamentals and the presence of feedback traders activity is a necessary condition to match the results obtained when analysing real survey data.

²¹We perform this analysis only on simulated expectations of rational agents because by construction we know that feedback traders expectations are a linear function of cumulated returns, i.e. time-series momentum.

1.5 Concluding Remarks

We study the links between expectations, fundamentals, and asset returns using an unexplored database on commodity survey price forecasts. This unique dataset offers us the opportunity to analyse survey-based expectations with less measurement error problems compared to previous studies.²² More importantly, the data provide us with the opportunity to explore how the forecasts predictive power and drivers change at different forecasting horizons and when the population surveyed consists of professional forecasters.²³

The results show that, in commodity markets, survey-based expectations of returns are strongly correlated with past returns, but not with fundamentals. In fact, we find that survey forecasts are largely explained by time-series momentum and value factors. Expectations have positive, but not significant correlation with future realised returns, which implies little predictive power.

We rationalise these findings by setting up a simulation experiment, which shows that such results are not incompatible with rational expectations. To the contrary, we show that the predictive power of rational forecasts purely depends on the predictability of the fundamentals and the asset demand of other types of agents. Additionally, the experiment shows that the correlation with past returns can also arise in rational expectations when rational agents take into account the speculative activity of momentum (feedback) traders, whose asset demand depends on past returns.

In terms of contribution, this Chapter enriches the strand of literature on asset pricing and survey forecasts by extending the studies of Koijsen et al. [70]; Greenwood and Shleifer [49] and Beber et al. [11] to the commodity markets and by adding the analysis on multiple forecasting horizons. Most importantly, our findings shed some light on the puzzling evidence of contrarian predictive power of survey forecasts and positive correlation with past returns found by Greenwood and Shleifer [50], by showing that when a population of professional forecasters is interviewed, the contrarian predictive power disappears and the sign of the correlation with past returns changes. Our simulation experiment also shows that both results

²²Most important advantages: clear forecast objective, quantitative nature, higher frequency.

²³Greenwood and Shleifer [50] use survey data input by non-professional forecasters.

can be obtained depending on which of the two populations, rational or extrapolative, is surveyed. Therefore, we hypothesize that our and the previous empirical evidence is driven by the population being interviewed, with professional forecasters being close to rational expectations agents and the non-professional forecasters close to the extrapolative type.

Finally, our analysis also suggests that survey-based expectations can have a crucial role to understand better the dynamics of trading flows and the drivers of the option implied volatility risk premium. We leave our preliminary results on these two extremely fascinating topics as starting points for avenues of future research.

Appendix A

Appendix and Tables

A.1 Model Description

Consider a futures market, where there is a well defined fundamental equal to the terminal value and no dividends are paid. The final condition is $p_T = f_T$ where f_T is the natural logarithm of the fundamental value of the future contract at maturity and p_T is the log of the settlement price of the contract. The fundamental price and the market price can be different at any date prior to maturity $t < T$. The fundamental value is assumed to be mean reverting, more specifically it is assumed to be a stationary autoregressive process of order one.¹

$$f_{t+1} = c + \alpha f_t + z_{t+1}; \quad 0 < \alpha < 1 \quad (\text{A.1})$$

where z_t is a martingale difference sequence, i.e. $E_t[z_{t+1}] = 0$. Given the stationarity assumption the fundamental price dynamics can be written in terms of deviation from its equilibrium price $\bar{f} \equiv c/(1 - \alpha)$.

$$f_{t+1} = (1 - \alpha) \bar{f} + \alpha f_t + z_{t+1} \quad (\text{A.2})$$

¹In commodity markets, the assumption of mean reversion can be justified by the existence of production reactions that drives the fundamental prices towards the marginal cost of production.

therefore, next period fundamental price can be interpreted as a weighted average of the fundamental price equilibrium and the current period price, plus a random shock. Without loss of generality, in order to simplify the algebra, we rescale the fundamental equilibrium price to 1, which implies $\bar{f} = 0$.² Therefore, the fundamental price process can be written as

$$f_{t+1} = \alpha f_t + z_{t+1} \quad (\text{A.3})$$

Furthermore, we assume the existence of N traders of two types, namely: N^R rational expectation and N^B feedback traders, such that $N = N^R + N^B$. We define the share of rational expectation traders as $\gamma \equiv N^R/N$ and the share of feedback traders as $\delta \equiv (1 - \gamma) = N^B/N$.

A.1.1 Rational Expectations Traders

Rational expectations traders demand function is positive and linear in next period expected returns

$$d_t^R = N^R E_t [R_{t+1}] \quad (\text{A.4})$$

where $R_{t+1} \equiv p_{t+1} - p_t$ is the log-return on the future contract. In terms of share of the traders population equation A.4 can be written as

$$d_t^R = \gamma N E_t [R_{t+1}] = \gamma N (E_t [p_{t+1}] - p_t) \quad (\text{A.5})$$

A.1.2 Feedback Traders

Feedback traders demand is positive and linear in past returns

$$d_t^B = N^B L^k R_t; \quad (\text{A.6})$$

where L^k is a lag polynomial of order k . In terms of share of the population equation A.6 can be written as

$$d_t^B = \delta N L^k R_t; \quad (\text{A.7})$$

²Because f denotes the log of the fundamental price.

For exposition purpose we assume $k = 3$

$$d_t^B = \delta N (R_t + R_{t-1} + R_{t-2}) \quad (\text{A.8})$$

applying the definition of log-returns in equation A.8 we obtain

$$d_t^B = \delta N (p_t - p_{t-3}); \quad (\text{A.9})$$

therefore, feedback trader base their demand on the returns cumulated over the previous three periods.

A.1.3 Market Clearing

In futures markets the contracts net supply must be zero, i.e. the number of long and short position must be equal at each point in time. Therefore, the market clearing requires

$$d_t^R + d_t^B = 0 \quad (\text{A.10})$$

plugging-in equations A.5 and A.9 into A.10 we obtain

$$\gamma N E_t [R_{t+1}] + \delta N (p_t - p_{t-3}) = 0 \quad (\text{A.11})$$

dividing by N and rearranging we obtain a rational expectations difference equation in prices

$$p_t = A E_t [p_{t+1}] - B p_{t-3} \quad (\text{A.12})$$

with $A \equiv \gamma / (\gamma - \delta) > 0$ and $B \equiv \delta / (\gamma - \delta) > 0$ for $\gamma > 0.5$. Therefore, prices react positively to an increase in expected prices and negatively to past prices. Equation A.12 can be solved by forward recursion imposing the final condition that the fundamental and observed future prices at maturity must be equal to the spot price, i.e. $f_T = p_T = s_T$ and recalling that by using the law of iterated expectations $E_t [f_{t+k}] = \alpha^k f_t$ we obtain the price dynamics of a future contract with 3-periods maturity as a function of contemporaneous and lagged

variables only

$$p_t^{(3)} = A^3 E_t[f_{t+3}] - A^2 B p_{t-1} - A B p_{t-2} - B p_{t-3} \quad (\text{A.13})$$

where $E_t[f_{t+3}] = \alpha^3 f_t$. Recalling that the rational expectation demand is given by $d_t^R = \gamma N(E_t[p_{t+1}] - p_t)$ and plugging-in A.13 we obtain that the rational expectation traders demand for a future contract expiring in 3 months is given by

$$d_t^{(3)R} = \gamma N[A^3(\alpha - 1)\alpha^3 f_t - B(A^2 p_t - A(A - 1)p_{t-1} - (A - 1)p_{t-2} - p_{t-3})] \quad (\text{A.14})$$

A.2 Tables

Table A.1 Summary Statistics

The Table reports the time-series average and standard deviation for: spot price returns r ; futures returns $r^{(m)}$ where m is the average time to maturity; expectations of returns $\hat{r}^{(m)}$ where m is the average forecasting horizon; the logarithm of inventories and $\Delta Net Long^*$ is the adjusted variation in net long position for hedgers (H) , large speculators (LS) and small speculators (SS) as defined in equation 1.13. The sample is monthly from variable specific start date to 2015M2.

	Crude Oil				Copper			
	Mean	St. Dev.	N	Start	Mean	St. Dev.	N	Start
r_t	0.43%	10.06%	374	1984M1	0.50%	7.63%	314	1989M1
$r_t^{(+2Q)}$	0.45%	8.27%	374	1984M1	0.53%	7.40%	314	1989M1
$r_t^{(+3Q)}$	0.43%	7.33%	374	1984M1	0.56%	6.87%	314	1989M1
$r_t^{(+4Q)}$	0.41%	6.57%	374	1984M1	0.56%	6.66%	314	1989M1
$\hat{r}_{t,Bloomberg}^{(+2Q)}$	0.10%	9.77%	102	2006M9	-1.02%	9.38%	102	2006M9
$\hat{r}_{t,Bloomberg}^{(+3Q)}$	1.50%	11.76%	105	2006M6	-0.89%	11.27%	105	2006M6
$\hat{r}_{t,Bloomberg}^{(+4Q)}$	3.02%	13.57%	106	2006M4	-1.19%	14.08%	106	2006M4
$\hat{r}_{t,Consensus}^{(+2Q)}$	-4.02%	13.33%	79	1995M8	-0.30%	8.57%	79	1995M8
$\hat{r}_{t,Consensus}^{(+3Q)}$	-3.96%	15.18%	79	1995M8	-0.27%	10.13%	79	1995M8
$\hat{r}_{t,Consensus}^{(+4Q)}$	-4.14%	17.00%	79	1995M8	-0.63%	12.12%	79	1995M8
$inventories_t$	12.71	0.08	374	1984M1	12.40	0.76	374	1984M1
$\Delta Net Long_{H,t}^*$	-0.18%	5.97%	265	1993M2	0.05%	13.78%	265	1993M3
$\Delta Net Long_{LS,t}^*$	0.14%	4.60%	265	1993M2	-0.05%	10.90%	265	1993M3
$\Delta Net Long_{SS,t}^*$	0.03%	1.92%	265	1993M2	-0.01%	4.17%	265	1993M3

Table A.2 Bloomberg's Survey Participants

The Table reports summary statistics on the number of survey participants per month.

Summary Statistics on Survey Participants

Horizon	Crude Oil (WTI)			Copper		
	+2 <i>Q</i>	+3 <i>Q</i>	+4 <i>Q</i>	+2 <i>Q</i>	+3 <i>Q</i>	+4 <i>Q</i>
<i>Mean</i>	12.5	12.1	11.3	8.2	8.0	7.3
<i>Median</i>	11.0	11.0	11.0	8.0	8.0	7.0
<i>Max</i>	26.0	26.0	26.0	18.0	18.0	14.0
<i>Min</i>	3.0	3.0	3.0	2.0	2.0	1.0
<i>Std.Dev.</i>	4.3	4.1	4.3	3.2	3.1	3.1
<i>Participants</i>	95	95	95	65	65	65

Table A.3 Multi-Horizon Expectations of Returns and Futures Curve

The Table reports pairwise sample correlation among expectations of returns based on survey forecasts as defined in equation 1.4 and the slope of the future curve at multiple horizons as defined in equation 1.5. Panel A reports the sample correlation for the Bloomberg survey over the sample 2006M9-2015M2. Panel B reports the sample correlation for the Consensus survey over the sample 1995M8-2015M2. The data frequency is monthly.

Panel A: Sample Correlations (Bloomberg Survey 2006M9-2015M2)

Crude Oil (WTI)										Copper					
Survey	Horizon	Survey			Slope			Horizon	Survey			Slope			
		+2Q	+3Q	+4Q	+2Q	+3Q	+4Q		+2Q	+3Q	+4Q	+2Q	+3Q	+4Q	
		+2Q	-	0.95	0.92	0.58	0.58		0.58	-	0.91	0.82	0.19	0.51	0.56
		+3Q	0.95	-	0.98	0.67	0.68		0.68	0.91	-	0.88	0.30	0.59	0.65
Slope	+4Q	0.92	0.98	-	0.67	0.69	0.70	+4Q	0.82	0.88	-	0.35	0.61	0.65	
	+2Q	0.58	0.67	0.67	-	0.99	0.97	+2Q	0.19	0.30	0.35	-	0.78	0.66	
	+3Q	0.58	0.68	0.69	0.99	-	1.00	+3Q	0.51	0.59	0.61	0.78	-	0.98	
	+4Q	0.58	0.68	0.70	0.97	1.00	-	+4Q	0.56	0.65	0.65	0.66	0.98	-	

Panel B: Sample Correlations (Consensus Survey 1995M8-2015M2)

Crude Oil (WTI)										Copper					
Survey	Horizon	Survey			Slope			Horizon	Survey			Slope			
		+2Q	+3Q	+4Q	+2Q	+3Q	+4Q		+2Q	+3Q	+4Q	+2Q	+3Q	+4Q	
		-	0.97	0.96	0.67	0.73	0.76		-	0.91	0.85	0.43	0.55	0.57	
		0.97	-	0.99	0.70	0.76	0.79		0.91	-	0.96	0.56	0.70	0.73	
		0.96	0.99	-	0.70	0.77	0.80		0.85	0.96	-	0.63	0.77	0.81	
Slope	+2Q	0.67	0.70	0.70	-	0.98	0.97	+2Q	0.43	0.56	0.63	-	0.91	0.88	
	+3Q	0.73	0.76	0.77	0.98	-	1.00	+3Q	0.55	0.70	0.77	0.91	-	0.99	
	+4Q	0.76	0.79	0.80	0.97	1.00	-	+4Q	0.57	0.73	0.81	0.88	0.99	-	

Table A.4 Correlations Between Surveys

The Table reports the (common sample) pairwise correlations between Bloomberg and Consensus Economics survey expectations of returns (as defined in equation 1.4). The sample is monthly from 2006M10 to 2015M2.

Sample Correlations (Common Sample 2006M10-2015M2)															
Crude Oil (WTI)															
Horizon	Bloomberg			Consensus			Horizon	Bloomberg			Consensus				
	+2Q	+3Q	+4Q	+2Q	+3Q	+4Q		+2Q	+3Q	+4Q	+2Q	+3Q	+4Q		
	+2Q	-	0.94	0.89	0.93	0.91		0.88	+2Q	-	0.91	0.75	0.80	0.83	0.81
	+3Q	0.94	-	0.96	0.89	0.95		0.95	+3Q	0.91	-	0.91	0.67	0.76	0.79
	+4Q	0.89	0.96	-	0.85	0.92		0.94	+4Q	0.75	0.91	-	0.53	0.67	0.76
Bloomberg						Copper									
Consensus	+2Q	0.93	0.89	0.85	-	0.96	0.92	+2Q	0.80	0.67	0.53	-	0.95	0.86	
	+3Q	0.91	0.95	0.92	0.96	-	0.99	+3Q	0.83	0.76	0.67	0.95	-	0.97	
	+4Q	0.88	0.95	0.94	0.92	0.99	-	+4Q	0.81	0.79	0.76	0.86	0.97	-	

Fig. A.1 Agreement

The charts show the time-series average of the agreement statistic across forecasting horizons. The agreement statistic is defined as $Agreement_t^{(m)} = \max [\%Bullish_t^{(m)}, \%Bearish_t^{(m)}] \in [0.5, 1]$. When the agreement measure takes the value of 100%, all the forecasters agree on an expected increase (or decrease) in prices at horizon m . Conversely, when it takes the value of 50%, only half of the forecasters expect the price to move in one specific direction. The sample is monthly from 2006M9 to 2015M2. The data source is Bloomberg.




Fig. A.2 Bloomberg Survey Forecasts: Crude Oil (WTI)




Fig. A.3 Consensus Survey Forecasts: Crude Oil (WTI)

Panel A shows the evolution of the Crude Oil (WTI) spot price in USD/Barrel and the corresponding Consensus average survey price forecasts at 2-3-4 quarters horizons. When the forecasts (blue dots) are below the corresponding spot price (red line), on average analysts expect a future decrease in the spot price and vice versa. Panel B shows the monthly spot price variations (log-returns) and the survey-based expected future spot price variations (log-returns) at 2-3-4 quarters horizons computed as per equation 1.4 . The sample is monthly from 1995M8 to 2015M2.




Fig. A.4 Bloomberg Survey Forecasts: Copper

Panel A shows the evolution of the Copper spot price in USD/Tonne and the corresponding Bloomberg average survey price forecasts at 2-3-4 quarters horizons. When the forecasts (blue dots) are below the corresponding spot price (red line), on average analysts expect a future decrease in the spot price and vice versa. Panel B shows the monthly spot price variations (log-returns) and the survey-based expected future spot price variations (log-returns) at 2-3-4 quarters horizons computed as per equation 1.4 . The sample is monthly from 2006M9 to 2015M2.




Fig. A.5 Consensus Survey Forecasts: Copper

Panel A shows the evolution of the Copper spot price in USD/Tonne and the corresponding Consensus average survey price forecasts at 2-3-4 quarters horizons. When the forecasts (blue dots) are below the corresponding spot price (red line), on average analysts expect a future decrease in the spot price and vice versa. Panel B shows the monthly spot price variations (log-returns) and the survey-based expected future spot price variations (log-returns) at 2-3-4 quarters horizons computed as per equation 1.4 . The sample is monthly from 1995M8 to 2015M2.




Table A.5 Predictive Regressions

This Table reports results for predictive regressions of realised future returns on survey expectations of returns: $r_{t+m:t} = c + \beta \hat{r}_t^{(m)} + \varepsilon_{t+m:t}$. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Panel A: Crude Oil WTI (Mean Forecast)						
Dep. Var.	Bloomberg (2006M9-2014M2)			Consensus (1995M8-2014M2)		
	$r_{t+2Q:t}$	$r_{t+3Q:t}$	$r_{t+4Q:t}$	$r_{t+2Q:t}$	$r_{t+3Q:t}$	$r_{t+4Q:t}$
<i>Constant</i>	0.00	-0.01	-0.04	0.02	0.04	0.03
	[0.011]	[-0.067]	[-0.399]	[0.470]	[0.841]	[0.622]
$\hat{r}_t^{(+2Q)}$	-0.13			-0.20		
	[-0.321]			[-0.928]		
$\hat{r}_t^{(+3Q)}$		0.52			0.07	
		[0.740]			[0.252]	
$\hat{r}_t^{(+4Q)}$			0.78			0.15
			[1.191]			[0.467]
Observations:	90	90	90	73	73	73
R-squared:	0.00	0.03	0.07	0.01	0.00	0.00
F-statistic:	0.18	2.59	6.37	0.75	0.07	0.28

Panel B: Copper (Mean Forecast)						
Dep. Var.	Bloomberg (2006M9-2014M2)			Consensus (1995M8-2014M2)		
	$r_{t+2Q:t}$	$r_{t+3Q:t}$	$r_{t+4Q:t}$	$r_{t+2Q:t}$	$r_{t+3Q:t}$	$r_{t+4Q:t}$
<i>Constant</i>	0.01	0.01	0.00	0.01	0.03	0.03
	[0.135]	[0.216]	[0.061]	[0.391]	[0.755]	[0.709]
$\hat{r}_t^{(+2Q)}$	0.25			-0.18		
	[0.846]			[-0.779]		
$\hat{r}_t^{(+3Q)}$		1.01			-0.16	
		[1.919]			[-0.403]	
$\hat{r}_t^{(+4Q)}$			0.85			-0.32
			[1.811]			[-0.752]
Observations:	90	90	90	73	73	73
R-squared:	0.01	0.15	0.16	0.01	0.00	0.02
F-statistic:	0.84	15.09	16.58	0.41	0.27	1.17

Table A.6 Inventories Predictive Regressions

This Table reports results for predictive regressions of realised future inventories on survey expectations of returns; a time trend and time-series momentum factor: $inventories_{t+m} = c + \beta_1 \hat{r}_t^{(m)} + \beta_2 t + \beta_3 momentum_{t:t-12} + \varepsilon_{t+m}$. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Panel A: Crude Oil WTI (Mean Forecast)						
Dep. Var.	Bloomberg (2006M9-2014M2)			Consensus (1995M8-2014M2)		
	$inv_{t+2Q:t}$	$inv_{t+3Q:t}$	$inv_{t+4Q:t}$	$inv_{t+2Q:t}$	$inv_{t+3Q:t}$	$inv_{t+4Q:t}$
<i>Constant</i>	12.14	12.16	12.03	12.49	12.47	12.46
	[162.295]**	[113.128]**	[109.640]**	[541.040]**	[433.674]**	[368.148]**
$\hat{r}_t^{(+2Q)}$	0.30			0.33		
	[4.680]**			[5.288]**		
$\hat{r}_t^{(+3Q)}$		0.22			0.32	
		[3.098]**			[5.602]**	
$\hat{r}_t^{(+4Q)}$			0.07			0.23
			[1.048]			[3.283]**
$momentum_{t:t-12}$	0.00	0.03	0.02	0.01	0.06	0.08
	[0.068]	[1.397]	[0.614]	[0.668]	[2.538]*	[2.380]*
Observations:	90	90	90	73	73	73
R-squared:	0.71	0.64	0.64	0.69	0.66	0.59
F-statistic:	69.06	50.2	51.23	51.33	43.97	32.68

Panel B: Copper (Mean Forecast)						
Dep. Var.	Bloomberg (2006M9-2014M2)			Consensus (1995M8-2014M2)		
	$inv_{t+2Q:t}$	$inv_{t+3Q:t}$	$inv_{t+4Q:t}$	$inv_{t+2Q:t}$	$inv_{t+3Q:t}$	$inv_{t+4Q:t}$
<i>Constant</i>	9.81	10.86	12.22	12.86	12.95	13.09
	[6.823]**	[6.989]**	[8.553]**	[41.336]**	[42.082]**	[47.008]**
$\hat{r}_t^{(+2Q)}$	0.18			3.30		
	[0.201]			[3.087]**		
$\hat{r}_t^{(+3Q)}$		0.25			3.49	
		[0.328]			[3.671]**	
$\hat{r}_t^{(+4Q)}$			0.57			3.35
			[1.067]			[5.405]**
$momentum_{t:t-12}$	0.08	-0.09	-0.18	-0.67	-0.55	-0.54
	[0.355]	[-0.305]	[-0.705]	[-2.696]**	[-2.020]*	[-2.376]*
Observations:	90	90	90	73	73	73
R-squared:	0.21	0.11	0.08	0.31	0.4	0.54
F-statistic:	7.63	3.45	2.45	10.56	15.47	27.18

Table A.7 Determinants of Analysts' Expectations

This Table reports results for time-series regressions of survey expectations of returns on time-series momentum, a proxy for value, the slope of the futures curve and fundamentals variables: $\hat{r}_t^{(m)} = c + \beta_1 momentum_{t:t-b} + \beta_2 value_{t:t-b} + \beta_3 slope_t^{(m)} + \beta_4 inventories_t + \beta_5 \Delta ind.prod_{t:t-b} + \epsilon_t$. Expectations of returns are defined in equation 1.4. The momentum and value variables are defined in equations 1.9 and 1.10 respectively. Slope is defined in equation 1.5. Inventories are the the commodity specific logarithm of inventory levels and $\Delta ind.prod_{t:t-b}$ are the log-differences in world industrial production calculated over the previous b months. The proxy for value is obtained by setting $B = 60$ as in Asness et al. [6]. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Panel A: Crude Oil WTI (Bloomberg Survey 2006M9-2015M2) $b = 3; B = 60$											
Dep. Var:	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$				3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$				4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$		
<i>Constant</i>	-0.01 [-0.756]	0.01 [0.683]	-0.00 [-0.195]	-0.58 [-0.275]	0.01 [0.543]	0.03 [1.751]	0.02 [2.468]*	-0.50 [-0.352]	0.02 [1.532]	0.05 [2.387]*	0.03 [4.637]**
<i>momentum_{t:t-3}</i>	-0.34 [-8.463]**		-0.33 [-10.365]**	-0.31 [-11.849]**	-0.42 [-9.757]**	-0.39 [-16.499]**	-0.31 [-7.553]**	-0.50 [-11.638]**		-0.47 [-22.694]**	-0.37 [-12.161]**
<i>value_{t:t-60}</i>		-0.07 [-2.499]*	-0.05 [-4.102]**	-0.05 [-3.086]**		-0.09 [-3.294]**	-0.08 [-6.283]**	-0.07 [-4.981]**		-0.10 [-3.402]**	-0.09 [-6.963]**
<i>inventories_t</i>				0.04 [0.272]				0.04 [0.361]			0.03 [0.335]
<i>$\Delta ind.prod_{t:t-3}$</i>				0.14 [0.355]				-0.35 [-1.049]			-0.59 [-1.787]
<i>slope_t^(+2Q)</i>				0.30 [1.149]							
<i>slope_t^(+3Q)</i>								0.44 [2.980]**			
<i>slope_t^(+4Q)</i>											0.38 [4.518]**
Observations:	102	102	102	102	102	102	102	102	102	102	102
R-squared:	0.61	0.19	0.74	0.74	0.63	0.26	0.82	0.85	0.68	0.25	0.85
F-statistic:	154.93	23.25	137.62	55.08	169.86	35.35	221.02	106.30	208.76	32.89	279.62
											141.16

Table A.7 Continued
 Panel B: Crude Oil WTI (Consensus Survey 1995M8-2015M2) $b = 3$; $B = 60$

Dep. Var.	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$			3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$			4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$			
<i>Constant</i>	-0.04 [-2.926]**	-0.01 [-0.637]	-0.02 [-1.489]	-0.13 [-0.076]	-0.04 [-2.493]*	0 [0.042]	-0.01 [-0.575]	-0.04 [-2.287]*	0 [0.458]	0.13 [0.078]
<i>momentum_{t,t-3}</i>	-0.49 [-5.619]**	-0.44 [-5.434]**	-0.44 [-5.319]**	-0.33 [-5.319]**	-0.57 [-6.994]**		-0.5 [-6.842]**	-0.64 [-7.586]**	-0.54 [-7.632]**	-0.3 [-4.935]**
<i>value_{t,t-60}</i>		-0.09 [-4.918]**	-0.06 [-4.603]**	-0.05 [-2.913]**		-0.12 [-5.983]**	-0.09 [-6.116]**	-0.07 [-3.979]**	-0.15 [-6.814]**	-0.1 [-5.099]**
<i>inventories_t</i>				0.01 [0.059]			0 [-0.004]			-0.01 [-0.080]
<i>Δind.prod_{t,t-3}</i>				1.36 [1.156]			1.02 [1.091]			0.73 [0.869]
<i>slope_t^(+2Q)</i>				1.06 [2.905]**						
<i>slope_t^(+3Q)</i>							0.93 [4.995]**			
<i>slope_t^(+4Q)</i>										0.86 [6.146]**
Observations:	79	79	79	79	79	79	79	79	79	79
R-squared:	0.47	0.20	0.56	0.66	0.50	0.27	0.63	0.49	0.34	0.82
F-statistic:	68.16	19.68	47.97	27.99	76.12	28.81	64.85	74.9	39.03	68.76

Table A.7 Continued
Panel C: Copper (Bloomberg Survey 2006M9-2015M2) $b = 3$; $B = 60$

Dep. Var:	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$				3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$				4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$			
<i>Constant</i>	-0.01 [-1.101]	0.05 [4.89]**	0.04 [4.590]**	-0.14 [-0.963]	-0.01 [-0.417]	0.07 [4.576]**	0.06 [4.404]**	-0.20 [-0.900]	-0.01 [-0.336]	0.08 [3.960]**	0.06 [3.955]**	-0.17 [-0.514]
<i>momentum_{t-3}</i>	-0.32 [-5.007]**		-0.28 [-7.406]**	-0.29 [-7.457]**	-0.35 [-6.135]**		-0.30 [-6.963]**	-0.22 [-5.711]**	-0.41 [-7.351]**		-0.35 [-5.390]**	-0.25 [-4.349]**
<i>value_{t-60}</i>		-0.07 [-6.393]**	-0.06 [-8.185]**	-0.06 [-7.406]**		-0.09 [-7.869]**	-0.08 [-8.679]**	-0.06 [-7.187]**		-0.10 [-6.680]**	-0.08 [-5.505]**	-0.05 [-4.031]**
<i>inventories_t</i>				0.01 [1.194]				0.02 [1.085]				0.02 [0.623]
<i>$\Delta \ln d.prod_{t,t-3}$</i>				0.02 [0.096]				-0.99 [-3.869]**				-0.92 [-1.905]
<i>slope_t^(+2Q)</i>				-0.28 [-0.232]								
<i>slope_t^(+3Q)</i>								1.08 [1.750]				
<i>slope_t^(+4Q)</i>												1.86 [2.149]*
Observations:	102	102	102	102	102	102	102	102	102	102	102	102
R-squared:	0.39	0.34	0.63	0.63	0.35	0.38	0.62	0.68	0.31	0.3	0.52	0.61
F-statistic:	64.19	52.42	83.55	32.92	53.14	62.23	82.24	41.21	45.14	43.59	54.71	29.62

Table A.7 Continued
Panel D: Copper (Consensus Survey 1995M8-2015M2) $b = 3$; $B = 60$

Dep. Var.	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$			3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$			4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$					
<i>Constant</i>	0.00 [-0.028]	0.02 [1.832]	0.01 [1.430]	-0.32 [-1.448]	0.00 [0.055]	0.03 [2.677]**	0.02 [2.345]*	-0.39 [-1.566]	0.00 [-0.205]	0.04 [2.874]**	0.03 [2.388]*	-0.56 [-2.271]*
<i>momentum_{t,t-3}</i>	-0.40 [-8.317]**		-0.36 [-8.024]**	-0.35 [-6.780]**	-0.47 [-10.286]**		-0.41 [-9.420]**	-0.32 [-6.394]**	-0.53 [-9.308]**		-0.44 [-8.632]**	-0.28 [-6.360]**
<i>value_{t,t-60}</i>		-0.05 [-4.084]**	-0.03 [-3.349]**	-0.01 [-0.981]		-0.07 [-5.852]**	-0.05 [-5.635]**	-0.02 [-1.983]		-0.10 [-6.589]**	-0.07 [-6.691]**	-0.03 [-2.618]*
<i>inventories_t</i>				0.03 [1.505]				0.03 [1.684]				0.05 [2.446]*
<i>$\Delta ind.prod_{t,t-3}$</i>				0.44 [0.891]				-0.17 [-0.324]				-0.61 [-1.347]
<i>slope_t^(+2Q)</i>				0.33 [0.586]								
<i>slope_t^(+3Q)</i>								0.64 [2.363]*				
<i>slope_t^(+4Q)</i>												0.80 [3.990]**
Observations:	79	79	79	79	79	79	79	79	79	79	79	79
R-squared:	0.47	0.15	0.51	0.57	0.47	0.25	0.58	0.75	0.41	0.30	0.57	0.85
F-statistic:	66.93	13.33	39.76	19.56	66.99	25.65	51.91	44.21	53.28	33.31	49.78	81.35

Table A.8 Survey Forecasts and Net Long Positions

The Table reports time-series regression results of variations in the Commitment of Traders (COT) Net Futures positions on contemporaneous expectations of returns at different horizons. We control for the 2008 financial crises by adding a dummy variable which interacts with the expectations of returns. The dummy variable is equal to one from January 2008 to December 2009. The variations in net futures positions are scaled by the previous period open interest as defined in equation 1.13. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator. $\Delta Net Long_{i,t}^* = c + \beta_1 \hat{r}_t^{(m)} + \beta_2 dummy_t \times \hat{r}_t^{(m)} + \epsilon_{i,t}$ $i \in \{H \equiv Hedgers; LS \equiv Large Speculators; SS \equiv Small Speculators\}$

Panel A: Crude Oil WTI (Bloomberg Survey 2006M9-2015M2)									
Dep. Var:	2 Quarters Forecasting Horizon			3 Quarters Forecasting Horizon			4 Quarters Forecasting Horizon		
	H	LS	SS	H	LS	SS	H	LS	SS
	$\Delta Net Long_{i,t}^*$			$\Delta Net Long_{i,t}^*$			$\Delta Net Long_{i,t}^*$		
<i>Constant</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[-0.831]	[0.842]	[0.341]	[-0.966]	[1.004]	[0.353]	[-1.090]	[1.099]	[0.547]
$\hat{r}_t^{(+2Q)}$	0.10	-0.09	-0.01						
	[2.111]*	[-2.069]*	[-1.330]						
$\hat{r}_t^{(+2Q)} \times dummy$	-0.12	0.09	0.03						
	[-1.891]	[1.419]	[1.985]*						
$\hat{r}_t^{(+3Q)}$				0.07	-0.06	-0.01			
				[2.043]*	[-1.891]	[-1.822]			
$\hat{r}_t^{(+3Q)} \times dummy$				-0.09	0.06	0.02			
				[-2.036]*	[1.472]	[2.074]*			
$\hat{r}_t^{(+4Q)}$							0.06	-0.05	-0.01
							[2.093]*	[-1.898]	[-2.078]*
$\hat{r}_t^{(+4Q)} \times dummy$							-0.08	0.06	0.02
							[-2.290]*	[1.732]	[2.180]*
Observations:	102	102	102	102	102	102	102	102	102
R-squared:	0.07	0.06	0.03	0.05	0.03	0.03	0.04	0.03	0.03
F-statistic:	3.5	2.93	1.36	2.4	1.7	1.38	2.32	1.53	1.33

Table A.8 Continued										
Panel B: Copper (Bloomberg Survey 2006M9-2015M2)										
Dep. Var:	2 Quarters Forecasting Horizon			3 Quarters Forecasting Horizon			4 Quarters Forecasting Horizon			$\Delta Net Position_{i,t}$
	H	LS	SS	H	LS	SS	H	LS	SS	
	$\Delta Net Position_{i,t}$			$\Delta Net Position_{i,t}$			$\Delta Net Position_{i,t}$			
<i>Constant</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[0.562]	[-0.793]	[0.273]	[0.096]	[-0.231]	[0.361]	[0.148]	[-0.240]	[0.198]	
$\hat{r}_t^{(+2Q)}$	0.45	-0.37	-0.08							
	[6.346]**	[-6.731]**	[-3.226]**							
$\hat{r}_t^{(+2Q)} \times dummy$	-0.12	0.03	0.10							
	[-1.390]	[0.393]	[2.829]**							
$\hat{r}_t^{(+3Q)}$				0.36	-0.29	-0.07				
				[6.105]**	[-6.331]**	[-2.836]**				
$\hat{r}_t^{(+3Q)} \times dummy$				-0.14	0.05	0.09				
				[-1.404]	[0.609]	[2.794]**				
$\hat{r}_t^{(+4Q)}$							0.24	-0.18	-0.06	
							[4.548]**	[-4.287]**	[-3.730]**	
$\hat{r}_t^{(+4Q)} \times dummy$							-0.04	-0.02	0.07	
							[-0.460]	[-0.308]	[2.503]*	
Observations:	102	102	102	102	102	102	102	102	102	
R-squared:	0.19	0.2	0.07	0.15	0.15	0.06	0.12	0.12	0.07	
F-statistic:	11.48	12.38	3.56	8.43	9.00	3.33	6.82	6.82	3.79	

Table A.8 Continued
Panel C: Crude Oil WTI (Consensus Survey 1995M8-2015M2)

Dep. Var:	2 Quarters Forecasting Horizon			3 Quarters Forecasting Horizon			4 Quarters Forecasting Horizon		
	H	LS	SS	H	LS	SS	H	LS	SS
$\Delta Net Position_{i,t}$									
<i>Constant</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[-0.423]	[0.450]	[0.268]	[-0.600]	[0.707]	[0.242]	[-0.628]	[0.751]	[0.233]
$\hat{r}_t^{(+2Q)}$	0.06	-0.06	-0.01						
	[1.627]	[-1.830]	[-0.436]						
$\hat{r}_t^{(+2Q)} \times dummy$	0.05	-0.11	0.06						
	[0.473]	[-1.480]	[1.329]						
$\hat{r}_t^{(+3Q)}$				0.04	-0.03	-0.01			
				[1.141]	[-1.235]	[-0.425]			
$\hat{r}_t^{(+3Q)} \times dummy$				0.08	-0.11	0.03			
				[1.120]	[-2.193]*	[0.709]			
$\hat{r}_t^{(+4Q)}$							0.04	-0.03	-0.01
							[1.106]	[-1.198]	[-0.434]
$\hat{r}_t^{(+4Q)} \times dummy$							0.06	-0.08	0.02
							[1.151]	[-1.991]	[0.527]
Observations:	79	79	79	79	79	79	79	79	79
R-squared:	0.03	0.05	0.01	0.02	0.03	0.01	0.02	0.03	0.00
F-statistic:	1.10	1.88	0.44	0.76	1.32	0.20	0.81	1.31	0.13

Dep. Var:	2 Quarters Forecasting Horizon			3 Quarters Forecasting Horizon			4 Quarters Forecasting Horizon		
	H	LS	SS	H	LS	SS	H	LS	SS
	$\Delta Net Position_{i,t}$			$\Delta Net Position_{i,t}$			$\Delta Net Position_{i,t}$		
<i>Constant</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	[-0.192]	[0.288]	[-0.121]	[-0.164]	[0.260]	[-0.149]	[-0.134]	[0.232]	[-0.183]
$\hat{r}_t^{(+2Q)}$	0.55	-0.49	-0.06						
	[2.964]**	[-3.323]**	[-1.089]						
$\hat{r}_t^{(+2Q)} \times dummy$	-0.33	0.25	0.07						
	[-1.552]	[1.404]	[1.209]						
$\hat{r}_t^{(+3Q)}$				0.34	-0.30	-0.04			
				[2.252]*	[-2.521]*	[-0.866]			
$\hat{r}_t^{(+3Q)} \times dummy$				-0.08	0.03	0.05			
				[-0.503]	[0.238]	[0.891]			
$\hat{r}_t^{(+4Q)}$							0.17	-0.15	-0.03
							[1.378]	[-1.455]	[-0.778]
$\hat{r}_t^{(+4Q)} \times dummy$							0.11	-0.15	0.04
							[0.781]	[-1.148]	[0.716]
Observations:	79	79	79	79	79	79	79	79	79
R-squared:	0.12	0.15	0.02	0.07	0.09	0.01	0.04	0.04	0.01
F-statistic:	5.29	6.75	0.63	3.05	3.82	0.43	1.46	1.77	0.34

Fig. A.6 Bloomberg Forecasts Dispersion: WTI

The chart shows the evolution of Crude Oil WTI price and analysts forecasts dispersion measured as the standard deviation of price forecasts divided by the spot price. The sample is monthly from 2006M10 to 2015M2.




Fig. A.7 Consensus Forecasts Dispersion: WTI

The chart shows the evolution of Crude Oil WTI price and analysts forecasts dispersion measured as the standard deviation of price forecasts divided by the spot price. The sample is monthly from 1995M8 2015M2.




Fig. A.8 Implied Volatility and Bloomberg Forecasts Dispersion: WTI

The chart shows the evolution of Crude Oil WTI implied volatility and analysts forecasts dispersion measured as the standard deviation of price forecasts divided by the spot price. The sample is monthly from 2006M10 to 2015M2.




Fig. A.9 Implied Volatility and Consensus Forecasts Dispersion: WTI

The chart shows the evolution of Crude Oil WTI implied volatility and analysts forecasts dispersion measured as the standard deviation of price forecasts divided by the spot price. The sample is monthly from 2006M10 to 2015M2.




Table A.9 Crude Oil WTI Implied Volatility and Forecasts Dispersion

This Table reports results for time-series regressions of WTI implied volatility on forecasts dispersion, realised volatility, the slope of the futures curve, lagged implied volatility and inventories levels:

$IV_t = c + \beta_1 \hat{\sigma}_t^{(+2Q)} + \beta_2 \sigma_t + \beta_3 slope_t^{(+2Q)} + \beta_4 IV_{t-1} + \beta_5 inventories_t + \varepsilon_t$ where IV_t represents the crude oil implied volatility index at time t of the WTI futures options, $\hat{\sigma}_t^{(+2Q)}$ represents the analysts forecasts dispersion measured as the standard deviation of price forecasts divided by the spot price, σ_t is the realised volatility measured as the monthly average of daily squared log-returns on the front month WTI contract and $slope_t^{(+2Q)}$ is the slope of the futures curve as defined in equation 1.5. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Dep. Var.	Crude Oil WTI Implied Volatility					
	Bloomberg (2006M9-2015M2)					
	IV_t	IV_t	IV_t	IV_t	IV_t	IV_t
<i>Constant</i>	0.09 [3.393]**	0.07 [4.716]**	0.08 [5.121]**	0.03 [1.716]	2.63 [3.150]**	0.14 [3.442]**
$\hat{\sigma}_t^{(+2Q)}$	3.05 [8.174]**	1.18 [6.054]**	1.11 [5.645]**	0.63 [4.585]**	0.54 [3.912]**	2.44 [4.975]**
σ_t		2.59 [7.871]**	2.43 [6.973]**	1.65 [4.802]**	1.77 [7.634]**	2.72 [4.602]**
$slope_t^{(+2Q)}$			0.38 [2.075]*	-0.1 [-0.694]	0.08 [0.464]	-0.05 [-0.237]
IV_{t-1}				0.44 [4.729]**	0.37 [5.371]**	0.51 [5.330]**
<i>inventories_t</i>					-0.2 [-3.126]**	-0.27 [-4.789]**
Observations:	102	102	102	102	102	61
R-squared:	0.66	0.89	0.9	0.93	0.94	0.47
F-statistic:	194.16	405.93	280.6	344.43	321.94	52.78
					111.5	73.16
					121.68	138.54

Table A.10 Implied and Realized Volatility Spread and Forecasts Dispersion

This Table reports results for time-series regressions of WTI implied-realised volatility spread on forecasts dispersion, the slope of the futures curve and lagged spread values:

$spread_t = c + \beta_1 \hat{\sigma}_t^{(+2Q)} + \beta_2 slope_t^{(+2Q)} + \beta_3 spread_{t-1} + \beta_4 inventories_t + \varepsilon_t$ where $spread_t$ represents the difference between crude oil implied volatility index at time t of the WTI futures options and the realised volatility on the front month contract. The realised volatility is measured as the monthly average of daily squared log-returns on the front month WTI contract. $\hat{\sigma}_t^{(+2Q)}$ represents the analysts forecasts dispersion measured as the standard deviation of price forecasts divided by the spot price, and $slope_t^{(+2Q)}$ is the slope of the futures curve as defined in equation 1.5. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Dep. Var.	Crude Oil WTI Implied - Realized Volatility Spread											
	Bloomberg (2006M9-2014M2)						Consensus (1995M8-2014M2)					
	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$	$spread_t$
<i>Constant</i>	0.08 [4.002]**	0.08 [4.002]**	0.11 [5.763]**	0.02 [1.096]	1.86 [2.353]*	0.12 [3.580]**	0.12 [3.580]**	0.13 [4.085]**	0 [0.183]	0 [0.183]	3.34 [4.686]**	3.34 [4.686]**
$\hat{\sigma}_t^{(+2Q)}$	2.33 [8.764]**	2.33 [8.764]**	1.86 [6.637]**	0.81 [3.807]**	0.77 [3.618]**	1.83 [4.439]**	1.83 [4.439]**	1.73 [4.581]**	0.8 [2.489]*	0.8 [2.489]*	0.99 [3.546]**	0.99 [3.546]**
$slope_t^{(+2Q)}$			0.94 [3.407]**	0.03 [0.208]	0.17 [1.115]			0.16 [0.482]	-0.27 [-1.974]	-0.27 [-1.974]	0.09 [0.742]	0.09 [0.742]
$spread_{t-1}$				0.69 [9.828]**	0.63 [9.524]**			0.74 [10.373]**	0.74 [10.373]**	0.74 [10.373]**	0.55 [7.020]**	0.55 [7.020]**
<i>inventories_t</i>					-0.14 [-2.323]*						-0.26 [-4.676]**	-0.26 [-4.676]**
Observations:	102	102	102	102	102	61	61	61	61	61	61	61
R-squared:	0.66	0.66	0.72	0.86	0.87	0.44	0.44	0.45	0.81	0.81	0.85	0.85
F-statistic:	196.46	196.46	126.96	207.29	165.38	46.64	46.64	23.33	79.58	79.58	80.09	80.09

Table A.11 Simulation: Predictive Regressions - Rational Traders

We regress simulated future realised returns on rational expectations of returns under different model specifications. The structural parameter γ governs the share of trading activity undertaken by rational traders ($1 - \gamma$ governs the feedback traders activity) and α the autoregressive coefficient of the fundamentals process. We run 1000 repetitions and for each repetition we simulate 200 observations burning the first 100 realizations. The fundamentals have AR(1) dynamics with zero mean and Gaussian white noise shocks with unit variance. The Table reports mean values across repetitions.

Calibration	<i>I</i>	<i>II</i>	<i>III</i>	<i>IV</i>	<i>V</i>
Agents	R.E.	R.E. + Feedback	R.E.	R.E. + Feedback	R.E. + Feedback
Fundamentals	Random Walk	Random Walk	Mean Reverting	Mean Reverting	Mean Reverting
Dep. Var.	R_{t+1}	R_{t+1}	R_{t+1}	R_{t+1}	R_{t+1}
$E_t[R_{t+1}]$	-	1.24	1.81	1.36	1.17
t-stat	-	1.14	2.10	2.38	2.78
R^2	0	0.02	0.05	0.06	0.08
Structural Parameters					
γ	1	0.95	1	0.95	0.90
α	1	1	0.95	0.95	0.98

Table A.12 Simulation: Predictive Regressions - Feedback Traders

We regress simulated future realised returns on positive feedback traders expectations of returns (coinciding with past cumulated returns) under different model specifications. The structural parameter γ governs the share of trading activity undertaken by rational traders ($1 - \gamma$ governs the feedback traders activity) and α the autoregressive coefficient of the fundamentals process. We run 1000 repetitions and for each repetition we simulate 200 observations burning the first 100 realizations. The fundamentals have AR(1) dynamics with zero mean and Gaussian white noise shocks with unit variance. The Table reports mean values across repetitions.

Calibration Agents Fundamentals	I		II		III		IV		V	
	R.E.	Random Walk	R.E. + Feedback	Random Walk	R.E.	Mean Reverting	R.E. + Feedback	Mean Reverting	R.E. + Feedback	Mean Reverting
Dep. Var.	R_{t+1}		R_{t+1}		R_{t+1}		R_{t+1}		R_{t+1}	
Cum. Ret.	-0.01		-0.07		-0.03		-0.09		-0.16	
t-stat	-0.19		-1.14		-0.50		-1.35		-2.42	
R^2	0.01		0.02		0.01		0.03		0.06	
Structural Parameters										
γ	1		0.95		1		0.95		0.90	
α	1		1		0.95		0.95		0.98	

Table A.13 Simulation: Expectations of Returns and Past Cumulated Returns

We regress simulated expectations of returns of rational agents under different model specifications on past simulated cumulated returns. The structural parameter γ governs the share of trading activity undertaken by rational traders ($1 - \gamma$ governs the feedback traders activity) and α the autoregressive coefficient of the fundamentals process. We run 1000 repetitions and for each repetition we simulate 200 observations burning the first 100 realizations. The fundamentals have AR(1) dynamics with zero mean and Gaussian white noise shocks with unit variance. The Table reports mean values across repetitions.

Calibration Agents	I		II		III		IV		V	
	R.E.	Random Walk	R.E. + Feedback	Random Walk	R.E.	Mean Reverting	R.E. + Feedback	Mean Reverting	R.E. + Feedback	Mean Reverting
Fundamentals										
Dep. Var.	$E_t[R_{t+1}]$		$E_t[R_{t+1}]$		$E_t[R_{t+1}]$		$E_t[R_{t+1}]$		$E_t[R_{t+1}]$	
$p_t - p_{t-3}$	-		-0.06		-0.01		-0.07		-0.14	
t-stat	-		-196.57		-3.06		-14.45		-48.61	
$p_t - p_{t-60}$	-		0.00		-0.02		-0.03		-0.01	
t-stat	-		-0.25		-12.42		-11.92		-10.66	
R^2	1		0.99		0.61		0.81		0.96	
Structural Parameters										
γ	1		0.95		1		0.95		0.90	
α	1		1		0.95		0.95		0.98	

Table A.14 Determinants of Analysts' Expectations: Look-Back

This Table reports results for time-series regressions of survey expectations of returns on the time-series momentum factor computed using different look-back periods: $\hat{p}_t^{(m)} = c + \beta_1 momentum_{t-1-b} + \varepsilon_t$. Expectations of returns are defined in equation 1.4. The momentum factor is defined in equation 1.9. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Panel A: Crude Oil WTI (Bloomberg Survey 2006M9-2015M2)				
Dep. Var:	2 Quarters Forecasting Horizon $\hat{p}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{p}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{p}_t^{(+4Q)}$	
<i>Constant</i>	-0.01 [-0.756]	0.00 [0.543]	0.01 [0.567]	0.02 [1.532]
<i>momentum_{t-3}</i>	-0.34 [-8.463]**	-0.42 [-9.757]**	-0.42 [-9.757]**	-0.50 [-11.638]**
<i>momentum_{t-6}</i>	-0.19 [-4.91]**	-0.26 [-7.368]**	-0.26 [-7.368]**	-0.3 [-6.385]**
<i>momentum_{t-12}</i>	-0.10 [-2.269]*	-0.10 [-2.269]*	-0.15 [-2.908]**	-0.16 [-2.505]*
Observations:	102	102	102	102
R-squared:	0.61	0.63	0.56	0.68
F-statistic:	154.93	73.56	127.15	208.76
				126.93
				28.27
Panel B: Crude Oil WTI (Consensus Survey 1995M8-2015M2)				
Dep. Var:	2 Quarters Forecasting Horizon $\hat{p}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{p}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{p}_t^{(+4Q)}$	
<i>Constant</i>	-0.04 [-2.926]**	-0.04 [-2.880]**	-0.03 [-2.189]*	-0.02 [-1.788]
<i>momentum_{t-3}</i>	-0.49 [-5.619]**	-0.57 [-6.994]**	-0.64 [-7.586]**	-0.64 [-7.586]**
<i>momentum_{t-6}</i>	-0.30 [-5.158]**	-0.37 [-6.758]**	-0.42 [-7.523]**	-0.42 [-7.523]**
<i>momentum_{t-12}</i>	-0.20 [-5.352]**	-0.20 [-5.352]**	-0.26 [-6.761]**	-0.29 [-7.316]**
Observations:	79	79	79	79
R-squared:	0.47	0.46	0.43	0.53
F-statistic:	68.16	66.76	58.08	76.12
				95.52
				85.36
				74.9
				99.05
				94.47
Panel C: Copper (Bloomberg Survey 2006M9-2015M2)				
Dep. Var:	2 Quarters Forecasting Horizon $\hat{p}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{p}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{p}_t^{(+4Q)}$	
<i>Constant</i>	-0.01 [-1.101]	-0.01 [-0.521]	-0.01 [-0.360]	-0.01 [-0.336]
<i>momentum_{t-3}</i>	-0.32 [-5.007]**	-0.35 [-6.135]**	-0.41 [-7.351]**	-0.41 [-7.351]**
<i>momentum_{t-6}</i>	-0.20 [-4.363]**	-0.24 [-7.666]**	-0.26 [-4.896]**	-0.26 [-4.896]**
<i>momentum_{t-12}</i>	-0.08 [-2.233]*	-0.14 [-3.300]**	-0.14 [-3.300]**	-0.18 [-2.725]**
Observations:	102	102	102	102
R-squared:	0.39	0.31	0.37	0.18
F-statistic:	64.19	45.04	57.85	45.14
				37.72
				22.17
Panel D: Copper (Consensus Survey 1995M8-2015M2)				
Dep. Var:	2 Quarters Forecasting Horizon $\hat{p}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{p}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{p}_t^{(+4Q)}$	
<i>Constant</i>	0.00 [-0.028]	0.00 [0.221]	0.00 [0.221]	0.01 [0.568]
<i>momentum_{t-3}</i>	-0.40 [-8.317]**	-0.47 [-10.286]**	-0.53 [-9.308]**	-0.53 [-9.308]**
<i>momentum_{t-6}</i>	-0.21 [-4.316]**	-0.27 [-6.029]**	-0.33 [-6.885]**	-0.33 [-6.885]**
<i>momentum_{t-12}</i>	-0.09 [-2.648]**	-0.16 [-4.717]**	-0.16 [-4.717]**	-0.21 [-5.985]**
Observations:	79	79	79	79
R-squared:	0.47	0.30	0.11	0.47
F-statistic:	66.93	33.35	9.75	66.99
				46.16
				24.78
				53.28
				49.36
				35.69

Table A.15 Determinants of Analysts' Expectations: Perfect Foresight

This Table reports the results for time-series regressions of survey expectations of returns on the same explanatory variables used in table A.7 but assuming that analysts have perfect foresight on future inventories and industrial production growth: $\hat{r}_t^{(m)} = c + \beta_1 \text{momentum}_{t-b} + \beta_2 \text{value}_{t-b} + \beta_3 \text{slope}_t^{(m)} + \beta_4 \text{inventories}_{t+m} + \beta_5 \Delta \text{ind. prod.}_{t+m} + \epsilon_t$. Expectations of returns are defined in equation 1.4. The momentum and value variables are defined in equations 1.9 and 1.10 respectively. Slope is defined in equation 1.5. Inventories are the the commodity specific logarithms of inventory levels at time $t + m$ and $\Delta \text{ind. prod.}_{t+m}$ are the log-differences in world industrial production between time $t + m$ and t . The proxy for value is obtained by setting $B = 60$ as in Asness et al. [6]. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Panel A: Crude Oil WTI (Mean Forecast) $b = 3; B = 60$						Panel B: Copper (Mean Forecast) $b = 3; B = 60$					
Bloomberg (2006M9-2015M2)			Consensus (1995M8-2015M2)			Bloomberg (2006M9-2015M2)			Consensus (1995M8-2015M2)		
Dep. Var:	$\hat{r}_t^{(+2Q)}$	$\hat{r}_t^{(+3Q)}$	$\hat{r}_t^{(+4Q)}$	$\hat{r}_t^{(+2Q)}$	$\hat{r}_t^{(+3Q)}$	$\hat{r}_t^{(+4Q)}$	$\hat{r}_t^{(+2Q)}$	$\hat{r}_t^{(+3Q)}$	$\hat{r}_t^{(+4Q)}$	$\hat{r}_t^{(+2Q)}$	$\hat{r}_t^{(+4Q)}$
<i>Constant</i>	-0.01 [-1.490]	0.00 [0.071]	0.02 [2.865]**	-0.02 [-1.509]	-0.01 [-0.490]	0.00 [0.333]	0.04 [4.697]**	0.05 [3.373]**	0.04 [2.057]*	0.01 [0.900]	0.02 [2.444]*
<i>momentum</i> _{t,t-3}	-0.25 [-7.122]**	-0.27 [-7.336]**	-0.36 [-1.5295]**	-0.24 [-2.341]*	-0.24 [-2.982]**	-0.23 [-3.147]**	-0.25 [-6.097]**	-0.28 [-5.844]**	-0.32 [-4.126]**	-0.31 [-5.689]**	-0.34 [-7.967]**
<i>value</i> _{t,t-60}	-0.04 [-3.094]**	-0.06 [-5.886]**	-0.06 [-6.525]**	-0.05 [-3.165]**	-0.08 [-5.031]**	-0.1 [-6.726]**	-0.07 [-6.834]**	-0.06 [-5.284]**	-0.05 [-3.380]**	-0.01 [-0.965]	-0.02 [-2.349]*
<i>slope</i> _t ^(+2Q)	0.63 [3.410]**			1.08 [3.066]**			0.46 [0.498]		0.45 [0.735]		
<i>inventories</i> _{t+2Q}	0.29 [3.919]**			0.15 [1.255]			0 [-0.046]		0.03 [2.068]*		
<i>Δind. prod.</i> _{t+2Q,t}	-0.25 [-1.326]			-0.36 [-1.101]			-0.33 [-1.821]		-0.35 [-1.260]		
<i>slope</i> _t ^(+3Q)	0.8 [5.734]**			1.01 [5.970]**			1.26 [2.115]*		0.65 [2.346]*		
<i>inventories</i> _{t+3Q}	0.28 [4.537]**			0.02 [0.166]			0.02 [0.930]		0.04 [2.689]**		
<i>Δind. prod.</i> _{t+3Q,t}	-0.14 [-1.131]			-0.24 [-1.236]			0.04 [0.216]		-0.11 [-0.587]		
<i>slope</i> _t ^(+4Q)	0.62 [8.577]**			0.96 [6.856]**			1.68 [1.738]		0.92 [4.442]**		
<i>inventories</i> _{t+4Q}	0.27 [3.436]**			-0.08 [-0.591]			0.03 [0.975]		0.04 [3.284]**		
<i>Δind. prod.</i> _{t+4Q,t}	-0.08 [-0.617]			-0.15 [-0.940]			0.19 [0.809]		0.08 [0.577]		
Observations:	102	102	102	79	79	79	102	102	102	79	79
R-squared:	0.81	0.87	0.88	0.65	0.76	0.82	0.63	0.66	0.59	0.59	0.85
F-statistic:	77.99	116.4	125.18	26.12	43.57	60.87	30.6	33.83	24.97	20.51	48.93

Table A.16 Determinants of Analysts' Expectations: Skipping Most Recent Returns

This Table reports a robustness test for the relationship between survey-based expectations of returns and past cumulated returns when the momentum and value factors are computed skipping the most recent return. More specifically the table reports results for time-series regressions of survey expectations of returns on momentum and value: $\hat{r}_t^{(m)} = c + \beta_1 \text{momentum}_{t-1:t-b-1} + \beta_2 \text{value}_{t-1:t-b-1} + \varepsilon_t$. t-Stat are reported in brackets. Standard errors are computed using the Newey-West HAC estimator.

Panel A: Crude Oil WTI (Bloomberg Survey 2006M9-2015M2) $b = 3; B = 60$									
Dep. Var:	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$
<i>Constant</i>	0.00 [-0.380]	0.01 [0.119]	0.00 [0.724]	0.01 [1.676]	0.03 [2.214]*	0.02 [1.638]	0.03 [2.331]*	0.05 [3.675]**	0.04 [3.675]**
<i>mom_{t-1:t-b-1}</i>	-0.24 [-4.963]**	-0.23 [-5.158]**	-0.34 [-6.769]**	-0.34 [-7.295]**	-0.32 [-7.295]**	-0.42 [-7.534]**	-0.42 [-7.534]**	-0.4 [-8.064]**	-0.4 [-8.064]**
<i>value_{t-1:t-b-1}</i>	-0.05 [-1.709]	-0.04 [-1.988]*	-0.07 [-2.396]*	-0.07 [-3.122]**	-0.06 [-3.122]**	-0.08 [-3.135]**	-0.08 [-3.135]**	-0.07 [-3.135]**	-0.07 [-3.135]**
Observations:	102	102	102	102	102	102	102	102	102
R-squared:	0.28	0.10	0.35	0.40	0.15	0.51	0.46	0.15	0.57
F-statistic:	39.72	10.77	27.1	67.11	18.16	52.3	86.73	17.24	65.1
Panel B: Crude Oil WTI (Consensus Survey 1995M8-2015M2) $b = 3; B = 60$									
Dep. Var:	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$
<i>Constant</i>	-0.04 [-2.721]**	-0.02 [-1.062]	-0.02 [-1.682]	-0.02 [-2.402]*	-0.04 [-0.369]	-0.02 [-1.008]	-0.04 [-2.244]*	-0.02 [-0.018]	-0.04 [-0.551]
<i>mom_{t-1:t-b-1}</i>	-0.29 [-3.804]**	-0.25 [-3.530]**	-0.39 [-4.828]**	-0.39 [-4.828]**	-0.33 [-4.565]**	-0.44 [-5.332]**	-0.33 [-4.565]**	-0.37 [-5.119]**	-0.37 [-5.119]**
<i>value_{t-1:t-b-1}</i>	-0.07 [-3.892]**	-0.05 [-3.024]**	-0.05 [-3.024]**	-0.05 [-3.024]**	-0.10 [-4.980]**	-0.07 [-4.152]**	-0.05 [-3.024]**	-0.12 [-5.693]**	-0.09 [-5.086]**
Observations:	79	79	79	79	79	79	79	79	79
R-squared:	0.21	0.12	0.26	0.29	0.18	0.38	0.31	0.24	0.43
F-statistic:	20.38	10.63	13.6	31.74	17.13	22.95	34.28	23.66	28.51
Panel C: Copper (Bloomberg Survey 2006M9-2015M2) $b = 3; B = 60$									
Dep. Var:	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$
<i>Constant</i>	-0.01 [-0.972]	0.04 [3.246]**	-0.01 [-0.388]	0.06 [4.254]**	0.05 [3.964]**	-0.01 [-0.297]	0.07 [3.804]**	0.03 [2.244]*	0.02 [1.791]
<i>mom_{t-1:t-b-1}</i>	-0.23 [-5.538]**	-0.19 [-5.297]**	-0.29 [-6.266]**	-0.25 [-3.902]**	-0.32 [-3.581]**	-0.27 [-2.265]*	-0.24 [-3.330]**	-0.30 [-4.410]**	-0.31 [-4.694]**
<i>value_{t-1:t-b-1}</i>	-0.06 [-4.984]**	-0.05 [-5.128]**	-0.08 [-5.903]**	-0.07 [-5.512]**	-0.08 [-5.512]**	-0.07 [-5.005]**	-0.04 [-1.491]	-0.06 [-3.693]**	-0.06 [-3.338]**
Observations:	102	102	102	102	102	102	102	102	102
R-squared:	0.19	0.24	0.38	0.24	0.28	0.45	0.19	0.36	0.42
F-statistic:	24.05	31.9	29.94	31.83	39.25	40.9	23.4	29.07	27.69
Panel D: Copper (Consensus Survey 1995M8-2015M2) $b = 3; B = 60$									
Dep. Var:	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$	2 Quarters Forecasting Horizon $\hat{r}_t^{(+2Q)}$	3 Quarters Forecasting Horizon $\hat{r}_t^{(+3Q)}$	4 Quarters Forecasting Horizon $\hat{r}_t^{(+4Q)}$
<i>Constant</i>	0.00 [0.036]	0.01 [1.281]	0.01 [0.906]	0.00 [0.145]	0.03 [2.082]*	0.02 [1.694]	0.00 [-0.065]	0.03 [2.244]*	0.02 [1.791]
<i>mom_{t-1:t-b-1}</i>	-0.24 [-3.330]**	-0.21 [-2.932]**	-0.21 [-2.932]**	-0.30 [-4.410]**	-0.26 [-3.840]**	-0.26 [-3.840]**	-0.38 [-5.201]**	-0.38 [-5.201]**	-0.31 [-4.694]**
<i>value_{t-1:t-b-1}</i>	-0.04 [-2.586]*	-0.02 [-1.491]	-0.02 [-1.491]	-0.02 [-1.491]	-0.06 [-3.693]**	-0.04 [-2.653]**	-0.04 [-2.653]**	-0.08 [-4.236]**	-0.06 [-3.338]**
Observations:	79	79	79	79	79	79	79	79	79
R-squared:	0.24	0.09	0.26	0.29	0.17	0.36	0.32	0.22	0.42
F-statistic:	24.57	7.24	13.67	31.64	16.19	21.53	35.96	21.76	27.09

Chapter 2

Expected Spot Prices and the Dynamics of Commodity Risk Premia¹

2.1 Introduction

The way in which investors form expectations about future commodity prices is of great interest to economists and market participants at least since Keynes [66]. Forward prices have been used extensively in economic models as an approximation of market beliefs.² However, the forward curve includes not only investors' expectations for the future, but also a component reflecting the compensation required by market participants for bearing the risk of uncertain fluctuations in spot prices, i.e. a risk premium.³ Whether this risk premium is positive, negative, or time-varying and driven by changes in economic fundamentals has

¹A research paper joint with Dr. D. Bianchi, entitled “Expected Spot Prices and the Dynamics of Commodity Risk Premia” is based on this Chapter and submitted for publication. The paper has been presented at the 2016 NBER Economics of Commodity Markets meeting, the 2016 European Winter Meeting of the Econometric Society, the Barcelona GSE Summer Forum, the Commodity and Energy Markets meeting at Oxford 2017 and the Workshop on the Financialization of Commodity Markets Workshop at Bolzano 2017.

²For instance, futures-based forecasts for the Oil price play a role in the policy decision making process at the ECB, see e.g. Svensson [92], at the Federal Reserve Board, see e.g. Bernanke [13], and at the International Monetary Fund, see e.g. IMF World Economic Outlook 2005.

³Throughout the Chapter we use the terms *risk premium* and *expected payoff* interchangeably. In fact, all these terms identify a payoff expected at time t as a compensation for a risk which materializes at maturity $t + h$. Differently, a realised payoff, or realised risk premium, couples the risk premium with any unanticipated deviation of the future spot price from the expected future spot price (see Section 2 for a more detailed discussion).

been controversial in the literature.⁴ This controversy stems from the fact that investors' expectations are not directly observable.

If in the first Chapter, we have shown that the dependence of expectations on past prices, can also arise under the limiting case of perfect rational expectations when extrapolative agents are present in the market; in the second Chapter we model price expectations to study the dynamics of commodity risk premia, by still taking into account past price dynamics but inside a more realistic framework of rational learning. More specifically, we assume that investors do form their expectations rationally according to a rational Muth's model, but they learn adaptively about the coefficients of the prices perceived law of motion from past prediction errors and as new fundamental data and prices become available.

Such model of learning, allows us to approximate the time-varying *ex-ante* risk premia – calculated as the spread between the futures price as of date t with maturity $t + h$ and expectations at time t on future spot prices over the same time-horizon – for a reasonably long sample period. Once the dynamics of ex-ante risk premia is extracted from futures prices, we investigate their determinants across investment horizons and commodities by using a dynamic linear regression framework, which features random-walk betas on a set of widely discussed economic risk factors. This approach allows us to shed some light on the dynamics and determinants of risk premia, that have been controversial in the literature.⁵

Our main results show that risk premia are time-varying, both across commodities and time-horizons, and their dynamics is predominantly driven by risks sharing mechanisms and the changing nature of market activity, as proxied by Open Interest (OI henceforth), Hedging Pressure (HP henceforth) and time-series Momentum (TSMOM henceforth). These

⁴See, e.g. Keynes [66], Hicks [55], Kaldor [61], Working [98], Brennan [17], Hsieh and Kulatilaka [60], Fama and French [40], Fama and French [42], Gorton et al. [47], Singleton [89], Szymanowska et al. [93] and Bakshi et al. [8] just to cite a few.

⁵See, e.g. Keynes [66], Hicks [55], Kaldor [61], Working [98], Brennan [17], Hsieh and Kulatilaka [60], Fama and French [40], Fama and French [42], Gorton et al. [47], Singleton [89], Szymanowska et al. [93] and Bakshi et al. [8] just to cite a few. The controversy stems from the fact that investors' expectations are not directly observable.

results hold after controlling for a variety of other commonly used proxies for risk factors, e.g. changes in inventories and realised volatility. Yet, we show that emerging markets, as proxied by the MSCI Emerging Market Index (MXEF), plays a sensible role for both WTI Oil and Copper, which is coherent with the increasing weight of emerging economies in the global economic growth and the presence of potential spillover effects to be associated with concerns about a worldwide economic slowdown.⁶ More generally, we provide evidence of heterogeneity in the dynamics of factor loadings in the time-series of commodity risk premia across both products and maturities.

Also, we compare the expected future spot prices obtained from our model with the cross-sectional average of survey forecasts provided by Bloomberg. In this respect, we show that, although with differences across commodities, our model generates conditional expectations which are broadly consistent with the average survey forecast from two to four quarters ahead.

Finally, we show that our model, in which individuals learn over time about the prices law of motion, compares favourably against alternative specifications for forecasting future spot prices. More precisely, an out-of-sample comparison of mean squared prediction errors against models in which expectations are based on either futures or current spot prices, or a spread of the two, demonstrates that the forecasts generated by our model reach a statistically significance 1% higher predictive R^2 on average across commodities and maturities. This result, possibly, rules out the concern that the model-implied ex-ante risk premia merely represent forecast errors which have nothing to do with investors' preferences or the actual expectations formation process. As a matter of fact, a further analysis clarifies that the expected payoffs extracted in this Chapter are highly correlated to the actual, realised, excess rolling returns in the same-maturity generic futures contract.

In terms of contribution, the second Chapter builds on a number of existing works such as Nerlove [80], Evans and Honkapohja [39], Sargent [87], Sargent and Williams [88], and Mal-

⁶China itself is the second largest economy and the second largest importer of both goods and commercial services.

mendier and Nagel [72], who consider a model of adaptive learning to explain the dynamics of expectations on inflation and more general macroeconomic outcomes. Also, our study is related to recent research that posits trading activity is the result of a learning process in which hedgers and speculators update their views as news about economic fundamentals, and prices become available (see, e.g. Singleton [89]). Finally, the second Chapter contributes also to the recent literature that aims at understanding the origins of unconditional realised commodity risk premia such as Carter et al. [22], Bessembinder [15], De Roon et al. [29], Acharya et al. [1], Hong and Yogo [59], Asness et al. [5], Basu and Miffre [10], Hamilton and Wu [53], Szymanowska et al. [93], and Bakshi et al. [8].

The rest of the Chapter is organized as follows. Section 2 discusses the motivation of this Chapter, while Section 3 introduces the model of learning as well as compares the implied expectations with the cross-sectional average of the Bloomberg's individual analysts forecasts. Section 4 represents the core of the Chapter and reports the empirical results. Section 5 concludes. We leave the details of the model derivation and further results to the Appendix.

2.2 Motivation

Let S_t denote the spot price of a given commodity at time t , and $F_t^{(h)}$ the price of a futures at time t with maturity $t + h$. The basis $F_t^{(h)} - S_t$ can be decomposed in two main components,

$$F_t^{(h)} - S_t = E_t [\Delta S_{t+h}] + \underbrace{F_t^{(h)} - E_t [S_{t+h}]}_{y_t^{(h)}} \quad (2.1)$$

with $E_t [S_{t+h}]$ the market aggregate expected spot price for time $t + h$, $y_t^{(h)}$ a risk premium component in dollar terms, and $E_t [\Delta S_{t+h}]$ the expected change in spot valuations between t and $t + h$. To the extent that one wants to investigate the origins of risk premia, equation (2.1) offers an ideal setting since directly isolates risk-related components in futures prices

conditioning on investors' expectations about the spot commodity.

One may argue that the ex-ante and realised payoff of a futures position are equivalent, such that we can indifferently use the spread between the spot price at maturity S_{t+h} and the futures price $F_t^{(h)}$ as a reliable proxy for risk premia. Unconditionally, this is indeed the case. Suppose S_t evolves according to a simple AR(1) process $S_{t+h} = \phi S_t + v_{t+1}$, the expectation at time t for the spot price at time $t+h$ is $E_t[S_{t+h}] = \phi^h S_t$, and the realised forecast error would be

$$S_{t+h} - E_t S_{t+h} = \sum_{i=0}^{h-1} \phi^i v_{t+h-i},$$

Note that by definition the forecast error is autocorrelated. Now, the *realised* payoff of a futures contract held until maturity can be decomposed as

$$F_t^{(h)} - S_{t+h} = y_t^{(h)} - \sum_{i=0}^{h-1} \phi^i v_{t+h-i}, \quad (2.2)$$

If expectations are unbiased the unconditional average of the forecasting error is zero. However, the persistence of price dynamics can make the conditional expectation errors sizeable for finite samples and horizons. Figures B.1 makes this case in point; the expectation errors $E_t[S_{t+h}] - S_{t+h}$ for two different horizons, i.e. $h = 2, 4$ quarters ahead, and two alternative commodities, i.e. WTI Crude Oil and Silver, tend to be time-varying and quite persistent.⁷

[Insert Figure B.1 about here]

Unsurprisingly, unexpected depreciation for crude oil occurred over the great financial crisis of 2008/2009 and the recent collapse of late 2014/beginning of 2015. Similarly, unexpected appreciation of Silver occurred in the recovery of financial markets after 2009, consistent with the idea that the value of precious metals tend to be negatively correlated with the

⁷The aggregate forecast $E_t[S_{t+h}]$ is proxied by the cross-sectional average of the Bloomberg's survey individuals forecasts. A complete discussion on how the survey is collected and structured, as well as a description of the data, is provided below.

business cycle.

As a whole, the assumption of either small or constant conditional unexpected price change turns out to be fairly restrictive. Figure B.2 makes a case in point, where to the extent that investors' misjudge the level of future spot prices over time, the ex-ante and realised risk premia differ.

[Insert Figure B.2 about here]

For instance, let us consider a simple situation in which the price of the commodity at time t is equal to 50\$ and market expectations for the future spot price at time $t + h$ are equal to 47\$, i.e. $E_t[S_{t+h}] = 47$. Also, let us assume that in order to make investors willing to enter the market the current price of a futures contract at time t for delivery at time $t + h$ is equal to 43\$, which means futures are sold at a discount. The difference between the futures price and $E_t[S_{t+h}]$ at time t implies that the expected payoff of a long position is equal to 4\$.

The top panel of Figure B.2 shows the case in which the commodity is indeed traded at 47\$ at maturity. Under no-arbitrage and given there are no unexpected price changes, the ex-ante and the realised payoffs are equivalent. Consider instead a situation in which investors make errors in forecasting future spot prices (see, e.g. Alquist and Kilian [4] for a complete discussion on the predictability of nominal spot prices). More specifically, let assume that the commodity is traded at a lower price of 45\$ at time $t + h$ on the spot market, which implies a forecast error equal to -2 (bottom panel). The realised payoff is now 2\$; the ex-ante and the realised risk premia differ by the amount of the unexpected price change. Figures B.1 and B.2 coupled, make clear that although expectations error can be zero asymptotically, they might have sizeable effects on investigating the origins of risk premia for reasonable sample sizes. In the following, we propose a reduced-form model of adaptive learning which allows to disentangle the actual, ex-ante, risk premium $y_t^{(h)}$.

2.3 Learning and Expectations

To set up an analytical framework, we start from an extended Muth [79]’s market model with the addition of both predictable changes in aggregate demand and the presence of a futures market (see, e.g. Turnovsky [95], Kawai [65], and Beck [12]). The market is characterized as an infinite horizon, discrete time model with both spot and futures market clearing conditions that hold in each time period. By including a futures market we assume that suppliers, buyers and inventory holders hedge their positions by trading on futures, and so we explicitly consider the effect of hedging in the decision-making process that leads to the Perceived Law of Motion (PLM henceforth) of spot prices. By allowing demand shocks to be predictable and possibly persistent we make explicit the effect of changes in aggregate demand in the dynamics of equilibrium spot prices.⁸ A unique reduced-form rational expectations equilibrium is defined as (see Appendix B.1)

$$S_{t+1} = \phi_0 + \phi_1 S_t + \phi_2 z_t + \eta_{t+1}, \quad (2.3)$$

with S_{t+1} the commodity price at date $t + 1$, z_t the change in aggregate demand at time t , and η_{t+1} an unobservable random shock.⁹ Notably, a similar solution would be obtained by assuming market segmentation between spot and futures as originally proposed in Muth [79]’s model.

We do not take a stand on the marginal relevance of supply vs. demand shocks in the dynamics of commodity stock prices, and assume that changes in aggregate supply are conditionally i.i.d. This assumption can be relaxed at the cost of having some reliable empirical

⁸In the original Muth [79] framework demand shocks that induce changes in inventories quickly revert to their long-run equilibrium values. In this respect, inventories adjustments are perceived to have a stabilizing effect on prices. However, as recently showed by Dvir and Rogoff [36] quick adjustments in inventories to demand shocks cannot explain the persistence in the time-series of commodity prices and volatilities.

⁹One may also specify a model in which expectations of future changes in aggregate demand rather than current values enter in the equilibrium outcome. As far the unique reduced-form solution in Eq. (2.3) is concerned, the two things are virtually equivalent. Aggregate demand is specified as an AR(1), i.e. $z_{t+1} = bz_t + e_{t+1}$. This implies that $E_t z_{t+1} = bz_t$, which means that the structural coefficient b of the actual law of motion, although cannot be identified, is embedded in the reduced-form parameter ϕ_2 of the perceived law of motion.

proxy for aggregate supply shocks for agriculturals, e.g. Corn, and precious metals, e.g. Silver, to be used as exogenous variables in the adaptive learning dynamics. Also, while the i.i.d. assumption for supply shocks can be restrictive for energy or industrial commodities, the same assumption possibly represents a fair approximation of supply shocks in agriculturals and precious metals, e.g. “harvest” can be thought as i.i.d. and storage of, say, corn is temporally limited.

A visual inspection of the relationship between (a proxy for) economic growth and spot prices confirms that changes in aggregate demand represent an important source of fluctuations in commodity prices. Figure B.3 shows the year-on-year changes in the (log of) commodity spot prices (blue line) and aggregate demand as proxied by an index of world industrial production published by the Netherlands Bureau for Economic and Policy Analysis (magenta line).¹⁰

[Insert Figure B.3 about here]

With the only partial exception of Corn (bottom-left panel), which is less sensitive to business cycles, changes in spot prices tend to align with changes in aggregate demand, especially after the beginning of the 2000’s. Similarly, Kilian and Hicks [67] show that unexpected economic growth sensibly affects the dynamics of spot prices in the Oil market. In our adaptive learning framework, beliefs are revised in line with past prediction errors based on available information, i.e. aggregate demand shocks affect investors’ expectations as well. This is consistent with Singleton [89], who argue that differences in beliefs can generate persistence in the dynamics of commodity spot prices.¹¹

¹⁰The index of world industrial production is published by the Netherlands Bureau for Economic and Policy Analysis and aggregate information from 81 countries worldwide, which account for about 97% of the global industrial production. The aggregate series starts in January 1991 and relate to import-weighted, seasonally adjusted industrial production.

¹¹Related to the Oil market, Singleton [89] pointed out that “Perhaps more plausible is the assumption that participants [...] learn about the true mapping between changes in fundamentals and prices by conditioning on past fundamentals and prices”.

Learning is introduced by assuming that agents do not know true values of the parameters of the PLM $\phi = (\phi_0, \phi_1, \phi_2)$ and expectations are instead formed on the basis of a weaker form of rational expectations that allow for model instability, uncertainty, and learning (see, e.g. Hsieh and Kulatilaka [60], Frenkel and Froot [44], Marcet and Sargent [73], Evans and Honkapohja [39], and Sargent [87], and Sockin and Xiong [90] relatively to commodity markets). Aggregate beliefs on the parameters are updated over time conditioning on current observations plus a constant $X_t = (1, S_t, z_t)$. More specifically, we follow Cho et al. [24], Sargent [87], and Sargent and Williams [88] and model the agents' recursive estimates in terms of a Bayesian prior that describes how coefficients drift at each time t ;¹²

$$\begin{aligned} S_{t+1} &= \phi'_{t+1} X_t + \eta_{t+1}, & \text{with } \omega_{t+1} &\sim N(0, \sigma^2), \\ \phi_{t+1} &= \phi_t + \xi_{t+1} & \text{with } \xi_{t+1} &\sim N(0, \Omega), \end{aligned} \quad (2.4)$$

with $\phi_t = (\phi_{0,t}, \phi_{1,t}, \phi_{2,t})'$ and $X_t = (1, S_t, z_t)'$. The shock ω_{t+1} is uncorrelated with ξ_{t+1} , and $\Omega \ll \sigma^2 I$. The innovation covariance matrix Ω governs the perceived volatility of increments to the parameters (see, Sargent and Williams [88]). Agents' recursive optimal estimate of ϕ_{t+1} conditional on information available at time t , $\gamma_{t+1} = \hat{\phi}_{t+1|t}$ are provided by a standard recursion;

$$\begin{aligned} \gamma_{t+1} &= \gamma_t + K_t (S_{t+1} - \gamma'_t X_t), \\ R_{t+1} &= R_t - \frac{R_t X_t X'_t R_t}{X'_t R_t X_t + 1} + \sigma^{-2} \Omega, \end{aligned} \quad (2.5)$$

where $K_t = R_t X_t (X'_t R_t X_t + \sigma^2)^{-1}$ determines the degree of updating of agents' beliefs when faced when an unexpected commodity spot price $S_t - \gamma'_t X_t$. This beliefs updating dynamics represents a generalization of recursive learning with constant gain. The recursive estimates (2.5) imply perpetual learning as they converge to a steady-state solution for a given initial condition of the state covariance matrix Ω (see, Hamilton [51] Proposition 13.1, pag. 390).

¹²This random walk specification for the evolution of the parameters is widely used in applied work in macroeconomics and finance, e.g. Frühwirth-Schnatter [45], West and Harrison [96], Stock and Watson [91], Primiceri [84], Hansen [54], and Leduc et al. [71].

We use the subscript $t+h|t$ to indicate a forecast for the $h > 0$ horizon made using information available to agents' at time t . The market price expected to prevail at time $t+1$ given the information available through the t -th period is obtained as

$$\hat{E}_t[S_{t+1}] = \gamma'_{t+1} X_t, \quad (2.6)$$

Multi-period forecasts $\hat{E}_t[S_{t+h}]$ are obtained by iterating forward the time- t estimates of the model parameters. Learning schemes as (2.5) are widely motivated in the macroeconomics literature by the fact that agents face constraints in cognitive abilities that limit their possibility to observe the true equilibrium parameters and use optimal forecasting rules (see, e.g. Carceles-Poveda and Giannitsarou [20], Adam and Marcet [3] and Malmendier and Nagel [72]). Conditional forecasts from Eq. (2.5) allows to extract risk premia across predictive horizons and commodities. More specifically, let $\hat{E}_t[S_{t+h}]$ be the model-implied expected future spot price of a given commodity at time t for the horizon $t+h$. The dollar value risk premium can be extracted from the price of a future contract at time t for delivery at time $t+h$, $F_t^{(h)}$, as;

$$\hat{y}_t^{(h)} = F_t^{(h)} - \hat{E}_t[S_{t+h}], \quad (2.7)$$

Eq. (2.7) implies that it is not necessary for the investors to have private information for their actions to affect commodity risk premia. As a consequence, the latter may depend on the nature of agents' learning mechanism based on common signals.

2.3.1 Comparison with Survey Expectations

We now compare the time-series of monthly expected future spot prices obtained from our model with the average forecast by professional analysts that operate in commodity markets. Individual price forecasts for different commodities and horizons are obtained from the Bloomberg's commodity price forecasts database. This database contains analysts' price expectations at multiple quarterly forecasting horizons and across diverse commodities from

2006 to 2016. The survey includes only operators highly specialised in commodity markets mainly from banks and consulting firms. Participants are asked to provide a point forecast on the average quarterly commodity price for a specified futures contract.

A deep knowledge of the peculiarities of commodity markets from the survey respondent, coupled with a clear objective of the survey, allows to reduce the effect of potential biases, quality homogeneity issues, and limited information processing, which generally characterises directional forecasts of non-specialised, or retail, cross-markets investors (see, e.g. Cutler et al. [27], Greenwood and Shleifer [50] and Koijen et al. [69]).¹³ There are two main objections on the use of survey expectations in empirical studies; first, the respondent may misunderstand the question which, for instance, can be posed in a simple directional way, e.g. do you expect prices increase, decrease or stay roughly constant. Second, a respondent may intentionally hide their true expectations for strategic purposes. Our survey mitigates the effect of both of these sources of error as (1) the question is about giving a clear point estimate for future spot prices, and (2) survey participants are professional market participants who possibly have payoffs that directly depends on the precision of their estimates.¹⁴ One comment is in order; the use of the Survey does not represent on itself the core of the Chapter, which is instead based on a model of adaptive learning. In this respect, we use the survey as an instrument to “validate” our model. As a matter of fact, although the survey represents the closest possible approximation of observable expectations, it still suffers from potential strategic biases and interactions among analysts.

The survey allows to retrieve for each analyst the historical price forecasts and the related publication date. Analysts provide their expectations for spot prices in different days for fixed common maturities that correspond to calendar quarters, i.e. they provide discontinued

¹³More specifically, the fact that only operators specialised in commodity markets are being surveyed increase the proportion of “truly informed” agents in the survey population compared to a case in which cross-market analysts are being surveyed.

¹⁴As we take the cross-sectional average of investors’ forecast as our proxy for market expectations, any non-coordinated strategic bias/error at the individual level is mitigated (see, e.g. Bernhardt and Kutsoati [14], Hong et al. [57], and Hong and Kubik [56])

fixed-calendar maturity quarterly expectations. Such feature makes the use of the survey for operational purposes quite challenging as the quarterly analysts' forecasts submission are recorded daily and not evenly spaced in time.

To perform a sensible time-series comparison with the model-implied expectations, we need to transform analysts' responses in continued constant-horizon price forecasts. We aggregate responses at the monthly frequency to reduce the difference in the market information available between early and late submitters within a month. Then, we compute the forecasting horizon with respect to the end of the month of the last month in the quarter which is the object of the prediction. More specifically, at each point in time, we stack the forecasts with residual life that belongs to the following groups: 4 to 6; 7 to 9 and 10 to 12 months, then finally we approximate the aggregate expectations as the cross-sectional average prediction across analysts and time-horizons.

Short-term moving average effects are reduced by discarding the horizon between one and three months as the analysts take into account what has been the realised price over the first part of the quarter generating nowcasting dynamics which makes hard to disentangle the role of expectations versus current information in the dynamics of short-term risk premia.¹⁵

For the ease of exposition, we report the results for dollar value expectations at maturities $h = 2, 4$ quarters. The sample period is from 12:2006 to 01:2016 for the survey, and is from 01:1995 to 01:2016 for the model-implied expectations. Figure B.4 reports the results for $h = 2$. The red circles represent the monthly observed survey forecast, and the light-blue circles show the expectations obtained from the adaptive learning model. The shaded area underlying the overlapping period between the survey and the model represents the difference between the two, i.e. a positive value means the model generates higher expected future spot

¹⁵Also, contracts close to expiration are typically illiquid in commodity markets as futures traders do not want to take the risk of a physical delivering of the underlying.

prices than the survey and vice versa.

[Insert Figure B.4 about here]

The survey forecasts and the adaptive learning expectations line up fairly well across the overlapping sample for WTI Crude Oil (top-left panel). This holds both during the dramatic rise and subsequent sharp fall in crude oil prices during the period 2008/2009, as well as during the market decline occurred since 2014. The “spread” between the model and the survey increases as high as 20\$ across the great financial crisis, although is sensibly reduced over the remaining sample. The top-right panel shows the results for Copper. Similar to Oil, adaptive learning can mimic the drop in expected spot prices in the period 2008/2009, the subsequent rapid price recovery, as well as the downward trend from 2011 until the end of the sample. Over a short-term horizon, the model still generates higher expected prices compared to the survey for a fraction of the sample, although the gap is small in magnitude after 2010.

A comparison with observable expectations for Corn (bottom-left panel) is limited by the few observations available from the survey, which does not provide opinions from analysts in the period 2011-2013. The divergence around the great financial crisis is non-negligible as indicated by an 80 cents/bushel negative gap. However, over the last part of the sample adaptive learning closely replicates average survey forecasts. Results are stronger for Silver. The gap is fairly small, with the partial exception of a negative “spread” during the dramatic rise in spot prices occurred in the aftermath of the great financial crisis of 2008/2009. In a separate calculation, we show that the sample correlation between the model- and survey-implied risk premia across commodities and horizons is 0.81, on average. Figure B.5 shows the results for a longer horizon,

[Insert Figure B.5 about here]

The expectations derived from our model of adaptive learning line up fairly closely with survey average forecasts over a four-quarter horizon, although the similarity between the model and the survey partly deteriorates as indicated by a more persistent gap throughout the sample. As a whole, the model performance tends to deteriorate in the longer-term, where the correlation between the model- and the survey-implied risk premia decreases to an average value of 0.64.

In the Appendix B.3 we further test the null hypothesis that average survey forecasts are consistent with a recursive learning framework. In this respect, we test for internal consistency between the model outlined to generate expectations and the observable proxy represented by the survey. We find evidence in support of adaptive learning in the expectations formation process across prediction horizons and commodity markets, meaning the elasticity of investors' expectations on future spot prices with respect to past forecasting errors is significant. However, notice that the evidence in favour of adaptive expectations does not rule out investors rationality (see Pesaran and Weale [82] for more details).

2.4 Empirical Analysis

We cover four main commodity futures which represent the energy, agricultural, industrial and precious metals markets. We focus on these commodities as they are the most traded consumption commodities with the most complete sample of survey data. The necessity to compare the model-implied expectations with the survey of professional analysts limits the possibility to increase the cross-section of commodities. In this respect, the choice of the commodity to be included in the analysis is mostly dictated by the length of the corresponding survey and the number of professional analysts responding. Including other commodities would come at the cost of using averages of few respondents or time-series with few observations.

2.4.1 Data

Data are obtained from different resources. Futures prices data on WTI Crude Oil are from the New York Mercantile Exchange (NYMEX), prices are in U.S. dollars per barrel. Futures on Silver and Copper are obtained from the Commodity Exchange (COMEX). Silver is quoted in U.S. dollars per troy ounce while Copper is quoted in U.S. cents/pound. We convert the price of Copper futures contracts to USD/tonne to match the measurement unit of the survey forecasts that instead refer to the London Metal Exchange (LME). Corn futures prices are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. As in Szymanowska et al. [93] the spot price for each commodity is approximated by using the nearest contract to maturity, and the futures price is the price of the next to the nearest futures contract for a given maturity.

We define the futures price at time t with average quarterly time to maturity h as $F_t^{(h)}$, where the definition of the average time to maturity is consistent with the average forecasting horizon for the survey expectations. For example, the price of a future for delivery four quarters ahead is computed interpolating the prices of the contracts between 10 and 12 months ahead. The sample period is monthly 01:1993-01:2016.

In order to study the sources of time variation in commodity risk premia, we collect diverse determinants that are considered to capture alternative sources of risk and/or economic fundamentals. Fluctuations in the global supply-demand imbalance for each commodity are captured by using inventory stocks. We collect data on Copper and Crude Oil inventories from the London Metal Exchange (LME) and Energy Information Administration (EIA) respectively. Copper inventory levels are recorded daily from June 1974 and relate the previous day closing stock of commodities held in LME. Crude Oil inventories are recorded weekly by the EIA and published monthly since January 1945. Stocks levels are measured in thousands of barrels and exclude strategic petroleum reserves.¹⁶ For Corn inventories, we use the U.S.

¹⁶We include in the level of inventories those domestic and Customs-cleared foreign stocks held at, or in transit to, refineries and bulk terminals, and stocks in pipelines. Stocks include an adjustment of 10,630

ending stocks reported in thousands of metric tons. The time-series is sampled at monthly frequency using the inventory level reported on the last business day of the month. Data are recorded from the United States Department of Agriculture (USDA) from January 1993. As far as Silver is concerned, we omit the inventory level variable as, similar to other precious metals, a considerable part of the existing reserves is privately held and therefore not reported in official statistics. In the regression specification we use the year-on-year growth rate of inventories as the levels are non-stationary and show the presence of a stochastic time trend.

Exchange rates are also a relevant risk factor as commodity trading takes place usually in U.S. Dollars, making FX a key factor for both producers and consumers that can directly affect profits and costs denominated in domestic currency. In order to account for the risk of appreciation and depreciation in the U.S. Dollar, we include the growth rate of Federal Reserve U.S. trade weighted exchange rate index, normalized to be equal to one hundred in March 1973.

Furthermore, we include a measure of time-series momentum among the risk factors in our analysis as it can be directly linked to asset demand by momentum traders as shown in Cutler et al. [27], Moskowitz et al. [76], and Kang et al. [63]. Momentum in commodity futures has been widely documented in the empirical finance literature, e.g. Erb and Harvey [37], Miffre and Rallis [75], Asness et al. [5] and Szymanowska et al. [93] among others. We construct time-series Momentum as the rolling return over the past 12 months skipping the most recent month on each commodity future. In addition, we include a Value factor which is assumed to be intimately interrelated to the dynamics of commodity risk premia, as it affects the propensity of market participants to trade in backwardation or in contango and can proxy the trading activity of speculators following mean-reversion type trading strategies. We follow Asness et al. [5] and define Value as the average of the log spot price from 4.5 to 5.5 years ago, divided by the most recent spot price, which is essentially the negative of the spot rolling return over the last 60 months. In addition to time-series Value and Momentum,

thousand barrels (constant since 1983) to account for incomplete survey reporting of stocks held on producing leases.

we also directly consider returns on the Standard and Poor's 500 and the MSCI Emerging Markets indexes as a proxy for financial risk. Beyond direct effects on financial flows, we incorporate stock indexes as they likely capture spillover effects to the real economy. As a measure of futures market uncertainty, we compute the Realised Volatility for a given maturity h as the sum of squared daily futures returns adjusted for roll-over and for delivery date $t + h$.

Finally, to capture market activity and risk sharing preferences in the economic mechanism that drives commodity risk premia we consider OI and HP (see e.g. Baker and Routledge [7] and Singleton [89]). OI is measured as the total number of outstanding contracts that are held by market participants at the end of the month. An outstanding contract is when a seller and a buyer combine to create a single contract. For each seller of a futures there must be a buyer of that contract, therefore to determine the total OI for any given market we need to know the totals from one side or the other, buyers or sellers, not the sum of both. Increasing OI means that new cash is flowing into the marketplace while declining activity means that the market is liquidating, which can be interpreted as a signal of a price turning point. As for inventories, we use the year-on-year growth rate of OI as the levels are non-stationary.

Hedging pressure represents a measure of net positions of hedgers in commodity futures markets which is the result of risks that market participants do not want, or cannot trade because of market frictions, information asymmetries and limited risk capacity (see, e.g. Hong and Yogo [59] and Kang et al. [63]). We compute the level of HP for different commodities as the net excess in short futures positions by commercial traders, i.e. short minus long positions, divided by the amount of outstanding contracts. The data on commercial traders futures positions are from the Commodity Futures Trading Commission (CFTC).

2.4.2 Dissecting Ex-Ante Risk Premia

The framework outlined in Section 3 allows to back out the dollar-valued time-varying ex-ante risk premia from our model of learning. In order to obtain the expected payoffs

as a returns quantity, which is more suitable for our regression analysis, we took a log transformation for both futures prices and the model-implied expected future spot prices in Eq.(2.7). This allows to approximate risk premia in percentage returns, up to a negligible Schwarz inequality term.¹⁷

Panel A of Table B.1 shows the in-sample descriptive statistics of the monthly risk premia (decimals). Unconditionally, the term structure of risk premia for Crude Oil and Copper is negatively sloped. Risk premia for these two commodities are negative and increasing over time in magnitude. This is consistent with the theory of Keynes [66] and Hicks [55], which posits that hedgers are net short and futures are set at a discount with respect to the future expected spot price. Conversely, for Corn and Silver average risk premia are positive and increasing as a function of time horizon. Hedging pressure theory states this is the result of hedgers predominantly being net-long with speculators willing to enter contracts with slightly negative payoff provided there are expectations of increasing future prices.

[Insert Table B.1 about here]

With the only exception in the short-term risk for Crude Oil and longer-term for Copper, the sample distribution of risk premia is far from Gaussian. Both Corn and Silver show a fairly large negative skewness and a substantial excess kurtosis. Departure from Normality is also mainly given to fatter tails in Oil and Copper. Overall, a Jarque-Bera test rejects the null hypothesis of Normality for nine out of twelve cases. Finally, the term structure of volatility for risk premia is positively sloped, i.e. the standard deviation of ex-ante risk premia increases with maturity. Panel B of Table B.1 shows the in-sample cross-sectional correlations of risk premia for each expectations horizon. Cross-sectional correlations are inversely related to the investment horizon; for instance, the correlation of WTI Crude Oil with Copper

¹⁷For the sake of completeness, we implement the empirical analysis in Section 4 by using Eq.(2.7) and rescale both futures prices and the model-implied expectations by current spot prices. The main results, available upon request from the authors, are in line with the log-transformation.

is 0.355 at $h = 2$ and decreasing to 0.243 at $h = 4$. A similar path is found across commodities.

We now investigate the determinants of the ex-ante risk premia through a static regression. Table B.2 shows the estimated standardized coefficients with the asymptotic t-statistics in parenthesis.¹⁸ Few comments are in order; first, there is significant heterogeneity in the significance of each factor across commodities. While emerging markets are strongly significant for the sample variation of WTI and Copper risk premia, the same are not relevant for Corn and Silver. In this respect, Copper and Oil are directly affected by the demand from, e.g. China, while food and precious metals are much less dependent on spillovers effects from emerging markets. Similarly, realised volatility seems to significantly affect futures expected payoff only for Silver, which is consistent with the fact that precious metals are safe-haven assets during market turmoil

[Insert Table B.2 about here]

Second, surprisingly hedging pressure and inventories are not significant determinants of risk premia sample variation. This somewhat contradicts early work by De Roon et al. [29], Basu and Miffre [10], and Szymanowska et al. [93]. However, we show below that once the dynamics of risk premia is fully considered, HP turns out to be a key component. Similarly, except for futures on Copper at a two- and three-quarter horizon, inventories are not significantly related to the ex-ante risk premia, after controlling for net supply-demand imbalances and spillover effects from emerging markets and currency fluctuations. Finally, sensitivity to past performances is significant and positive across commodities and horizons, with the only exception of short-term futures for Copper. This is consistent with Asness et al. [5], and can be possibly rationalized by a “bandwagon” effect in market activity and trading behaviour which increase the persistence of futures returns.

¹⁸The explanatory factors in the regression are pre-whitened, i.e. orthogonalized to each other and standardized. Pre-whitening helps to reduce the spurious effect of cross-factor correlations, which can be arguably relevant in a linear model with many factors like ours, e.g. HP and OI or S&P500 and MXEF. We estimate the model by OLS with GMM corrected standard errors.

The regression results of Table B.2 suggest that, unconditionally, risk sharing mechanism and market activity possibly explain the sample variation of the ex-ante risk premia. However, the fact that expected payoffs have their own dynamics could be the consequence of a heterogeneous exposure to different risk factors on a time scale. In this sense, the results of a static regression might be potentially incomplete, at best. For instance, the so-called financialization of commodity markets arguably increases the sensitivity of risk premia to market activity which is, by definition, contingent and time-varying and not necessarily linked to economic fundamentals. In the following, we use a dynamic regression modeling framework that explicitly allows for a time variation in the relationship between the risk premia $\hat{y}_{t+1}^{(h)}$ over the interval $[t, t+1)$ and the realizations of the explanatory factors observed at time t .

More specifically, we assume that the exposure of risk premia to each specific factor is a random walk (see, e.g. West and Harrison [96], Kilian and Taylor [68], and Ferreira and Santa-Clara [43]). Risk factors have been orthogonalized to avoid spurious effects due to cross-correlations in the explanatory variables. Methodologically, we opt for a Bayesian estimation framework, which allows to obtain robust finite-sample estimates that flexibly and explicitly accounts for different sources of uncertainty: uncertainty in the relative importance of predictors, uncertainty in the estimated coefficients and their degree of time-variation. Appendix B.2 provides a detailed explanation of the regression design and model estimation strategy. One comment is in order; assuming regression betas evolve as a random walk implies that the elasticity of risk premia to a given factor drift to deterministic high or low values of $\hat{y}_t$, hence generating non-stationarity. However, an alternative more general AR(1) specification for the dynamics of the regression betas shows that the state parameters are highly persistence with low conditional variance. In this respect, the random walk assumption represents an attractive approximation because of its parsimony, ease of computation and the

smoothness it induces in the estimated sensitivities over time.¹⁹

For the ease of exposition we first investigate what is the actual amount of explanatory power that can be associated to each of these factors within our dynamic regression exercise, and then we show the time-varying betas for the sub-set of risk factors that show most of the significance. In particular, we first decompose the overall R^2 to measure the improvement resulting from including covariate k in a dynamic regression model that already contains the other covariates (see Genizi [46] for more details). Given the regression covariates have been previously orthogonalized, this boils down to compute the ratio between the sum of explained residuals from the regressor k and the total sum of squares, as is done for a univariate regression with the single regressor k . Figure B.6 shows the marginal contribution of each risk factor for the total R^2 of the regression.²⁰

[Insert Figure B.6 about here]

The top-left panel confirms that most of the explanatory power for the dynamics of WTI Oil risk premia comes from three key variables, namely MXEF, HP and TSMOM, especially for short maturities. Indeed, HP alone contributes to around 20% of the explained variation for $h = 2$, proportion than shrink to around 10% for $h = 3, 4$. As far as Copper is concerned, inventories contributed to a large fraction of the explained in-sample variation (around 15%) especially for longer-term maturities. Similar to WTI, we attribute most of the explanatory power to time-series momentum, particularly in the short-term (around 20% for $h = 2, 3$). Bottom-left panel shows the same decomposition for Corn. Most of the R^2 is attributed to open interests. Also, time-series momentum and USDTW carry a significant explanatory power, with the latter contributing to around 10% of the explained sample variation. Finally, bottom-right panel shows that most of the R^2 obtained by the dynamic regression model

¹⁹We share these findings with a large literature on returns predictability that assumes time variation in the predictive coefficients. Similar to our argument they find that assuming parameters are random walks in predicting excess returns we benefit from a substantial reduction of estimation error without effectively increasing the precision in the estimated dynamics in a finite sample.

²⁰Notice that the percentages in the graph do not sum to one as we left aside the amount of sample variation explained by the intercept.

for Silver is due to market activity, past performances and uncertainty, as proxied by HP, momentum and realised volatility.

We now focus on those factors which shows most of the significance. Figure B.7 shows the time-varying betas for each risk factor on WTI Crude Oil ex-ante risk premia. For the ease of exposition we report the results for $h = 2$ (blue line) and $h = 4$ (dark yellow line). We report both the posterior medians (solid marked line) and the 95% credibility intervals (dashed-dot lines). Results on the intermediate horizon $h = 3$ are available upon request.

[Insert Figure B.7 about here]

The empirical evidence shows that the impact of emerging markets has become increasingly important in the aftermath of the great financial crisis of 2008/2009. A possible explanation is the presence of spillover effects due to the increasing weight of emerging economies in the global economic outlook.²¹ Indeed, although the direct impact of shocks in stock valuations in emerging markets is relatively low due to moderate foreign investments, financial turbulence in this area is often perceived as a signal of a slowdown in global economic growth, and thus aggregate demand.

Betas on HP show that risk sharing/appetite preferences partly explain the dynamics of risk premia in the period that coincides with the dramatic rise in oil prices between 2001 to the end of 2005, and in the aftermath of the great financial crisis of 2008/2009. During this period the propensity to buy futures by consumers to lock in oil prices increased substantially. Pressure on the demand side of futures possibly decreased the premium required by speculators to take the short side of the contract. The period 2001-2005 is more difficult to rationalize as hedging pressure was widely fluctuating around zero during this period. A possible explanation relates to the scarce risk-bearing capacity of investors during a period characterized by overall higher uncertainty in the aftermath of 9/11 attacks and the following

²¹For instance, the IMF Economic Outlook 2016 states that growth in developing economies accounted for over 70 percent of global growth in 2016.

Iraq invasion of March 20th, 2003. In this respect, e.g. Acharya et al. [2], Cheng et al. [23], Etula [38] and Hong and Yogo [59] showed that when there are limits to the risk-bearing capacity of investors and/or constraints on the amount of capital different investor categories are willing to commit, large changes in market liquidity possibly affect prices both in the futures and spot markets and ultimately affect risk premia. As a whole, the strong relevance of HP for the dynamics of expected payoffs confirms the primary relevance of futures as a risk insurance market place, as postulated by Keynes [66] and Hicks [55].

A substantial, positive, effect is also played by TSMOM, which can be generated by psychological biases of market participants and informational frictions that delay their learning about fundamentals (see, e.g. Cutler et al. [27], Greenwood and Shleifer [50], and Singleton [89]). More importantly, time-series momentum aims at capturing the changes in trading activity of feedback traders. In fact, as shown by Cutler et al. [27], the demand for futures contracts by feedback (momentum) traders depend on past market performances. By the same token, our results confirm the findings of Kang et al. [63], that show the importance of speculators following momentum strategies in determining the market demand for liquidity, and ultimately risk premia. Indeed, as shown by Kang et al. [63], momentum traders increase the demand for liquidity, which needs to be absorbed by risk-averse market makers and hedgers who will require therefore appropriate risk compensation. Surprisingly, other economic fundamentals such as Inventories, Exchange rates, and Value do not play a sensible role in the dynamics of crude oil risk premia. Figure B.8 shows the time-varying betas for each risk factor on Copper ex-ante risk premia.

[Insert Figure B.8 about here]

Except for few differences, much of the results of Oil also holds for Copper, which is not surprising as industrial metals and energy commodities are commonly sensitive to fluctuations over the business cycle and share most of the risk factors exposures and similar storage costs (see, e.g. Bhardwaj et al. [16]). The impact of emerging markets is increasing in the

aftermath of the great financial markets. As for crude oil, this is possibly due to the increasing impact of demand of Copper from Asian markets, and China in particular.

The positive effect of OI on risk premia is consistent with the idea that increasing market activity signals changes in economic conditions, which, in turn, increases the marginal propensity of hedgers to take a net long/short position, generating price pressure on futures. This result is in line with Hong and Yogo [59] that showed how OI has a significant predictive power for realised risk premia in futures markets in the presence of hedging demand and limited risk capacity. Also, the significant betas $\beta_{OI,t}$ provide some indirect evidence on the financialization of commodity markets whereby commodity risk premia are no longer simply determined by their supply-demand but are also affected by aggregate investment behaviour (see, e.g. Tang and Xiong [94]).

While the positive effect of TSMOM is similar to crude oil, the effect of inventories and realised volatility is much different. Indeed, in the longer-term, changes in inventories negatively affect risk premia. A possible explanation lies in the fact that inventories proxy supply-demand imbalances; a positive shock in stockpiles negatively correlates with prices, which in turn increases the risk premium required by speculators to take the long side of futures contracts. Figure B.9 shows the time-varying betas for Corn.

[Insert Figure B.9 about here]

Figure B.6 shows that, unlike WTI Oil and Copper, risk premia on Corn are not affected by possible shocks from emerging markets. Similarly, except few nuances, betas on realised volatility, value and hedging pressure are not significant across the sample. Most of the explanatory power is limited to open interests, time-series momentum and USD TW index. Momentum in agriculturals is mostly generated by irregular production. Taking Corn as our example, consumer demand remains fairly stable throughout the year whilst production is seasonal and can vary hugely. For instance, a bad harvest in October/November in the U.S. (which represents around 40% of the global production) cannot be rebalanced until a good

harvest occurs in the south hemisphere the next production cycle or in the U.S. the next year, increasing prices and possibly generating positive momentum as supply expectations are revised downward, and stockpiles decrease. The corresponding time-varying betas tell us that, except for the great financial crisis of 2008/2009, fluctuations in production make futures contracts more expensive on average. This is partly confirmed by the negative effect of changes in inventories, which becomes negative and significant towards the end of the sample.

Also, Figure B.9 shows that risk premia on Corn turn out to be related to the exchange rate. Time-varying betas show a positive and significant effect of FX shocks mostly during the first decade of the 2000's., while for $h = 4$ major positive effects appear across 2011/2012. The positive effect of USD TW is somewhat expected as the U.S. represents on itself 40% of the global production for Corn. A strong dollar generally leads to lower exports for the U.S. as a consequence of lower demand given less competitive prices but also means that the production of Corn will become more profitable (see, e.g. Hamilton [52]). These effects combined make overall more expensive to take the short side of a futures contract, therefore increasing the premium required by, for instance, speculators to sell contracts to hedgers. Another possible explanation relies on the increasing financialization of the agricultural commodity markets. As shown by Tang and Xiong [94], after 2004, agricultural commodities included in financial indexes such as the Goldman Sachs Commodity Index (GSCI) and the Dow Jones (DJ)-AIG, became much more responsive to shocks to the U.S. dollar exchange rate. Finally, Figure B.10 shows the time-varying betas for Silver.

[Insert Figure B.10 about here]

Figure B.6 shows that, similarly to Oil and Copper, betas on S&P500, Value and USD TW are not significant throughout the sample. On the other hand, β_{HP} tend to be negative for short-term contracts for the period 2003-2013. This period coincides with a massive increase in futures prices. The imbalance between short and long contracts was consistently positive during this period, i.e. hedgers were net short, although slightly decreasing. Given prices

were constantly increasing, a further positive change in HP would make cheaper to take the long side of the contract, which means a lower premium is required to bear the risk of decreasing prices. Across the same period, β_{RVol} are negative and significant. The fact that the effect of uncertainty is opposite than Copper is no surprise. In fact, a closer look at expected price dynamics (see, e.g. Figures B.5) shows that while uncertainty is associated with declining prices for Copper, the opposite holds for Silver.

Forecasting Performance

One may argue that the ex-ante risk premia extracted from equation (2.7) merely represent expectations errors which have nothing to do with investors' preferences and/or the actual expectation formation process. In this Section, we address this concern by directly testing the consistency of our model output, i.e. expected spot prices and ex-ante risk premia, with the observable realised payoffs and future spot prices.

We first investigate whether $\hat{E}_t[S_{t+h}]$ from (2.6) can effectively approximate latent expectations $S_{t+h|t}$. More specifically, we compare the forecasting performance of the model-implied expectations against alternative specifications which are mostly used in the forecasting literature to predict future spot prices S_{t+h} . As a performance metric, we use the out-of-sample R^2 statistics, R_{OS}^2 , suggested by Campbell and Thompson [19] to compare our benchmark forecast with alternative predictions. The R_{OS}^2 is akin to the standard in-sample R^2 statistics and is computed as one minus the ratio of the Mean Squared Prediction Error (MSPE) obtained from the alternative model and the one obtained from our benchmark (see, e.g. Rapach et al. [85]). In this respect, the R_{OS}^2 measures the improvement in forecasting future spot prices using the model-implied expectations relative to the competing predictors. Thus, $R_{OS}^2 > 0$ implies that adaptive learning is best performing according to the MSPE metric. Following Rapach et al. [85], statistical significance for the R_{OS}^2 statistic is based on the p-value for the Clark and West [25] out-of-sample MSPE; the statistics corresponds to a one-side test of the null hypothesis that the competing specification for the expected future

spot prices has equal forecasting performance that our model-implied expectations against the alternative that the competing model has a lower average square prediction error. We use the first ten years of data, i.e. 01:1995-12:2005 to train the model of learning, such that the out-of-sample evaluation period on which R_{OS}^2 is computed is 01:2006-01:2016. Table B.3 reports the results; a number is highlighted in grey when the null hypothesis is rejected at least at a 5% significance level.

[Insert Table B.3 about here]

The first row shows the model performance with respect to $\hat{E}_t[S_{t+h}] = F_t^{(h)}$, i.e. using futures as proxy for expectations. Adaptive learning compares favourably in seven out of twelve cases, four of which are significant at 5% level. Except for Corn, simply using futures to predict future spot prices seem beneficial for Corn. This is in line with Alquist and Kilian [4]. A similar result is found by assuming expectations are restricted to be equal to current spot prices, i.e. $\hat{E}_t[S_{t+h}] = S_t$. However, for Corn there is a significant under-performance, although small in magnitude (from -0.8% at $h = 2$ to -2% at $h = 4$). Finally, we compare the forecasting ability of our model-implied expectations against a baseline futures spread indicator, i.e. $\hat{E}_t[S_{t+h}] = S_t \left(1 + \ln \left(F_t^{(h)} / S_t \right) \right)$, (see Alquist and Kilian [4]). Our model compares favourably in eight out of twelve cases, i.e. $R_{OS}^2 > 0$; the improvement is significant in five out of eight cases.

A second check should be made is to investigate the correlation between expected and realised payoffs. Expected payoffs are extracted from our model according to equation (2.7), and realised returns are computed as the excess rolling return in the generic contract for the same maturity of the corresponding model-based expectations. Figure B.11-B.12 show the scatter plots of ex-ante vs realised risk premia for $h = 2$ and $h = 4$ maturities, respectively. The red line represents the regression line; betas and asymptotic t-statistics are reported within the graphs.

[Insert Figures B.11-B.12 about here]

Few comments are in order; first, the scatter plots make clear that there is a significant positive correlation between the model-implied risk premia and the observable rolling returns, across maturities. The correlation is higher for Silver and lower for WTI Crude Oil. Second, as we would expect, the correlation between expected and realised payoffs becomes lower as the contracts maturity increases. This is possibly due to the fact that as the maturity of the contract increases, it is more likely that investors make mistakes in forecasting future spot prices, therefore increasing the gap between ex-ante and ex-post returns (see, eq. (2.2) and Section 2 for a full discussion).

2.5 Concluding Remarks

Our empirical analysis shows that investors' expectations of future commodity spot prices can be approximated by a rational learning scheme in which expected future spot prices are revised in line with past price prediction errors and changes in aggregate demand. In fact, we show that, although with differences across commodities, our model-implied price expectations are broadly consistent with analysts' survey forecasts. As a second step, we contribute to the study of the dynamics of risk-premia by exploiting this expectations formation mechanism to extract time-varying (ex-ante) risk premia from futures across different commodities and maturities.

By using a dynamic linear regression in which we accommodate uncertainty in the estimated coefficients and their degree of time-variation, we show that the dynamics of commodity risk premia is predominantly driven by market activity and the changing nature of market participants, as proxied by open interests, hedging pressure and time-series momentum. Furthermore, we show that our model of learning compares favourably to other commonly used specifications in forecasting future spot prices and generates expected payoffs that are consistently linked to the actual, observable, returns on same-horizon futures contracts.

Appendix B

Appendix and Tables

B.1 A Simple Model of Learning

We start from a simple rational expectations model which is closely related to the Muth [79] market model with inventory speculation except demand shocks are predictable and not i.i.d. The market behaviour is characterized by an infinite horizon, discrete time model with a market clearing condition that holds in each period, $t + 1$;

$$C_{t+1} + I_{t+1} = Q_{t+1} + I_t, \quad (\text{B.1})$$

where Q_{t+1} represents the output produced for a commodity in a period lasting as long as the production lag, C_{t+1} is the amount of commodity consumed in the same time period, and I_{t+1} the commodity inventories at the end of period $t + 1$. The standard Muth [79] market model posits there are three categories of economic agents active in the market for commodities; the buyers, the producers and the inventory holders. The latter can capture speculation effects. The utility of price-taking consumers is declining in the current market price S_{t+1} and affected by an aggregate persistent demand shock z_t . On the other hand, the utility of risk-averse producers is positively related to expected spot prices $E_t S_{t+1}$, while inventories decisions depend on the expected capital gain of holding a unit of commodity. As a result, aggregate demand, supply and holding functions are defined as

$$C_{t+1} = A - \delta S_{t+1} + z_{t+1}, \quad (\text{B.2})$$

$$Q_{t+1} = \lambda E_t S_{t+1} + u_{t+1}, \quad (\text{B.3})$$

$$I_t = v(E_t S_{t+1} - S_t), \quad (\text{B.4})$$

with ν be a rescaled risk-aversion parameter. We extend the standard market model with inventory speculation assuming exogenous factors that affect aggregate demand are predictable and potentially persistent;

$$z_{t+1} = bz_t + e_{t+1}, \quad (\text{B.5})$$

with e_{t+1} and u_{t+1} zero-mean i.i.d. disturbance terms. Storage costs are assumed to be zero to simplify the model. These equations and assumptions are the same of the original Muth [79] model, except for the predictability of demand shocks. Substituting (B.2)-(B.5) in the equilibrium condition (B.1), the spot market equilibrium can be expressed in terms of prices, price expectations, demand shocks and disturbances;¹

$$\begin{aligned} A - (\nu + \delta)S_{t+1} + bz_t + e_{t+1} &= \lambda E_t S_{t+1} + u_{t+1} - \nu S_t, \\ (\nu + \delta)S_{t+1} &= A + bz_t + e_{t+1} - \lambda E_t S_{t+1} - u_{t+1} + \nu S_t, \end{aligned} \quad (\text{B.6})$$

which can be rewritten as a simple linear model as follows

$$S_{t+1} = \mu + \beta E_t S_{t+1} + \theta S_t + \omega z_t + \eta_{t+1}, \quad (\text{B.7})$$

By taking expectations on both sides and substituting back in (B.7), we can obtain a unique reduced-form Rational Expectations Equilibrium (REE) as

$$S_{t+1} = \phi_0 + \phi_1 S_t + \phi_2 z_t + \eta_{t+1}, \quad (\text{B.8})$$

with $\phi_0 = (1 - \beta)^{-1} \mu$, $\phi_1 = (1 - \beta)^{-1} \theta$, $\phi_2 = (1 - \beta)^{-1} \omega$ and $\eta_{t+1} = e_{t+1} - u_{t+1}$. This solution is the same as the original Muth [79]'s model except that future commodity spot prices now depends on aggregate demand. Notice that for a given level of commodity prices, Eq. (B.8) implies that a positive (negative) shock to aggregate demand increases (decreases) future prices, while a positive (negative) shock in aggregate supply decreases (increases) prices.

B.1.1 Introducing a Futures Market

We now introduce a futures market upon the process of price formation and show that the functional form of the perceived law of motion under rational expectations is observationally

¹We assume there is a period distance in the future where the forward expectations are equivalent, i.e. $E_t S_{t+1} \equiv E_{t+1} S_{t+2}$, (see, e.g. Beck [12]).

equivalent to Eq.(B.8). By including a futures market we explicitly consider the effect of hedging in the decision-making process of a representative investor. We assume that suppliers, buyers and inventory holders now hedge their commodity positions by trading on futures. Following Turnovsky [95], Kawai [65], and Beck [12], we start from the assumption that agents are making production, storage and hedging decisions simultaneously, depending on current futures and expectations of spot prices.² All agents are assumed to act as hedgers as well as speculators in the futures market. Assuming now futures prices and spot price expectations are linearly linked to each other, aggregate demand, supply and holding functions are defined as

$$C_{t+1} = A - \delta F_t + z_{t+1}, \quad (\text{B.9})$$

$$Q_{t+1} = \lambda F_t + u_{t+1}, \quad (\text{B.10})$$

$$I_{t+1} = \xi (F_t - S_t), \quad (\text{B.11})$$

$$-X_t^b = \chi^b [F_t - E_t S_{t+1}] - \tilde{C}_{t+1}, \quad (\text{B.12})$$

$$X_t^p = \chi^p [F_t - E_t S_{t+1}] + \tilde{Q}_{t+1}, \quad (\text{B.13})$$

$$X_t^i = \chi^i [F_t - E_t S_{t+1}] + \tilde{I}_t, \quad (\text{B.14})$$

where ξ represents the inverse of storage cost per unit of commodity, and X_t^b, X_t^p and X_t^i represent the speculative positions, i.e. excess supply of futures contracts, by buyers, producers and inventory holders, respectively. Planned levels of consumption, production and inventories denoted as $\tilde{C}_{t+1}, \tilde{Q}_{t+1}$ and $\tilde{I}_{t+1}$, indicate that commodity positions are completely hedged in the futures market. The market clearing condition on the futures market states that the aggregate excess supply of futures contract should be zero, i.e.

$$X_t^p + X_t^i - X_t^b = 0, \quad (\text{B.15})$$

Substituting (B.9)-(B.11) in the market clearing condition (B.1), the spot market equilibrium can be expressed in terms of both futures and spot prices, demand shocks and disturbances, i.e.

$$A - \delta F_t + bz_t + e_{t+1} + \xi (F_{t+1} - S_{t+1}) = aF_t + u_{t+1} + \xi (F_t - S_t), \quad (\text{B.16})$$

²More specifically, we assume buyers are intermediate producers, which therefore as well willing to reduce risk hedging their positions participating in the futures contract.

Similarly, by substituting (B.12)-(B.14) in (B.15) we obtain;

$$\begin{aligned}\chi^f [F_t - E_t S_{t+1}] + \tilde{Q}_{t+1} + \chi^i [F_t - E_t S_{t+1}] + \tilde{I}_t + \chi^b [F_t - E_t S_{t+1}] - \tilde{C}_{t+1} &= 0, \\ \chi^f [F_t - E_t S_{t+1}] + \lambda F_t + \chi^i [F_t - E_t S_{t+1}] + \xi (F_t - S_t) + \chi^b [F_t - E_t S_{t+1}] - A + \delta F_t &= 0,\end{aligned}$$

Where $\tilde{Q}_{t+1}$, $\tilde{I}_t$ and $\tilde{C}_{t+1}$ are defined as (B.9)-(B.11) without the error terms. Solving for F_t we have that

$$F_t = \bar{A} + \bar{\chi} E_t S_{t+1} + \bar{\xi} S_t, \quad (\text{B.17})$$

with $\bar{A} = A/a$, $\bar{\chi} = \chi/a$ and $\bar{\xi} = \xi/a$, where $a = (\chi + \lambda + \xi - \delta)$ and $\chi = \chi^p + \chi^b + \chi^i$. Similarly, F_{t+1} can be obtained as a function of $E_{t+1} S_{t+2}$ and S_{t+1} , and substitute these values into (B.16) to obtain;

$$\begin{aligned}\xi (\bar{\xi} - 1) S_{t+1} &= \delta \bar{A} + \bar{\chi} (a + \delta) E_t S_{t+1} + \bar{\chi} (\chi + \delta) S_t + b z_t + e_{t+1} - u_{t+1}, \\ S_{t+1} &= \mu + \beta E_t S_{t+1} + \theta S_t + \omega z_t + \eta_{t+1},\end{aligned} \quad (\text{B.18})$$

with $\mu = \delta \bar{A} / \xi (\bar{\xi} - 1)$, $\beta = \bar{\chi} (a + \delta) / \xi (\bar{\xi} - 1)$, $\theta = \bar{\chi} (\chi + \delta) / \xi (\bar{\xi} - 1)$, and $\omega = b / \xi (\bar{\xi} - 1)$, respectively. Equation (B.18) is analogous to (B.29) and can be solved in the same way. From (B.18), the solution procedure described above yields the same Perceived Law of Motion (PLM);

$$S_{t+1} = \phi_0 + \phi_1 S_t + \phi_2 z_t + \eta_{t+1}, \quad (\text{B.19})$$

with $\phi_0 = (1 - \beta)^{-1} \mu$, $\phi_1 = (1 - \beta)^{-1} \theta$, $\phi_2 = (1 - \beta)^{-1} \omega$ and $\eta_{t+1} = e_{t+1} - u_{t+1}$. To summarize, we show that by introducing a futures market in which different type of investors hedge their positions in physical commodities, the reduced form PLM has the same functional form of the case without a futures market. In the following Section we introduce recursive learning on the reduced-form parameters ϕ_0 , ϕ_1 and ϕ_2 in Eq.(B.19).

B.1.2 Learning Dynamics

The key assumption to introduce learning is that the expectations of economic agents $E_t [S_{t+1}]$ are rational, but not necessarily perfectly rational as agents do not know the structural parameters. Expectations are instead formed on the basis of current observations and predictions of parameters which are updated over time. There are two key building blocks

to explicit the agents' learning dynamics. First, agents' beliefs are described by means of a dynamic model. We assume the PLM as the same functional form of the REE (B.19), where the true values $\phi = (\phi_0, \phi_1, \phi_2)$ are not known. Second, we need to describe how agents obtain estimates for the parameters of the PLM. We explicit agents' recursive estimates in terms of a Bayesian prior that describes how coefficients in the PLM drift at each time t ;

$$\begin{aligned} S_{t+1} &= \phi'_{t+1} X_t + \eta_{t+1}, & \text{with } \eta_{t+1} &\sim N(0, \sigma^2), \\ \phi_{t+1} &= \phi_t + \varepsilon_{t+1} & \text{with } \varepsilon_{t+1} &\sim N(0, \Omega), \end{aligned} \quad (\text{B.20})$$

with $\phi_t = (\phi_{0,t}, \phi_{1,t}, \phi_{2,t})$ and $X_t = (1, S_t, z_t)$. The shock η_{t+1} is uncorrelated with ε_{t+1} , and $\Omega \ll \sigma^2 I$. The innovation covariance matrix Σ governs the perceived volatility of increments to the parameters, and is a key component of the model (see Sargent and Williams [88]). Agents' recursive optimal estimate of ϕ_{t+1} conditional on information available up to time t . $\gamma_{t+1} = \hat{\phi}_{t+1|t}$ are provided by the Kalman filter recursion;

$$\begin{aligned} \gamma_{t+1} &= \gamma_t + K_t (S_t - \gamma'_t X_t), \\ R_{t+1} &= R_t - \frac{R_t X_t X'_t R_t}{X'_t R_t X_t + 1} + \sigma^{-2} \Omega, \end{aligned} \quad (\text{B.21})$$

where $K_t = R_t X_t (X'_t R_t X_t + \sigma^2)^{-1}$ determines the degree of updating of agents' beliefs when faced when an unexpected commodity spot price $S_t - \gamma'_t X_t$, i.e. Kalman gain. The recursive learning dynamics (B.20) represents a generalization of a recursive learning with constant gain as specified in Evans and Honkapohja [39], Sargent [87], Cho et al. [24], and Williams [97], among others.

B.2 Econometric Design

In the following we specify the dynamic regression model used to capture the time-varying linkages between the ex-ante risk premia and the corresponding explanatory variables. Denoting these by Z_t the set of economic predictors a dynamic regression can be specified as a state-space model;

$$y_t = Z'_t \theta_t + v_t, \quad v_t \sim N(0, H), \quad (\text{B.22})$$

$$\theta_t = \theta_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, W), \quad (\text{B.23})$$

The vector θ_t consists of unobservable, time-varying, regression coefficients (see West and Harrison [96] for more details on dynamic linear models). The observational H and state variances W are estimated using the whole sample of observations of risk premia and factors. As such, although the “betas” are time-varying, the structural variances are considered constant over time.³

The sequential model description in (B.22)-(B.23) requires that the defining quantities at time t be known at that time. Let D_0 contains the initial prior information about the elasticities and structural variances. We assume prior information about θ_0 is vague and centered around the initial hypothesis of no effect of risk factors on premia, i.e. $\theta_0|D_0 \sim N(c_0, C_0)$, with $c_0 = 0$ and $C_0 = 10,000$. Also, we assume that the impact of risk factors is highly uncertain and volatile, as captured by an Inverse-Wishart distribution with small degrees of freedom and large scale parameter, i.e. $W|D_0 \sim IW(a_0, A_0)$ with $a_0 = 3$ and $A_0 = 10,000$. As a result we assume that when no historical information on expected risk premia and factors is available, elasticities are mainly driven by idiosyncratic risk as proxied by an Inverse-Gamma distribution with uninformative hyper-parameters, i.e. $H|D_0 \sim IG(n_0/2, n_0 N_0/2)$ with $n_0 = 0.001$ and $N_0 = 0.001$. Notice priors are constant for all maturities $h = 2, 3, 4$ quarters.

In the following we provide details of the Gibbs sampler we use for the estimation of the dynamic linear model (B.22)-(B.23). For the ease of exposition, we disregard the maturity super-script h . Let us denote $\mathbf{x}_{s:t} = (\mathbf{x}_s, \dots, \mathbf{x}_t)$, $s \leq t$, the set of vectors $\mathbf{x}_u$. The collections of parameters is defined as $\Theta = (\theta_{1:T}, W, H)$, respectively, where $\theta_{1:T}$ represents the $(T \times N)$ matrix of state parameters. Let θ_0 represents the initial value of the dynamic sensitivity to the k -dimensional vector of regressors. The complete likelihood function can be defined as

$$p(\mathbf{y}_{1:T}, \theta_{1:T} | \mathbf{Z}_{1:T}, W, H) = \prod_{t=2}^T p(\mathbf{y}_t | \mathbf{Z}_t' \theta_t, H) p(\theta_t | \theta_{t-1}, W), \quad (\text{B.24})$$

with $p(\mathbf{y}_t | \mathbf{Z}_t' \theta_t, H) = N(\mathbf{Z}_t' \theta_t, H)$ and $p(\theta_t | \theta_{t-1}, W) = N_k(\theta_{t-1}, W)$ two univariate and multivariate Gaussian distributions, respectively. Conditional on priors and the latent states $\theta_{1:T}$

³However, the framework could be easily extended by using an exponential weighted moving average recursion to obtain dynamic estimates for $H_{k,t}$ and $W_{k,t}$. We leave this for future research.

the complete likelihood can be factorized as

$$\begin{aligned} p(\theta_{1:T}, W, H | \mathbf{y}_{1:T}, \mathbf{Z}_{1:T}) &\propto p(\mathbf{y}_{1:T}, \theta_{1:T} | \mathbf{Z}_{1:T}, W, H) p(\theta_0, W, H), \\ &= p(\mathbf{y}_{1:T} | \theta_{1:T}, \mathbf{Z}_{1:T}, H) p(\theta_{1:T} | W) p(\theta_0, W, H), \end{aligned}$$

The joint posterior distribution of the states and parameters is not tractable analytically such that the estimator for the parameters cannot be obtained in closed form. The latent variables $\theta_{1:T}$ are simulated alongside the model parameters H and W . At each iteration, the sampler sequentially cycles through the following steps:

1. Draw $\theta_{1:T}$ conditional on H , W and the data $\mathbf{y}_{1:T}, \mathbf{Z}_{1:T}$.
2. Draw W conditional on $\theta_{1:T}$.
3. Draw H conditional on $\mathbf{y}_{1:T}, \mathbf{Z}_{1:T}$, and $\theta_{1:T}$.

In what follows we provide details of each step of the Gibbs sampler.

B.2.1 Step 1. Sampling the Conditional Factor Sensitivities $\theta_{1:T}$

The full conditional posterior density for the time-varying factor loadings is computed using a Forward Filtering Backward Sampling (FFBS) approach as in Carter and Kohn [21]. The initial prior are sequentially updated via the Kalman filtering recursion. Conditionally on idiosyncratic risk H , state variance W , and assuming an initial distribution $\theta_0 | y_0 \sim N(m_0, C_0)$, it is straightforward to show that the (see West and Harrison [96] for more details)

$$\begin{array}{ll} \theta_t | Z_{1:t-1}, W \sim N(a_t, R_t) & \text{Propagation Density} \\ Y_t | Z_{1:t-1}, H \sim N(f_t, Q_t) & \text{Predictive Density} \\ \theta_t | Z_{1:t} \sim N(m_t, C_t) & \text{Filtering Density} \end{array}$$

with

$$\begin{array}{ll} a_t = m_{t-1} & R_t = C_{t-1} + W \\ f_t = Z_t' a_t & Q_t = Z_t R_t X_t' + H \\ m_t = a_t + K_t e_t & C_t = R_t - K_t Q_t K_t' \end{array} \quad (\text{B.25})$$

and $K_t = R_t X_t Q_t^{-1}$ and $e_t = y_t - f_t$. Conditional thetas are drawn from the posterior distribution which is generated by backward recursion (see Frühwirth-Schnatter [45], Carter and Kohn [21], and West and Harrison [96]), i.e. $p(\theta_t | \mathbf{y}_{1:T}) = N_k(m_t^b, C_t^b)$, with

$$\begin{aligned} m_t^b &= (1 - B_t) m_t + B_t m_{t+1}^b, \\ C_t^b &= (1 - B_t) C_t + B_t^2 C_{t+1}^b, \quad \text{with} \quad B_t = \frac{C_t}{C_t + W}, \end{aligned}$$

B.2.2 Step 2. Sampling the State Variance Parameters W

Conditional on the risk exposures, the estimate of the state variance covariance matrix coincide with the update of an Inverse-Wishart distribution. Posterior estimates are obtained by updating the prior structure as

$$W | \theta_{1:T} \sim IW(a_1, A_1) \quad (\text{B.26})$$

with

$$\begin{aligned} a_1 &= a_0 + T \\ A_1 &= A_0 + \hat{\varepsilon} \hat{\varepsilon}' \end{aligned}$$

where $\hat{\varepsilon}' = (\hat{\varepsilon}_1, \dots, \hat{\varepsilon}_T)$ and $\hat{\varepsilon}_t = \hat{\theta}_t - \hat{\theta}_{t-1}$ given $\hat{\theta}_t = m_t^b$.

B.2.3 Step 3. Sampling the Idiosyncratic Risk H

For the posterior estimates of the idiosyncratic risk we exploit the fact that the prior and the likelihood are conjugate. The updating scheme is easily derived as

$$H | \theta_{1:T}, \mathbf{Z}_{1:T}, \mathbf{y}_{1:T} \sim IG(v_1/2, v_1 N_1/2) \quad (\text{B.27})$$

with

$$\begin{aligned} v_1 &= v_0 + T \\ v_1 N_1 &= v_0 N_0 + \hat{v} \hat{v}', \end{aligned}$$

where $\hat{v}' = (\hat{v}_1, \dots, \hat{v}_T)$ and $\hat{v}_t = y_t - Z_t' \hat{\theta}_{t-1}$ given $\hat{\theta}_{t-1} = m_{t-1}^b$

B.3 Forecasts and Learning

At the outset of the Chapter we argue that our model of rational learning closely track the average forecasts of professional analysts, which in turn represents an approximation of investors' expectations. In this Section we test for the null hypothesis that average survey forecast is consistent with a rational adaptive learning framework.

We test for a general rule of updating by estimating the impact of current prices on expectations. Let $E_t[\Delta S_{t+h}]$ represents the investors' expectations at time t for a change in the future spot price from t to $t+h$ (see, e.g. Frenkel and Froot [44] and Pesaran and Weale [82] for more details on testing rationality and adaptivity on survey forecasts). To test adaptivity we first estimate the following regression model;⁴

$$E_t[\Delta S_{t+h}] = \alpha + \beta \Delta S_t + e_t, \quad \text{for } h = 2, 3, 4, \quad \text{quarters}, \quad (\text{B.28})$$

with $\Delta S_t = (S_t - S_{t-h})$ representing past changes in spot prices. The regression equation (B.28) states that if a commodity has been recently depreciated, then it will be expected to depreciate in the near future as well. As we have shown in Chapter one, strong rationality, coupled with no feedback traders activity and no mean reversion in prices, would imply the null hypothesis that there is no "learning" from past information, i.e. $H_0 : \beta = 0$. Panel A of Table B.4 shows the results.

[Insert Table B.4 about here]

Consistently with the results presented in the first Chapter, the slope coefficients are all negative and strongly significant meaning that a recent depreciation of a commodity leads to an optimistic view on future spot prices, and vice versa. Such dynamics does not rule out the possibility of having positive autocorrelation in investors' expectations. Building on this result, we now test the further restriction that individuals are learning adaptively. Under adaptive learning investors adjust their expectations in line with past prediction errors (see, e.g. Nerlove [80], Evans and Honkapohja [39], Cho et al. [24], Sargent [87], Williams [97], Sargent et al. [86], Sargent and Williams [88] and Malmendier and Nagel [72], to cite a few). We test the adaptive learning hypothesis by regressing the expected price change on

⁴We estimate the model by OLS with GMM corrected standard errors to account for autocorrelation and heteroschedasticity in the residuals.

the lagged survey prediction error;

$$E_t [\Delta S_{t+h}] = \mu + \delta (E_{t-h} S_t - S_t) + v_t, \quad \text{for } h = 2, 3, 4, \quad \text{quarters}, \quad (\text{B.29})$$

Panel B of Table B.4 shows the results. The slope coefficients are positive and statistically significant across forecasting horizons and commodity markets. This implies that investors, on average, place positive weight on previous prediction errors. To summarize, investors' expectations on future spot prices are not static; in fact, the elasticity of the expected future spot prices with respect to past forecasting errors is positive and significant. Notably, the support for a form of learning in the expectations formation process does not depend on the prediction horizon and the specific commodity market.

Fig. B.1 Expectations Errors for Future Spot Prices

This figure shows the unexpected changes in spot prices $E_t [S_{t+h}] - S_{t+h}$ for two different horizons, i.e. $h = 2, 4$. Expectations are proxied by the cross-sectional average of the individual Bloomberg's survey of professional analysts, i.e. **Panel A:** Shows the unexpected price changes for WTI Crude Oil (USD/Barrel). **Panel B:** Shows the unexpected price changes for Silver (USD/Ounce). Data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX) and on Silver are obtained from the Commodity Exchange (COMEX). The sample period for the Survey is 12:2006-01:2016, aggregated monthly.




Table B.1 Summary Statistics

This table reports the descriptive statistics for the risk premia for WTI Oil Crude, Copper, Corn and Silver. The ex-ante risk premia are obtained by subtracting from the futures prices the model-implied expected future spot prices for the same maturity, $h = 2, 3, 4$ quarters ahead. WTI Crude Oil prices are in U.S. Dollars per barrel, whereas Copper prices are transformed from USD Cents/Pound USD/Tonne to match the measurement unit used in the survey forecasts. Data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX) and Copper are obtained from the Commodity Exchange (COMEX). Data on Silver are obtained from the Commodity Exchange (COMEX), and data for Corn are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. The sample period is 01:1995-01:2016, monthly.

Panel A: Descriptive Statistics

	WTI			Copper			Corn			Silver		
	$h = 2$	$h = 3$	$h = 4$	$h = 2$	$h = 3$	$h = 4$	$h = 2$	$h = 3$	$h = 4$	$h = 2$	$h = 3$	$h = 4$
Mean	-0.006	-0.012	-0.018	-0.013	-0.019	-0.024	0.020	0.023	0.027	0.019	0.022	0.027
Median	-0.015	-0.020	-0.029	-0.010	-0.012	-0.014	0.031	0.038	0.044	0.015	0.023	0.027
St. Dev.	0.100	0.111	0.123	0.090	0.097	0.101	0.119	0.128	0.137	0.086	0.087	0.096
Skewness	-0.021	0.008	0.022	0.068	-0.360	-0.250	-0.900	-0.785	-0.902	-0.455	-0.476	-0.610
Kurtosis	0.810	0.336	0.027	1.738	1.509	0.249	1.230	0.405	1.045	1.607	1.356	1.641
Jarque-Bera	7.286	1.252	0.029	33.68	30.99	3.47	52.66	29.11	48.19	37.83	30.40	46.33
p-value	0.026	0.535	0.985	0.000	0.000	0.176	0.000	0.000	0.000	0.000	0.000	0.000

Panel B: Correlations

Commodity	h=2				h=3				h=4			
WTI	1				1				1			
Copper	0.355	1			0.280	1			0.246	1		
Corn	0.149	0.139	1		0.139	0.129	1		0.116	0.114	1	
Silver	0.252	0.267	0.220	1	0.216	0.248	0.207	1	0.190	0.198	0.179	1

Table B.2 Static Regression Analysis

This table shows the results of a static regression analysis. The set of predictors Z_t is pre-whitened, i.e. regressors are orthogonal to each other and have standard deviation equal to one, to improve the signal informativeness about the unconditional ex-ante risk premia. The ex-ante risk premia are obtained by subtracting from the futures prices the model-implied expected future spot prices for the same maturity, $h = 2, 3, 4$ quarters ahead. Futures prices data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX), prices are in U.S. dollars per barrel. Futures on Silver and Copper are obtained from the Commodity Exchange (COMEX). S&P500 and MXEF represent monthly returns for the Standard and Poor's 500 and the MSCI Emerging Markets indexes. Hedging pressure (HP) is defined as the net excess in short futures positions by commercial traders, i.e. short minus long positions, divided by the amount of outstanding contracts. Open Interest (OIN) is defined as the total number of outstanding contracts that are held by market participants at the end of the month. The data on commercial traders futures positions are from the Commodity Futures Trading Commission (CFTC). Inventories for Copper and Crude Oil are from the London Metal Exchange (LME) and Energy Information Administration (EIA) respectively. For Corn inventories, we use the U.S. ending stocks reported in thousands of metric tonnes. USD TW stands for the Federal Reserve U.S. trade-weighted exchange rate index, normalized to be equal to one hundred in March 1973. Realised volatility (RVol) is computed as the sum of squared daily returns adjusted for roll-over. For both open interests and inventories, we take the year-on-year growth as explanatory variable. A number is highlighted in grey when the null hypothesis is rejected at least at a 5% significance level. The sample period is 01:1995-01:2016, monthly.

Panel A: Regression Analysis

		S&P500	MXEF	HP	OI	TSM	Value	Inv	USDTW	RVol	R_{adj}^2
WTI	$h = 2$	-0.207	0.319	0.138	0.075	0.144	-0.006	0.214	-0.938	-0.116	0.306
		(-1.060)	(2.001)	(1.673)	(1.917)	(5.880)	(-0.483)	(1.516)	(-3.095)	(-0.768)	
	$h = 3$	-0.182	0.317	0.118	0.084	0.109	-0.004	0.253	-0.855	-0.056	0.376
		(-0.890)	(2.084)	(1.151)	(1.813)	(4.396)	(-0.256)	(1.605)	(-2.595)	(-0.305)	
	$h = 4$	-0.147	0.304	0.086	0.102	0.078	0.003	0.270	-0.754	-0.038	0.486
		(-0.736)	(2.132)	(0.774)	(1.890)	(2.676)	(0.136)	(1.717)	(-2.163)	(-0.185)	
Copper	$h = 2$	-0.021	0.412	0.037	0.060	0.050	0.004	-0.311	-1.021	-0.260	0.345
		(-0.155)	(3.871)	(1.224)	(2.229)	(1.910)	(0.492)	(-1.890)	(-2.153)	(-1.171)	
	$h = 3$	-0.118	0.453	0.036	0.091	0.160	0.010	-0.380	-0.971	-0.336	0.360
		(-0.692)	(3.924)	(1.029)	(2.678)	(2.340)	(1.160)	(-1.986)	(-1.659)	(-1.237)	
	$h = 4$	-0.013	0.395	0.000	0.133	0.287	0.036	-0.413	-0.087	-0.051	0.370
		(-0.085)	(3.257)	(0.001)	(3.002)	(3.078)	(3.815)	(-3.170)	(-1.881)	(-0.198)	
Corn	$h = 2$	-0.008	0.052	0.208	0.098	0.093	0.072	-0.076	-1.008	0.147	0.232
		(-0.035)	(0.348)	(1.137)	(2.403)	(1.994)	(1.686)	(-1.695)	(-1.479)	(0.574)	
	$h = 3$	-0.103	0.083	0.151	0.016	0.120	0.084	-0.084	-1.095	0.372	0.301
		(-0.473)	(0.589)	(1.405)	(2.155)	(2.278)	(1.830)	(-1.065)	(-1.007)	(1.209)	
	$h = 4$	-0.078	0.088	0.089	0.027	0.133	0.094	0.087	-0.962	0.350	0.332
		(-0.321)	(0.566)	(1.796)	(2.876)	(2.509)	(1.996)	(0.800)	(-1.744)	(1.033)	
Silver	$h = 2$	-0.312	0.380	0.105	0.037	0.058	0.021		-0.918	-0.633	0.290
		(-1.386)	(2.288)	(2.713)	(1.449)	(2.740)	(1.025)		(-2.955)	(-3.375)	
	$h = 3$	-0.287	0.417	0.118	0.007	0.044	0.002		-1.130	-0.230	0.256
		(-1.409)	(1.564)	(2.731)	(1.312)	(2.158)	(0.323)		(-3.341)	(-2.482)	
	$h = 4$	-0.233	0.398	0.155	0.007	0.041	-0.005		-1.112	-0.436	0.257
		(-1.173)	(1.811)	(2.721)	(1.273)	(1.980)	(-0.596)		(-3.296)	(-3.765)	

Fig. B.2 Ex-Ante vs Realised Risk Premia

This figure sketches the differences between the *expected* payoff, i.e. ex-ante risk premium, and the *realised* payoff of a futures position. **Panel A:** shows the payoff structure of a futures position keeping the contract until maturity under no unexpected changes in spot prices. In this case, the expected and the realised risk premia coincide. **Panel B:** shows the payoff structure of a futures position keeping the contract until maturity under a negative unexpected fluctuation in spot prices. In this case, the ex-ante and the realised risk premia diverge.




Fig. B.3 Spot Prices and Aggregate Demand

This figure shows the year-on-year growth rates for commodity spot prices (blue line) and the index of world industrial production (magenta line). Top panels compare the changes in world industrial production to the changes in WTI Crude Oil (top-left) and Copper (top-right) spot prices. Bottom panels compare the changes in world industrial production to the changes in Corn (top-left) and Silver (top-right) spot prices. Spot prices data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX), prices are in U.S. dollars per barrel. Spot prices on Silver and Copper are obtained from the Commodity Exchange (COMEX). Silver futures are quoted in U.S. dollars per troy ounce while Copper is quoted in U.S. cents/pound. The index of world industrial production published by the Netherlands Bureau of Economic and Policy Analysis, and contains aggregate information on industrial production from 81 countries worldwide, which account for about 97% of the global industrial production. The sample period is 01:1995-01:2016.




Fig. B.4 Model-Implied vs Survey Expectations ($h = 2$)

This figure compares the expected future spot prices implied by our model with the Survey Price Forecasts from Bloomberg's Commodity Price Forecasts Database for $h = 2$ quarters ahead. WTI Crude Oil prices are in U.S. Dollars per barrel, whereas Copper prices are transformed from USD Cents/Pound USD/Tonne to match the measurement unit used in the survey forecasts. Data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX) and Copper are obtained from the Commodity Exchange (COMEX). Data on Silver are obtained from the Commodity Exchange (COMEX), and data for Corn are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. The shaded area at the bottom shows the difference between expectations from adaptive learning and the survey for the overlapping periods. The sample period for the Survey is 12:2006-01:2016, aggregated monthly. The sample period of the model-implied expected future spot price is 01:1995-01:2016.




Fig. B.5 Model-Implied vs Survey Expectations ($h = 4$)

This figure compares the expected future spot prices implied by our model with the Survey Price Forecasts from Bloomberg's Commodity Price Forecasts Database for $h = 4$ quarters ahead. WTI Crude Oil prices are in U.S. Dollars per barrel, whereas Copper prices are transformed from USD Cents/Pound USD/Tonne to match the measurement unit used in the survey forecasts. Data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX) and Copper are obtained from the Commodity Exchange (COMEX). Data on Silver are obtained from the Commodity Exchange (COMEX), and data for Corn are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. The shaded area at the bottom shows the difference between expectations from adaptive learning and the survey for the overlapping periods. The sample period for the Survey is 12:2006-01:2016, aggregated monthly. The sample period of the model-implied expected future spot price is 01:1995-01:2016.




Fig. B.6 R^2 Decomposition

This figure shows a decomposition of the R^2 obtained from each explanatory variable in the dynamic regression of ex-ante risk premia on pre-whitened risk factors, across different commodities and expectations horizons. The ex-ante risk premia are extracted from the futures prices by using the model-implied expectations for $h = 2, 3, 4$ quarters ahead. Data are obtained from different resources. Futures prices data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX), prices are in U.S. dollars per barrel. Futures on Silver and Copper are obtained from the Commodity Exchange (COMEX). Silver is quoted in U.S. Dollars per troy ounce while Copper is quoted in U.S. cents/pound. We convert the price of Copper futures contracts to USD/tonne to match the measurement unit of the survey forecast, that instead refer to the London Metal Exchange (LME). Corn futures prices are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. Details on the explanatory variables are outlined in Section 4. The sample period 1995:01-2016:01.




Fig. B.7 Time-Varying Betas for WTI Crude Oil

This figure shows the posterior median and credibility intervals for each explanatory variables on the ex-ante risk premia for futures on WTI Crude Oil. The solid blue (dark-yellow) line represents the estimated median beta for the two-quarter (four-quarter) ahead risk premia. The dashed-dot lines show the credibility intervals at the 5% significance level. The sample period 1995:01-2016:01.



Fig. B.8 Time-Varying Betas for Copper

This figure shows the posterior median and credibility intervals for each explanatory variables on the ex-ante risk premia for futures on Copper. The solid blue (dark-yellow) line represents the estimated median beta for the two-quarter (four-quarter) ahead risk premia. The dashed-dot lines show the credibility intervals at the 5% significance level. The sample period 1995:01-2016:01.




Fig. B.9 Time-Varying Betas for Corn

This figure shows the posterior median and credibility intervals for each explanatory variables on the ex-ante risk premia for futures on Corn. The solid blue (dark-yellow) line represents the estimated median beta for the two-quarter (four-quarter) ahead risk premia. The dashed-dot lines show the credibility intervals at the 5% significance level. The sample period 1995:01-2016:01.



Table B.3 Future Spot Prices Out-of-Sample Forecasting Comparison

This table reports the out-of-sample goodness-of-fit statistics R_{OS}^2 computed as in Campbell and Thompson [19]. Statistical significance for the R_{OS}^2 statistic is based on the p-value for the Clark and West [25] out-of-sample Mean Squared Prediction Error (MSPE); the test statistics corresponds to a one-side test of the null hypothesis that the competing specification for the expected future spot prices has equal forecasting performance to our benchmark rational learning specification against the alternative that the competing model has a lower average square prediction error. A number is highlighted in grey when the null hypothesis is rejected at least at a 5% significance level. The sample period is 01:1995-01:2016, monthly.

Panel A: Out-of-sample R^2 statistics ($R_{OS}^2\%$)

Predictor	WTI			Copper			Corn			Silver		
	$h = 2$	$h = 3$	$h = 4$	$h = 2$	$h = 3$	$h = 4$	$h = 2$	$h = 3$	$h = 4$	$h = 2$	$h = 3$	$h = 4$
Futures	0.042	0.616	-0.111	-0.157	0.629	0.999	-0.076	-0.106	-0.595	0.070	0.281	0.042
Current Spot	-0.095	0.227	0.772	-0.266	0.410	0.616	-0.873	-1.195	-2.035	0.040	0.229	-0.043
Spread	0.059	0.669	1.332	-0.143	0.665	1.069	-0.027	-0.089	-0.632	0.069	0.279	0.038

Fig. B.10 Time-Varying Betas for Silver

This figure shows the posterior median and credibility intervals for each explanatory variables on the ex-ante risk premia for futures on Silver. The solid blue (dark-yellow) line represents the estimated median beta for the two-quarter (four-quarter) ahead risk premia. The dashed-dot line shows the credibility intervals at the 5% significance level. The sample period 1995:01-2016:01.




Fig. B.11 Ex-Ante vs Realised Risk Premia (Horizon $h = 2$)

This figure shows the scatter plot of ex-ante vs. realised risk premia. The ex-ante risk premia are extracted from the futures prices by using the model-implied expectations for $h = 2$ quarters ahead. Realised returns are computed as the excess rolling return in the generic contract for the same maturity of the corresponding model-implied expectations. Data are obtained from different resources. Futures prices data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX), prices are in U.S. dollars per barrel. Futures on Silver and Copper are obtained from the Commodity Exchange (COMEX). Silver is quoted in U.S. dollars per troy ounce while Copper is quoted in U.S. cents/pound. We convert the price of Copper futures contracts to USD/tonne to match the measurement unit of the survey forecasts, that instead refer to the London Metal Exchange (LME). Corn futures prices are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. The red line represents the fitted value from a linear regression of the realised returns on the ex-ante risk premia. The sample period 1995:01-2016:01.




Fig. B.12 Ex-Ante vs Realised Risk Premia (Horizon $h = 4$)

This figure shows the scatter plot of ex-ante vs. realised risk premia. The ex-ante risk premia are extracted from the futures prices by using the model-implied expectations for $h = 4$ quarters ahead. Realised returns are computed as the excess rolling return in the generic contract for the same maturity of the corresponding model-implied expectations. Data are obtained from different resources. Futures prices data on Crude Oil (WTI) are from the New York Mercantile Exchange (NYMEX), prices are in U.S. dollars per barrel. Futures on Silver and Copper are obtained from the Commodity Exchange (COMEX). Silver is quoted in U.S. dollars per troy ounce while Copper is quoted in U.S. cents/pound. We convert the price of Copper futures contracts to USD/tonne to match the measurement unit of the survey forecasts, that instead refer to the London Metal Exchange (LME). Corn futures prices are from the Chicago Board of Trade (CBOT) with price quotation in USD cents per bushel. The red line represents the fitted value from a linear regression of the realised returns on the ex-ante risk premia. The sample period 1995:01-2016:01.




Table B.4 Forecasts and Learning

This table shows the results of a test for learning on the cross-sectional average of individual Bloomberg survey price forecasts. The sample period for the survey is 12:2006-01:2016, aggregated monthly and collected for alternative commodities and time-horizons. We exclude from the analysis the survey for Corn as the survey has lots of missing data which would make the sample size subject to small-sample biases. Regressions are estimated by GMM correcting standard errors to account for autocorrelation and heteroskedasticity in the residuals. **Panel A:** shows the results for the null hypothesis that expectations depend on past prices. **Panel B:** shows the results for the null hypothesis that expectations are revised in line with past prediction errors on future spot prices, i.e. learning. Robust standard errors are in parenthesis. A number is highlighted in grey when the null hypothesis is rejected at least at a 5% significance level.

Horizon (Quarters)	Commodity	Panel A: Past Prices				Panel B: Learning			
		α	β	adj R^2	DW	μ	δ	adj R^2	DW
$h = 2$	Crude Oil (WTI)	-0.002 (0.017)	-0.668 (0.095)	0.466	0.567	0.003 (0.009)	0.244 (0.038)	0.558	1.261
	Copper	-0.019 (0.018)	-0.491 (0.193)	0.186	1.019	-0.004 (0.013)	0.226 (0.056)	0.349	1.441
	Silver	0.027 (0.012)	-0.695 (0.069)	0.513	1.123	0.020 (0.009)	0.290 (0.042)	0.396	1.240
$h = 3$	Crude Oil (WTI)	0.020 (0.023)	-0.691 (0.127)	0.324	0.391	0.0216 (0.009)	0.269 (0.031)	0.669	1.199
	Copper	0.002 (0.019)	-0.497 (0.199)	0.166	0.876	0.006 (0.015)	0.197 (0.031)	0.334	1.140
	Silver	0.041 (0.013)	-0.723 (0.079)	0.487	0.932	0.036 (0.013)	0.173 (0.038)	0.201	1.013
$h = 4$	Crude Oil (WTI)	0.041 (0.027)	-0.802 (0.131)	0.321	0.441	0.041 (0.013)	0.221 (0.041)	0.512	0.677
	Copper	0.009 (0.021)	-0.537 (0.161)	0.156	0.982	0.013 (0.019)	0.168 (0.057)	0.197	1.051
	Silver	0.059 (0.015)	-0.769 (0.059)	0.465	1.212	0.054 (0.017)	0.119 (0.038)	0.106	1.251

Chapter 3

From Prices to Fundamentals: A Round-Trip

3.1 Introduction

Over the last decades, different theories have been proposed to explain the puzzling empirical evidence of the persistence in asset returns over short horizons (up to twelve months) and reversal at longer horizons. The seminal works by Cutler et al. [27] and Hong and Stein [58] propose the existence of delays in the processing of fundamental news by rational traders and the activity of momentum traders as the main explanation of this phenomenon.

In their framework, the persistence in returns is caused by an initial underreaction of prices to fundamental news and an overreaction induced by momentum traders. Unfortunately, it is not easy to test directly these predictions empirically by studying the correlation between news and returns.¹ This complexity stems from the fact that it is difficult, not only to agree on what can be considered a stable fundamental pricing factor for a given asset class, but also to agree on what can be considered a news.² Even if these obstacles can be overcome, it remains complicated to find a time-series of fundamental news sufficiently long

¹When it is not possible to study directly the correlation between news and returns, the empirical analysis must rely on indirect effects. For instance studying the autocorrelation of returns itself is an indirect way to test for the possibility that news are not incorporated instantaneously into prices. Unfortunately, such indirect test cannot rule out the possibility of alternative explanations to the phenomenon.

²The stability in terms of the importance of the fundamental pricing factor is crucial for time-series analysis. If the pricing factors follow switching regimes, the analysis of the relationship between news and returns becomes much more difficult in small samples.

to perform a meaningful analysis.

In this Chapter, we exploit the peculiarities of commodity markets to test this mechanism more directly and in an innovative way. In fact, our approach does not consist in collecting news on supply or demand to test if this is reflected into prices, but we do the opposite. We devise an identification strategy, which extracts news from prices of 48 different commodities, and then we test if such information can predict future official statistical publications on industrial production, which is a highly regarded indicator of global demand growth. In this way, we first test if prices carry information about fundamentals, and then we use the filtered news to study the speed at which such information is reflected into different speculative assets.

The news identification strategy is based on the very simple assumptions that commodity prices increase (decrease) across the board after a positive (negative) news about global growth, but supply shocks are uncorrelated and affect a single commodity or substitutes.³ For instance, an unexpected acceleration in China's economic activity will increase commodities demand, which in turn will lead to an increase in commodity prices (with price variations that depend on each commodity sensitivity to economic growth). To the contrary, a drought in South America will affect agricultural commodity prices, but not metals and energy commodities. Given these assumptions, we show that price co-movements are caused by news about global growth, and that news can be filtered using Principal Component Analysis (PCA).⁴

Our empirical analysis shows that prices do reflect fundamental information about global growth. The news extracted from prices correlates well with the OECD publications on industrial production and the level of correlation changes with the basket of commodities used as predicted by the theoretical model we use to guide us in the empirical analysis. Furthermore, we show that prices incorporate also non-publicly available information by showing that the news extracted from prices helps in predicting the official statistical publications that are usually released with three to six months of delay. In fact, in a out-of-sample exercise, the inclusion of the filtered news in the forecasting model reduces the RMSE by 18% compared

³The role of news about global growth on commodity demand is undisputed. Deaton and Laroque [34] explicitly model the effect of global income dynamics on commodity demand and show how it generates co-movements in commodity prices. The role of fundamental factors in commodity markets has been analysed by: Alquist and Kilian [4]; Bakshi et al. [9]; Bhardwaj et al. [16]; Brennan [17]; Deaton and Laroque [31]; Deaton and Laroque [32]; Deaton and Laroque [34]; Fama and French [40]; Gorton et al. [47]; Hamilton [52]; Kilian and Hicks [67]; Szymanowska et al. [93]; Turnovsky [95]; Working [98]

⁴See Section 3.3.1 for a detailed discussion of the assumptions.

to a benchmark forecasting model.

In Section 3.4 we investigate the existence of delays in the information processing as the possible cause of the persistence and predictability of returns by studying if the news extracted from commodity prices are incorporated instantaneously into different speculative assets. To this end, we study the delay in the reaction of returns to fundamental news.

Our findings show that such news are reflected into speculative assets but not instantaneously and with a delay of several months. More specifically, we demonstrate the existence of a delay in information processing, by showing that the correlation between news and asset returns is positive for several periods and goes gradually to zero to become negative only after fifteen months. Furthermore, to quantify the economic significance of such delay in the information processing, we simulate an out-of-sample trading strategy in which we use news with different delays as a trading signal. The results show that also in economic terms the absorption of news about global growth into prices is rather slow, leading to positive excess returns of 12% after one month, and yet 5% six months after the news has occurred.⁵ Furthermore, similar, but negative performances obtained with a delay between fifteen and eighteen months support the existence of positive feedback traders causing overreaction to news as proposed by Cutler et al. [27] and by Hong and Stein [58].

In the last part of the Chapter, we contribute to the literature by rationalising our empirical findings in light of the theories of prices under and overreaction to news by analysing the data generated from a modified version of the Cutler et al. [27] model. The model, which features fully rational traders, latecomers (who are rational but observe fundamental information with a delay), and positive feedback traders (who engage in momentum strategies) allows us to study quantitatively how the interaction among the different types of traders can alter the speed of information processing leading to returns persistence and the time-series momentum anomaly. The results of the Monte-Carlo simulation exercise show that the predictability found in the data is consistent with a delay in the information processing generated by the presence of rationally-bounded traders. More specifically we find that our empirical results can be obtained when around 70% of the traders process news with a delay of several periods. Finally, we show that also momentum traders, but to a lesser extent (20%), are necessary to generate the over-reaction mechanism to news we have found in the data.

⁵This refers to the case of commodity futures.

The Chapter is organised as follows: in Section 3.2 we present the data; in Section 3.3 we introduce the identification strategy of fundamental global growth shocks and we show the predictive power of filtered news; in Section 3.4 we study returns predictability and in Section 3.5 we rationalise our findings by analysing the simulated data. We leave the model derivation and further results to the Appendix.

3.2 Data

We obtain spot prices on 48 commodities from the International Monetary Fund commodity prices database. The sample consists of monthly observations from January 1980 to November 2016.⁶ Commodities are classified, according to the IMF definition, in five groups, namely: Agricultural; Beverages; Food; Metals and Energy commodities. In the analysis, we define returns as monthly log-returns. Year-on-Year returns are defined as the log-returns over the last twelve months. Summary statistics and a detailed description of the commodities are reported in the Appendix.

[Insert Table C.1 about here]

To test how fast the fundamental information on global growth is incorporated into speculative commodity markets we use data on the S&P Goldman Sachs Commodity Index. The index is used to track the performance of commodity markets over time and includes 24 liquid commodities futures contracts from all commodity sectors. The weight of each constituent is based on the average production quantity of each commodity included in the index. Also, to test how fast such information flows into other speculative markets that have a different risk exposure to global growth we, use data on the S&P 500, FTSE All-Share Basic Material Index and the MSCI World Index.⁷ All the data are from Bloomberg. The data frequency and sample size coincide with the commodity spot prices sample. To use the same sample size and data frequency, we proxy global economic activity with the OECD monthly total industrial production, which is usually published with at least three months

⁶In order to exploit the longest homogeneous time-series dimension possible, contracts on natural gas (starting in 1985) and on European sugar (starting in 1991) are not included in the sample.

⁷The MSCI World Index is a stock index including more than 1'600 stocks from 23 countries. Emerging and frontiers economies are not included in the index. The FTSE All-Share Basic Materials Index is a capitalization-weighted index measuring the performance of the basic materials sector of the FTSE Index.

delay and subsequently revised after several months.⁸ Also in this case the sample is monthly from January 1980 to November 2016.

3.3 From Prices to Fundamentals

In this Section, we provide evidence that prices reflect fundamental news about global growth, by showing that such information can be extracted from prices. More specifically, we show that commodity prices contain information about global economic activity, which has not been released yet by statistical agencies, and that it can be used to predict the delayed OECD publications on total industrial production.

3.3.1 Identification

The reason why commodity markets offer a very appealing setting to test if prices reflect news about fundamentals, more specifically about global growth, are two. First, supply and demand are the key fundamentals pricing factors, which importance is indisputable and not model dependent (this reduces the set of fundamentals factors affecting prices, making the analysis easier). Second, the structural heterogeneity of commodity markets gives us the unique opportunity of devising an identification strategy, which relies mainly on three straightforward and credible assumptions: 1) an increase in global economic activity increases commodity demand across all commodity sectors via an income effect; 2) an unexpected increase in demand (which is what we call a news about fundamentals) leads to an increase in prices; 3) supply shocks are uncorrelated across commodities.⁹ The first assumption says that if world economic activity increases, for instance, because China is growing, the demand for commodities will increase across-the-board. The second assumption states that an unexpected increase in commodity demand drives prices up. The third assumption affirms that a drought in South America will affect the supply of corn, but not the amount of copper and oil produced.¹⁰

From a theoretical point of view, these three assumptions together allow us to make two restrictions on the generalised Muth's model price dynamics, shown in Equation (3.1): 1)

⁸ The delay is even longer for global industrial production data. Unfortunately, we cannot use global level data as they start only from 1990.

⁹ Deaton and Laroque [34] show how global income shocks can affect commodity demand.

¹⁰ Of course, these assumptions can be relaxed, allowing for some degree of correlation across supply shocks and the existence of some other common fundamental factor. This will worsen the signal-to-noise problem in the filtering of fundamentals from prices, but it will not alter the prediction of our theoretical framework.

It exists an exogenous common fundamental state variable $x_{j,i,t} = y_t$ for any commodity i , which represents the global income level with innovations $u_{m,i,t+1} = \varepsilon_{t+1}$; 2) The reduced form coefficients $\phi_{\varepsilon i}$ are strictly positive, meaning that a positive news about global growth increases commodity prices across-the-board; 3) the supply shocks $u_{m,i,t+1} = z_{i,t+1}$ are uncorrelated, i.e. $\text{corr}(z_{i,t+1}, z_{j,t+1}) = 0$ for any $i \neq j$.¹¹

$$p_{i,t+1} = \phi_0 i + \sum_{j=1}^J \phi_{x j,i} x_{j,i,t} + \sum_{m=1}^M \phi_{u m,i} u_{m,i,t+1} \quad (3.1)$$

As shown in our model presented in the Appendix C.1, by applying these assumptions to the Muth's model, we obtain the following price dynamics,

$$p_{i,t+1} = \phi_0 i + \phi_{y i} y_t + \phi_{\varepsilon i} \varepsilon_{t+1} - \phi_{z i} z_{i,t+1} \quad (3.2)$$

which implies that co-movements in commodity prices are driven by common demand shocks. In the next Subsection we use Principal Component Analysis to extract news about global growth from commodity returns.

[Insert Table C.2 about here]

3.3.2 Filtering

We perform the filtering exercise to extract the news about global growth applying Principal Component Analysis on Year-on-Year log-returns on commodity spot prices.¹²

[Insert Figure C.1 about here]

The price dynamics introduced in Equation (3.2), predicts that the first principal component captures common demand news about the global economic activity. Figure C.1 and C.2 confirm that the correlation between the first principal component and changes in industrial production is indeed positive and high. Correlation tends to increase when PCA is performed

¹¹Both demand and supply shocks are not autocorrelated and not cross-correlated. See Muth [79] for the original model.

¹²Using Log-Returns instead of prices reduces the spurious correlation problem that arises when analysing highly persistent time-series. Furthermore, we use Year-on-Year returns to filter-out seasonality effect, and as shown in Section 3.5, longer look-back periods improve the signal-to-noise when fundamentals are persistent.

on portfolios including industrial commodities and metals. This result is intuitive and in line with the theoretical framework, which predicts that highly pro-cyclical commodities, such as industrial goods, should be more sensitive to global demand shocks and therefore more informative.¹³

[Insert Figure C.2 about here]

A further prediction of the model (see Appendix C.1) is that the first principal component loadings should be positive and that they should be higher for commodities that are more sensitive to global economic activity variations, such as industrial and energy products. Indeed, Figure C.3 and Table C.2 show that the weight is higher for energy and metals, then decreasing for agricultural industrial products, food and beverage. All loadings as expected are positive, with the only exception of soft sawn-wood, which is negative, but very close to zero. Furthermore, copper price, which is traditionally considered a barometer of global industrial activity, receives one of the largest weight in the first component.

[Insert Figure C.3 about here]

3.3.3 Predicting Fundamental News

If prices do reflect fundamental news about global economic growth, then looking at commodity prices should give us real-time information about the state of the economy and therefore help us in predicting the official statistical publication on industrial production. In fact, Equation (3.2) shows that spot prices, besides idiosyncratic changes in supply, reflect unexpected variations in common commodity demand driven by global growth. Intuitively, an increase in economic activity should increase first demand of raw material, which *ceteris paribus* increases spot prices, and only later it should be recorded in statistical publications. In practice, this means that using information derived from prices at the end of March should help us in predicting the official statistics (referring to March), that will be published with a delay at the end of June 2016.¹⁴

¹³The price dynamics in Equation (3.2) shows that when commodities are not sensitive to world income variations, i.e. $\phi_{y,i} = 0$ and $\phi_{e,i} = 0$, they do not carry any information about global growth.

¹⁴Usually, official publications occur with a three months delay. As shown in the model, prices informativeness increases with commodity price sensitivity to global economic activity and decreases with the size of supply shocks.

To test the informative content of commodity prices, we compare the out-of-sample forecasts Root Mean Squared Error (RMSE) of a predictive regression, which includes real-time information extracted from commodity returns, against the ones obtained using univariate benchmark models. More specifically, the augmented model includes as an exogenous predictor the first principal component of commodity returns, which as shown in Subsection 3.3.1 filters news about the global growth process. To make the out-of-sample exercise as realistic as possible, the first principal component is estimated using a rolling sample, which does not include prices that would not be available in a real-life application.¹⁵ Likewise, the predictive regressions parameters are calculated using the corresponding data sample.

[Insert Table C.3 about here]

Table C.3 shows that the predictive regression, which includes the real-time information extracted from commodity returns (model $m = 1$), has lower RMSE compared to the benchmark Random Walk and AR models. The results are also robust to different lag specifications, with the models of order one performing better than less parsimonious lag specifications. To understand, if the forecasting performance improves because we are actually extracting real-time information about global industrial production, we perform the same out-of-sample exercise using the information obtained from different portfolios of commodities. If what we are capturing is truly fundamental information about economic growth and not just a collection of unspecified news, we should obtain a lower RMSE when extracting information from industrial commodities compared to food and beverages. Indeed, Table C.3 shows that the lowest RMSE is obtained when the group of Metals and Industrial commodities are used. On the opposite, the highest RMSE is obtained when we attempt to extract information about global growth from Food and Beverages commodities, which is intuitive since they are expected to carry less information about economic cycles. Therefore, also the cross-sectional evidence supports the model's prediction that the first principal component does capture news about global growth.

¹⁵We use a rolling sample to account for possible trends or structural changes in the parameters.

3.4 Returns Predictability

3.4.1 Simulated Trading Strategy

In the previous Section, we showed that spot prices do reflect news about global growth, but as shown in our Model presented in Section C.1, this does not necessarily imply that futures prices do not deviate from the rational expectations price. Indeed, prices can reveal information about non-publicly available information even if it exists some level of mispricing. In our model, futures price deviations occur when latecomers, who observe news with some delay, or momentum traders exist.¹⁶ In both cases, the presence of rationally bounded traders causes mispricing and returns predictability opportunity arise for fully informed traders.

In this Section, we test if commodity futures prices are informationally efficient with respect to fundamental news on global economic growth. To do so, we simulate trading strategies that entail using real-time information about fundamental news extracted from commodity returns to generate trading signals. In fact, the principal component of commodity returns should approximate fairly well the information set on global demand growth of fully informed traders.¹⁷ If such trading strategy did not generate extra profits compared to long-only strategies, this would be an indication that mispricing does not exist or if exists it is too little to matter in small samples.¹⁸

[Insert Figure C.4 about here]

To the contrary, Figure C.4 shows that the simulated trading strategy in which we buy (sell) a portfolio of commodity futures using the news about fundamentals extracted in real-time from prices to generate trading signals, performs better than the long-only strategy (see Figure C.4 for a detailed description of the trading strategies). Also, such performance mimics quite well the one obtained with a perfect foresight test, which entails using the

¹⁶In general, prices deviate from the fully rational equilibrium price when bounded rational traders are present. Bounded rationality, intended as limited access to information, can be more or less severe, with traders having access to lagged fundamental information and traders engaging in momentum strategies because either they have no access to data or because the analysis is too costly.

¹⁷In Section 3.3 we show in detail how to identify news about industrial production from commodity returns and we show that the industrial metals portfolio gives, as predicted, the first principal component with the highest correlation with the industrial production. In this Section we use the industrial metals subgroup to estimate the rolling first principal component.

¹⁸Obviously, it could be the case that we do not obtain any extra profit because our proxy for the news about global growth is not good enough. However, this makes our test even tougher with respect to the existence of mispricing.

future releases of the industrial production statistics, that would have not been available when the trading decision was made. This suggests that our proxy does not differ significantly from what should be the information set on the global growth of fully informed (rational) traders. We explain the slightly better performance of the trading strategy on the S&P Goldman-Sachs Commodity Index based on the filtered information, compared to the future official statistics, as the result of the discrepancy from the OECD industrial production statistics and what it would be the global figure including all countries if available. In fact, whereas the official OECD statistics track only the developed countries industrial production levels, commodity prices also reflect the commodity demand of important economies, like China, making our measure possibly a more accurate metrics of global economic activity.

Figure C.5 shows that our results are statistically significant when we compute confidence intervals obtained by applying the same trading signals to ten-thousands random walk processes that by construction contain no predictability. This experiment gives us confidence that the performances obtained are unlikely the result of chance (i.e. that the performances could also be obtained by chance when applying our trading signal to a price process that in reality is not predictable at all). As a further robustness check we perform the same exercise but assuming that we cannot observe up-to-date commodity spot prices and we can only observe prices with a three months delay. Since in our framework returns predictability is generated by latecomers that observe fundamentals information with a delay, by using lagged prices we should observe lower trading strategy performances compared with ones obtained by extracting fundamental information from live prices. However, such performances should still beat the long-only strategy, provided that on average fundamentals are observed and processed with more than three months of delay. Figure C.6 indeed shows that, as expected, when we introduce a three months delay, we obtain lower trading strategy performance that however still beat long-only strategies. The only exception is the performance of the trading strategy on the S&P500, which after a long period of over-performance, obtains at the end of the sample the same cumulated return of the long-only strategy. This difference is mainly due to the fact that by delaying commodity prices by three months, the industrial metals prices that started decreasing during the summer of 2008 before the financial crisis, cannot provide a timely sell signal.

We also apply the same strategy to the S&P 500, the Basic Material Stocks Index and the MSCI World to test if, in other markets, the information about global growth is better incorporated into prices compared to the commodity market. Figure C.4 shows that also in

these cases we obtain extra profits, but especially during periods in which global growth mattered the most, i.e. prior and post the 2008 financial crises and after the end of the "Dot-Com Bubble". This evidence can therefore solve the puzzling evidence that time-series momentum strategies performance tend to be correlated across different markets especially during periods of high market volatility. This is because, when news about global growth becomes a crucial pricing factor in all markets, the existence of delays in processing such news creates the persistence in returns, which makes time-series momentum strategies profitable.

In terms of cross-sectional evidence, the fact that the trading performance is higher for the assets that are supposed to be more exposed to global growth, like commodity derivatives and basic material sector stocks, compared to general stocks market indices, offers additional supporting evidence that the predictability is not the result of chance. Intuitively, global growth is constantly an important pricing factor for commodities and stocks related to the commodity sector, and therefore trading strategies based on the news on global growth should work more consistently in these markets compared to other markets where global growth becomes the most important pricing factor only in certain periods of time.

3.4.2 Returns Response to News

In this Subsection, we quantify how long it takes for news about fundamentals to be absorbed into prices. As shown by Cutler et al. [27] and Hong and Stein [58], predictability arises when latecomers, who process news with a delay, are present in the market. In these models, the delay in news processing generates underreaction to news, which implies a positive correlation between returns and past news. Also our model (see Appendix C.1) generates data with a positive and statistically significant correlation between returns and past fundamental news when a sufficiently large speculative activity of non-fully rational traders take place. Furthermore, as in the models mentioned above, the predictability horizon depends on the speed of news processing.

To study the speed at which news is incorporated into prices, we analyse the predictability of returns using the news on fundamentals with different delays. In an entirely rational model, only the contemporaneous correlation between returns and fundamental news should be significantly different from zero. Otherwise, also lagged news should be significantly correlated with returns, showing evidence of predictability due to informational frictions and speculative activity of rationally bounded traders.

[Insert Figure C.7 about here]

Figure C.7 shows that on average news on global growth tend to be absorbed gradually, and only after three months the correlation between the news and returns becomes not significantly positive. This is coherent with the existence of latecomers, who update their pricing model and their portfolio positions only after the news has become publicly available on the OECD website (which usually takes three months). After three months, the decline in correlation continues, except a non-statistically significant bump, which peaks with a lag of four quarters. This subsequent increase in positive correlation can be explained, as in Cutler et al. [27], by a group of latecomers which processes the information only after four quarters or/and by the existence of negative feedback traders. Finally, the negative correlation, which occurs after fifteen months, can be explained by positive feedback traders with a long memory (i.e. with a long look-back period in past returns) that can generate negative correlations at long horizons.

[Insert Figure C.8 about here]

To quantify the economic significance of such evidence of returns predictability, we perform an out-of-sample trading strategy in which we use news with different delays as a trading signal on the S&P Goldman Sachs Commodity Index. If the correlation reported in Figure C.7 is truly evidence of predictability of returns due to slow absorption of news into prices, then we should observe positive trading performances that will first decrease to become negative after fifteen months, mimicking the correlation pattern. Indeed, Figure C.8 shows that annualised returns (in excess of the long-only performance) are positive and decrease progressively to plateau after three months. The performance starts becoming negative after one year and bottoming after fifteen months, which corresponds to the beginning of the negative correlations reported in Figure C.7. In conclusion, the simulated out-of-sample trading exercise shows that also in economic terms the absorption of news about global growth into prices is rather slow, leading to positive excess returns of 5% also after six months the news has occurred. Furthermore, similar, but negative performances obtained with a delay between fifteen and eighteen months are further supporting evidence of the existence of positive feedback traders with a long memory as proposed by Cutler et al. [27] and by Hong and Stein [58].

3.5 Simulation

In this last Section, we rationalise our findings by analysing simulated data from a modified version of the Cutler et al. [27] model, which features both commodity spot and futures markets. Futures prices are determined by the interaction of fully rational traders with latecomers, who are rational but observe fundamental information with a delay, and positive feedback traders who engage in momentum strategies. The Model is presented in the Appendix C.1.

In Subsection 3.5.1 we investigate how returns respond to news on global growth when the share of different types of traders changes. In Subsection 3.5.2, we show that using monthly data, the extraction of global growth news from prices improves when log-returns are overlapping and reaches a plateau with a look-back period close months. In Subsection 3.5.3 we show how the effectiveness of PCA changes when the number of commodities changes in the portfolio for different levels of persistence in the global growth dynamics.

3.5.1 Returns Response to News: Interpretation

As shown in Figure C.7, on average news on global growth tend to be absorbed gradually and only after three months the correlation between the news and returns becomes not significantly positive. After three months, the decline in correlation continues, except a non-statistically significant bump, which peaks with a lag of four quarters. In this Subsection, we analyse the returns response to news from simulated data from the model presented in the Appendix C.1.

The model allows us to change the share of trading activity of rational, latecomers and feedback traders respectively, enabling us to understand the role of each category in the diffusion process of fundamental news on global growth and therefore explaining the correlation between returns and lagged news.

[Insert Figure C.9 about here]

Panel (a) and (b) of Figure C.9 analyse the case in which the market is populated only by rational traders and latecomers. The two Figures show that the presence of latecomers generates a positive correlation between returns and lagged news on global growth. In fact, latecomers who observe and process information with a one-period delay can generate significantly positive correlation at the first lag. This result is because it takes only one-

period before the information is fully available. Therefore, to produce a more persistent positive correlation between past news and returns, as shown in the data, it is necessary to include latecomers who process information with a larger delay. Panel (b) indicates that by including latecomers with an information processing delay of twelve months, we can obtain significantly positive correlation up to the 12th lag. This suggests that the gradual decline in correlation observed in the data could be explained by the presence of different groups of latecomers who process the information with various delays. However, latecomers can generate only underreaction to news and cannot account for the existence of a negative correlation between past news and returns. To reproduce such feature, we need to allow for the presence of positive feedback traders that generate price overreaction as in Cutler et al. [27] and Hong and Stein [58].

Panel (c) and (d) show that the presence of positive feedback traders generates first a price overreaction, which is adjusted at a later period. Positive feedback traders, who base their asset demand on prior period return, produce negative correlation between a positive fundamental news and returns in the following period as rational traders who accommodate their demand taking a short position need to be compensated on average by negative returns in the next period. Panel (d) shows that when the look-back period of feedback traders is larger than one-period, the overreaction to news persist for longer periods, to eventually disappear when the news fall out of the look-back period. The existence of feedback traders with a long memory can, therefore, explain the negative correlation between news and returns at distant lags found in the data.

Panels (e) shows that to reproduce a similar correlation pattern between lagged fundamental news and returns as shown in Figure C.7, we need to allow for a significant presence of latecomers (see Figure C.9 for calibration) and a minority of positive feedback traders with a long memory. It is important also to notice that to have such pattern, rational traders who observe fundamental shock instantaneously need have a minor role in the trading activity otherwise the impact of the bounded traders becomes negligible from a statistical point of view given our calibration.

3.5.2 Look-Back Period

In this Subsection, we use Monte-Carlo simulation, to show how the extraction of news on global growth from prices improves when log-returns are overlapping and reaches a plateau

with a look-back period close to twelve months.

We define the price dynamics as

$$p_{i,t+1} = \alpha y_t + \varepsilon_{t+1} - z_{i,t+1} \quad (3.3)$$

which is a simplified version of the price process defined in Equation (3.2), and y_t is an AR(1) process with autoregressive coefficient α with unit variance white-noise innovations ε_t . Also, supply shocks z_t are unit variance white-noise.

[Insert Figure C.10 about here]

Figure C.10 shows that the signal extraction improves with the look-back period. As the look-back period increases, also the correlation between log-returns and log-differences in the demand process y_t , which is the true signal, increases. Intuitively, whereas supply shocks are transitory, demand shocks are persistent and sum-up with an exponentially decreasing weight given by the AR(1) coefficient; therefore increasing the Signal-to-Noise. Proposition C.4, presented in the Appendix C.4, shows this result analytically.

3.5.3 Portfolio Size

In this Subsection, we show with a Monte-Carlo experiment, that the effectiveness of Principal Component Analysis, in extracting the signal on common demand fundamentals increases rapidly with the portfolio size, and increases more quickly for less persistent demand processes. Figure C.11 shows changes in average correlation between the first principal component and log-differences of the signal process as a function of the number of assets in the portfolio. In this case, the returns look-back period is one. Correlations and confidence intervals are estimated on 10000 simulated dynamics of the demand and the price process defined by Equation (3.3).

[Insert Figure C.11 about here]

3.6 Concluding Remarks

Our analysis exploits the peculiarities of commodity markets to show that fundamental news about global growth is reflected into prices, but not instantaneously. The news can, therefore, be filtered in real-time from commodity prices, but such news takes several months before being fully incorporated into prices, leading to returns predictability.

Coherently with the theories of overreaction and underreaction to news, we show using simulated data that the results obtained can be explained by the existence of latecomers, who process information with a delay, and momentum traders.

Our analysis assumes therefore that not all investors have access to the same information and learn from past experience. We are confident that such assumptions, not only allow to give a simple explanation to the existence of asset pricing anomalies such as prices under and overreaction to news and the time-series momentum but also provide a more realistic representation of investors' behaviour. Asymmetric information, high costs of data analysis, unstable time-series relationships and limited data samples, force often investors to use simpler expectations formation rules that are far from being perfectly rational and should be taken into account more seriously in future research.

Appendix C

Appendix and Tables

C.1 Model

In this Section, we present a model, which builds on Muth [79] and Cutler et al. [27], to illustrates how the presence of traders with different levels of information about fundamental news affects returns predictability and the diffusion process of information into prices.

We consider a futures market on a non-storable commodity, in which contracts have one-period maturity and are cash-settled against the spot market price. The spot price follows a fully rational Cobb-Web model in which commodity demand is affected positively by shocks (news) to global growth.¹ Futures contracts are in zero net supply, and the futures price is determined by the interaction of three different categories of traders, namely: rational, latecomers and feedback traders. Rational traders make their investment choices using all the existing information, whereas latecomers act rationally but observing global growth with different delays.² Feedback traders instead, have no access to fundamental information and form their future returns expectations extrapolating future returns from price trends, i.e. they engage in momentum strategies.

In the following two Subsections we describe the spot and futures market dynamics.

¹Persistent global income shocks are introduced to make the model more realistic and to create a role for forecasting demand. Furthermore, it allows us to show formally how global shocks can affect the cross-section of commodity prices.

²This is a realistic assumption since statistics on global GDP, and industrial production are published with several months of delay

C.1.1 Spot Market

Consider a partial equilibrium Muth's model for a non-storable commodity, in which producers are rational and can observe the current state of the economy.

$$D_{t+1} = a + by_{t+1} - dp_{t+1}, \quad \text{with } a, b, d > 0 \quad (\text{Demand}) \quad (\text{C.1})$$

$$y_{t+1} = \alpha y_t + \varepsilon_{t+1}, \quad \text{with } 0 < \alpha < 1, \quad \varepsilon_{t+1} \sim iid(0, \sigma_\varepsilon^2) \quad (\text{Global Income}) \quad (\text{C.2})$$

$$S_{t+1} = \lambda E_t p_{t+1} + z_{t+1}, \quad \text{with } \lambda > 0 \quad z_{t+1} \sim iid(0, \sigma_z^2) \quad (\text{Supply}) \quad (\text{C.3})$$

$$D_{t+1} = S_{t+1} \quad (\text{Market Clearing}) \quad (\text{C.4})$$

Where D_{t+1} is the aggregated demand at time $t + 1$, which is log-linear and increasing function of global income y and decreasing in prices p . Global income is assumed to be exogenous and follows a stationary autoregressive process of order one. Supply S_{t+1} is log-linear in producers price expectations.

By plugging Equations (C.1) and (C.3) in the market clearing condition (C.4), we obtain a standard rational expectations dynamics for the spot price.

$$p_{t+1} = \frac{1}{d} [a + b\alpha y_t - \lambda E_t p_{t+1} + b\varepsilon_{t+1} - z_{t+1}] \quad (\text{C.5})$$

By taking expectations on both sides, we obtain producers price expectations as function of future expected global income

$$E_t p_{t+1} = \frac{a + b \overbrace{\alpha y_t}^{E_t y_{t+1}}}{d + \lambda} \quad (\text{C.6})$$

By plugging price expectations in the market equilibrium equation (C.5) we obtain the following RE price dynamics.

$$p_{t+1} = \frac{a}{\lambda + d} + \frac{b\alpha}{\lambda + d}y_t + \frac{1}{d}(b\varepsilon_{t+1} - z_{t+1}) \quad (\text{C.7})$$

Equation (C.7) can be written also be written in reduced form as

$$p_{t+1} = \phi_0 + \phi_1 y_t + \eta_{t+1}, \quad \text{with } \eta_{t+1} \sim iid(0, \sigma_\eta^2) \quad (\text{C.8})$$

with $\phi_0 = a/(\lambda + d)$; $\phi_1 = b\alpha/(\lambda + d)$ and $\eta_{t+1} = (b\varepsilon_{t+1} - z_{t+1})/d$

C.1.2 Futures Market

Consider a futures market on a non-storable commodity in which futures have one-period maturity and are cash-settled against the spot price. The futures price is determined by the interaction of three types of traders: rational, latecomers and feedback traders. Rational traders' futures demand depends on expected returns given all the existing information and perfect knowledge of the model and its parameters.

$$q_{R,t} = N_R \gamma_R E_t[R_{t+1} - \rho | y_t], \quad (\text{Rational}) \quad (\text{C.9})$$

Latecomers are rational traders who observe information about global growth with a delay and therefore create their returns expectations based on multiple periods forecasts on the growth process y .

To analyse the effect of heterogeneous information diffusion, we subdivide latecomers into two groups: $\delta_{L|1}$ share of the latecomers observe the fundamental y with one-period delay and $\delta_{L|l} = 1 - \delta_{L|1}$ share with l lags being a delay larger than one-period.

$$q_{L,t} = N_L \left\{ \delta_{L|1} \gamma_{L|1} E_t[R_{t+1} - \rho | y_{t-1}] + (1 - \delta_{L|1}) \gamma_{L|l} E_t[R_{t+1} - \rho | y_{t-l}] \right\}, \quad (\text{Latecomers}) \quad (\text{C.10})$$

Feedback traders have no access to any information but past spot prices (or it is too costly for them to analyse other data) and engage in time-series momentum strategies. Also in this case we introduce two sub-groups to analyse the effect of different look-back periods l_F on returns.

$$q_{F,t} = N_F \left\{ \delta_{F|1} \gamma_{F|1} [(p_t - p_{t-1}) - \rho] + (1 - \delta_{F|1}) \gamma_{F|l_F} [(p_t - p_{t-l_F}) - \rho] \right\}, \quad (\text{Feedback}) \quad (\text{C.11})$$

where $N_i \geq 0$ defines the number of traders of type i (which sum is rescaled to one $N_R + N_L + N_F = 1$ to be interpreted as a share of the total population); ρ is the required rate of return on the risky asset; and $\gamma_i > 0$ is the responsiveness of traders to a change in expected returns (or price trends in the case of feedback traders). Returns are defined as $R_{t+1} \equiv f_{t+1}^0 - f_t^1$, where f_t^M is the log-price at time t of futures with maturity M . Since futures at maturity are cash settled, i.e. by no-arbitrage $f_{t+1}^0 = p_{t+1}$, returns can be written also as $R_{t+1} = p_{t+1} - f_t^1$.

By plugging the definition of returns, the reduced form spot price dynamics defined in equation (C.8) and setting the required rate of return to zero ($\rho = 0$), and traders responsiveness equal to one across groups ($\gamma = 1$), traders demand can be written as

$$q_{R,t} = N_R (\phi_0 + \phi_1 y_t - f_t^1), \quad (\text{Rational}) \quad (\text{C.12})$$

$$q_{L,t} = N_L \left\{ \delta_{L|1} (\phi_0 + \phi_1 \alpha y_{t-1} - f_t^1) + (1 - \delta_{L|1}) (\phi_0 + \phi_1 \alpha^{l_L} y_{t-l_L} - f_t^1) \right\}, \quad (\text{Latecomers}) \quad (\text{C.13})$$

$$q_{F,t} = N_F \left\{ \delta_{F|1} (p_t - p_{t-1}) + (1 - \delta_{F|1}) (p_t - p_{t-l_F}) \right\}, \quad (\text{Feedback}) \quad (\text{C.14})$$

From the above equations we can see that, whereas rational traders demand depends on the current state of the economy y_t , latecomers' demand reflects not only predictions about the future state of the economy, but also predictions about the current state of the economy as they cannot observe global growth in real time. Conversely, feedback traders' demand depends only on past price changes over a different look-back period depending on the group they belong.

By plugging traders' demand into the zero-net supply condition of futures market

$$q_{R,t} + q_{L,t} + q_{F,t} = 0, \quad (\text{Market Clearing}) \quad (\text{C.15})$$

we obtain the futures price dynamics reduced form

$$f_t^1 = \phi_0 + \psi_1 y_t + \psi_2 y_{t-1} + \psi_3 y_{t-l_L} + \psi_4 p_t - \psi_5 p_{t-1} - \psi_6 p_{t-l_F} \quad (\text{C.16})$$

with

$$\psi_1 = \frac{N_R \phi_1}{k} \quad (\text{C.17})$$

$$\psi_2 = \frac{N_L \delta_{L|1} \phi_1 \alpha}{k} \quad (\text{C.18})$$

$$\psi_3 = \frac{N_L (1 - \delta_{L|1}) \phi_1 \alpha^{l_L}}{k} \quad (\text{C.19})$$

$$\psi_4 = \frac{N_F}{k} \quad (\text{C.20})$$

$$\psi_5 = \frac{N_F \delta_{F|1}}{k} \quad (\text{C.21})$$

$$\psi_6 = \frac{N_F (1 - \delta_{F|1})}{k} \quad (\text{C.22})$$

$$k = N_R + N_L \quad (\text{C.23})$$

from the future price dynamics we can see that futures prices are increasing in the current and past state of the economy and in current price levels, but decreasing in past prices over different look-back period.

C.2 Tables and Figures

Fig. C.1 Filtered Fundamentals vs. Actuals

The Figure shows the rescaled filtered and actual global growth news. Filtered fundamentals are given by the first principal component, which is computed on the correlation matrix of Year-on-Year returns on monthly commodity prices from January 1980 to November 2016. Actual fundamental variations are the Year-on-Year log-differences on the monthly OECD total industrial production. The commodity portfolio contains 48 commodities.




Table C.1 Summary Statistics

Commodity	Mean	Median	Std	Skew	Kurt	Max	Min	Range	Obs.
Aluminium	0.000	-0.002	0.056	-0.454	6.329	0.180	-0.326	0.507	442
Copper	0.002	0.003	0.063	-0.381	7.068	0.248	-0.354	0.602	442
Iron Ore	0.004	0.000	0.061	3.030	28.018	0.539	-0.191	0.730	442
Lead	0.002	0.002	0.070	-0.261	5.318	0.307	-0.284	0.591	442
Nickel	0.001	-0.006	0.083	0.751	9.173	0.581	-0.381	0.962	442
Tin	0.000	0.000	0.056	-0.482	5.721	0.159	-0.252	0.411	442
Uranium	-0.002	0.000	0.066	0.370	10.337	0.374	-0.313	0.687	442
Zinc	0.003	0.003	0.062	-0.330	4.719	0.234	-0.292	0.526	442
Cotton	0.000	-0.002	0.055	-0.051	6.142	0.206	-0.269	0.475	442
Hides	0.001	0.000	0.066	-1.515	16.998	0.257	-0.547	0.804	442
Soft Logs	0.001	-0.002	0.068	0.005	5.385	0.287	-0.344	0.631	442
Hard Logs	0.001	0.000	0.063	-0.152	11.591	0.348	-0.415	0.763	442
Rubber	0.000	0.000	0.067	-0.432	6.574	0.196	-0.394	0.590	442
Hard Sawnwood	0.002	0.000	0.050	0.682	14.005	0.348	-0.271	0.618	442
Soft Sawnwood	0.002	0.001	0.071	-0.066	16.881	0.502	-0.538	1.040	442
Wool (coarse)	0.001	-0.002	0.048	0.082	5.846	0.199	-0.252	0.450	442
Wool (fine)	0.001	-0.002	0.063	-0.150	7.673	0.240	-0.398	0.638	442
Cocoa	-0.001	-0.003	0.059	0.175	3.516	0.210	-0.188	0.398	442
Arabica	0.000	-0.003	0.077	0.586	6.679	0.423	-0.361	0.784	442
Robusta	-0.001	-0.006	0.067	0.446	6.368	0.374	-0.251	0.625	442
Tea	0.001	0.000	0.075	-0.007	4.080	0.290	-0.280	0.570	442
Bananas	0.002	-0.002	0.160	0.377	4.450	0.593	-0.456	1.049	442
Barley	0.002	0.000	0.070	0.032	5.392	0.287	-0.274	0.561	442
Beef	0.001	0.000	0.039	0.028	5.711	0.180	-0.180	0.360	442
Rapeseed oil	0.001	-0.002	0.077	0.705	16.170	0.576	-0.455	1.031	442
Fishmeal	0.001	0.000	0.049	-0.006	7.041	0.266	-0.240	0.506	442
Groundnuts	0.001	0.000	0.075	0.475	10.769	0.352	-0.451	0.803	442
Lamb	0.000	-0.001	0.041	-0.053	4.506	0.132	-0.180	0.312	442
Corn	0.001	0.001	0.058	-0.160	6.419	0.287	-0.252	0.538	442
Olive Oil	0.001	0.001	0.044	0.253	8.050	0.258	-0.207	0.466	442
Oranges	0.003	0.005	0.126	-0.319	4.077	0.413	-0.467	0.880	442
Palm oil	0.000	0.003	0.078	-0.085	4.914	0.290	-0.316	0.606	442
Swine	-0.001	-0.002	0.109	0.124	6.093	0.602	-0.492	1.094	442
Poultry	0.003	0.002	0.022	0.782	5.535	0.105	-0.056	0.160	442
Rice	0.000	-0.001	0.059	1.210	12.253	0.412	-0.281	0.693	442
Fish	0.000	0.002	0.058	-0.131	4.283	0.189	-0.210	0.399	442
Shrimp	0.000	0.000	0.045	-1.250	11.449	0.170	-0.322	0.491	442
Soybean Meal	0.001	-0.003	0.063	0.032	5.573	0.269	-0.316	0.585	442
Soybean Oil	0.001	0.000	0.059	0.362	5.944	0.343	-0.253	0.597	442
Soybeans	0.001	0.001	0.057	-0.002	6.167	0.250	-0.256	0.507	442
Sugar	0.000	0.002	0.092	0.290	4.059	0.376	-0.273	0.648	442
Sugar U.S.	0.001	0.000	0.042	0.781	19.210	0.341	-0.260	0.601	442
Sunflower oil	0.001	0.000	0.074	2.170	23.988	0.661	-0.422	1.083	442
Wheat	-0.001	-0.002	0.059	0.154	5.219	0.247	-0.253	0.500	442
Coal	0.002	0.000	0.054	0.076	12.569	0.364	-0.329	0.692	442
Dated Brent	0.000	0.002	0.087	-0.053	5.790	0.466	-0.313	0.780	442
Dubai Crude Oil	0.000	0.006	0.086	-0.152	7.787	0.521	-0.335	0.857	442
WTI	0.000	0.001	0.082	-0.396	6.269	0.391	-0.395	0.786	442

Table C.2 1st Principal Component Loadings on Commodity Returns

Portfolios	Metals	Agri	Beverage	Food	Energy	Industrial	Non Energy	All
Aluminium	0.31					0.28	0.21	0.19
Copper	0.39					0.33	0.24	0.23
Iron Ore	0.09					0.09	0.08	0.09
Lead	0.35					0.30	0.21	0.20
Nickel	0.58					0.47	0.31	0.29
Tin	0.28					0.24	0.20	0.20
Uranium	0.25					0.19	0.14	0.12
Zinc	0.37					0.28	0.17	0.16
Cotton		0.44				0.21	0.19	0.18
Hides		0.28				0.17	0.12	0.12
Soft Logs		0.04				0.01	0.00	0.00
Hard Logs		0.24				0.09	0.07	0.07
Rubber		0.54				0.33	0.28	0.25
Hard Sawnwood		0.22				0.09	0.07	0.07
Soft Sawnwood		-0.01				-0.01	-0.02	-0.01
Wool (coarse)		0.34				0.19	0.16	0.14
Wool (fine)		0.45				0.28	0.21	0.20
Cocoa			0.23				0.07	0.05
Arabica			0.71				0.16	0.13
Robusta			0.66				0.16	0.13
Tea			0.08				0.04	0.03
Bananas				0.08			0.04	0.04
Barley				0.30			0.19	0.18
Beef				0.03			0.03	0.04
Rapeseed oil				0.34			0.21	0.19
Fishmeal				0.13			0.11	0.09
Groundnuts				0.22			0.12	0.12
Lamb				0.06			0.06	0.06
Corn				0.30			0.17	0.15
Olive Oil				0.02			0.04	0.04
Oranges				0.07			0.05	0.05
Palm oil				0.38			0.26	0.23
Swine				-0.04			0.00	0.03
Poultry				0.03			0.00	0.00
Rice				0.21			0.10	0.09
Fish				0.03			0.07	0.07
Shrimp				0.00			0.03	0.03
Soybean Meal				0.23			0.13	0.12
Soybean Oil				0.35			0.21	0.19
Soybeans				0.29			0.17	0.15
Sugar				0.18			0.17	0.14
Sugar U.S.				0.05			0.05	0.04
Sunflower oil				0.27			0.10	0.09
Wheat				0.25			0.16	0.15
Coal					0.24			0.15
Dated Brent					0.57			0.23
Dubai Crude Oil					0.57			0.22
WTI					0.54			0.22

Table C.3 Out-of-Sample Fundamentals Forecasts

The Table reports the percentage decrease in the RMSE for different predictive models with different lag specifications with respect to the RMSE obtained using as a benchmark the Random Walk model (10% means a reduction of 10% with respect to the RW RMSE). The predictive content of the news on global growth extracted from commodity prices is tested estimating the following predictive regression $y_t = \gamma + \sum_{i=1}^b \phi_i y_{t-i} + \sum_{i=1}^b \beta_i x_{t-i+m} + \varepsilon_t$, where y_t is the log-difference in the OECD industrial production publications between time t and time $t - 1$, which will be officially released at time $t + 1$ and X_t is the real-time information extracted from commodity returns using the first principal component estimated using the information available up to time t . With $m = 1$ the predictive regression exploits the real-time information that can be extracted at time t to predict the industrial production growth that will be reported at $t + 1$. The sample is quarterly from 1980 to 2016. The RMSE is computed out-of-sample from 2000 to 2016. The model is re-estimated at the end of each quarter using a rolling sample of ten years. Each column reports the RMSE obtained using different portfolios of commodities to estimate the first principal component. The table also reports the performance of univariate AR models.

Portfolio	Metals	Agri	Beverage	Food	Energy	Industrial	Non Energy	All
AR Order: 1								
RW	-	-	-	-	-	-	-	-
AR	-3%	-3%	-3%	-3%	-3%	-3%	-3%	-3%
m = 0	0%	-5%	-4%	-5%	-5%	-2%	-4%	-4%
m = 1	17%	15%	-2%	1%	11%	18%	12%	18%
AR Order: 2								
RW	-	-	-	-	-	-	-	-
AR	-11%	-11%	-11%	-11%	-11%	-11%	-11%	-11%
m = 0	-8%	-10%	-12%	-3%	-7%	-3%	0%	-5%
m = 1	14%	9%	-10%	-7%	6%	17%	8%	14%
AR Order: 3								
RW	-	-	-	-	-	-	-	-
AR	-13%	-13%	-13%	-13%	-13%	-13%	-13%	-13%
m = 0	-15%	-11%	-16%	-3%	-7%	-5%	-3%	-8%
m = 1	13%	7%	-11%	-3%	9%	19%	11%	14%
AR Order: 4								
RW	-	-	-	-	-	-	-	-
AR	-16%	-16%	-16%	-16%	-16%	-16%	-16%	-16%
m = 0	-25%	-20%	-20%	-7%	-14%	-10%	-9%	-15%
m = 1	4%	2%	-15%	-5%	3%	12%	2%	6%

Positive numbers represent a percentage decrease in RMSE compared to the RW model.

m = 1 uses real time information from commodity prices, m = 0 previous quarter information

Fig. C.2 Filtered - Actual Fundamentals Correlation

The Figure reports the correlation between filtered and actual global growth fundamentals when information is extracted using different portfolios. Filtered fundamentals are given by the first principal component, which is computed on the correlation matrix of Year-on-Year returns on monthly commodity prices from January 1980 to November 2016. Actual fundamentals are the Year-on-Year log-differences on the monthly OECD total industrial production. Portfolios composition is described in the Appendix C.3.




Fig. C.3 1st Principal Component Weights

The Figure reports the average weight, by commodity group, in the first principal component. Principal components are computed on the correlation matrix of Year-on-Year returns on monthly commodity prices from January 1980 to November 2016. The portfolio contains 48 commodities.




Fig. C.4 Trading Strategy Performance

The Figure presents the simulated trading performance (cumulated log-returns) of long/short strategies based on the evolution of news about global economic activity. Specifically, the trading signal relies on the first principal component of metal commodities returns, which identifies global growth news. The first principal component is estimated using a rolling window of ten years of data at monthly frequency. The portfolio position is updated monthly based on the sign of the latest value of the first the principal component. At time t a position is long if the sign of the latest value of the first principal component (estimated using data from $t-120$ to t) is positive, and vice versa. The holding period is monthly. The trading strategy is simulated from January 2000 to October 2016. Performances of long-only strategies and strategies based on future statistical publication on OECD industrial production are reported as benchmark performances. The trading signal of the strategy based on future statistical publications generates a buy (sell) signal when the Year-on-Year growth is positive (negative).




Fig. C.5 Trading Strategy Statistical Significance

The Figure presents the simulated trading performance (cumulated log-returns) of the long/short strategies based on the evolution of news about global economic activity together with simulated confidence intervals. Specifically, the trading signal relies on the first principal component of metal commodities returns, which identifies global growth news. The first principal component is estimated using a rolling window of ten years of data at monthly frequency. The portfolio position is updated monthly based on the sign of the latest value of the first the principal component. At time t a position is long if the sign of the latest value of the first principal component (estimated using data from $t-120$ to t) is positive, and vice versa. The holding period is monthly. The trading strategy is simulated from January 2000 to October 2016. The confidence intervals are obtained by applying the same trading strategy (i.e. same signals) on 10000 simulated random walk processes that by construction contain no predictability. The random walk processes of each panel are calibrated with the same volatility of the corresponding traded index. The top and the bottom black dashed lines represent the 95th and 5th percentile respectively.




Fig. C.6 Trading Strategy and Lagged Prices

The Figure presents a robustness test for the simulated trading strategy which accounts for the possibility that the spot prices are only available with a three months delay. Specifically, in this case, the trading signal relies on a three months lagged first principal component of metal commodities returns, which identifies global growth news. The first principal component is estimated using a rolling window of ten years of data at monthly frequency. The portfolio position is updated monthly. The holding period is monthly. The trading strategy is simulated from January 2000 to October 2016. Performances of long-only strategies and long-short based on up-to-date prices are reported as benchmark performances.




Fig. C.7 Returns Response to News about Fundamentals

The Figure presents the correlation between monthly returns on the S&P Goldman Sachs Commodity Index and lagged values of the first principal component of metal commodity returns, which approximates the news about global growth (See Section 3.3.1 for the news identification.). The sample is monthly from January 1980 to November 2016. Confidence intervals at 95% probability are reported in red.




Fig. C.8 Delayed Information and Trading Performance

The Figure presents the annualised returns (in excess of the long-only performance) of the simulated long/short strategy on the S&P Goldman Sachs Commodity Index when trading signals are used with increasing time lag. The trading signals are based on the filtered news about global growth extracted from commodity prices using PCA. The first principal component is estimated using a rolling window of ten years of data at monthly frequency. The portfolio position is updated monthly based on the sign of the latest value of the first principal component. At time t a position is long if the sign of the value of the first principal component at time $t - \text{Lag}$ is positive, and vice versa. The holding period is monthly. The trading strategy is simulated from January 2000 to October 2016.




Fig. C.9 Simulated Returns Response to News

The Figures present the correlation between returns and lagged values of global growth news for different model parametrizations. Above each panel the share of traders types is reported together with the percentage of latecomers and feedback traders with different delays and look-back periods. Correlations and confidence intervals are estimated on 10000 simulated dynamics of the model presented in the Appendix C.1. Innovations are unit variance white-noise, with: $\alpha = 0.99$; $a = 1$; $\lambda = 0.5$; $d = 0.5$; $b = 1$. The time-series dimension of each process (after burn-in) is 500 observations.




Fig. C.10 Signal Extraction and Look-Back Period

The Figures show changes in average correlation between log-returns and news (innovations) of the signal process as a function of the look-back period used to compute log-returns. Correlations and confidence intervals are estimated on 10000 simulated dynamics of the demand and the price process defined by Equation (3.3). Innovations are unit variance white-noise. The time-series dimension of each process is (after burn-in) 500 observations.





Fig. C.11 Signal Extraction and Portfolio Size

The Figures show changes in average correlation between the first principal component and news (innovations) of the signal process as a function of the number of assets in the portfolio. In this case, the returns look-back period is one. Correlations and confidence intervals are estimated on 10000 simulated dynamics of the demand and the price process defined by Equation (3.3). Innovations are unit variance white-noise. The time-series dimension of each process (after burn-in) is 500 observations.



C.3 Spot Prices Data

Metals: **Aluminium**, 99.5% minimum purity, LME spot price, CIF UK ports, US\$ per metric ton; **Copper**, grade A cathode, LME spot price, CIF European ports, US\$ per metric ton; China import **Iron Ore** Fines 62% FE spot (CFR Tianjin port), US dollars per metric ton; **Lead**, 99.97% pure, LME spot price, CIF European Ports, US\$ per metric ton; **Nickel**, melting grade, LME spot price, CIF European ports, US\$ per metric ton; **Tin**, standard grade, LME spot price, US\$ per metric ton; **Uranium**, NUEXCO, Restricted Price, Nuexco exchange spot, US\$ per pound; **Zinc**, high grade 98% pure, US\$ per metric ton.

Agri: **Cotton**, Cotton Outlook 'A Index', Middling 1-3/32 inch staple, CIF Liverpool, US cents per pound; **Hides**, Heavy native steers, over 53 pounds, wholesale dealer's price, US, Chicago, fob Shipping Point, US cents per pound; **Soft Logs**, Average Export price from the U.S. for Douglas Fir, US\$ per cubic meter; **Hard Logs**, Best quality Malaysian meranti, import price Japan, US\$ per cubic meter; **Rubber**, Singapore Commodity Exchange, No. 3 Rubber Smoked Sheets, 1st contract, US cents per pound; **Hard Sawnwood**, Dark Red Meranti, select and better quality, C&F U.K port, US\$ per cubic meter; **Soft Sawnwood**, average export price of Douglas Fir, U.S. Price, US\$ per cubic meter; **Wool**, coarse, 23 micron, Australian Wool Exchange spot quote, US cents per kilogram; **Wool**, fine, 19 micron, Australian Wool Exchange spot quote, US cents per kilogram.

Beverage: **Cocoa Beans**, International Cocoa Organization cash price, CIF US and European ports, US\$ per metric ton; **Coffee**, Other Mild Arabicas, International Coffee Organization New York cash price, ex-dock New York, US cents per pound; **Coffee, Robusta**, International Coffee Organization New York cash price, ex-dock New York, US cents per pound; **Tea**, Mombasa, Kenya, Auction Price, US cents per kilogram, From July 1998, Kenya auctions, Best Pekoe Fannings. Prior, London auctions, c.i.f. U.K. warehouses.

Food: **Bananas**, Central American and Ecuador, FOB U.S. Ports, US\$ per metric ton; **Barley**, Canadian no.1 Western Barley, spot price, US\$ per metric ton; **Beef**, Australian and New Zealand 85% lean fores, CIF U.S. import price, US cents per pound; **Rapeseed oil**, crude, fob Rotterdam, US\$ per metric ton; **Fishmeal**, Peru Fish meal/pellets 65% protein, CIF, US\$ per metric ton; **Groundnuts** (peanuts), 40/50 (40 to 50 count per ounce), cif Argentina, US\$ per metric ton; **Lamb**, frozen carcass Smithfield London,

US cents per pound; **Maize** (corn), U.S. No.2 Yellow, FOB Gulf of Mexico, U.S. price, US\$ per metric ton; **Olive Oil**, extra virgin less than 1% free fatty acid, ex-tanker price U.K., US\$ per metric ton; **Oranges**, miscellaneous oranges CIF French import price, US\$ per metric ton; **Palm oil**, Malaysia Palm Oil Futures (first contract forward) 4-5 percent FFA, US\$ per metric ton; **Swine** (pork), 51-52% lean Hogs, U.S. price, US cents per pound; **Poultry** (chicken), Whole bird spot price, Ready-to-cook, whole, iced, Georgia docks, US cents per pound; **Rice**, 5 percent broken milled white rice, Thailand nominal price quote, US\$ per metric ton; **Fish** (salmon), Farm Bred Norwegian Salmon, export price, US\$ per kilogram; **Shrimp**, No.1 shell-on headless, 26-30 count per pound, Mexican origin, New York port, US\$ per kilogram; **Soybean Meal**, Chicago Soybean Meal Futures (first contract forward) Minimum 48 percent protein, US\$ per metric ton; **Soybean Oil**, Chicago Soybean Oil Futures (first contract forward) exchange approved grades, US\$ per metric ton; **Soybeans**, U.S. soybeans, Chicago Soybean futures contract (first contract forward) No. 2 yellow and par, US\$ per metric ton; **Sugar**, Free Market, Coffee Sugar and Cocoa Exchange (CSCE) contract no.11 nearest future position, US cents per pound; **Sugar**, U.S. import price, contract no.14 nearest futures position, US cents per pound (Footnote: No. 14 revised to No. 16); **Sunflower oil**, Sunflower Oil, US export price from Gulf of Mexico, US\$ per metric ton; **Wheat**, No.1 Hard Red Winter, ordinary protein, FOB Gulf of Mexico, US\$ per metric ton.

Energy: **Coal**, Australian thermal coal, 12,000- btu/pound, less than 1% sulphur, 14% ash, FOB Newcastle/Port Kembla, US\$ per metric ton; Crude Oil (petroleum), **Dated Brent**, light blend 38 API, fob U.K., US\$ per barrel; Oil; Dubai, medium, Fateh 32 API, fob **Dubai Crude Oil** (petroleum), Dubai Fateh Fateh 32 API, US\$ per barrel; Crude Oil (petroleum), **West Texas Intermediate** 40 API, Midland Texas, US\$ per barrel.

Industrial: Includes Metals and Agri commodities.

Non-Energy: Includes all commodities but Energy commodities defined above.

All: Includes: Metals; Agri; Beverage; Food; Energy.

C.4 Look-Back Period and Signal Extraction

Proposition 1. *The Signal-to-Noise of the price dynamics increases with the returns look-back period.*

Proof. Consider the price process defined in equation (3.3). We define log-returns with look-back period b as $r_{t:t-b} = p_t - p_{t-b}$. Plugging into this definition the price dynamics (3.3), log-returns with look-back period b can be written as

$$p_t - p_{t-b} = \alpha(\alpha^b - 1)y_{t-b} + \sum_{j=0}^b \alpha^j \varepsilon_{t-j} - \varepsilon_{t-b} - z_t + z_{t-b} \quad (\text{C.24})$$

Defining the signal as $S = \sum_{j=0}^b \alpha^j \varepsilon_{t-j} - \varepsilon_{t-b}$ and the noise as $N = 1 - z_t + z_{t-b}$, the Signal-to-Noise can be written as

$$\frac{S}{N} = \frac{\sum_{j=0}^b \alpha^j \varepsilon_{t-j} - \varepsilon_{t-b}}{1 - z_t + z_{t-b}} \quad (\text{C.25})$$

and for simplicity, with the assumption that the innovations ε_t and z_t have unit mean, we can show that³

$$\frac{E[S|B]}{E[N|B]} \geq \frac{E[S|b]}{E[N|b]}, \quad \text{with } B > b \quad (\text{C.26})$$

since

$$\frac{E[\sum_{j=0}^B \alpha^j \varepsilon_{t-j} - \varepsilon_{t-B}]}{E[1 - z_t + z_{t-B}]} = \frac{\sum_{j=0}^B \alpha^j - 1}{1} \geq \frac{\sum_{j=0}^b \alpha^j - 1}{1} = \frac{E[\sum_{j=0}^b \alpha^j \varepsilon_{t-j} - \varepsilon_{t-b}]}{E[1 - z_t + z_{t-b}]} \quad (\text{C.27})$$

□

³Assuming a positive mean in both income and supply shocks is equivalent to assume that demand and supply grow over time.

Chapter 4

Concluding Remarks

This Thesis has contributed to the study of expectations formation and their relationship with fundamentals and asset returns along different directions.

In the first Chapter, we have studied the links between expectations, fundamentals, and asset returns by using an unexplored database on commodity survey price forecasts. This unique dataset has offered us the opportunity to analyse survey-based expectations with less measurement error problems compared to previous studies.¹ More importantly, the data have provided us with the opportunity to explore how the forecasts predictive power and drivers change at different forecasting horizons and when the population surveyed consists of professional forecasters.²

The empirical findings have shown that, in commodity markets, survey-based expectations of returns are strongly correlated with past returns, but not with fundamentals. In fact, we have found that survey forecasts are largely explained by time-series momentum and value factors. Expectations have positive, but not significant correlation with future realised returns, which implies little predictive power.

We have rationalised these findings by setting up a simulation experiment, which has shown that such results are not incompatible also with the limiting case of perfect rational expectations. To the contrary, we have demonstrated that the predictive power of rational forecasts purely depends on the predictability of the fundamentals and the asset demand of other types of agents. Additionally, the experiment has shown that the correlation with past returns can also arise in rational expectations when rational agents take into account the

¹Most important advantages: clear forecast objective, quantitative nature, higher frequency.

²Greenwood and Shleifer [50] use survey data input by non-professional forecasters.

speculative activity of momentum (extrapolative) traders, whose asset demand depends on past returns.

In terms of contribution, this Chapter has enriched the strand of literature on asset pricing and survey forecasts by extending the studies of Koijsen et al. [70]; Greenwood and Shleifer [49] and Beber et al. [11] to the commodity markets and by adding the analysis on multiple forecasting horizons. Most importantly, our findings have shed some light on the puzzling evidence of contrarian predictive power of survey forecasts and positive correlation with past returns found by Greenwood and Shleifer [50], by showing that when a population of professional forecasters is interviewed, the contrarian predictive power disappears and the sign of the correlation with past returns changes. Our simulation experiment has also shown that both results can be obtained depending on which of the two populations, rational or extrapolative, is surveyed. Therefore, we have hypothesized that our and the previous empirical evidence is driven by the population being interviewed, with professional forecasters being close to rational expectations agents and the non-professional forecasters close to the extrapolative type.

Finally, our analysis has also suggested that survey-based expectations can have a crucial role to understand better the dynamics of trading flows and the drivers of the option implied volatility risk premium. We have left our preliminary results on these two extremely fascinating topics as starting points for avenues of future research.

In the second Chapter, our empirical analysis has shown that investors' expectations of future commodity spot prices can be approximated by a rational learning scheme in which expected future spot prices are revised in line with past prediction errors and changes in aggregate demand. In fact, we have shown that, although, with differences across commodities, our model-implied price expectations are broadly consistent with analysts' survey forecasts. As a second step, we have contributed to the study of the dynamics of risk-premia by exploiting this expectations formation mechanism to extract time-varying (ex-ante) risk premia from futures across different commodities and maturities. By using a dynamic linear regression in which we accommodate uncertainty in the estimated coefficients and their degree of time-variation, we have provided evidence that the dynamics of commodity risk premia is predominantly driven by market activity and the changing nature of market participants, as proxied by open interests, hedging pressure and time-series momentum. Furthermore, we have shown that our model of learning compares favourably to other commonly used specifications in forecasting future spot prices and generates expected payoffs that are consistently

linked to the actual, observable, returns on same-horizon futures contracts.

In the third Chapter of the thesis, we have exploited the peculiarities of commodity markets to show that fundamental news about global growth is reflected into prices, but not instantaneously. The news can, therefore, be filtered in real-time from commodity prices, but such news takes several months before being fully incorporated into prices, leading to returns predictability. Coherently with the theories of overreaction and underreaction to news, we have shown using simulated data that the results obtained can be explained by the existence of latecomers, who process information with a delay, and momentum traders. Our analysis has assumed therefore that not all investors have access to the same information and learn from past experience. We are confident that such assumptions, not only allow to give a simple explanation to the existence of asset pricing anomalies such as prices under and overreaction to news and the time-series momentum but also provide a more realistic representation of investors' behaviour. Asymmetric information, high costs of data analysis, unstable time-series relationships and limited data samples, often force investors to use simpler expectations formation rules that are far from being perfectly rational and should be taken into account more seriously in future research.

In conclusion, the three Chapters constituting this thesis have shown that simple departures from strict rationality are sufficient to explain a series of empirical evidences that are at odds with the predictions of standard rational expectations models. More specifically, our results have shown that the combination of frictions in the information flow, and the existence of agents using simpler expectations formation rules based on extrapolation, constitutes a promising avenue for novel theoretical models aiming at explaining better asset prices dynamics.

References

- [1] Acharya, V., Lochstoer, L. A., and Ramadorai, T. (2010). Does hedging affect commodity prices? the role of producer default risk. *Working Paper, London Business School*.
- [2] Acharya, V., Lochstoer, L. A., and Ramadorai, T. (2013). Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, 109(2):441–465.
- [3] Adam, K. and Marcet, A. (2011). Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory*, 146(3):1224–1252.
- [4] Alquist, R. and Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4):539–573.
- [5] Asness, C., Moskowitz, T., and Pedersen, L. (2013a). Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985.
- [6] Asness, C. S., Moskowitz, T. J., and Pedersen, L. H. (2013b). Value and Momentum Everywhere. *Journal of Finance*, 68(3):929–985.
- [7] Baker, S. and Routledge, B. (2011). The price of oil risk.
- [8] Bakshi, G., Gao, X., and Rossi, A. G. (2015a). Understanding the sources of risk underlying the cross-section of commodity returns. *Management Science*, Forthcoming.
- [9] Bakshi, G., Rossi, A., and Gao, X. (2015b). Understanding the sources of risk underlying the cross-section of commodity returns. *Working Paper*.
- [10] Basu, D. and Miffre, J. (2013). Capturing the risk premium of commodity futures: The role of hedging pressure. *Journal of Banking & Finance*, 37(7):2652–2664.
- [11] Beber, A., Breedon, F., and Buraschi, A. (2010). Differences in beliefs and currency risk premiums. *Journal of Financial Economics*, 98(3):415–438.
- [12] Beck, S. E. (1993). A rational expectations model of time varying risk premia in commodities futures markets: Theory and evidence. *International Economic Review*, pages 149–168.
- [13] Bernanke, B. S. (2004). Oil and the economy. *Speech presented at Darton College, Albany, Ga.*
- [14] Bernhardt, D. and Kutsoati, E. (1999). Can relative performance compensation explain analysts' forecasts of earnings? Technical report, Department of Economics, Tufts University.

- [15] Bessembinder, H. (1992). Systematic risk, hedging pressure, and risk premiums in futures markets. *Review of Financial Studies*, 5(4):637–667.
- [16] Bhardwaj, G., Gorton, G., and Rouwenhorst, K. (2015). Facts and fantasies about commodity futures ten years later. *NBER Working Paper*, (w21243).
- [17] Brennan, M. (1958). The supply of storage. *American Economic Review*, 48:50–72.
- [18] Buraschi, A. and Jiltsov, A. (2006). Model uncertainty and option markets with heterogeneous beliefs. *The Journal of Finance*, 61(6).
- [19] Campbell, J. and Thompson, S. (2008). Predicting the equity premium out of sample: Can anything beat the historical average?, forthcoming. *Review of Financial Studies*.
- [20] Carceles-Poveda, E. and Giannitsarou, C. (2008). Asset pricing with adaptive learning. *Review of Economic Dynamics*, 11(3):629–651.
- [21] Carter, C. and Kohn, R. (1994). On gibbs sampling for state-space models. *Biometrika*, (81):541–553.
- [22] Carter, C., Rausser, G., and Schmitz, A. (1983). Efficient asset portfolios and the theory of normal backwardation. *Journal of Political Economy*, 91(2):319–331.
- [23] Cheng, H., Kirilenko, A., and Xiong, W. (2015). Convective risk flows in commodity futures markets. *Review of Finance*, 19(5):1733–1781.
- [24] Cho, I.-K., Williams, N., and Sargent, T. J. (2002). Escaping nash inflation. *The Review of Economic Studies*, 69(1):1–40.
- [25] Clark, T. and West, K. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of econometrics*, 138(1):291–311.
- [26] Cootner, P. H. (1960). Returns to speculators: Telser versus keynes. *Journal of Political Economy*, 68(4):pp. 396–404.
- [27] Cutler, D., Poterba, J., and Summers, L. (1990a). Speculative dynamics and the role of feedback traders. *American Economic Review*, 80(2):63–68.
- [28] Cutler, D. M., Poterba, J. M., and Summers, L. H. (1990b). Speculative Dynamics and the Role of Feedback Traders. *American Economic Review*, 80(2):63–68.
- [29] De Roon, F., Nijman, T., and Veld, C. (2000). Hedging pressure effects in futures markets. *The Journal of Finance*, 55(3):1437–1456.
- [30] de Roon, F. A., van den Goorbergh, R. W. J., and Nijman, T. E. (2014). An Anatomy of Futures Returns: Risk Premiums and Trading Strategies. WO Research Memoranda (discontinued) 757, Netherlands Central Bank, Research Department.
- [31] Deaton, A. and Laroque, G. (1992). On the behaviour of commodity prices. *The Review of Economic Studies*, 59(1):1–23.
- [32] Deaton, A. and Laroque, G. (1996). Competitive storage and commodity price dynamics. *Journal of Political Economy*, 104(5):896–923.

- [33] Deaton, A. and Laroque, G. (2003a). A model of commodity prices after Sir Arthur Lewis. *Journal of Development Economics*, 71(2):289–310.
- [34] Deaton, A. and Laroque, G. (2003b). A model of commodity prices after sir arthur lewis. *Journal of Development Economics*, 71(2):289–310.
- [35] Dusak, K. (1973). Futures trading and investor returns: An investigation of commodity market risk premiums. *Journal of Political Economy*, 81(6):1387–1406.
- [36] Dvir, E. and Rogoff, K. (2010). Three epochs of oil. Technical report, National Bureau of Economic Research.
- [37] Erb, C. B. and Harvey, C. R. (2006). The strategic and tactical value of commodity futures. *Financial Analysts Journal*, 62(2):69–97.
- [38] Etula, E. (2013). Broker-dealer risk appetite and commodity returns. *Journal of Financial Econometrics*, 11(3):486–521.
- [39] Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.
- [40] Fama, E. and French, K. (1987a). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60:55–73.
- [41] Fama, E. and French, K. (1987b). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *The Journal of Business*, 60(1):55–73.
- [42] Fama, E. and French, K. (1988). Business cycles and the behavior of metals prices. *Journal of Finance*, 43:1075–1093.
- [43] Ferreira, M. and Santa-Clara, P. (2011). Forecasting stock market returns: The sum of the parts is more than the whole. *Journal of Financial Economics*, 100(3):514–537.
- [44] Frenkel, J. and Froot, K. (1987). Using survey data to test standard propositions regarding exchange rate expectation. *American Economic Review*, 77:133–153.
- [45] Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis*, 15:183–202.
- [46] Genizi, A. (1993). Decomposition of r^2 in multiple regression with correlated regressors. *Statistica Sinica*, pages 407–420.
- [47] Gorton, G., Hayashi, F., and Rouwenhorst, K. G. (2013a). The fundamentals of commodity futures returns. *Review of Finance*, 17:35–105.
- [48] Gorton, G. B., Hayashi, F., and Rouwenhorst, K. G. (2013b). The Fundamentals of Commodity Futures Returns. *Review of Finance*, 17(1):35–105.
- [49] Greenwood, R. and Shleifer, A. (2014a). Expectations of Returns and Expected Returns. *Review of Financial Studies*, 27(3):714–746.
- [50] Greenwood, R. and Shleifer, A. (2014b). Expectations of returns and expected returns. *Review of Financial Studies*, 27(3):714–746.

- [51] Hamilton, J. (1994). *Time Series Analysis*, volume 2. Princeton university press Princeton.
- [52] Hamilton, J. (2009). Understanding crude oil prices. *The Energy Journal*, pages 179–206.
- [53] Hamilton, J. and Wu, J. (2014). Risk premia in crude oil futures prices. *Journal of International Money and Finance*, 42:9–37.
- [54] Hansen, L. (2007). Beliefs, doubts and learning: Valuing macroeconomic risk. *The American Economic Review*, 97(2):1–30.
- [55] Hicks, C. (1939). *Value and Capital*. Oxford University Press, Cambridge.
- [56] Hong, H. and Kubik, J. (2003). Analyzing the analysts: Career concerns and biased earnings forecasts. *The Journal of Finance*, 58(1):313–351.
- [57] Hong, H., Kubik, J., and Solomon, A. (2000). Security analysts’ career concerns and herding of earnings forecasts. *The Rand Journal of Economics*, pages 121–144.
- [58] Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of finance*, 54(6):2143–2184.
- [59] Hong, H. and Yogo, M. (2012). What does futures market interest tell us about the macroeconomy and asset prices? *Journal of Financial Economics*, 105(3):473–490.
- [60] Hsieh, D. A. and Kulatilaka, N. (1982). Rational expectations and risk premia in forward markets: primary metals at the london metals exchange. *The Journal of Finance*, 37(5):1199–1207.
- [61] Kaldor, N. (1939a). Speculation and economic stability. *Review of Economic Studies*, 7:1–27.
- [62] Kaldor, N. (1939b). Speculation and economic stability. *The Review of Economic Studies*, 7(1):1–27.
- [63] Kang, W., Rouwenhorst, K. G., and Tang, K. (2014). The role of hedgers and speculators in liquidity provision to commodity futures markets. *Yale International Center for Finance Working Paper*, (14-24).
- [64] Kang, W., Rouwenhorst, K. G., and Tang, K. (2015). The Role of Hedgers and Speculators in Liquidity Provision to Commodity Futures Markets. *Yale International Center for Finance Working Paper*, 5(4):637–67.
- [65] Kawai, M. (1983). Price volatility of storable commodities under rational expectations in spot and futures markets. *International Economic Review*, pages 435–459.
- [66] Keynes, J. (1930). *A treatise on money*.
- [67] Kilian, L. and Hicks, B. (2013). Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *Journal of Forecasting*, 32(5):385–394.

- [68] Kilian, L. and Taylor, M. (2003). Why is it so difficult to beat the random walk forecast of exchange rates? *Journal of International Economics*, 60(1):85–107.
- [69] Koiien, R., Schmeling, M., and Vrugt, E. (2015a). Survey expectations of returns and asset pricing puzzles. *Working Paper*.
- [70] Koiien, R. S. J., Schmeling, M., and Vrugt, E. B. (2015b). Survey Expectations of Returns and Asset Pricing Puzzles. *Working Paper*.
- [71] Leduc, S., Moran, K., and Vigfusson, R. (2015). A decade of learning: The role of beliefs in oil futures markets during the 2000s. *Working Paper*.
- [72] Malmendier, U. and Nagel, S. (2015). Learning from inflation experiences. *The Quarterly Journal of Economics*, pages 02–23.
- [73] Marcet, A. and Sargent, T. (1989). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic theory*, 48(2):337–368.
- [74] Meese, R. A. and Rogoff, K. (1983). Empirical exchange rate models of the seventies : Do they fit out of sample? *Journal of International Economics*, 14(1-2).
- [75] Miffre, J. and Rallis, G. (2007). Momentum strategies in commodity futures markets. *Journal of Banking & Finance*, 31(6):1863–1886.
- [76] Moskowitz, T., Ooi, Y., and Pedersen, L. (2012a). Time series momentum. *Journal of Financial Economics*, 104(2):228–250.
- [77] Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H. (2012b). Time series momentum. *Journal of Financial Economics*, 104(2):228–250.
- [78] Moskowitz, T. J., Ooi, Y. H., and Pedersen, L. H. (2012c). Time series momentum. *Journal of Financial Economics*, 104(2):228–250.
- [79] Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pages 315–335.
- [80] Nerlove, M. (1958). Adaptive expectations and cobweb phenomena. *The Quarterly Journal of Economics*, pages 227–240.
- [81] Newey, W. and West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–08.
- [82] Pesaran, M. and Weale, M. (2006). Survey expectations. *Handbook of economic forecasting*, 1:715–776.
- [83] Pindyck, R. S. (2004). Volatility and commodity price dynamics. *Journal of Futures Markets*, 24(11):1029–1047.
- [84] Primiceri, G. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72:821–852.

- [85] Rapach, D., Strauss, J., and Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies*, 23(2):821–862.
- [86] Sargent, T., Williams, N., and Zha, T. (2004). Shocks and government beliefs: The rise and fall of american inflation. Technical report, National Bureau of Economic Research.
- [87] Sargent, T. J. (2002). *The Conquest of American Inflation*. Princeton University Press.
- [88] Sargent, T. J. and Williams, N. (2005). Impacts of priors on convergence and escapes from nash inflation. *Review of Economic Dynamics*, 8(2):360–391.
- [89] Singleton, K. (2014). Investor flows and the 2008 boom/bust in oil prices. *Management Science*, 60(2):300–318.
- [90] Sockin, M. and Xiong, W. (2015). Informational frictions and commodity markets. *The Journal of Finance*, 70(5):2063–2098.
- [91] Stock, J. H. and Watson, M. W. (1998). Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association*, 93(441):349–358.
- [92] Svensson, L. E. (2005). Oil prices and ecb monetary policy. *Committee on Economic and Monetary Affairs*, pages 1–4.
- [93] Szymanowska, M., De Roon, F., Nijman, T., and Van Den Goorbergh, R. (2014). An anatomy of commodity futures risk premia. *Journal of Finance*, 69(1):453–482.
- [94] Tang, K. and Xiong, W. (2012). Index investment and the financialization of commodities. *Financial Analysts Journal*, 68(5):54–74.
- [95] Turnovsky, S. (1983). The determination of spot and futures prices with storable commodities. *Econometrica*, 51(5).
- [96] West, M. and Harrison, J. (1997). *Bayesian forecasting and dynamics models*. Springer.
- [97] Williams, N. (2003). Adaptive learning and business cycles. *Manuscript, Princeton University*.
- [98] Working, H. (1949). The theory of the price of storage. *American Economic Review*, 39:1254–1262.