Marko Melolinna

What explains risk premia in crude oil futures?

SUOMEN PANKKI
FINLANDS BANK

EUROJÄRJESTELMÄ
EUROSYSTEMET

Marko Melolinna*

What explains risk premia in crude oil futures?

What explains risk premia in crude oil futures?

Marko Melolinna
Monetary Policy and Research Department

Abstract

This paper studies the existence of risk premia in crude oil futures prices with simple regression and Bayesian VAR models. It also studies the importance of three main risk premia models in explaining and forecasting the risk premia in practice. Whilst the existence of the premia and the validity of the models can be established at certain time points, it turns out that the choice of sample period has a considerable effect on the results. Hence, the risk premia are highly time-varying. The study also establishes a model, based on speculative positions in the futures markets, which has some predictive power for future oil spot prices.

Mikä selittää raakaöljyfutuurien riskipreemioita?

Suomen Pankin keskustelualoitteita 2/2011

Marko Melolinna
Rahapolitiikka- ja tutkimusosasto

Tiivistelmä

Tässä työssä tutkitaan öljyfutuurien riskipreemioiden olemassaoloa yksinkertaisilla regressioilla ja bayesilaisilla VAR-malleilla. Samalla tutkitaan kolmen keskeisen riskipreemiomallin kykyä selittää ja ennustaa riskipreemioita käytännön sovelluksessa. Riskipreemioiden olemassaolo ja mallien validiteetti voidaan vahvistaa tiettyinä ajankohtina, mutta tutkimuksessa ilmenee, että nämä seikat riippuvat hyvin selvästi valitun otoksen ajanjaksosta. Näin ollen riskipreemiot vaihtelevat vahvasti ajassa. Tutkimuksessa myös löydetään futuurimarkkinoiden spekulatiivisiin positioihin perustuva malli, jolla on kykyä ennustaa tulevia öljyn hintoja.

Avainsanat: ennustaminen, öljyfutuurit, riskipreemiot, bayesilaiset VAR-mallit

JEL-luokittelu: C01, C32, C53

4

Contents

1 Introduction

The relationship between commodity futures prices and spot prices has been the subject of a vast amount of literature for a long time. The relevance of this research effort has been proved during the past decade, especially in the oil market, as oil price volatility has been one of the major forces behind global macroeconomic developments. The literature on risk premia - that is, the difference between futures price and expected future spot price (or also defined as the biasedness of futures prices as forecasts of future spot prices) - dates back to The Theory of Normal Backwardation, introduced by Keynes (1930), which asserts that in order to induce storage, futures price and expected future spot prices have to rise over time to compensate storage holders for the costs of storage. Kaldor (1939), on the other hand, introduced The Theory of Storage and the concept of convenience yield. The Theory of Storage establishes a link between contemporaneos spot and futures prices, and rectifies The Theory of Normal Backwardation by introducing the convenience yield. Keynes' original theory has also been used as a foundation to a so called "Market Pressure" Theory, which relates risk premia to the actions of hedgers and speculators. More recently, Pindyck (2001) proposed a definition for the risk premium based on a link to general market risk.

In recent years, a number of studies have found evidence for the existence of risk premia in crude oil futures prices, and many have attempted to exploit this to predict future spot prices. Evidence of the relationship between futures prices and risk premia is provided, for example, by Gorton and Rouwenhorst (2006), who find that a portfolio (including crude oil) of commodity futures risk premia have been equal in size to the historical risk premium on stocks. Moosa and Al-Lougani (1994) find evidence of a varying risk premium crude oil futures that can be modelled by a GARCH process, while Considine and Larson (2001) suggest that crude oil assets contain a risk premium that rises with higher price volatility. Alquist and Kilian (2007) show that futures-based forecasts are biased in the crude oil market.

Another strand of literature has related the variation in risk premia to macro-economic factors. For example, Coimbra and Esteves (2004) find evidence of a correlation between oil futures forecast errors and market expectation errors on world economic activity, and Pagano and Pisani (2009) show that the forecast error on oil futures could have been partly explained by US business-cycle indicators.

In contrast, there are also studies that have failed to find evidence of a risk premium in crude oil futures markets. For example, Chinn et al. (2005) find futures prices to be unbiased predictors of future spot prices over the period of 1999 to 2004, whilst Chernenko et al. (2004) find mixed evidence on the unbiasedness using the same methodology but a longer sample. A number of earlier studies, like Peroni and McNown (1998) and Kellard et al. (1999), also find little evidence of

a risk premium, especially for futures of short maturities.

As the above discussion indicates, risk premia in crude oil futures have been studied with various methods and theories. However, the results of these efforts are somewhat mixed. The outcome of these studies seems to be dependent on, for example, the futures maturity examined and the methodology used. Furthermore, the risk premia models used have not always been grounded in the seminal theories introduced by Keynes, Kaldor and Pindyck.

The aim of the current study is to shed light on factors explaining risk premia in crude oil futures, and to use this model to predict futures prices at different horizons. The first task is to establish the existence of the risk premia. Then, simple regression models are used to study the validity of the three main risk premia models above (The Theory of Storage, The "Market Pressure" Theory and Pindyck's Theory), and finally, Bayesian vector autoregression models are used for an out-of-sample prediction experiment.

This study finds - in line with most previous literature - that there exist negative risk premia in crude oil futures. These premia also tend to be time-varying. However, the statistical significance of such premia is highly dependent on the sample period studied. Furthermore, the factors that are relevant for explaining these risk premia change over time, but overall, the results support the validity of each of the main three models at different points in time. Some of the models introduced in the current paper can be used to predict future oil prices in a fairly satisfactory manner, but, as ever, consistently being able to outperform simple random walk models remains challenging.

The structure of this paper is as follows. The next section introduces the theoretical foundations of the risk premia models as well as describes the econometric methodology used in the current study. Section 3 presents the empirical results of the study, and Section 4 concludes.

2 Methodology

2.1 Theoretical foundations

To fully appreciate the difficulty of modelling and forecasting oil prices, it is useful to consider Figure 1, which depicts a long time series of the log of oil price, as well as the change of the log oil price. In the crude oil markets - like in any segment of financial markets - prices are constantly adjusted to take into account market participants' views of demand and supply shocks as well as other factors. Hence, it is extremely difficult to try and model the oil price purely as a function of any slow-moving macroeconomic variable or a traditional macroeconomic model. This problem is exacerbated by the constant changes in the volatility of the oil price,

Figure 1: Oil price and oil price change (West Texas Intermediate Grade)

illustrated by the strong time-variation in the change of the price. Specifically, it seems that the price has been, on average, more volatile from around the year 2000 compared to the previous two decades.

Based on the apparent difficulty in trying to model oil prices and forecasts of them, it is necessary to ground explanations of oil price risk premia on theories which allow for simplified, time-varying relationships between key variables in the markets to be modelled. The literature on risk premiums in commodities prices dates back to The Theory of Normal Backwardation, introduced by Keynes (1930), which compares futures prices to *expected* future spot prices. This theory is based on a definition of the basis, which is defined as the difference between the current fufures price maturing at time T and the current spot price. The theory divides the basis to the difference between the spot price expected to prevail at time T and the current spot price, minus a risk premium. Thus, the risk premium is defined as the difference between the expected future spot price at T and the current futures price for maturity T:

$$F_{t,T} - S_t = [E_t(S_T) - S_t] - \pi_{t,T} \tag{1}$$

where $F_{t,T}$ is the futures price at t for a futures contract expiring at T, S_t is the spot price at time t, and where the basis is divided into the difference between the expected spot price at T ($E_t(S_T)$) and the current spot price, and a risk premium ($-\pi_{t,T}$). Rearranging equation (1), the risk premium can be defined:

$$-\pi_{t,T} = F_{t,T} - E_t(S_T) \tag{2}$$

which is the definition used in this paper as well. When the risk premium can be studied after time T, I will talk of the *ex post* risk premium. In its earliest form, The Theory of Normal Backwardation asserts that in order to induce storage, futures price and expected future spot prices have to rise over time to compensate storage holders for the costs of storage. This cost of carry principle, however, had difficulties explaining downward sloping futures curves.

Kaldor (1939) introduced The Theory of Storage and the concept of convenience yield. The Theory of Storage establishes a link between contemporaneous spot and futures prices, and rectifies The Theory of Normal Backwardation by introducing the concept of convenience yield. The theory defines the basis as the cost of carry (the interest foregone to borrow to buy the commodity at the spot price plus the marginal storage cost), minus a convenience yield. The existence of the convenience yield allows for the basis to be negative, i.e. the futures price to be below the current spot price. Hence, according to this theory:

$$F_{t,T} - S_t = S_t r_t + w_t - c_t \tag{3}$$

where r_t is the risk-free interest rate at time t for the duration of the futures contract, w_t is marginal storage cost and c_t is the so-called "convenience yield". Intuitively, equation (3) is a no-arbitrage condition, which states that the return from buying oil in the spot market and selling it in the futures market (the left-hand side) must equal the interest return foregone when buying the oil, the storage cost of the oil, minus a "convenience yield" measuring the tightness of the physical spot oil market. The higher the convenience yield, the tighter the spot market, which allows the basis to be negative.

Plugging the right-hand side of equation (1) into the left-hand side of equation (3), one can derive another definition for the risk premium:

$$-\pi_{t,T} = (1 + r_t)S_t - E_t(S_T) + w_t - c_t \tag{4}$$

which can be used to define a relationship between the change in inventories and the risk premium. The change in inventories can be thought to have a negative relationship with the convenience yield, because the higher the inventories, the

10

lower the tightness of the spot market is (i.e. $\partial c_t / \partial I_t < 0$, where I_t is the level of inventories at time t). Thus, $\partial(-\pi_{t,T}) / \partial I_t > 0$. Furthemore, the risk premium can also be related to the interest rate variable r_t so that $\partial(-\pi_{t,T}) / \partial r_t > 0$.

Keynes' theory has also given rise to another theory, developed by Cootner (1960) and Deaves and Krinsky (1995) According to this so called "Market Pressure" Theory, the risk premium is determined by the actions of hedgers and speculators in the marketplace. Hedgers are willing to pay for the reduction in risk that the futures provide, whereas speculators demand compensation for the risk they are taking. Hence, if hedgers hold (net) short positions and speculators (net) long positions in the market, the price of futures is lower than the expected future spot price (and thus the risk premium $-\pi_{t,T} < 0$), since the speculators net long in the markets require compensation (i.e. a lower futures price) to enter the market. Conversely, if hedgers hold (net) long positions and speculators (net) short positions, the price of futures is higher than the expected future spot price. Therefore, the risk premium and the net long speculative positions (*spec*) have a negative relationship $(\partial(-\pi_{t,T}) / \partial(spec) < 0)$.

As Fama and French ((1987) and (1988)) point out, the two theories described above (The "Market Pressure" Theory and The Theory of Storage) are not necessarily mutually exclusive; the existence of basis can be defined in either way. The authors find evidence for the Theory of Storage for several commodities (they do not study crude oil), and they also find evidence for certain metal commodities of the implication of the theory that the convenience yield falls at decreasing rate when inventory increases. However, the evidence for the existence of a risk premium in the commodities the authors study is mixed.

The two theories have been the subject of more recent research as well. For example, Gorton et al. (2007) build a theoretical model to allow for risk-averse agents and a hedging motive of the commodity producers. Given that futures provide insurance against price volatility, which is negatively correlated with the level of inventories, the level of inventories should be negatively related to the required risk premium. In fact, the authors find evidence of this being the case for a wide range of commodities, including crude oil. On the other hand, they find no support for the "Market Pressure" Theory, since there is no evidence that the positions of futures traders help predict risk premiums on commodity futures.

Pindyck (2001) proposes a definition for the risk premium based on the Theory of Storage, as well as an equilibrium that connects the crude oil futures markets with the cash and the storage market. In particular, the author shows that the risk premium is related to the risk associated with holding the commodity. If the spot price co-varies positively with the overall economy (or in terms of the Capital Asset Pricing Model, the 'beta' of the commodity is positive), the holder of the commodity will expect to be rewarded for the risk by the spot price (on average)

rising above the current futures price over the holding period. Thus, this implies that the risk premium, as I define it, is, on average, negative. The author also finds evidence of this being the case.

Mathematically, according to Pindyck (2001), the return for the risky investment in spot crude oil must equal the current spot price with a discount factor $(\rho_T S_t)$:

$$E_t(S_T) - S_T + c_t - w_t = \rho_T S_t \tag{5}$$

Solving equation (3) for the term $c_t - w_t$ and plugging this into equation (5) yields

$$-(r_t - \rho_T)S_t = E_t(S_T) - F_{t,T} \equiv \pi_{t,T} \tag{6}$$

where the term $-(r_t - \rho_T)$ is defined as the "beta" of the investment.

2.2 Econometric modelling of the risk premia

As the above discussion indicates, the evidence on the existence and justification for risk premia on commodity futures is somewhat mixed. The outcome of the studies seems to be dependent on the commodity being studied, the sample period examined and the methodology used. There does not seem to be an agreement in the literature on whether risk premia exist and if they exist, what explains their existence.

The aim of this paper is to shed light on a number of issues related to risk premia in crude oil futures. The first aim is to study whether risk premia that are statistically significantly different from zero over time can be found. If the existence of such risk premia can be established, exogenous explanatory variables will be attached to the model that can potentially help shed light on the justification for the existence. This is carried out with simple OLS regressions. Finally, the ability of the different exogenous variables in helping predict future spot oil prices is tested with unrestricted VAR models.

The selection of the relevant explanatory variables requires some justification. As regards The Theory of Storage, based on the equations above, it is intuitively clear that the level of inventories as well as market interest rates can help explain the relevance of the theory. For the "Market Pressure" Theory, the level of net speculative positions offers a natural explanatory variable. For Pindyck's Theory, no such obvious explanatory variable exists. However, as a proxy, the correlation between the crude oil spot price and a measure of stock market performance will provide a variable that by definition, is proportional to the beta of the oil price. A variable of this type is used as an explanatory variables in the models.

2.3 Data

To calculate the risk premia, information on both spot prices as well as futures prices at different horizons is needed. Furthermore, for explanatory variables, I need data on crude oil stocks, speculative positions, interest rates and a relevant stock market index. All this data is available from the Bloomberg data service.

The risk premia are calculated for the main benchmark oil type, which is the West Texas Intermediate (WTI), as futures data (traded in NYMEX) as well as stocks and futures positions data for this is readily available. The horizon studied is restricted by the availability of the futures data to be 1989:1 to 2008:12, but this is considered to be long enough for reliable inference. The data frequency studied is monthly. The aim is to study short-term dynamics, so the analysis is restricted to 1-month, 3-month, 6-month and 12-month risk premia[1].

As regards the explanatory variables, US crude oil stocks (published weekly by The U.S. Energy Information Administration) are deemed to be a relevant variable, since I am studying a benchmark priced in the US. Hence, no information on oil stocks in other countries is taken into account (and would not be readily available either). For the speculative positions, US Commodity Futures Trading Commission (CFTC) data published on Bloomberg for the net long positions of speculators (non-commercial traders) is used. It is worth noting that while by no means a complete measure of speculative activity in the market (due, for example, to the fact that it excludes OTC derivatives not traded as standardised exchange), it is the only data available and it is widely used in the literature. To measure the "beta" for Pindyck's model, I use the rolling 12-month correlation coefficient between the spot oil price and the S&P 500 stock index, which is the main U.S. stock index. The market interest rates used are interbank US dollar interest rates of the relevant maturities (1,3,6 and 12 months).

3 Results

3.1 Forecast biasedness

Before modelling the risk premia in crude oil futures, it is obviously necessary to establish the existence of such premia. The ex post risk premia for the for maturities are illustrated in Figure 2. The risk premia seem to have gyrated without any clear trend or pattern, and there are especially strong movements

---

[1]In principle, it might be interesting to look at other maturities as well. However, the choice of maturities is dictated by the availability of market interest rate data for the entire sample period. While other interest rate maturities could of course be constructed using, for example, term structure models, this is not considered vital for the current study as relevant results can be obtained with the four maturities used.

Figure 2: Ex-post risk premia in crude oil futures
Note: CL1, CL3, CL6 and CL12 refer to 1-, 3-, 6- and 12-month risk premia.

14

towards the end of the sample. This is due to the rapid changes seen in the oil price, when prices increased during the global economic expansion and then dropped sharply when the financial market turbulence set in during 2007 and 2008.

More formal analysis of the risk premia proceeds by estimating the average forecast error over the sample period, and then testing whether this error is statistically significantly different from zero. Hence, to test this bias, I run the following regressions for different futures horizons T:

$$-\pi_{t.T} = \alpha_T + \epsilon_{t,T} \tag{7}$$

where α_T is a constant measuring the ex post realised forecast error and $\epsilon_{t,T}$ is an error term.

The dependent variable is computed by using daily futures quotations at different horizons during a one-week window around the middle of the month. This is done to avoid the effect of possible daily outliers, and conforms to choices made in previous literature (see Pagano and Pisani 2009). Furthemore, as the overlap present in the futures contracts can be expected to cause autocorrelation in the error term, the standard errors are computed using Newey-West autocorrelation- and heteroscedasticity-consistent (HAC) standard errors. The lag truncation parametre is chosen to be $2(T-1)$.

The results of the biasedness regressions are reported in Table 1. In line with most recent literature, like Pagano and Pisani (2009), the regressions suggest that the forecast error (and hence the risk premium) is negative for all horizons. This also confirms the observation made by Pindyck about the negative sign of the risk premium. However, unlike most authors, I do not find this biasedness to be statistically significant. The reason for this discrepancy is the sample period. The steep decrease in crude oil prices that resulted from the onset of the financial market crisis in 2007-2008 has had a dramatic effect on the biasedness of crude oil futures-based forecasts. This is illustrated in figures 3 and 4, which depict the rolling constant coefficients in the regressions (i.e., the sample period starts at 1989:1 and the end point is rolled forward month-by-month), as well as the corresponding P-value of the constant coefficients in these rolling regressions. Both the absolute values and their statistical significance gyrate strongly over time, and especially during the year 2008. Hence, the risk premia are highly time-varying[2]. This underscores an important point often overlooked in the literature; the conclusion on whether the futures prices are deemed to be biased predictors or not is highly dependent on the sample period studied.

---

[2]These type of results are commonplace in financial markets literature. See, for example, Fama (1984) and Frankel (1988) for seminal work on time-varying risk premia in the foreign exchange markets.

In particular, the results here suggest that it may be worthwhile to also look at the sample excluding the year 2008 when modelling the risk premia. However, a priori there is no justification for excluding the period of the financial market turbulence from the analysis altogether. A case can be made for the rapid drop in the oil price being "the flip side of the coin"; it may have merely corrected the potentially excessive oil price increases seen in the preceding years. Therefore, while some analysis below is also carried out for a sample that excludes the years 2007 and 2008, the overall analysis takes account of all available data in the full sample.

T	constant
1	-0.18 (-0.5)
3	-0.59 (-0.5)
6	-1.41 (-0.7)
12	-1.46 (-0.4)

Note: Sample 1989-2008. Estimation by OLS, t-statistic (HAC) in parantheses.

Table 1: Average futures-based forecast errors

3.2   Risk premia models

The previous section suggests that statistically significant risk premia have indeed existed in crude oil futures markets for most time periods.To ground the analysis on theoretical aspects, and to shed light on the importance of the different variables in explaining the risk premia, simple OLS regressions of the following form were carried out:

$$-\pi_{t,T} = \alpha + \beta x_t + \epsilon_{t,T} \tag{8}$$

where α is a constant, x_t is alternatively the level of crude oil stocks (in thousands of barrels, stockl), the market interest rate (r), the correlation between share prices and crude oil prices (corr) or net speculative positions (in thousands of contracts, spec) with the coefficient β, and $\epsilon_{t,T}$ is an error term Thus, in these regressions, the risk premia at different horizons were explained by variables which could be expected to have explanatory power based on the theories discussed above. The idea is to gauge the importance of the individual variables in explaining the risk premia. Hence, I will restrict the analysis to univariate one-equation models at this stage.

The results of these regressions are reported in Table 2. Finding statistically significant explanatory variables for the risk premia is challenging, especially for

Figure 3: Constant coefficients in rolling regressions
Note: clx-coeff refers to the coefficient for the x-month risk premium.

Figure 4: Constant coefficient P-values in rolling regressions
Note: P-clx refers to the P-value for the x-month risk premium.

the shorter maturities. Nevertheless, some inference is possible, and the results that are statistically significant, also bear the correct signs in accordance with the theories presented above. The results suggest that, when one excludes the years 2007 and 2008, the stock levels are statistically very significant explanatory variables for the 6- and 12-month risk premia. This finding is in line with Pagano and Pisani (2009). Furthermore, the fact that the market interest rate variable is likewise statistically significant for the longer maturities lends strong support for The Theory of Storage in the shorter sample. There is also some weaker evidence of the correlation variable having explanatory power in the full sample in the 6-month contract, and the speculative positions having statistical significance in the short sample in the 12-month contract.

		Explanatory variable							
		stockl		corr		spec		r	
Dep. variable		full	short	full	short	full	short	full	short
cl1	coefficient	0.002	0.013	-0.379	-0.461	-0.015	0.006	0.010	0.078
	R^2	0.00	0.01	0.00	0.01	0.01	0.00	0.00	0.00
cl3	coefficient	-0.007	0.028	-2.124	-0.968	-0.051	-0.004	-0.030	0.232
	R^2	0.00	0.02	0.02	0.00	0.02	0.00	0.00	0.01
cl6	coefficient	0.024	0.084 **	-5.067 *	-1.520	-0.059	-0.025	-0.047	0.642 *
	R^2	0.00	0.10	0.04	0.03	0.01	0.01	0.00	0.05
cl12	coefficient	-0.058	0.122 ***	-2.933	-1.878	-0.046	-0.088 **	-0.010	1.555 ***
	R^2	0.00	0.11	0.00	0.01	0.00	0.07	0.00	0.14
Note: Full sample 1989:1-2008:12, short sample 1989:1-2007:1. Estimation by OLS,									
statistical significance (HAC standard errors) in indicated by * 10 %, ** 5 %, *** 1 %.									

Table 2: Results of risk premia regressions

Once again, to take into account the time variation in the risk premia, it is interesting to look at how the significance changes when the regressions are rolled forward in time. The t-statistics[3] for the different explanatory variables are presented in Appendix A. The graphs indicate large changes in significance both over time and over different contract horizons. There seem to have been different time periods during which the different models to explain the risk premia have been more valid. Pindyck's model, as proxied by the corr variable, tended to have more statistical significance when measured in the early 2000's The spec variable has also been a relevant variable, especially for the longer contract maturitites, suggesting the "Market Pressure" Theory is also a potential explanation. The Theory of Storage seems to have been most relevant when studied between 2004 and 2007.

Overall, the simple regressions performed here offer no conclusive evidence on preferring one model over another. This is not surprising, given the irregular and

---

[3]The critical t-values are 1.66 (10 %), 1.98 (5 %) and 2.62 (1 %).

time-varying nature of the risk premia time series. However, it is interesting to notice that the variables grounded on simple theoretical models do indeed carry significance at different points in time when explaining the risk premia. This is especially true for the longer maturities. The movements in the risk premia ignited by the finacial markets turbulance seems to have blurred almost all the links between the risk premia and the explanatory variables. Hence, at this point in time, it is difficult to claim any of the fundamental models as particularly relevant in explaining risk premia in crude oil futures.

3.3   Forecasting models

As the risk premia models above indicate, single-equation univariate models fail to capture much of the movement in risk premia, so these models are unlikely to be very useful for trying to forecast the risk premia. Also, the risk premia time series exhibit a large amount of built-in autoregression, which supports the use of more dynamic models. For this, vector autoregressive (VAR) models offer a simple, flexible tool widely used for forecasting in the literature, and these kinds of models are also used in the current study. In particular, an unrestricted VAR(p) model of the following form is used:

$$X_t = C + \sum_{l=1}^{p} A_l X_{t-l} + u_t \qquad (9)$$

where X_t is an $Nx1$ vector of N variables, C is an $Nx1$ vector of constants, A_l is an NxN coefficient matrix, u_t is an error term and p is the number of lags. One of the variables in X_t will always be the risk premium, whilst the other variables vary in an effort to find the best forecasting model.

The strategy here is to construct a number of VAR models with the different explanatory variables and test their forecasting ability. For all the four different forecasting horizons, a "general-to-specific" approach to the experiment is adopted by first including in the models as many variables as possible, and then gradually dropping the variables that do not improve the forecasting performance. The model with the best forecasting ability is then used in a Bayesian VAR (BVAR) estimation framework. BVAR models have consistently been proved to perform better than pure VAR models in forecasting experiments (for early evidence of this, see Litterman (1986)). A grid search for two of the BVAR hyperparametres (tightness parameter gamma as well as weight parameter w) is then carried out to maximise the forecasting performance (for details of the structure of BVAR models, see Appendix B).

To evaluate the forecasting ability of the models, an out-of-sample rolling simulation to forecast the risk premia at different horizons was carried out. All the

models were first estimated for the time period January 1989 to December 1994, and the out-of-sample forecasts were then obtained for January 1995. This division of the sample period is believed to strike a good balance between having enough observations for the initial model estimation on the one hand and having enough observations to calculate a reliable forecast error on the other. The forecasts were then rolled forward each month to take into account the new information available, e.g., in January 1995, the model is estimated on the sample January 1989 to January 1995, and the forecast is made for February 1996, and so on until the end of the sample[4].

One important practical point to consider whilst carrying out the rolling forecasting simulation is the data set available for forecasting each month. Due to the way the ex-post risk premium is calculated, the value for a particular month will not be known until the time when the futures contract for the particular horizon has expired, and the ex post expected spot price can be recovered (i.e., the term $E_t(S_T)$ in equation (2) is known). Hence, when calculating the forecast for the risk premium at horizon i at time t, only the risk premium up to month t-i is known. Obviously, for the other endogenous variables in the model, the values will be known up to time t (see Table 3). Thus, the data set used at each point in time is unbalanced (this is known as the so-called "ragged edge" problem). This is taken into account in the experiment carried out in this study by making a conditional forecast for the risk premium based on the other endogenous variables in the model from time point t-i+1 up to time t. The forecast is then made for the risk premium at time t+i to recover the implied spot price forecast i months forward.

	Time	Data			
		CI for mat i	Other endogenous variables		
	t-i	x	x	x	x
	.		x	x	x
	.		x	x	x
	t-1		x	x	x
	t		**x**	**x**	**x**
	t+1				
	.				
	.				
	t+i				
	Note: CI for mat i indicates the risk premium at horizon i.				

Table 3: An illustration of an unbalanced data set in the forecast experiment

[4]WinRats econometric package was used for the estimation. The code is available from the author on request.

To analyse the forecasting performance in the experiment, I need to define a loss function for the forecast errors. As is common in the literature, the loss function considered and evaluation of the different forecast models is based on the Theil U statistic:

$$U_t = \frac{\sqrt{(\sum_{i=1}^{N_t} e_{it}^2)/N_t}}{\sqrt{(\sum_{i=1}^{N_t} (y_{it} - y_{i0})^2)/N_t}} \tag{10}$$

where N_t is the sample size at time i, e_{it}^2 is the forecast error for horizon t at time i, and $(y_{it} - y_{i0})^2$ is the forecast error for horizon t at time i of a naive random walk forecast, where the dependent variable is assumed to stay constant. In other words, Theil U is the root mean square error of the forecast, normalised by the root mean square error of a random walk forecast. Theil U statistic that gets a value lower then 1 indicates a forecasting model that outperforms a random walk model.

The results of the forecasting experiment are presented in Table 4. Apart from the BVAR models detailed above, two other forecasting tools are used for comparison; the pure futures forecast (Futures) as well as the single-equation model based on capacity utilisation in the U.S. introduced by Pagano and Pisani (2009) (Caputl). Several points are worth highlighting. First, as has been pointed out by a number of previous studies, the forecasting ability of a pure futures forecast is not statistically significantly different from a pure random-walk forecast. Second, the forecasting ability of neither Caputl models nor BVAR models is statistically significantly[5] better than random walk at the shorter horizons, even though the Theil U numbers are clearly below one at the one-month horizon and in the case of the BVAR model, also at the 3-month horizon. However, the BVAR model clearly outperforms the random walk model at the 12-month horizon, as does the Caputl model. The BVAR model has a lower Theil U statistic than the Caputl model in all the cases considered, although the difference is not statistically significant[6].

It is worth noting that the results presented in Table 4 for the Caputl model differ from those reported in Pagano and Pisani (2009). Based on replications of their methodology, it is not clear that the "ragged edge" problem detailed above has been considered in their paper. When carrying out the forecasting experiment, it seems that Pagano and Pisani (2009) assume that the current period risk premia are known when forecasting the next period. In particular, they use

---

[5]I use the traditional Diebold and Mariano (1995) test to measure the statistical significance.

[6]The lag lengths in the BVAR models are chosen to maximise the forecasting performance, however, the Theil U values are relatively robust to changes in the lag length.

the following regression for prediction of the Caputl model (see Pagano and Pisani (2009), methodology 4., page 16):

$$\widehat{p}_{t+n} = f_t^{(n)} - \widehat{\alpha}^{(n)} - \widehat{\beta}^{(n)} UCap_{t-1} \qquad (11)$$

where $\widehat{p}_{t+n}$ is the model forecast at time t for the spot oil price at horizon n, f_t^n is the futures price of maturity n at time t, $\widehat{\alpha}^{(n)}$ is the constant coefficient in a regression explaining the risk premium (which is equal to $f_t^{(n)} - p_{t+n}$), $\widehat{\beta}^{(n)}$ is the coefficient of capacity utilisation in the same regression, and $UCap$ is the degree of capacity utilisation in US manufacturing using real-time data.

The problem with this model is that the regression coefficents are not known at time t, because the regression explaining the risk premium cannot be carried out before time $t + n$. Hence, the regression coefficients α and β need to be estimated conditional on information available at time t, which results in the following model:

$$\widehat{p}_{t+n} = f_t^{(n)} - \widehat{\alpha}^{(n|t-n)} - \beta^{(n|t-n)} UCap_{t-1} \qquad (12)$$

where the regression coefficients take into account the fact that at time t, only the risk premium at time $t - n$ is known. This problem is taken into account in Table 4, and the Caputl model forecast errors are based on equation (12). As it turns out, the difference between the models in equations (11) and (12) is highly significant.

	T			
Model	1	3	6	12
Futures	1.00	1.02	1.02	0.98
Caputl	0.94	0.99	1.03	0.92 **
BVAR	0.93	0.96	1.01	0.86 **
variables	*spec*	*spec*	*spec*	*spec, r*
lags	*1*	*3*	*14*	*10*
hyperparametres	*0.22, 0.01*	*0.28, 0.07*	*1, 0.22*	*0.4, 0.19*

Note: Forecast simulation for period 1995:1-2008:12.
Statistical significance (HAC standard errors) in indicated by * 10 %, ** 5 %, *** 1 %.
Hyperparametres list the tightness (gamma) and weight numbers of the best performing model.

Table 4: Theil U statistics for different forecasting models

The best performing BVAR models are quite small, including only one or two explanatory variables in addition to the risk premium. In fact, the best performing BVAR models for all horizons include the net speculative long positions (spec), and only the 12-month model includes another variable (the interest rate). This is an important finding and suggests that, when using theory-based models to forecast the risk premia, the speculative positions are much more informative than,

for example, changes in stocks (note that this result differs from Gorton et al. (2007)). The interest rate variable is also an important factor at the 12-month horizon, which perhaps suggests it adds important information about the state of the business cycle.

The results of the grid search for the hyperparametres are presented in Appendix C for the four different horizons considered. These graphs prove that the forecasting performance is not particularly sensitive to the hyperparametre values, as the areas close to the minimum Theil U values are large flat areas rather than sharp points. In general however, the lower the overall tightness of the prior (the higher the gamma), the better the forecasting performance tends to be.

4 Conclusions

Risk premia in crude oil futures have provided an ample ground for research without an overwhelming consensus on whether such risk premia exist, what explains them and whether such premia can be used to forecast oil prices. The majority view tends to support the existence of the risk premia, but attempts to explain them and use them for forecasting purposes have varied in their success.

The current study attempted to shed light on the issues related to risk premia in crude oil futures, and makes a number of important findings. It confirmed the result of many previous studies by finding time-varying risk premia that tend to be negative; in other words, on average, pure futures price forecasts have tended to underestimate the future spot price. However, after the financial markets turbulance caused a rapid decrease in the price of oil in recent years, the biasedness of the futures forecast is no longer obvious. Furthermore, there is support for the traditional theories, especially the Theory of Storage and the "Market Pressure" Theory, having explanatory power for the risk premia. These results, again, were weaker when the turbulence of recent years was taken into account.

The true test of any risk premia model is its ability to forecast future events. A number of studies have claimed to have found such forecasting ability for crude oil prices with various models, some of them based on risk premia. The current study also found a model, based on speculative positions in the fufures markets, which does a good job in forecasting the oil price one year forward. The fact that speculative positions can have predictive power for future oil prices is a significant finding, which few, if any, previous studies have claimed.

However, the perhaps most important finding of the current study - and a point usually overlooked in the literature- is the extent to which the results vary depending on the sample period considered. It is relatively easy to find desired results if one chooses the correct sample. This is an especially important point for the forecasting experiment. The appropriate model that produces the best forecasting

results has no doubt changed many times in the past and it will keep changing in the future. Forecasting future developments in a segment of the financial markets which can be expected to be fairly competitive - like the crude oil futures market - remains as elusive as ever.

The current study suggests a number of avenues for future research. As regards the forecasting experiment, more advanced models, like, for example, dynamic factor models might be used to try improve on the forecasting performance. Furthermore, the predictive power of speculative positions deserves further exploration, when more data - and hopefully more detailed data - becomes available. Finally, this study, along with other recent research, suggests that concentrating on traditional spot market demand and supply shocks when studying the effects of oil in marcoeconomic models might be insufficient. We require a deeper understanding of the role of financial markets and other expectational channels through which oil shocks can affect the economy.

APPENDIX A
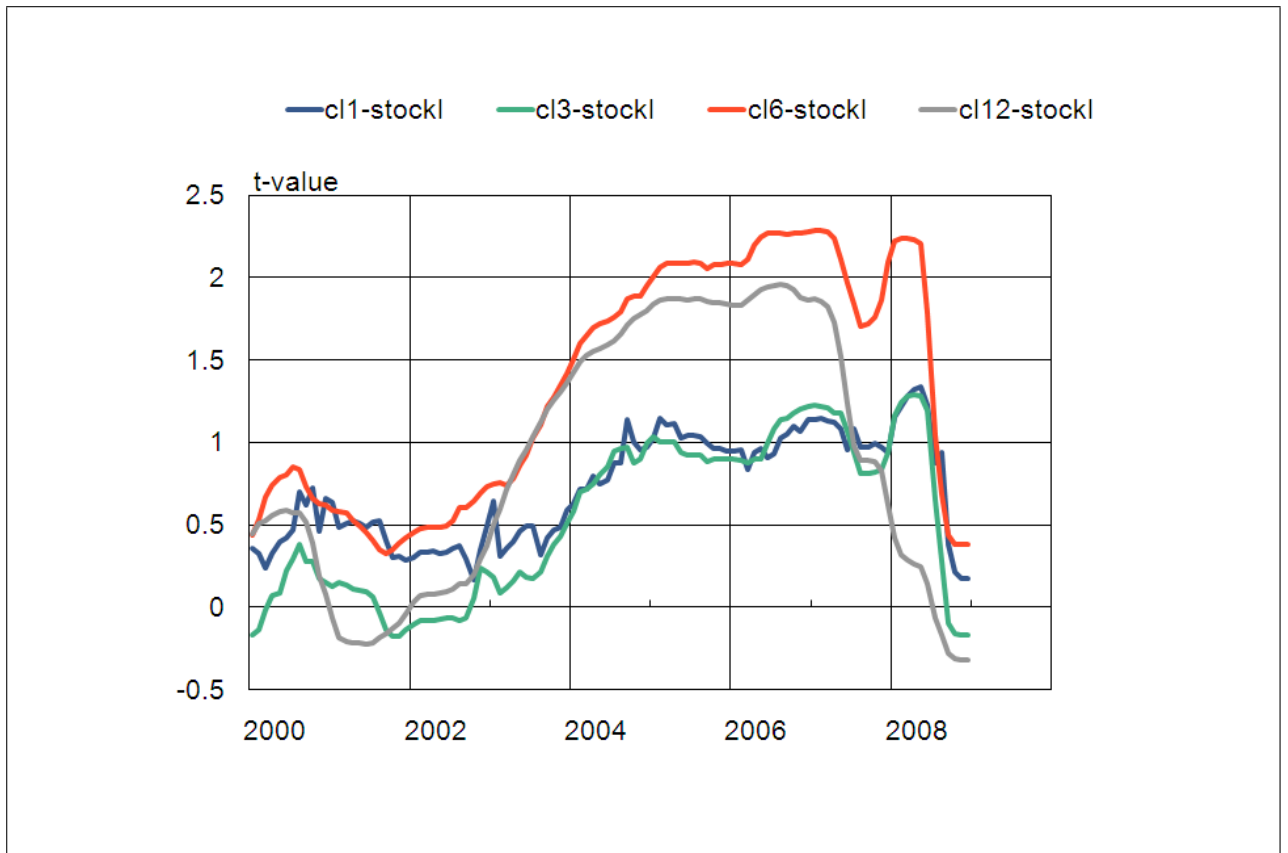
T-statistics in the rolling risk premia regressions for stockl
Note: the graph depicts t-values of regressing risk premia on crude oil stocks clx
indicates the x-month risk premium.

T-statistics in the rolling risk premia regressions for corr

Note: the graph depicts t-values of regressing risk premia on correlation between share prices and crude oil prices. clx indicates the x-month risk premium.

T-statistics in the rolling risk premia regressions for spec
Note: the graph depicts t-values of regressing risk premia on speculative oil positions.
clx indicates the x-month risk premium.

T-statistics in the rolling risk premia regressions for r

Note: the graph depicts t-values of regressing risk premia on short-term interest rates.
clx indicates the x-month risk premium.

APPENDIX B

In many forecasting applications, a Bayesian vector autoregression (BVAR) approach is used. This can be expected to lead to better forecasts than the standard unrestricted VAR approach, which suffers from overparameterisation. BVAR models allow for applying "fuzzy" restrictions on the parameters of the model, instead of sharp exclusion restrictions implied by, for example, SIC lag length criterion in standard VARs. The approach is based originally on the so-called Minnesota (or Litterman) priors introduced by Doan et. al. (1984). The approach has been used extensively since then (for an example on forecasting euro area inflation, see Benalal et. al. (2004)).

In BVARs with a Minnesota prior (used in the current paper), the standard priors have the following characteristics:

- The prior distributions on the lags of the endogenous variables are independent Normal.

- The means of the prior distributions for all coefficients are zero. The only exception is the first lag of the dependent variable in each equation, which has a prior mean of one (which centers the prior around a random walk process).

These standard priors restrict the standard error of the coefficient estimate for lag l of variable j in equation i to be of the form:

$$S(i,j,l) = \frac{\{\gamma g(l) f(i,j)\} s_i}{s_j} \tag{13}$$

where $f(i,j) = 1$ if $i = j$ and w_{ij} otherwise. s_i is the standard error of a univariate autoregression on equation i (scaling by standar errors is done to correct for the different magnitudes of the variables in the model).The hyperparameter γ and the functions $g(l)$ and $f(i,j)$ determine the 'tightness' of the prior distribution. γ is often seen as the overall tightness of the prior (the lower the number, the tighter the prior). The function $g(l)$ determines the tightness on lag one relative to lag l. The function $f(i,j)$ determines the tightness of the prior on variable j to variable i in the equation for variable i. If this is the same across all equations, $w_{ij} = w$ is a constant and the prior is said to be symmetric.

In the current study, a grid search for the best-performing forecasting model is performed for each risk premium time horizon (i.e., 1,3,6 and 12 months). For γ, the range between 0 and 1 is considered. It is assumed that $g(l)$ decays harmonically with decay factor 1, so that $g(l) = l^{-d}$ and $d = 1$. The prior is assumed to be symmetric, and w is allowed to range between 0 and 1. In other words, the two hyperparametres that are allowed to change in the grid search are γ and w.

Theil U for cl1
Note: cl1 is the 1-month risk premium, w (weight) and gamma are the two
hyperparametres used in the grid search.

Theil U for cl3

Note: cl3 is the 3-month risk premium, w (weight) and gamma are the two
hyperparametres used in the grid search.

Theil U for cl6

Note: cl6 is the 6-month risk premium, w (weight) and gamma are the two
hyperparametres used in the grid search.

Theil U for cl12

Note: cl12 is the 12-month risk premium, w (weight) and gamma are the two
hyperparametres used in the grid search.

References

Alquist, R. and Kilian, L.: 2007, What do we learn from the price of crude oil futures? CEPR Discussion Paper No. 6548, November.

Benalal, N., Hoyo, J. L. D. D., Landau, B., Roma, M. and Skudelny, F.: 2004, To aggregate or not to aggregate? euro area inflation forecasting. ECB Working Paper No. 374 (July).

Chernenko, S. V., Schwarz, K. B. and Wright, J. H.: 2004, The information content of forward and futures prices: Market expectations and the price of risk. Board of Governors of the Federal Reserve System, International Finance Discussion Papers, No. 808.

Chinn, M., LeBlanc, M. and Coibion, O.: 2005, The predictive content of energy futures: An update on petroleum, natural gas, heating oil and gasoline. NBER Working Paper No. 11033, February.

Coimbra, C. and Esteves, P. S.: 2004, Oil price assumptions in macroeconomic forecasts: Should we follow futures markets expectations?, *OPEC Review* **28**, 87–106.

Considine, T. J. and Larson, D. F.: 2001, Risk premiums on inventory assets: The case of crude oil and natural gas, *The Journal of Futures Markets* **21**, 109–126.

Cootner, P. H.: 1960, Returns to speculators: Telser vs. keynes, *Journal of Political Economy* **68**, 396–404.

Deaves, R. and Krinsky, I.: 1995, Do futures prices for commodities embody risk premiums?, *Journal of Futures Markets* **15**, 637–648.

Diebold, F. and Mariano, R.: 1995, Comparing predictive accuracy, *Journal of Business and Economic Statistics* **13**, 253–263.

Doan, T., Litterman, R. and Sims, C.: 1984, Forecasting and conditional projection using realistic prior distributions, *Econometric Reviews* **3**, 1–100.

Erb, C. and Harvey, C.: 2006, The strategic and tactical value of commodity futures, *Financial Analysts Journal* **62**, 69–97.

Fama, E. F.: 1984, Forward and spot exchange rates, *Journal of Monetary Economics* **14**, 319–338.

Fama, E. F. and French, K. R.: 1987, Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage, *Journal of Business* **60**, 55–73.

Fama, E. F. and French, K. R.: 1988, Business cycles and the behavior of metals prices, *Journal of Finance* **43**, 1075–1093.

Frankel, J. A.: 1988, Recent estimates of time-variation in the conditional variance and in the exchange risk premium, *Journal of International Money and Finance* **7**, 115–125.

Gorton, G. and Rouwenhorst, K. G.: 2006, Facts and fantasies about commodity futures, *Financial Analysts Journal* **62**, 47–68.

Horton, G. B., Hayashi, F. and Rouwenhorst, K. G.: 2007, The fundamentals of commodity futures returns. NBER Working Paper No. 13249, July.

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

Kellard, N., Newbold, P., Rayner, T. and Ennew, C.: 1999, The relative efficiency of commodity futures markets, *The Journal of Futures Markets* **19**, 413–432.

Keynes, J. M.: 1930, *A Treatise on Money, Volume II*, London: MacMillan.

Litterman, R.: 1986, Forecasting with bayesian vector autoregressions - five years of experience, *Journal of Business and Economic Statistics* **4**, 25–38.

Moosa, I. A. and Al-Loughani, N. E.: 1994, Unbiasedness and time varying risk premia in the crude oil futures market, *Energy Economics* **16**, 99–105.

Pagano, P. and Pisani, M.: 2009, Risk-adjusted forecasts of oil prices, *The B.E. Journal of Macroeconomics* **9: Issue 1 (Topics)**, 1–26.

Peroni, E. and McNown, R.: 1998, Noninformative and informative tests of efficiency in three energy futures markets, *The Journal of Futures Markets* **18**, 939–964.

Pindyck, R.: 2001, The dynamics of commodity spot and futures markets: A primer, *The Energy Journal* **22**, 1–29.

BANK OF FINLAND RESEARCH
DISCUSSION PAPERS

1/2011   Hanna Freystätter **Financial factors in the boom-bust episode in Finland in the late 1980s and early 1990s.** 2011. 39 p. ISBN 978-952-462-653-8, print; ISBN 978-952-462-654-5, online.

2/2011   Marko Melolinna **What explains risk premia in crude oil futures?** 2011. 37 p. ISBN 978-952-462-658-3, print; ISBN 978-952-462-659-0, online.

Suomen Pankki
Bank of Finland
P.O.Box 160
FI-00101 HELSINKI
Finland

. 2343

Suomen Pankki
Bank of Finland
P.O.Box 160
FI-00101 HELSINKI
Finland