



Trading time seasonality in commodity futures: An opportunity for arbitrage in the natural gas and crude oil markets?

Christian-Oliver Ewald^{a,b,*}, Erik Haugom^{b,1}, Gudbrand Lien^{b,1}, Ståle Størdal^{b,1},
Yuexiang Wu^{c,1}

^a Adam Smith Business School - Economics, University of Glasgow, Glasgow, G12 8QQ, United Kingdom

^b Department of Economics, Inland Norway University of Applied Sciences, Lillehammer, Norway

^c Rayliant Kingston, 2705, East Tower, Raffles City, Dongdaming Road, Hongkou District, Shanghai, China

ARTICLE INFO

JEL classification:

Q49
Q41
G12
G13

Keywords:

Futures
Crude oil
Natural gas
Seasonality
Arbitrage

ABSTRACT

In this paper we investigate energy futures contracts and the presence of a type of seasonality, that has been given very little to no attention in the literature — we call it trading time seasonality. Such seasonality is exposed through the futures trading time, not its maturity time, nor the underlying spot price. As we show, it can be linked to seasonality in the pricing kernel, but the latter cannot explain it fully. Its relationship to arbitrage and CAPM violation is investigated, and its presence is confirmed for natural gas and crude oil futures markets using descriptive analysis, Kruskal–Wallis testing and CAPM methodology. We provide an informal discussion around possible reasons for the effect and identify seasonal hedging pressure and market sentiments as such.

1. Introduction

Seasonality in commodity prices refers to periodical fluctuations in the distribution of spot or futures prices. This most commonly results from seasonal shifts in demand and supply but can also be the result of seasonal shifts in preferences. A large body of academic literature is dedicated to the topic of seasonality in commodities. Some of the early work includes Fama and French (1987), who confirm that seasonality exists in the convenience yield, which is closely related to the inventory level of the specific commodity, which in turn is usually subject to seasonal changes of demand and supply. Kramer (1994) studies the January effect in the stock market and argues that its source could be seasonality relating to the macro-economy. More recently, Sørensen (2002) extended the Schwartz (1997) and Schwartz and Smith (2000) model by adding a deterministic seasonal factor, governed by a linear combination of trigonometric functions, and evaluated the new model using agricultural commodity futures prices. Elsewhere, Lucia and Schwartz (2002) and Cartea and Figueroa (2005), among others, endeavour to describe spot and forward prices in the electricity market, where seasonality plays a crucial role. In other work, Lucia and Schwartz (2002) propose one-factor and two-factor models with seasonal components, while Cartea and Figueroa (2005) introduce a

mean-reverting model with jump diffusion. Finally, a large body of literature specifically focuses on energy commodities. For example, Mirantes et al. (2012, 2013) propose a number of pricing models including seasonality in the form of a stochastic factor, while Borovkova and Geman (2006) examine the seasonal pattern in the forward curves of commodity prices. All these studies focus on seasonality in the spot price, which then naturally transcends to seasonal patterns in futures prices.

In addition to the above literature, studies such as Suenaga and Smith (2011), Back et al. (2013), Koekebakker and Lien (2004), Aris-mendi et al. (2016) and Ewald and Zou (2021) attempt to model seasonality in the volatility of commodity prices. With risk aversion, this can lead to time varying and in fact seasonal risk premia. Shao et al. (2015) develop a model that includes time-varying and seasonal risk premia in the US natural gas market.

In this paper, we focus on natural gas and crude oil spot markets and their respective futures markets. All aforementioned papers, with the exception of Shao et al. (2015) and to some extent also Koekebakker and Lien (2004) and Ewald and Zou (2021), argue that seasonality in futures prices primarily relates to the maturity dates of the relevant

* Corresponding author at: Department of Economics, Inland Norway University of Applied Sciences, Lillehammer, Norway.

E-mail address: christian.ewald@inn.no (C.-O. Ewald).

¹ All authors contributed equally to this work.

futures contracts and that the likely cause of seasonality in the futures prices is the seasonality in the spot prices. In this paper, however, we emphasize other forms of seasonality which are present in commodity futures markets, particularly in natural gas and crude oil. These are exhibited through the trading time of futures contracts, not their maturities. We show that within an arbitrage-free framework, seasonality in the spot price is strongly tied to seasonality in the forward curves leaving the so called backward curves unaffected, while seasonality in preferences is tied to seasonality in both backward and forward curves, and these two cannot be uncoupled. In conclusion, it appears impossible within an arbitrage free framework to have seasonality in the backward curves only, without having seasonality in the forward curves.

In our analysis of futures prices of Henry Hub natural gas and West Texas Intermediate (WTI) crude oil, we find that in both cases futures prices reveal evidence to suggest that the trading time of the contracts influences the prices in a seasonal manner, independent of the maturity dates. More specifically, we obtain statistical evidence showing that the futures prices of natural gas, irrespective of the maturity time, reach their annual peak when traded in June, and bottom out when traded in February. Conversely, the highest and lowest futures prices for crude oil futures irrespective of the time of maturity usually occur with contracts traded in July and December, respectively. This is compelling evidence for seasonality in the preferences and supports the findings in [Shao et al. \(2015\)](#). However, we also find that the seasonal patterns in the backward curves are far more pronounced than in the forward curve, with crude oil featuring almost no seasonal patterns in the spot price and forward curves at all. This suggests the possible presence of an arbitrage.

In detail, we argue that there may exist an arbitrage opportunity in the market through exploiting this seasonal pattern by trading long/short the relevant futures contracts. Note that unlike in the spot market, futures do not carry a convenience yield, so there is no cost of storage. We design a simple trading strategy of the type “buy low sell high” to exploit this opportunity and find that the identified strategy produces excessive profits with very low risk. To benchmark this strategy in a more scientific manner, we assess its returns in the context of the capital asset pricing model (CAPM). We use the S&P 500 stock index over the same period as the market benchmark and investigate whether our trading strategy can produce significant and positive alphas for both commodities, which indeed it does. This part of our analysis is related to interesting work by [Han et al. \(2016\)](#) who assess whether a simple moving average strategy is able to generate superior performance when applied to a range of commodities. However, they only consider the Sharpe ratio, which is an inferior performance measure as compared with the CAPM alpha.

Finally, we provide a discussion about the possible origins of trading time seasonality. We identify seasonal changes in hedging pressure, seasonal risk aversion, market sentiments and natural factors as such. We discuss the role of the put–call ratio, which is connected to both hedging pressure and sentiments in more detail and provide an empirical analysis. Our discussion of the other factors remains informal, a detailed analysis would be beyond the scope of this paper, but will be included in a future research.

The remainder of the paper is organized as follows. Section 2 discusses in more detail the issue of seasonality in futures prices and the two channels through which it is created. In Section 3 we briefly describe the data, while in Section 4 we present our primary empirical findings related to trading time seasonality, including graphic and statistical evidence for trading time seasonality for both of the two commodities. We then introduce a suitable long/short strategy seeking to exploit the seasonal patterns in the backward curve in Section 5. We investigate the design, execution and profitability from this trading strategy. In particular we investigate the question as to whether our trading strategy can produce significant and positive alphas in the context of the CAPM. In Section 6 we investigate possible reasons for the effect of trading time seasonality. Our main conclusions are summarized in Section 7.

2. The two channels for seasonality

Following the no-arbitrage principle, futures prices are determined through

$$F(t, T) = \mathbb{E}_t^{\mathbb{Q}}(P(T)) = \mathbb{E}_t^{\mathbb{P}} \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{t, T} P(T) \right), \quad (1)$$

where \mathbb{Q} denotes the risk-neutral pricing measure,² \mathbb{P} the physical (real world) measure, $\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_T$ the corresponding Radon–Nikodym derivative and $\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{t, T} = \mathbb{E}_t^{\mathbb{P}} \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_T \right)^{-1} \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_T$ the corresponding pricing kernel for the interval $[t, T]$. The expectations $\mathbb{E}_t^{\mathbb{Q}}$ and $\mathbb{E}_t^{\mathbb{P}}$ are the conditional expectations under the respective measures and $P(T)$ denotes the spot price at time of maturity T .

The classical notion of seasonality is reflected in the functions $T \mapsto F(t, T)$ for fixed t , also called the forward curves. In this paper, however, we focus on a different type of seasonality, as reflected in seasonal statistical patterns in the realizations of the stochastic processes $t \mapsto F(t, T)$, the backward curves.

We refer to this form of seasonality as trading time seasonality. To the best of our knowledge, there is no literature which systematically studies trading time seasonality.

Looking at the right hand side of Eq. (1) we can immediately identify the two main channels that create seasonality: The spot price $P(T)$ and the pricing kernel $\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{t, T}$.

$P(T)$ only depends on T and any seasonal patterns in the spot price directly transcend into seasonal patterns in the forward curves. The pricing kernel on the other hand depends on both t and T and can therefore cause seasonal patterns in both the forward and backward curves, in fact it connects the two. Further, while the spot price is not explicitly affected by the pricing kernel, its convenience yield is, see [Casassus and Collin-Dufresne \(2005\)](#). Hence seasonal patterns in the pricing kernel can be linked to the spot price as well.

Both economic and asset pricing theory suggest that in equilibrium the pricing kernel $\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{t, T}$ is determined through the marginal utilities of the agents. It is therefore accounting for current and future risk preferences, and it makes sense to connect this expression to concepts such as hedging pressure and market sentiments. However, this connection has so far not been established in a formal theoretical model, even though [Hirshleifer \(1990\)](#) has provided an important contribution toward reaching this goal. We leave this as future research.

3. The data

We use daily futures prices for two energy commodities,³ namely Henry Hub natural gas and WTI crude oil. Since our focus is on annual seasonality in terms of monthly price changes, we take the average price for each month of all daily observed prices within this month as the monthly price. These prices are used in this paper unless stated otherwise. This removes any short-term patterns and abnormalities and does not affect the reliability of our results. The notation for the futures contracts in most of the paper follows the classical literature, i.e., F1, F2, F3, ..., where F1 indicates future contracts that mature in the next month, F2 in two months, and so on. In the case of natural gas, we have data from 1997 to 2020. However, from 1997 until 2002, only shorter-term futures contracts with maturity dates expanding for the next 36 months from trading dates (F1~F36) were traded. Since 2002, F37~F72 have been added. In conclusion, we divide the natural gas data into two groups accordingly. The first group includes all futures prices F1~F36 from April 1997 till July 2020. The second group

² We do not assume market completeness at this point. If the market is incomplete, we assume that the risk neutral measure is chosen by the market, which is why we refer to this measure as risk-neutral pricing measure.

³ We refer to the closing price one day ahead as the spot price.

Table 1
Average prices and standard deviations of selected contracts.

Commodity	F1	F4	F10	F18	F30	F42	F60
Natural gas, Group 1 (4/1997~7/2020)	4.2810 (2.2808)	4.5039 (2.3324)	4.6309 (2.3191)	4.6320 (2.2937)	4.5763 (2.1684)		
Natural gas, Group 2 (1/2002~7/2020)	4.5849 (2.3174)	4.8642 (2.4247)	5.0490 (2.3733)	5.0739 (2.3409)	5.0213 (2.1922)	4.9930 (2.0703)	4.9948 (1.8815)
Crude oil, Group 1 (9/1995~7/2020)	53.9105 (28.7582)	54.3511 (28.9701)	54.1549 (29.1098)	53.5834 (28.9392)	52.9350 (28.5990)		
Crude oil, Group 2 (3/2006~7/2020)	71.7816 (23.1105)	72.9415 (22.0872)	73.4422 (21.0016)	73.1917 (20.0166)	72.7226 (19.0791)	72.5273 (18.5915)	72.7690 (18.2789)

Numbers in parenthesis refer to the observation period in the first column, and the standard deviation in the remaining columns.

consists of futures contracts from January 2002 until July 2020, with maturity dates expanding for the next 60 months.⁴ For the crude oil contracts, we use data from 1995 to 2020. Between 1995 and 2006 only futures contracts with maturity dates expanding for the next 32 months (F1~F32) were traded. Therefore we divided the sample in two groups. The first group includes all futures prices F1~F32 from September 1995 till July 2020. The second group consists of futures contracts from March 2006 until July 2020, with maturity dates up to 60 months ahead. Table 1 provides a brief statistical summary of the data set. For the technical analysis, programming, and the interpretation of the results in the following section, it is good to think of all the data as being arranged in four-dimensional arrays of the type $F(i, j, k, l)$, with (i, j) the trading month and year and (k, l) the maturity month and year, with monthly averaging as discussed.⁵

4. Empirical findings using statistical analysis

In this section, we present our first set of empirical findings. Initially this involves a descriptive analysis that helps us to identify the presence of a seasonal pattern relating to trading time in addition to the classical seasonal pattern relating to maturity dates. We then formally confirm these findings in the second part of this section using statistical testing, specifically the Kruskal–Wallis test (Kruskal and Wallis, 1952). Our results complement the empirical findings obtained in Shao et al. (2015) for natural gas.

4.1. Preliminary findings

Fig. 1 plots the average $\bar{F}_{mat}(k)$ over all futures prices in the relevant data set for futures maturing in a particular month k (left)⁶ as well as the average $\bar{F}_{trd}(i)$ over all futures prices traded in a particular month (right).⁷ Averaging across years eliminates any long-term trends and emphasizes any systematic seasonal patterns.

Mathematically we have

$$\bar{F}_{mat}(k) = \text{avg}\{F(i, j, k, l) \mid i, j, l \text{ s.t. } F(i, j, k, l) \text{ is not void}\} \quad (2)$$

$$\bar{F}_{trd}(i) = \text{avg}\{F(i, j, k, l) \mid j, k, l \text{ s.t. } F(i, j, k, l) \text{ is not void}\}. \quad (3)$$

It is easy to observe from Fig. 1 that the futures prices for natural gas for both groups 1 and 2 tend to be higher when maturing during wintertime as in December and January, and remain relatively low during the rest of the year. This is the classical pattern of seasonality with respect to the maturity month and consistent with the main

⁴ We do not include F61~F72 owing to a large amount of missing data for these longer-term contracts early in the period.

⁵ Entries in the array where there are no corresponding futures prices are set void and ignored in all algebraic operations.

⁶ Here the average is taken over all trading months for contracts maturing in that particular month (but possibly different years).

⁷ Here the average is taken for a particular trading month over all contracts traded in that month (but possibly in different years) for all maturities.

Table 2
Results of Kruskal–Wallis test.

Commodity	Chi-Square (P-value)	Sample size (N)
Natural gas, Group 1	116.897 (0.0001)***	18 740
Natural gas, Group 2	597.267 (0.0001)***	29 570
Crude oil, Group 1	441.301 (0.001)***	16 160
Crude oil, Group 2	610.434 (0.001)***	19 380

Degrees of freedom: 11

10% significance *, 5% significance **, 1% significance ***.

literature on seasonality. Looking at the right-hand side of Fig. 1, however, we also notice an effect associated with the trading month. Irrespective of maturity, the prices of natural gas futures reach their annual peak when traded around summertime in May and June and bottom out when traded at the beginning of a calendar year.

The most interesting case here seems to be that of crude oil, which according to conventional theory, entails no seasonality. The graphs in the two bottom left panels of Fig. 1 are in line with this hypothesis. The curves are almost completely flat, indicating practically no differences in prices across the different maturity months. The two bottom right graphs of Fig. 1 reveal a different story however. Clearly, crude oil futures display a very similar seasonal pattern as natural gas with respect to the trading time, with the highest prices occurring when traded in the summer, and the lowest prices during the winter months.

We now perform the Kruskal–Wallis test to formalize our earlier results and provide a statistical foundation for the analysis to follow. We first investigate the raw data of daily observations before averaging monthly. The Kruskal–Wallis test involves a nonparametric one-way analysis of variance (ANOVA) for testing two or more groups of data if they are statistically different. Given there are 12 months in a year, our samples are split into 12 groups, with a degree of freedom of 11. The null hypothesis is that there is no seasonality relating to trading dates in our sample. The results of the Kruskal–Wallis test are presented in Table 2. The null hypothesis is clearly rejected for natural gas groups 1 and 2 and crude oil groups 1 and 2.

The observed seasonality patterns can be emphasized in the following way. In each year, there are 12 months, and in each month, contracts can be either traded or they mature. Hence, there is a maturity month and a trading month for every year for every contract where the highest price of the year and, respectively, the lowest price is obtained, both for trading date and maturity. We count how many times each maturity has the highest or lowest price over all trading months, and similarly, how many times each trading month has the highest or lowest price over all maturities. Mathematically, the max and min price counts $\#_{mat}^{max}(k)$ resp. $\#_{mat}^{min}(k)$ for a particular maturity month k are given as

$$\#_{mat}^{max}(k) = \text{card}\{(i, j, l) \mid F(i, j, \bar{k}, l) \neq \text{“void” for all } \bar{k} \in \{1, \dots, 12\} \text{ and } F(i, j, k, l) = \max_{\bar{k} \in \{1, \dots, 12\}} F(i, j, \bar{k}, l)\}$$

$$\#_{mat}^{min}(k) = \text{card}\{(i, j, l) \mid F(i, j, \bar{k}, l) \neq \text{“void” for all } \bar{k} \in \{1, \dots, 12\} \text{ and } F(i, j, k, l) = \min_{\bar{k} \in \{1, \dots, 12\}} F(i, j, \bar{k}, l)\}$$

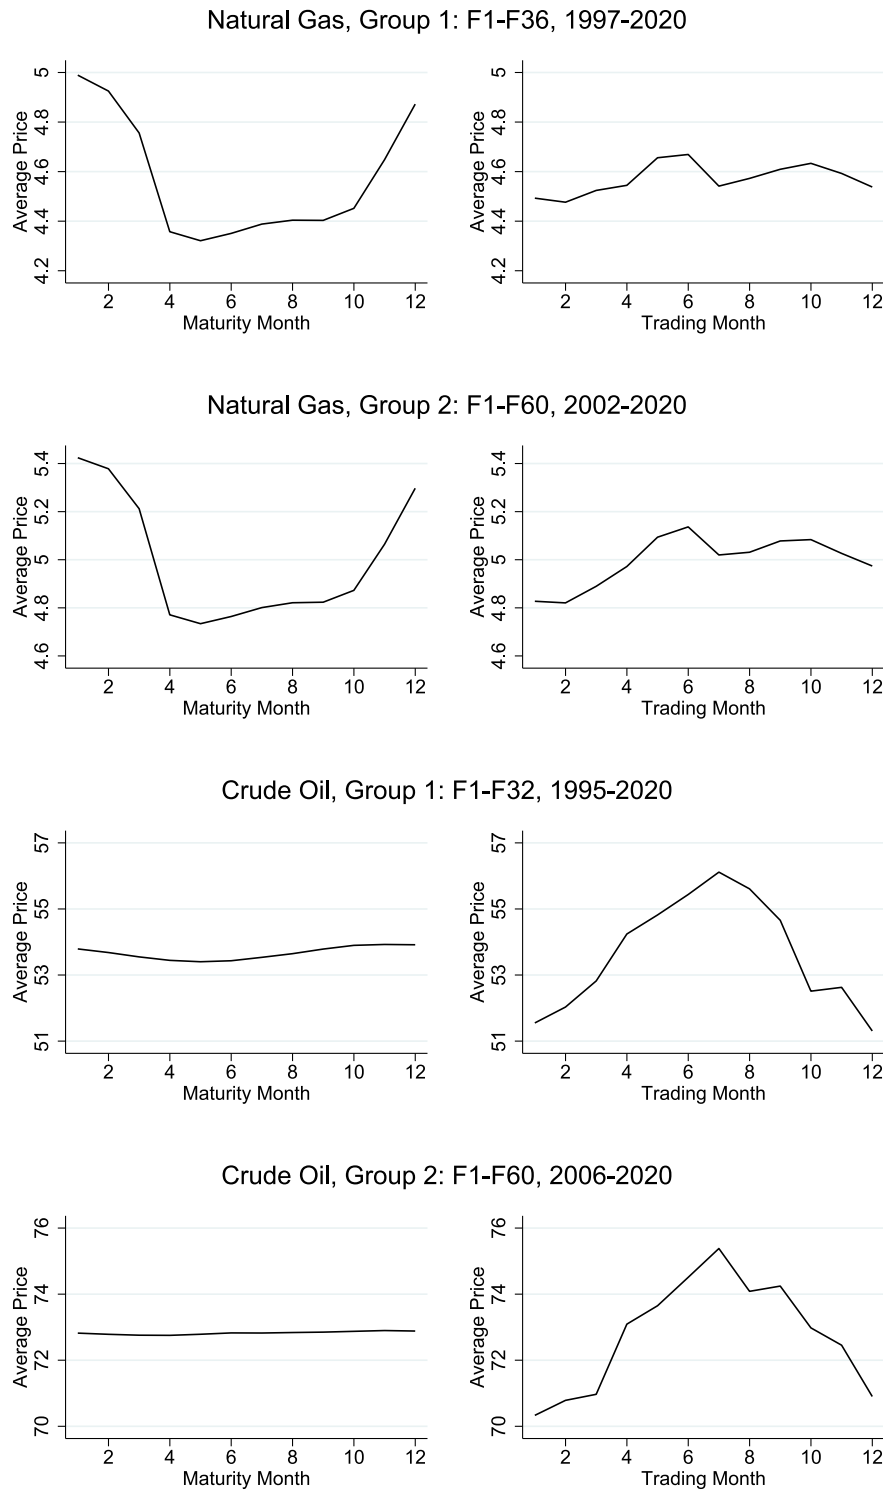


Fig. 1. Average monthly price vs. Maturity and trading months.

and by analogy, the max and min counts $\#_{trd}^{max}(i)$, respectively, $\#_{trd}^{min}(i)$ for a particular trading month i are given as

$$\#_{trd}^{max}(i) = \text{card}\{(j, k, l) | F(\tilde{i}, j, k, l) \neq \text{"void"} \text{ for all } \tilde{i} \in \{1, \dots, 12\} \text{ and}$$

$$F(i, j, k, l) = \max_{\tilde{i} \in \{1, \dots, 12\}} F(\tilde{i}, j, k, l)\}$$

$$\#_{mat}^{min}(i) = \text{card}\{(j, k, l) | F(\tilde{i}, j, k, l) \neq \text{"void"} \text{ for all } \tilde{i} \in \{1, \dots, 12\} \text{ and}$$

$$F(i, j, k, l) = \min_{\tilde{i} \in \{1, \dots, 12\}} F(\tilde{i}, j, k, l)\}$$

Here, the mathematical operator *card* denotes the cardinality of the set, i.e. the number of elements. The definition of the respective sets for the max and min counts guarantees that there will be an equal number of contracts considered for all maturities and trading dates. The values are illustrated in Figs. 2 to 5. We first look at the two groups of natural gas contracts. The upper left graphs in Figs. 2 and 3 show that almost all of the maximum prices appear when the contracts mature in the wintertime (December or January), consistent with the conventional belief that natural gas prices are higher in the winter. However, the




Fig. 2. Number of counts by maturity and trading month, natural gas, group 1 (F1~F36, 1997~2020).

upper right graphs in all three figures suggest that the minimum prices with respect to the maturity months frequently occur in April as well as in October, instead of May and June as we would conventionally expect. In fact, the summer months of July and August usually do not produce the smallest prices of the year, which may be due to increased energy use for cooling.

When it comes to trading months, the evidence is less clear, as no single month seems to have a dominant number to produce the highest prices of that year, as illustrated in the lower left graphs of Figs. 2 and 3. Nevertheless, it is obvious from the lower right graph that the contracts traded in January and December often have the lowest price of that year over all maturities for both natural gas groups. This is slightly inconsistent with our observations in Fig. 1, where the lowest average prices for contracts usually occur when traded in February. However, this effect is merely the result of the different ways of accounting for seasonality, i.e., averaging vs. counting.

The case of crude oil is more interesting. From the upper panels of Figs. 4 and 5, it seems that these contracts do not exhibit any seasonal patterns over maturity months. For the trading months, there appears to be no discernible pattern for max prices, as shown by the lower panels of these Figs. 4 and 5 while the pattern for min prices is present, but less prominent.

4.2. The forward and backward curves

To look for further evidence of trading-date seasonality for the two energy commodities, we adopt the idea of forward and backward curves. The forward curve is a classical concept and has been used

for a long time in the literature in order to describe the evolution of the expected future spot price, see Borovkova and Geman (2006). Denoting with $F(t_i, T_j)$, the futures price that is traded at time t_i and matures at time T_j , a forward curve depicts the function $F(t^*, T_j)$, where t^* indicates a fixed date of trade, and T_1, T_2, \dots, T_N the various maturity dates after t^* , with $j = 1$ being the closest to t^* , $j = 2$ the second closest, and so on. In other words, a forward curve includes all contracts with different maturity dates traded on the same date, and (at least under the no-arbitrage assumption) shows the expected spot prices under the risk-neutral market measure at the futures maturity dates. If the underlying asset possesses any seasonality that relates to the maturity dates, the forward curve should in principle present a noticeable periodic pattern. Figs. 6 and 7 illustrate two examples of forward curves for some randomly selected trading months for either of our two commodities. One can easily observe an annual seasonal pattern for natural gas and a lack of a seasonal pattern for crude oil.

Nevertheless, forward curves are unable to capture seasonality relating to trading dates, at least explicitly, and we therefore introduce what we call backward curves. The only technical difference between the forward and backward curves is that now the maturity is fixed and the trading time varies, i.e., the backward curve includes all contracts traded on different dates in the past that mature on the same date. To be more specific, the series of contracts that appear in the backward curve can be identified as $F(t_i, T^*)$, where t_1, t_2, \dots, t_N indicates the different trading dates prior to the fixed maturity date T^* , with $i = 1$ being the closest to the maturity date, $i = 2$ the second closest, and so on. The statistics of the backward curve thus illustrates a backward-looking view of the prices of those future contracts traded in different

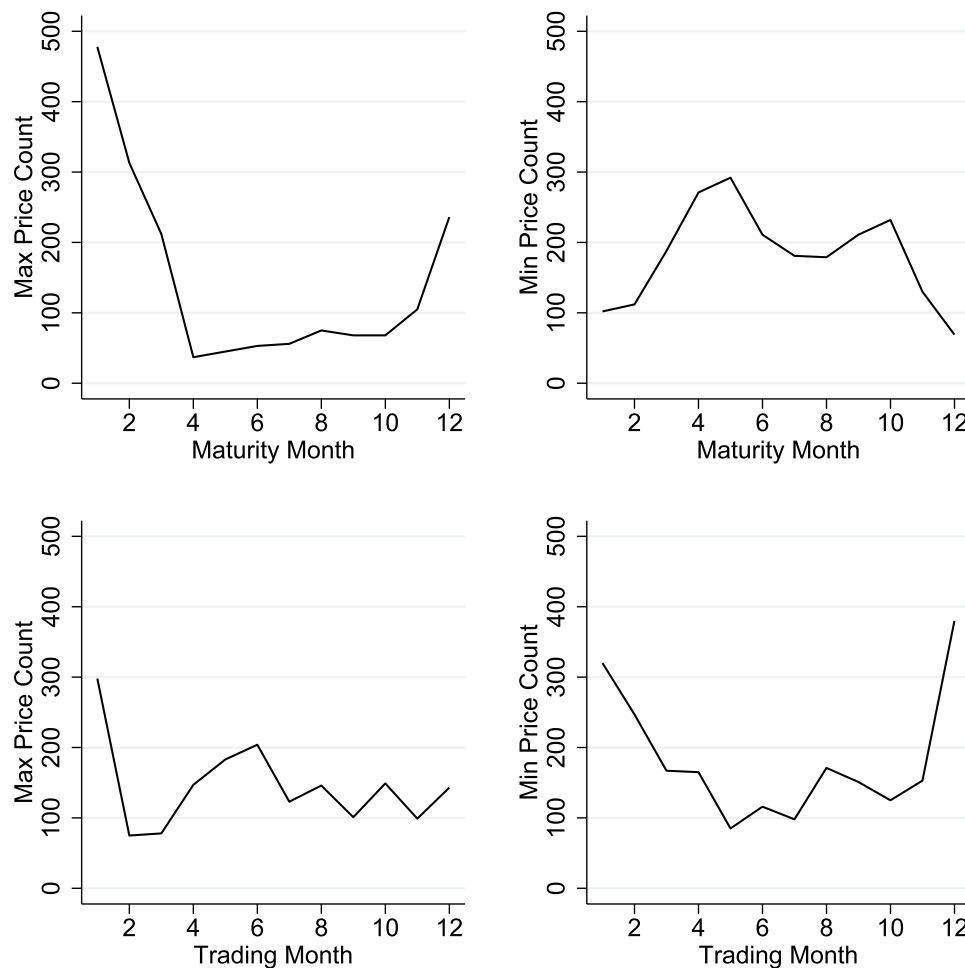


Fig. 3. Number of counts by maturity and trading month, natural gas, group 2 (F1~F60, 2002~2020).

months but maturing on the same date. If there is seasonality relating to the trading date, the backward curve should capture any and show a periodic statistical pattern.

Although the preliminary findings in the last subsection have shown evidence of trading time seasonality in the two energy commodities, individual (single realizations of) backward curves appear to fail to present any discernible seasonal pattern as volatility and the natural fluctuations in the prices obscure any such pattern. To reduce noise, we average for each maturity month and time to maturity over all maturity years. Mathematically, denoting with $F(\tau, k, l)$ the futures price for a contract maturing in month k of year l with a time to maturity τ , the average backward curve for month k is provided by

$$\tau \mapsto \bar{F}(\tau, k) = \frac{1}{N} \sum_{l=1}^N F(\tau, k, l), \quad (4)$$

where N denotes the number of years considered in the respective data set.⁸

Figs. 8 to 9 plot the aggregated backward curves for the two commodities. Each figure consists of six graphs, indicating 6 of the 12 months in a year. We chose $k = 1, 3, 5, 7, 9, 11$ for convenience. The figures present strong visual evidence of trading time seasonality in most of the cases, including crude oil, previously believed a nonseasonal commodity. In all the cases, the seasonal pattern repeats itself on an annual basis, and the peaks and bottoms coincide with our observations in Fig. 1.

5. A possible arbitrage opportunity in the future market

If trading time seasonality exists and future contracts with the same maturity are being traded systematically at higher prices in some specific months of the year and at lower prices in some others, then this pattern may be exploited to create an arbitrage-like trading strategy of the “buy low sell high” type, which can produce excessive returns.

We have demonstrated in Section 2 that under the physical measure trading time seasonality can only arise from seasonal changes in preferences and that this affects both, backward and forward curves, in a comparable manner. However, at least for crude oil, we see a much stronger seasonal effect in the backward curve than in the forward curve.

As a consequence, the trading time seasonality discovered in this paper in the context of natural gas and crude oil provides evidence of the existence of possible arbitrage, and in this section, we attempt to formally identify a possible arbitrage strategy and assess it in terms of the capital asset pricing model (CAPM). Our strategy follows a very simple approach, inspired by the naive idea of “buy low sell high” in exploiting the newly identified seasonal pattern. Nonetheless, we show that our strategy can produce consistent and significant positive alphas, the gold standard in the hedge fund literature. Nevertheless, we would like to point out that the main aim of this section is to present further evidence for trading time seasonality, reflected by a positive alpha, rather than creating a money-making machine.

Although the CAPM has been criticized in the past for many good reasons, the majority of hedge fund operators still consider the ability of a trading strategy to produce a positive alpha as the industry gold

⁸ Years are conveniently numbered as $l = 1, \dots, N$.

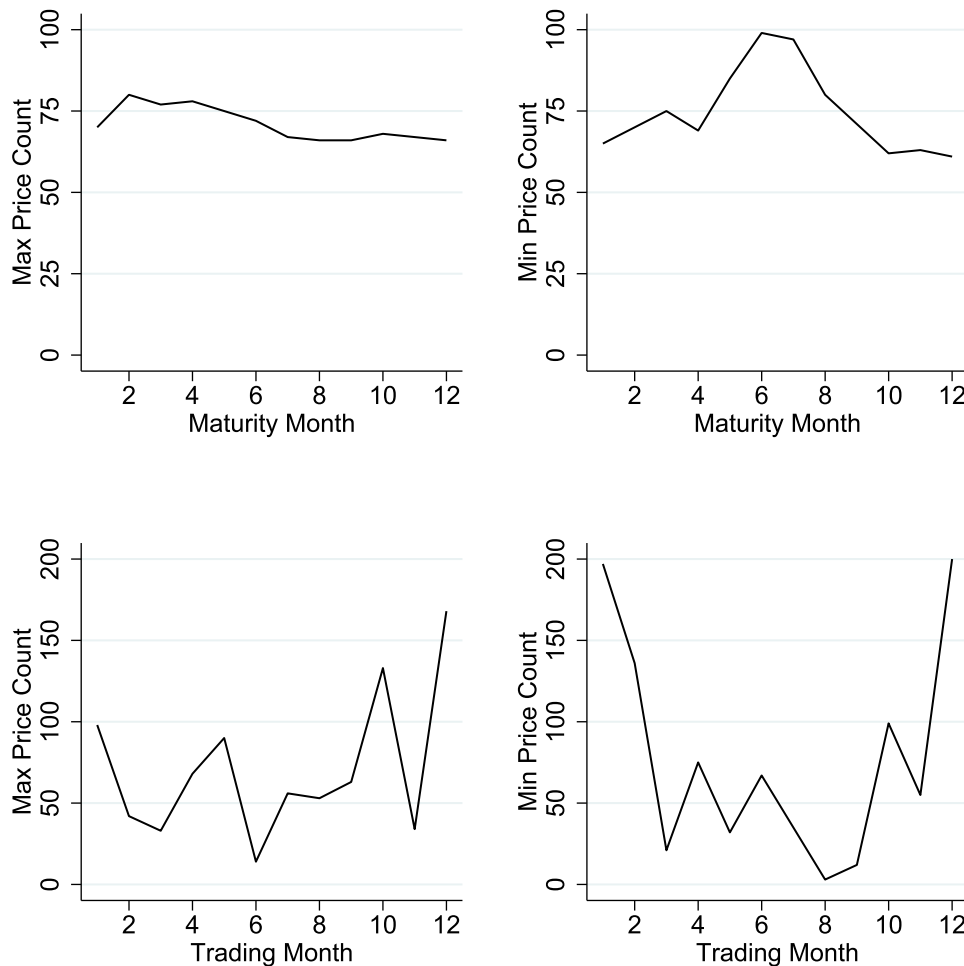


Fig. 4. Number of counts by maturity and trading month, crude oil, group 1 (F1~F32, 1995~2020).

standard. In consequence, our paper does not aim to defend the CAPM, and nor do we want to provide evidence of any CAPM violation; rather, our rationale is simply to show how trading time seasonality is reflected in the existence of a trading strategy that produces significantly positive alphas and that our strategy therefore meets the investment gold standard.⁹

5.1. A simple trading strategy, “buy low sell high”

The implementation of the strategy is purposefully simplified, in particular we assume abundant liquidity for all contracts in the market and no regulatory obstacles against any trading attempts. Our strategy is based on the unique annual trading time seasonality pattern in our two commodities. First, to simplify the strategy further, we decide to trade in only two months of each calendar year. We buy in the month of the lowest price and sell short in the month of the highest price, according to the identified seasonal pattern in each year. As a result, we view one year as a single period during which we conduct two transactions in opposite directions. For natural gas groups 1 and 2, we buy in February and then sell in June. For crude oil groups 1 and 2, we short first in July, and then take a corresponding long position in December. These choices obviously reflect our observations in the previous section.

⁹ Aiming to explain abnormal returns by including additional factors, such as in Fama and French (1992), would be an interesting avenue for future research. However this would also distract from the original question to benchmark our strategy, which is why we defer from this here.

Table 3
Details of trades in each period.

	First trading month			Second trading month		
	Position	Month	Contracts	Position	Month	Contracts
Natural gas group 1	Long	Feb.	F5~F36	Short	June	F1~F32
Natural gas group 2	Long	Feb.	F5~F60	Short	June	F1~F56
Crude oil group 1	Short	July	F6~F32	Long	Dec.	F1~F27
Crude oil group 2	Short	July	F6~F60	Long	Dec.	F1~F55

Second, each contract traded in the first trading month must match another in the second trading month with the same maturity date so that the initial position can be perfectly hedged. In other words, the contracts we trade in the first trading month must not mature before their position can be closed in the second trading month of the same year. Table 3 lists the details of the traded contracts for each commodity in one period.

Entering a futures contract is per se costless, i.e., it does not require an initial investment to set up a position in futures contracts other than a relatively small amount of money. This amount is usually proportional to the value of the open position and to be deposited at the beginning of each trade in the margin account (as collateral to settle any gain or loss when the position is marked to market on a daily basis). This complicates the computation of relative returns of any futures strategy and formulae which are typically expressed in relative returns, such as CAPM, which need to be adjusted accordingly, as we show later.

We assume for illustration that the initial value of the contracts that we buy or sell in the first trading month of each year will be 1

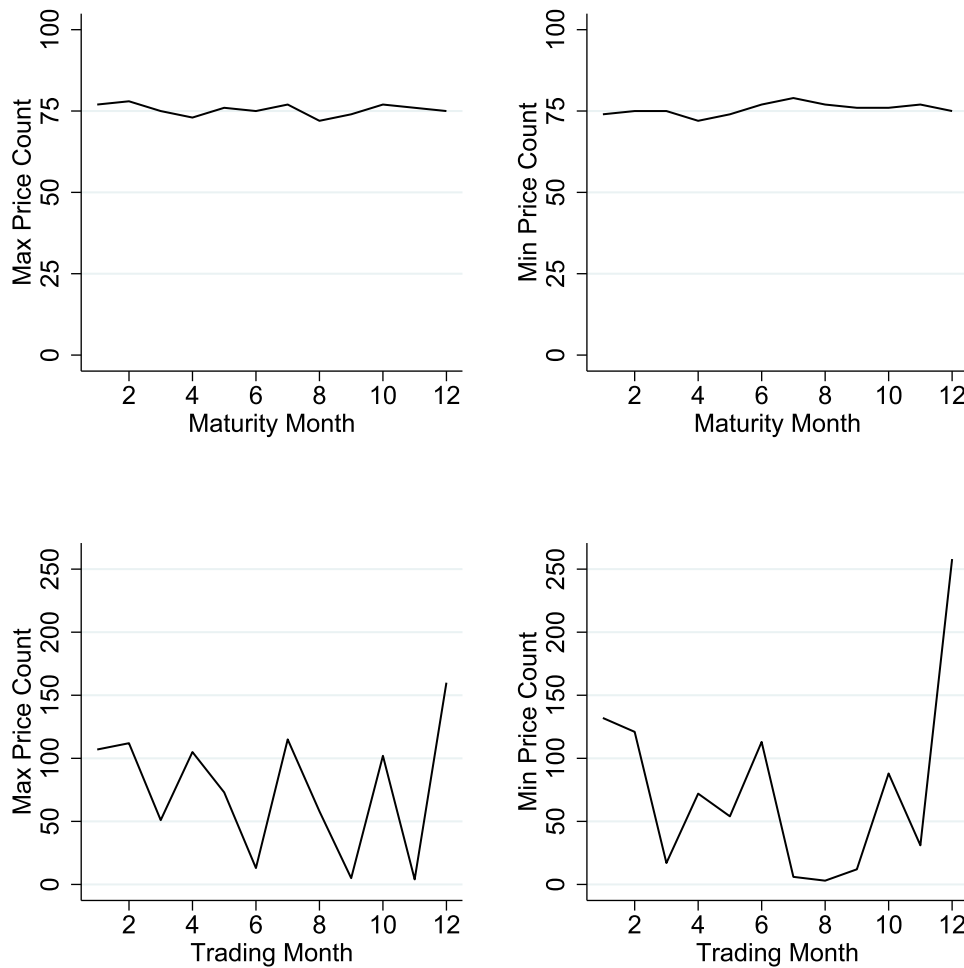


Fig. 5. Number of counts by maturity and trading month, crude oil, group 2 (F1~F60, 2006~2020).

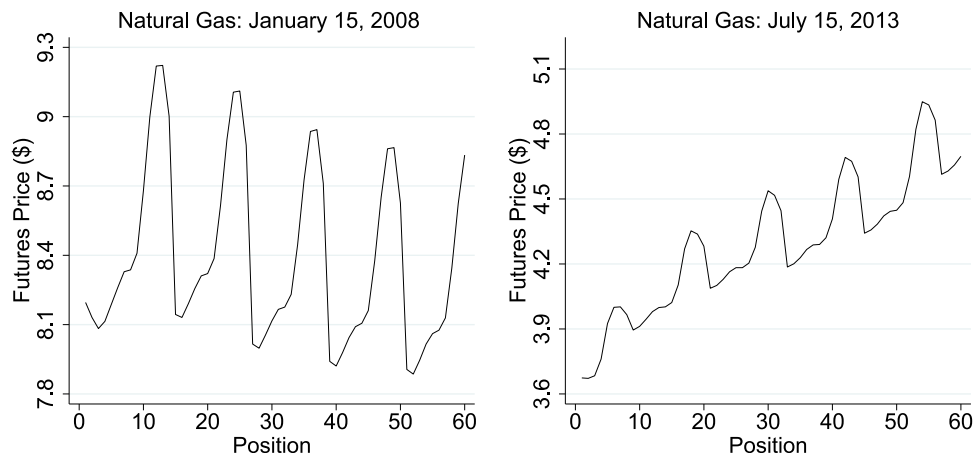


Fig. 6. Forward curve for two specific dates, natural gas.

million dollars. As for the transaction costs, we divide these into two parts according to the usual conventions for some trading outlets such as the CME Group. The first part involves the bid-offer spread, S , and the contract units, U . The cost per contract per transaction will be calculated as $C = S * U$. For natural gas, $S = \$0.005$ per million British thermal units (MMBtu), and $U = 10,000$ MMBtu per contract. In the case of crude oil, $S = \$0.02$ per barrel, and $U = 1000$ barrels per contract. The second part of the transaction cost consists of the

exchange fee, which is set at \$1.5 per contract per transaction for both commodities.¹⁰

Fig. 10 illustrates the pay-off of our trading strategy by year of operation for the two commodities. First, we can see that the strategy does

¹⁰ We assume that the investor has enough money to settle all margin calls, and in consequence we abstract from any form of margin risk. An assessment of this assumption is presented further below.

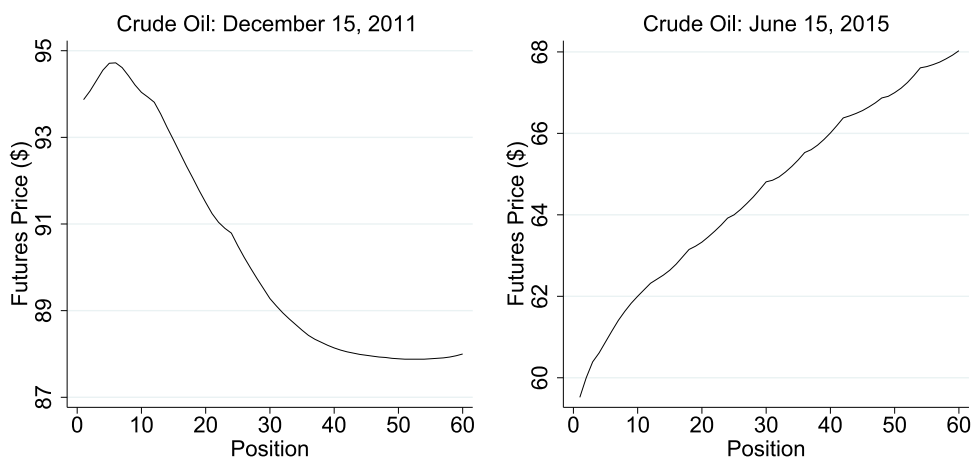


Fig. 7. Forward curve for two specific dates, crude oil.

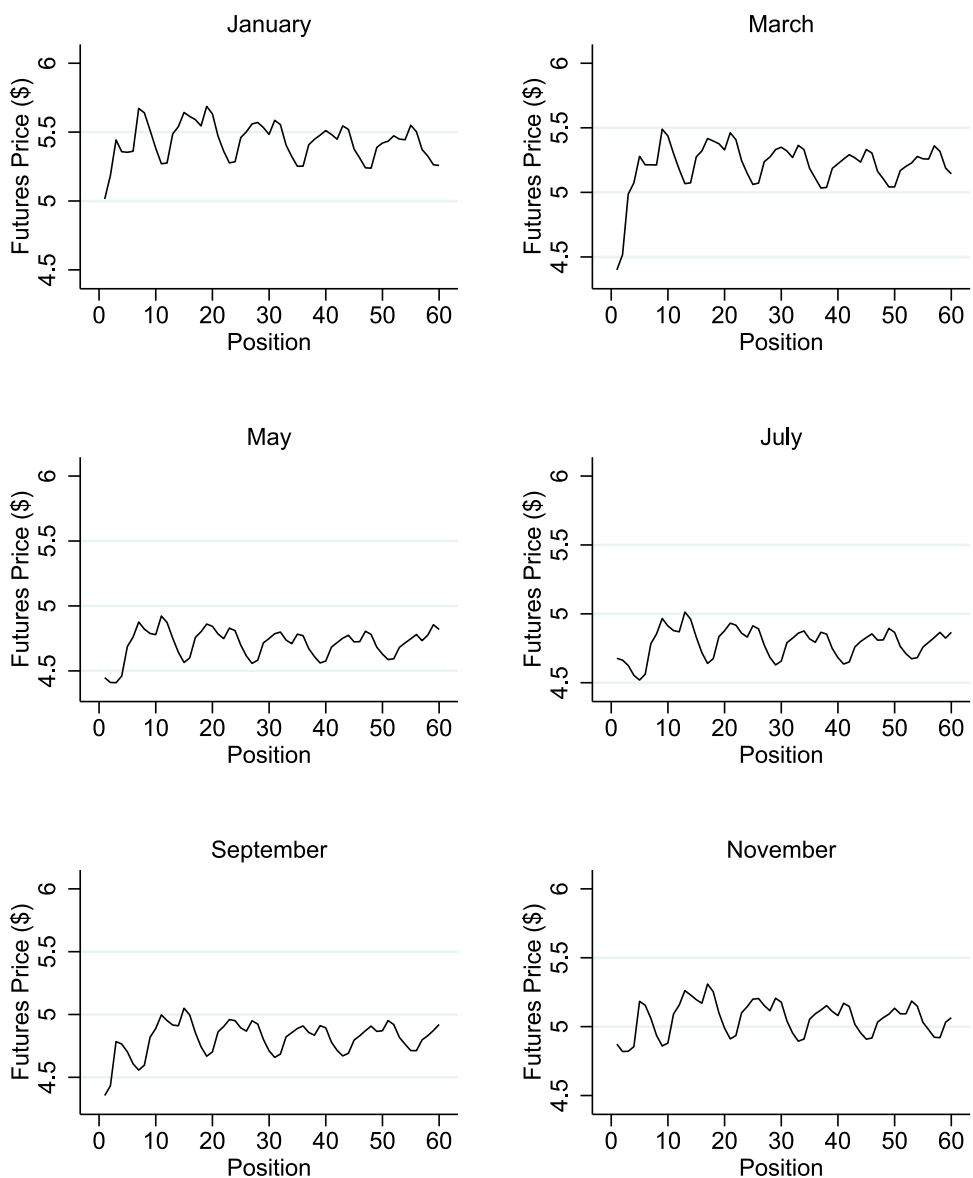


Fig. 8. Aggregated backward curves $\tau \mapsto \bar{F}(\tau, k)$ for months $k = 1, 3, 5, 7, 9, 11$, natural gas, group 2 (F1~F60, 2002~2020).

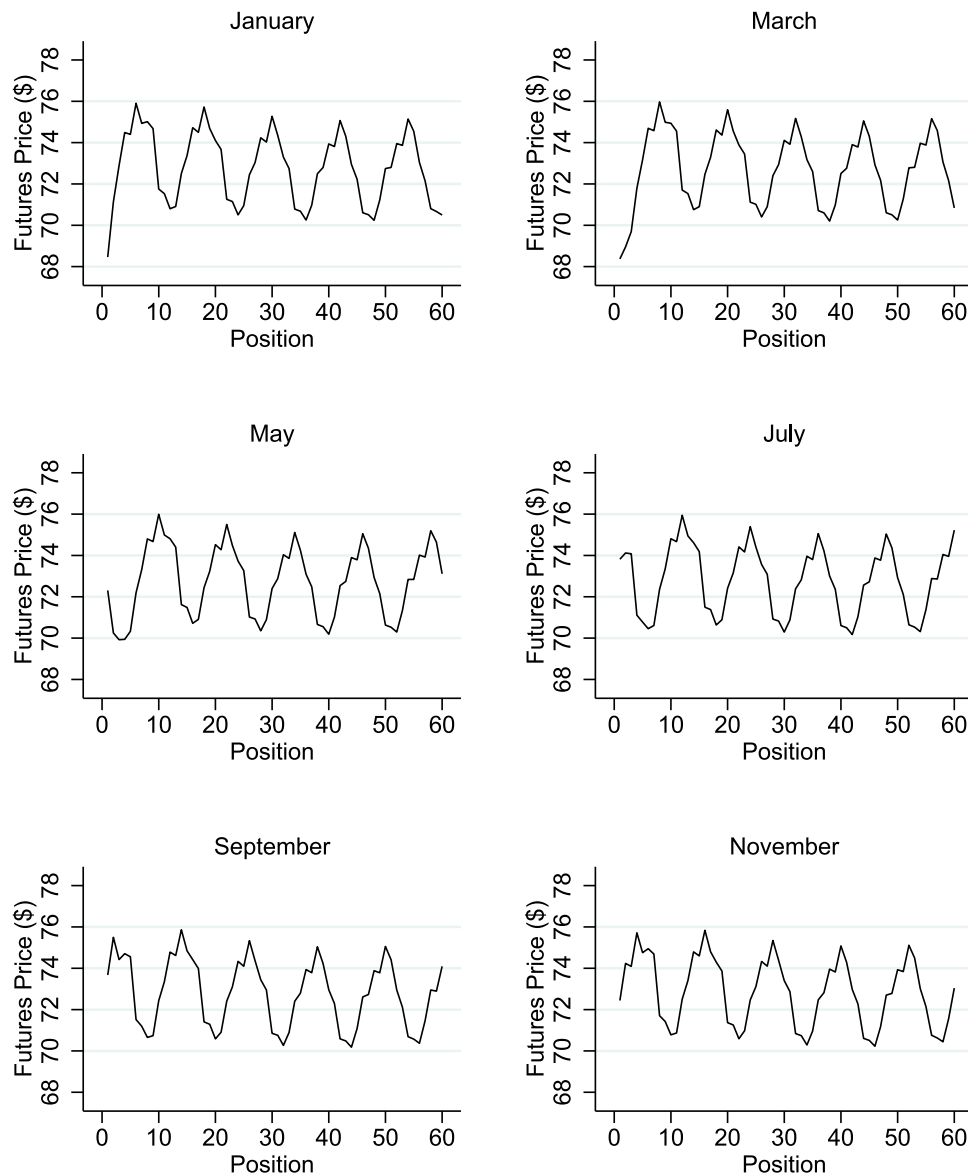


Fig. 9. Aggregated backward curves $\tau \mapsto \bar{F}(\tau, k)$ for months $k = 1, 3, 5, 7, 9, 11$, crude oil, group 2 (F1~F60, 2006~2020).

not perform positively in every year during our observation period. It appears to generate excessive profits more consistently for natural gas, but less so for crude oil. Nevertheless, the total pay-off of the entire operation period for both commodities results in profits, as shown in the right panels of Fig. 10, which means that the earnings from the profitable years more than compensate for all the losses in the losing years.

Second, it is worth noting that the annual pay-off of both commodities seems to fluctuate significantly. The most profitable year for all four groups is 2008, when the profits dwarf any other periods, contributing to a sizeable percentage of the total pay-off at the end of the observed period, especially for crude oil. However, even if we were to exclude 2008 as an exceptional year, profits would still be high.

Next, we calculate the average profit generated from the different contracts over the entire observation period. The results are illustrated in Fig. 11, where the x -axis represents time to maturity since the second trade of every year. First, it is easy to see that some contracts tend to generate more profits than others. In both groups of natural gas contracts, the contracts with around 8 to 20 months to maturity tend

to generate higher profits than the other contracts. In the case of crude oil, on the other hand, the contracts with the shortest time to maturity significantly yield more profit than the others. Accordingly, in practice a trader can build a more efficient strategy at lower cost by only trading those contracts generating the highest overall profits.¹¹

¹¹ In addition to studying the profitability of our strategy, it would be appropriate to account for its risks. One noticeable feature of our strategy is that the two opposing trades do not take place at the same time, but several months apart in every year. This may imply a certain level of risk, as for a short window of time (4 months for natural gas, and 5 months for crude oil), our initial positions after the first trading month are not hedged by any trade in the opposite direction. Given the unique feature of the futures market that each position is marked to market everyday, we also assessed the net fluctuation of the margin account during the risky period. The main finding is that even the largest floating losses we would ever have to bear during the risky months are relatively small, compared with the pay-off reported in Fig. 10. Detailed results are available from the authors upon request. Nevertheless, our CAPM analysis in the next subsection takes full account of market risk.

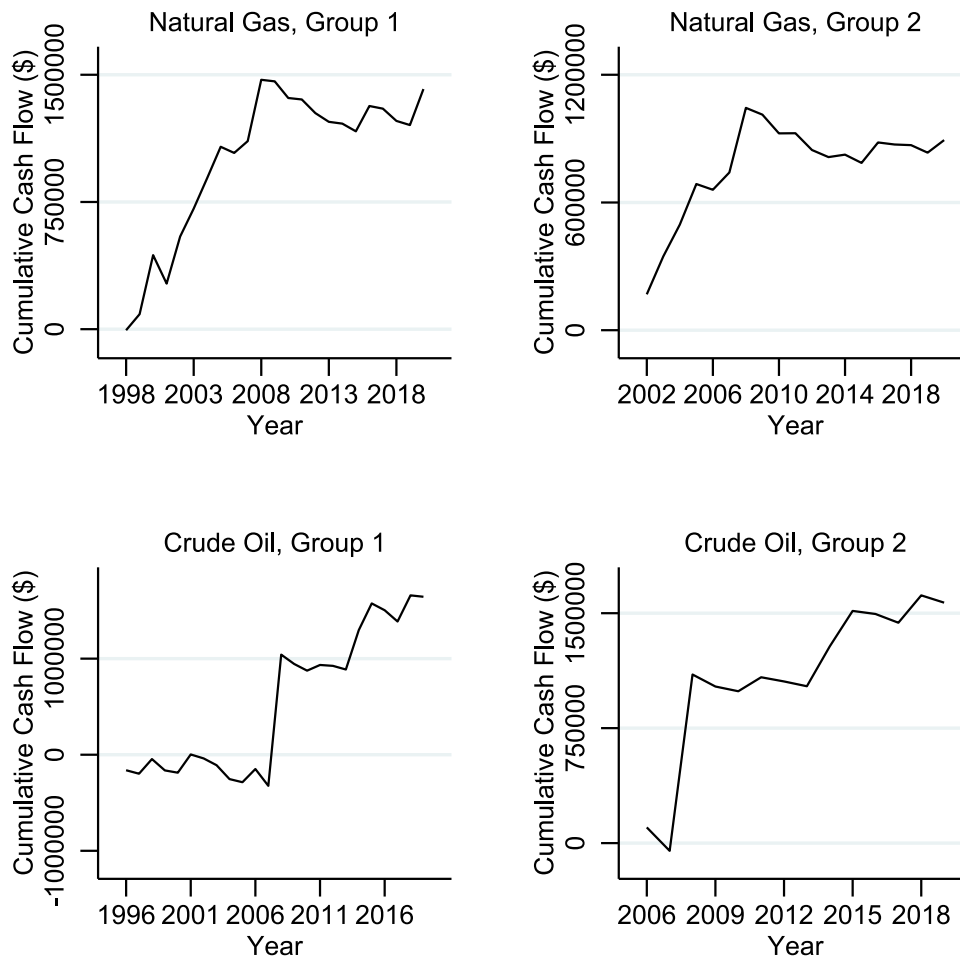


Fig. 10. Results of the trading strategy by year.

5.2. Our strategy in the context of the CAPM

Let us now benchmark our trading strategy within the framework of the CAPM. In order to do this, it is important to briefly review some issues around the applicability of CAPM in the context of commodity futures, which are due to net-zero supply and issues in computing returns for commodity futures, see [Dusak \(1973\)](#).

According to the CAPM, assuming there exists a market portfolio that can represent the performance of the entire market paying returns of R_m and denoting the risk-free rate as r_f , the expected return of a risky asset R_i satisfies

$$\mathbb{E}(R_i) - r_f = \beta_i(\mathbb{E}(R_m) - r_f), \tag{5}$$

where

$$\beta_i = \frac{cov(R_i, R_m)}{\sigma^2(R_m)}$$

represents the measure of exposure to market risk and \mathbb{E} denotes the expectation operator. Hence, the term on the right-hand side of equation (5) represents the risk premium that asset i should earn over the risk-free rate, based on its sensitivity to market performance.

As discussed above, the unique feature of trading futures contracts presents a challenge when applying the CAPM, as buying and selling future contracts does not actually require any initial investment. When a position is opened, the buyer and the seller sign a contract with a broker, in which the buyer promises to purchase an amount of the underlying asset from the seller at a future date for a pre-arranged price that both parties will honour by depositing a certain amount of money in the margin account as collateral and settling any fluctuation in the

value of the position with cash on a daily basis. Therefore, the value of the contract when opened is zero, and its fluctuations afterward are marked to market on a daily basis until the position is closed. This makes calculating the returns on investments in the futures market more difficult. It is also worth noting that unlike stocks, commodity futures contracts are not explicitly included in the market portfolio as they are in net-zero supply. For each long position of a futures contract, there must be a respective short position. Therefore, the overall net position of all futures contracts for the same commodity must equal zero.¹²

In light of this complication, we decide to follow the classic approach in [Black \(1976\)](#) and [Baxter et al. \(1985\)](#), who argue that although the relative return on the futures contract investment cannot be calculated given the initial value is zero, we can use the change in the value of the open position. Specifically, if $\mathbb{E}(R_i) = \frac{\mathbb{E}(P_i^1) - P_i^0}{P_i^0}$, where P_i^0 and P_i^1 are the values of the asset i at times 0 and 1, the CAPM equilibrium condition can be transformed into the following form:

$$\mathbb{E}(P_i^1) - P_i^0 = r_f P_i^0 + \beta_i^* (\mathbb{E}(R_m) - r_f), \tag{6}$$

where $\beta_i^* = \frac{cov(P_i^1 - P_i^0, R_m)}{\sigma^2(R_m)}$. As no initial investment is required, we have that P_i^0 is zero, while P_i^1 can be calculated as the change in the futures prices over the period, denoted as Δp_i . Hence,

$$\mathbb{E}(\Delta p_i) = \beta_i^* (\mathbb{E}(R_m) - r_f), \tag{7}$$

¹² However commodities are implicitly included in the market portfolio through various companies whose earnings reflect commodity prices.

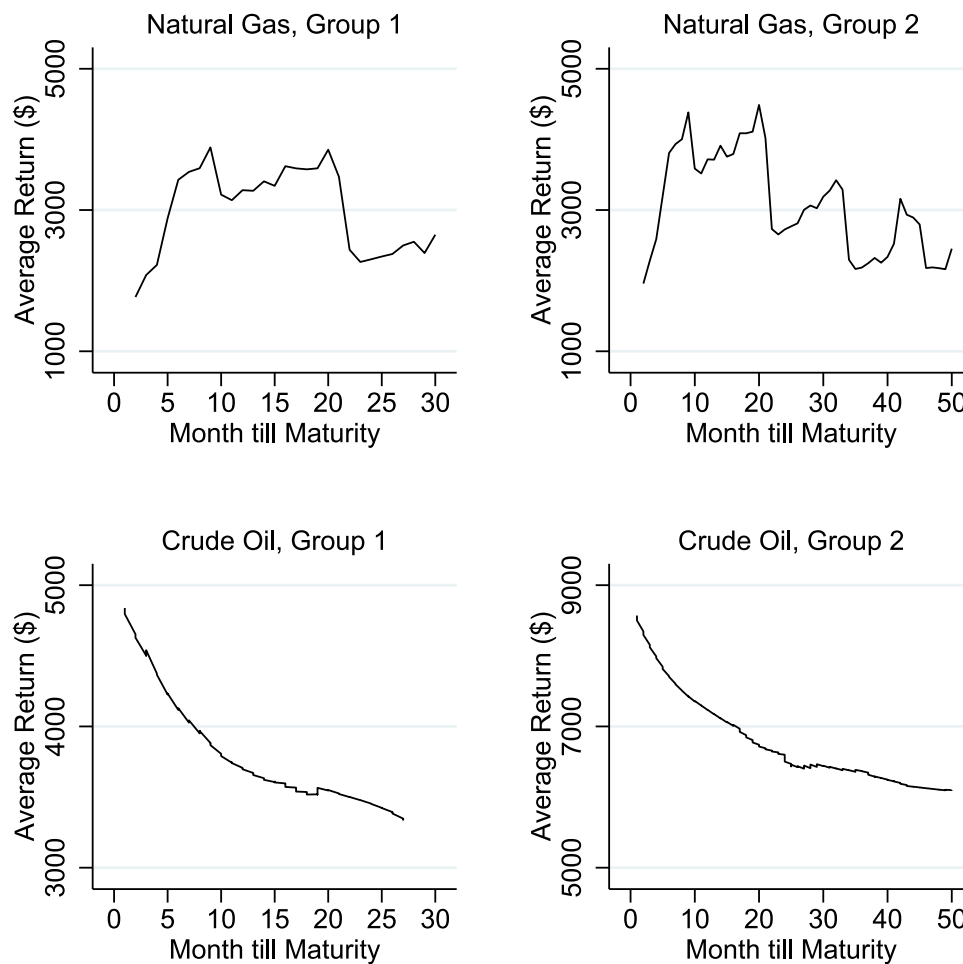


Fig. 11. Results of the trading strategy by contract.

and $\beta_i^* = \frac{cov(\Delta p_i, R_m)}{\sigma^2(R_m)}$. In other words, the expectation of the price change of a futures contract is equivalent to the product of the expected market return over the risk-free rate and the specific β_i^* that measures the sensitivity of the futures prices change to the market portfolio.

The empirical model reflecting the adaptation of the CAPM to the futures market is as follows:

$$r_i^t = \alpha_i^* + \beta_i^*(r_m^t - r_f^t) + \epsilon_i^t, \tag{8}$$

where r_i^t captures the change in the futures price during a discrete time interval starting at t . The expression r_m^t is the return of the market portfolio over the same period, and α_i captures the excessive return that trading a specific futures contract can generate. If there exists an arbitrage opportunity that can be exploited by a trading strategy, the trades should in theory generate a positive and statistically significant α_i . According to our trading strategy, commodity i corresponds to one of the three groups of the two commodities, and the time period covers the time between the two trades we conduct in each year. We select the S&P 500 index as our proxy for the market portfolio and the 10 year US Treasury bond yield as the risk-free rate, r_f^t .

The data used within our CAPM framework are divided into different panels according to their time to maturity. This takes account of our previous observation that contracts with different terms of maturity appear to carry different levels of profitability, see Fig. 10, and we would like to gather further evidence of such disparity in profitability in the context of the CAPM. For natural gas group 1, we split the data into three panels. The first panel includes the three contracts with the shortest times to maturity, or F1~F3 of the second trade (or F5~F7 if referring to the first trade), the second F16~F18, and the

last one F30~F32. In the case of natural gas group 2, the first panel includes F1~F3, the second one F19~F21, the third F37~F39, and the last F54~F56. For crude oil group 1, the first panel includes F1~F3, the second F16~F18, and the third F25~F27. For crude oil group 2, the first panel includes F1~F3, the second F18~F20, the third F35~F37, and the last F53~F55. For all panels, we enter a given position on all possible trading dates in the first trading month and close it out on the corresponding trading dates in the second trading month.

The results for the OLS regression of the CAPM are presented in Table 4. If we look at the two natural gas groups, it is easy to observe that although β_i^* is never significant (except for group 2, panel 3) in our tests, α_i in most cases is. This is especially true when trading contracts with longer times to maturity. The results for crude oil show significant α_i 's and β_i^* 's in all four panels, for both groups.

To illustrate the development in the markets over time, we also run recursive estimations of the CAPM and report the α_i 's with 95% confidence bands in Fig. 12. For expository purposes, we only report the results for the groups with shortest (left panel) and longest (right panel) time to maturity. The first estimation results in all these figures are based on historical data up until the given date on the x-axis. Then the models are re-run including one more observation at a time. For natural gas futures, the general pattern is that the α_i 's have slowly declined as the market has matured. For crude oil group 1, the strategy does not create any positive significant alphas until about 2008. After that however, the alpha's are highly significant and mainly increasing until 2020. For group 2 the alphas are highly significant over the observed period. At this point, one may question whether the positive alpha's for crude oil are possibly created by a single cataclysmic event such as the 2008 Global Financial Crisis. However, this is not the case. The Fig. 13

Table 4
CAPM regression parameters.

Panel	Natural gas group 1		Natural gas group 2		Crude oil group 1		Crude oil group 2	
	α_i^*	β_i^*	α_i^*	β_i^*	α_i^*	β_i^*	α_i^*	β_i^*
1	0.10 (0.08)	-0.42 (0.91)	0.19** (0.08)	1.66* (0.91)	6.86*** (0.93)	-138.62*** (13.98)	11.63** (1.30)	-163.88*** (12.37)
2	0.30*** (0.05)	0.15 (0.56)	0.39*** (0.05)	0.66 (0.54)	5.36*** (0.72)	-116.42*** (11.16)	8.89*** (0.95)	-136.96*** (9.35)
3	0.22*** (0.04)	-0.63 (0.48)	0.17** (0.04)	-1.09** (0.44)	4.80*** (0.67)	-106.58*** (10.51)	8.18*** (0.82)	-120.35*** (8.31)
4			0.24*** (0.04)	-0.68 (0.44)			7.72*** (0.79)	-109.93*** (8.09)

Number of observations in each panel: Natural gas group 1(2) 323(266) ; Crude oil group 1(2) 386(224).

The symbols *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Robust standard errors in parenthesis.

shows the alpha's obtained through a recursive regression in the same way as before, but this time for a period starting after the financial crisis. The alphas shown are still highly significant.¹³

These findings indicate that our trading strategy is able to generate significant alphas and therefore “beat the market”, even though some of this ability has declined in natural gas as the market has matured, while for crude oil it has become more pronounced. Yet, in the context of the CAPM, we can confirm the profitability of the proposed trading strategy. Positive and significant alphas are a strong indicator for the existence of arbitrage, or at least some form of market anomaly.

6. What are the potential factors that create trading time seasonality

Candidates are in principle all factors that have been shown to affect risk premia. Our analysis will focus on the most important ones and those where we believe time varying and seasonal patterns are most likely to be found. These are: hedging pressure, ESG indicators and natural factors. All three can be tied to sentiment indicators.

6.1. Hedging pressure

A futures price above the expected spot price indicates that the long-position (buyer) is willing to pay a premium for hedging the spot, on the other hand a futures price below the expected spot price indicates that the short position (seller) is willing to pay a premium for hedging the spot. As actual futures prices reflect supply and demand the question of whether risk premia in futures markets are positive or negative depends on what in the literature is referred to as hedging pressure, the relative difference between (reported) short hedge positions and (reported) long hedge positions in the futures market, see [De Roon et al. \(2000\)](#) and [Hirshleifer \(1990\)](#).¹⁴ Systematic hedging pressure is found to be a significant determinant of commodity futures risk premia, see also [Basu and Miffre \(2013\)](#). This literature argues that supply and demand and, hence, futures prices are determined through hedgers' and speculators' preferences, the size of inventories, access to hedging, and possibilities for diversification.

In difference to [De Roon et al. \(2000\)](#) we argue that risk is much better reflected in options contracts than futures contracts and that therefore hedging pressure originating from the aim to eliminate risk may be better reflected in the so called put–call ratio.

¹³ This specific result might point toward a regime shift in the crude oil futures market at the time of the Global Financial Crisis. Such a shift has also been reported by [Nikitopoulos et al. \(2017\)](#).

¹⁴ For futures there is of course one long position for every short position and the positions need to be classified as short hedge or long hedge by other means and different authors have proceeded differently. This problem is alleviated by considering the put–call ratio as a measure for hedging pressure, but this has limitations as well.

The put–call ratio is the relative difference between volume of traded put options vs. the volume of traded call options.¹⁵ Intuitively, someone buying a put (sell) option is to hedge the price risk when selling an asset (specifically when she/he believes that the asset-price may decrease) while someone buying a call (buy) option is to hedge the price risk when buying an asset (specifically when she/he believes that the asset-price may increase). Put–call ratios therefore contains relevant information as to whether the market believes prices are likely to go down or up and have been investigated in the past for their potential use as sentiment indicators, see for example [Pan and Poteshman \(2006\)](#) and [Bathia and Bredin \(2013\)](#). More recently put–call ratios have even gained more prominence by being included in a key sentiment indicator, the so called Market Sentiment Meter published by the CME Group, see [Kownatzki et al. \(2022\)](#) and [Putnam \(2020\)](#) for details. Considering someone who is buying a put as a short hedger and someone buying a call as a long hedger the put–call ratio can also be used as a proxy for hedging pressure. [Fig. 14](#) shows the put–call ratios for crude oil and natural gas reflecting volume (left) and open interest (right) aggregated over each month and years 2012–2021.¹⁶ Volume corresponds to daily trading volume first averaged for each month of each year, and then averaged for each month over the years, 2012–2021, the same process as applied for [Fig. 1](#). Open interest corresponds to the total number of option positions held at a day, then averaging is done in the same way as for volume. One can argue that as volume reflects trading at a given day, it is a better indicator for hedging pressure and hence the risk premia on that day and month. However, put–call ratios based on open interest have been considered in the literature, e.g. [Pan and Poteshman \(2006\)](#), so we include both.

As can be seen from the upper part of [Fig. 14](#), the put–call ratios for crude oil are very stable over the calendar year and hence cannot be linked to the observed trading time seasonality of crude oil, evidenced in [Fig. 1](#) lower right. For natural gas, the lower left part of [Fig. 14](#) seems to indicate that the put–call ratio (volume) is slightly higher in the summer months from June till August and lowest in the winter months from November till January. Dunn's test for pairwise comparison confirms this and shows that the difference is indeed statistically significant, see [Table 5](#). We therefore conclude that there is a seasonal periodic pattern in put–call ratios for natural gas.¹⁷ Furthermore, the

¹⁵ Some authors, e.g. [Pan and Poteshman \(2006\)](#) and [Bathia and Bredin \(2013\)](#) instead use the ratio of put volume to the total option volume (put and calls) as put–call ratio, but this is not consistent with the name giving and most major markets, including the CME and CBOE, indeed report the put–call ratio in the way as we use it, i.e. volume of puts to volume of calls.

¹⁶ We gratefully acknowledge support from the CME Group, providing us with an extensive data set of key market parameters including volumes of put and call options for the relevant commodities. Our analysis in this section is based on this data-set.

¹⁷ For open interest the same holds, but to a lesser degree. However, as noted above we consider put–call ratios based on volume as the more appropriate ratio.

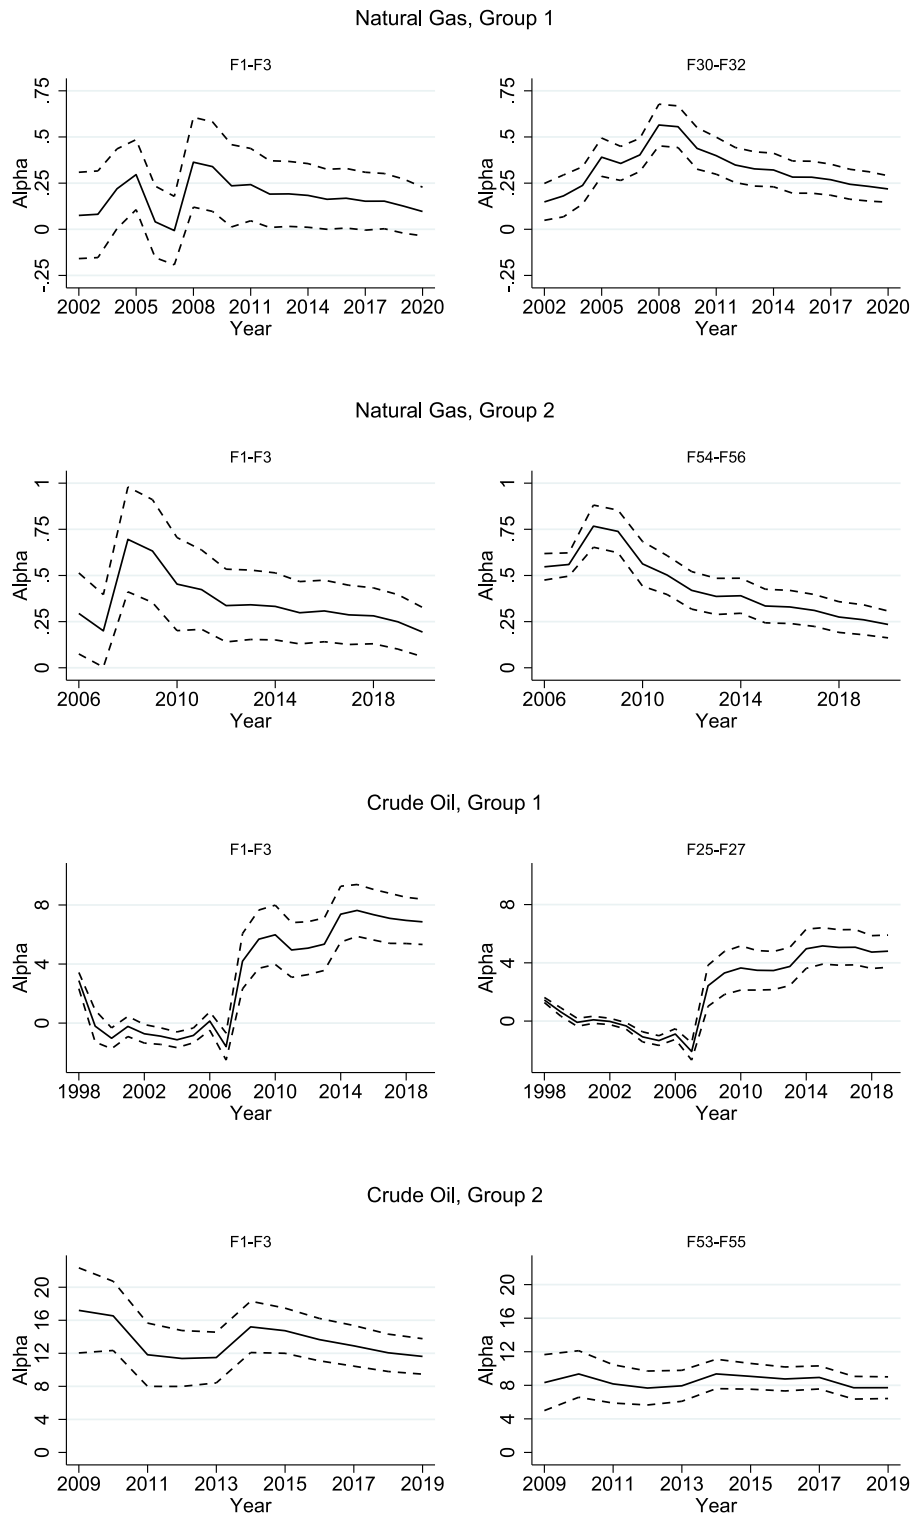


Fig. 12. Alpha for the shortest (F1-F3) and longest Contracts of natural gas (groups 1 and 2) and crude oil (groups 1 and 2).

timing of the seasonal pattern is consistent with the observed trading time seasonality in natural gas, Fig. 1 bottom right as well as our results in Sections 4.1 and 4.2. However, with the interpretation of the put-call ratio as a measure for hedging pressure presented as above, the two affects actually oppose each other. A high put-call ratio in the summer months would mean a higher demand in protecting sales, which should shift the risk premium to the sell side and hence cause lower futures prices in the summer months. Similarly, a relatively low put-call ratio

in the winter months should lead to higher futures prices in the winter months. However, this just shows that the interpretation of the put-call ratio as a proxy for hedging pressure in such simple manner is flawed. As for futures contracts, one has to recognize for each long put there is a short put, and for each long call there is short call. It is reasonable to assume that the long positions are held predominantly by hedgers while most of the short positions are held by speculators. In terms of risk premia, short positions are obviously inversely affected as compared to

Crude Oil

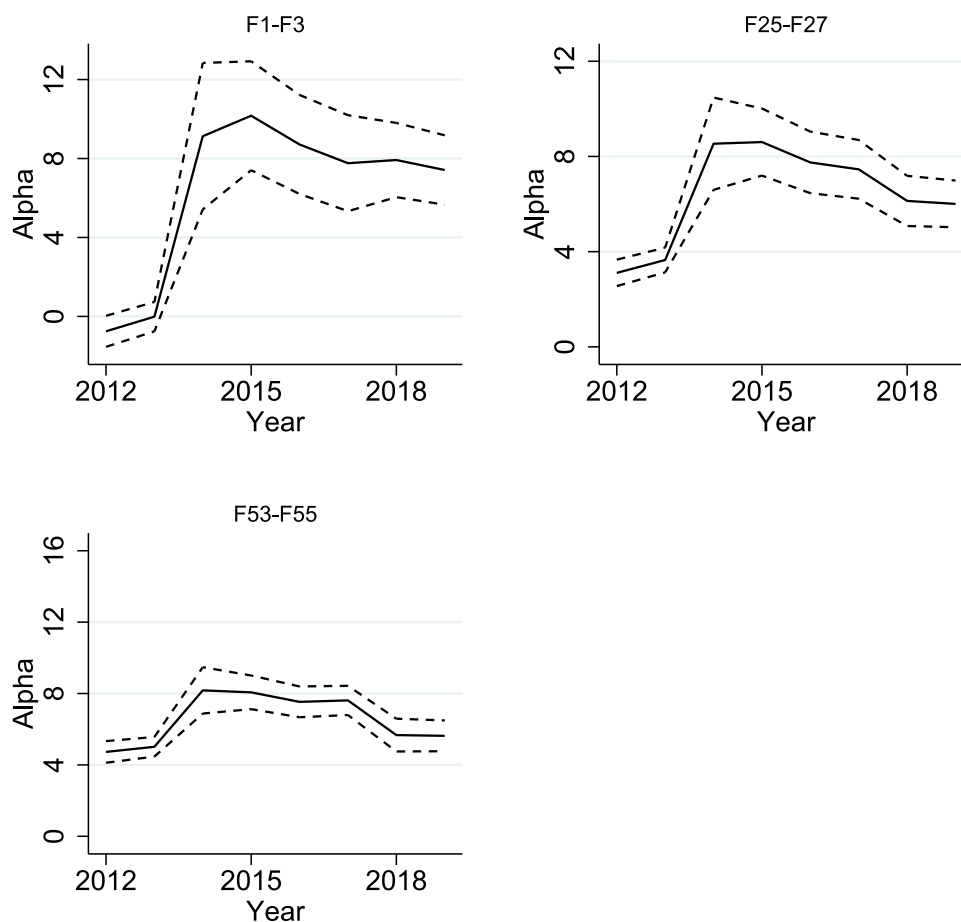


Fig. 13. Alpha for F1–F3, F25–F27 and F53–F55 contracts of crude oil for the period 2009–2019.

Table 5

Pairwise comparisons of put–call ratio (volume) for natural gas.

	Mean	Period: Jun.–Aug.	Period: Nov.–Jan.
Period: Jun.–Aug.	1.167		
Period: Nov.–Jan.	0.955	39.400 (0.000)	
Period: Other months	1.097	–0.0001 (0.000)	31.583 (0.000)

Dunn’s test (Dunn, 1964): Pairwise z-test statistics with p-values in parentheses.

the corresponding long positions.¹⁸ If one now assumes that speculators are more affected by seasonal sentiments than hedgers, then Fig. 14 can indeed contribute to explaining trading time seasonality in natural gas as observed in (1). In fact, there is literature that supports this argument: Deeney et al. (2015) report a negative loading for the put–call ratio within their WTI crude oil sentiment indicator (equations (2.1), (3.1) and Table (5)). In any case, the observation of a seasonal pattern in the put–call ratio in natural gas is highly interesting in itself. There is scope for future research in this field.

¹⁸ A short position in a call is effectively a bet on the underlying to go down and ideally stay below the strike and similarly a short position in a put is a bet for the underlying to go up and ideally stay above the strike.

6.2. ESG risk

Avramov et al. (2022) investigate how ESG rating uncertainty affects risk premia in the context of the CAPM and Cao et al. (2022) use company share options in order to reflect on ESG premia (not necessarily risk related). Preferences for good ESG ratings may well vary in a seasonal manner and in fact Pavlova et al. (2022) provide evidence for this hypothesis, but it is highly unclear how this cascades down to risk premia and more specifically to energy futures contracts. A possible avenue of investigation is to look at ESG index futures and options that have been introduced to various international exchanges since 2020. ISDA (2021) has recently published an overview about ESG related derivatives products. Investigating this further is beyond the scope of the current article and is left for future research.

6.3. Natural factors

Huisman and Kilic (2012) found that electricity futures prices in markets in which electricity is predominantly produced with perfectly storable fuels contained time varying risk premia. This applies in particular to the Nordic countries, and in fact Haugom et al. (2018) measured the effect of water inflow on the risk premium in electricity futures (among other things). While the two commodities considered in our paper are not directly affected by hydro levels, there is some evidence of cross-over. We argue that knowledge about hydro conditions is far lower during the first quarter than in the third quarter because of uncertainty about for example water content in the reservoirs and how snow melting will develop during the spring. Thus, risk-averse

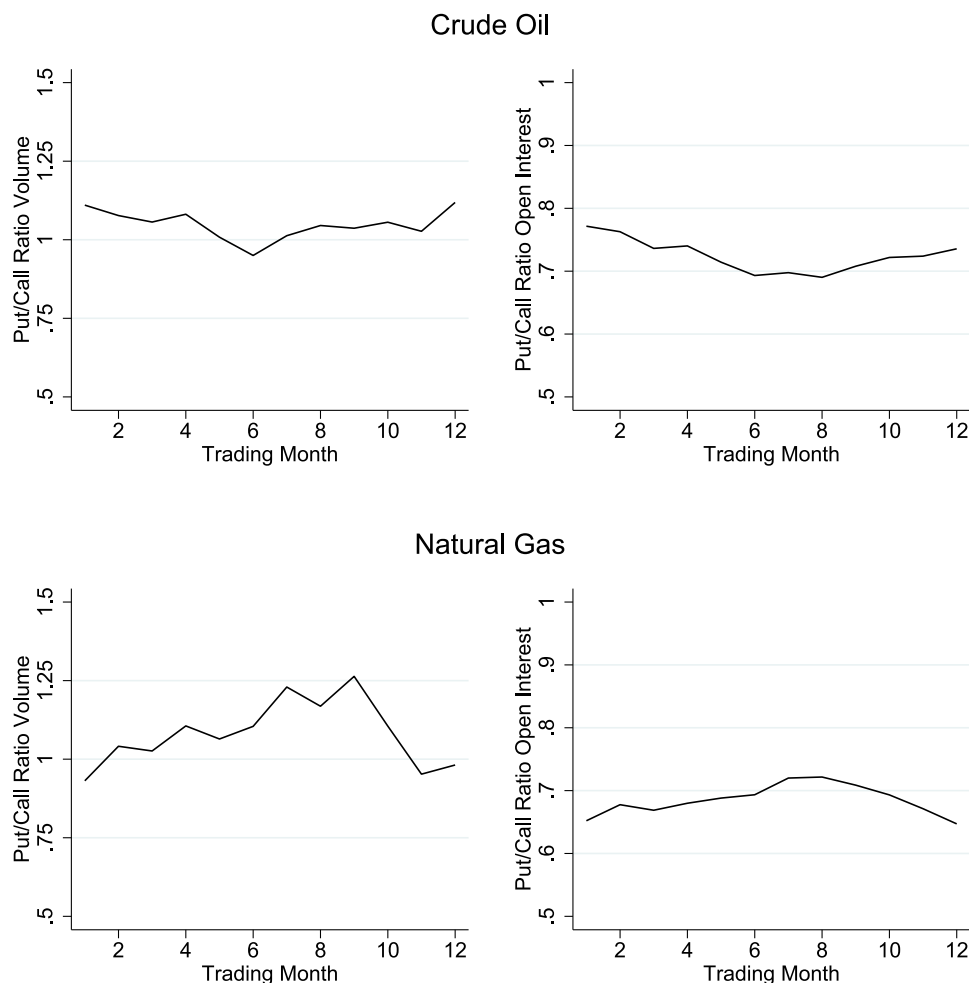


Fig. 14. Put-Call ratios for crude oil and natural gas reflecting volume (left) and open interest (right) aggregated over each month and years 2012–2021.

producers might be more inclined toward hedging during the first quarter and less so in the other quarters. This pattern corresponds with our observations of trading time seasonality in natural gas and crude oil. A detailed empirical analysis is left for future research.

7. Conclusions

In this paper, we present our findings on a new seasonal pattern reflected in futures' price paths with trading time, for futures contracts with fixed maturities. We investigate two commodities in particular, natural gas and crude oil. A preliminary study suggests that the futures prices for different trading times (in months) differ significantly during the year, but with some regularity. The concept of backward curves is introduced, which reveals that on average, the futures prices of both commodities present an obvious seasonal pattern that is both consistent with the preliminary findings and supplies further support to the possible existence of trading time seasonality. This type of seasonality is distinct from seasonal patterns in the spot prices. The latter are also reflected in futures prices, but in the corresponding forward curves; they depend on the time of maturity.

On the basis of empirical evidence, we argue that the no-arbitrage principal which lays the foundation of classic futures pricing models may be violated by potential arbitrage opportunities generated through trading time seasonality. We further emphasize our findings by constructing a relatively simple trading strategy of the “buy low sell high” type. Our results demonstrate a positive expected final payoff at the end of the operational period, with very low risk exposure. Formalizing this further, we test our trading strategy in the context of

the CAPM and we obtain positive and statistically significant alphas for both commodities. A theoretical discussion of the possible sources of both types of seasonality is provided and market sentiments linked to hedging pressure, ESG risk and natural factor are identified.

Our analysis is not without its limits. The precise source of trading time seasonality remains unknown to us, and while candidates have been identified, a more rigorous empirical analysis would go beyond the scope of the current paper and is therefore left for future research. However, we identify a seasonal pattern in the put-call ratio of natural gas by statistical means and are able to link it at least intuitively to the observed pattern of trading time seasonality. On the other hand, our theoretical discussion shows clearly how seasonal preference structures reflected in a seasonal pricing kernel can simultaneously generate seasonal patterns in the forward and backward curves, but not independent of each other. This contradicts our observation for crude oil, where there is no seasonal pattern in the forward curve but an exposed seasonal pattern in the backward curve, indicating some form of market anomaly.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Arismendi, J.C., Back, J., Prokopczuk, M., Paschke, R., Rudolf, M., 2016. Seasonal stochastic volatility: Implications for the pricing of commodity options. *J. Bank. Financ.* 66, 53–65.
- Avramov, D., Cheng, S., Lioui, A., Tarelli, A., 2022. Sustainable investing with ESG rating uncertainty. *J. Financ. Econ.* 145 (2), 642–664.
- Back, J., Prokopczuk, M., Rudolf, M., 2013. Seasonality and the valuation of commodity options. *J. Bank. Financ.* 37 (2), 273–290.
- Basu, D., Miffre, J., 2013. Capturing the risk premium of commodity futures: The role of hedging pressure. *J. Bank. Financ.* 37 (7), 2652–2664.
- Bathia, D., Bredin, D., 2013. An examination of investor sentiment effect on G7 stock market returns. *Eur. J. Finance* 19 (9), 909–937.
- Baxter, J., Conine Jr., T.E., Tamarkin, M., 1985. On commodity market risk premiums: Additional evidence. *J. Futures Mark.* 5 (1), 121–125.
- Black, F., 1976. The pricing of commodity contracts. *J. Financ. Econ.* 3 (1–2), 167–179.
- Borovkova, S., Geman, H., 2006. Seasonal and stochastic effects in commodity forward curves. *Rev. Deriv. Res.* 9 (2), 167–186.
- Cao, J., Goyal, A., Zhan, X., Zhang, W.E., 2022. Unlocking ESG premium from options. In: *Swiss Finance Institute Research Paper*, (21–39).
- Cartea, A., Figueroa, M.G., 2005. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Appl. Math. Finance* 12 (4), 313–335.
- Casassus, J., Collin-Dufresne, P., 2005. Stochastic convenience yield implied from commodity futures and interest rates. *J. Finance* 60 (5), 2283–2331.
- De Roon, F.A., Nijman, T.E., Veld, C., 2000. Hedging pressure effects in futures markets. *J. Finance* 55 (3), 1437–1456.
- Deeney, P., Cummins, M., Dowling, M., Bermingham, A., 2015. Sentiment in oil markets. *Int. Rev. Financ. Anal.* 39, 179–185.
- Dunn, O.J., 1964. Multiple comparisons using rank sums. *Technometrics* 6 (3), 241–252.
- Dusak, K., 1973. Futures trading and investor returns: An investigation of commodity market risk premiums. *J. Polit. Econ.* 81 (6), 1387–1406.
- Ewald, C., Zou, Y., 2021. Analytic formulas for futures and options for a linear quadratic jump diffusion model with seasonal stochastic volatility and convenience yield: Do fish jump? *European J. Oper. Res.* 294 (2), 801–815.
- Fama, E.F., French, K.R., 1987. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *J. Bus.* 60 (1), 55–73.
- Fama, E.F., French, K.R., 1992. Common risk factors in the returns on stocks and bonds. *J. Financ. Econ.* 33, 3–56.
- Han, Y., Hu, T., Yang, J., 2016. Are there exploitable trends in commodity futures prices? *J. Bank. Financ.* 70, 214–234.
- Haugom, E., Hoff, G.A., Molnár, P., Mortensen, M., Westgaard, S., 2018. The forward premium in the Nord Pool Power market. *Emerg. Mark. Finance Trade* 54 (8), 1793–1807.
- Hirshleifer, D., 1990. Hedging pressure and futures price movements in a general equilibrium model. *Econometrica* 58, 411–428.
- Huisman, R., Kilic, M., 2012. Electricity futures prices: Indirect storability, expectations, and risk premiums. *Energy Econ.* 34 (4), 892–898.
- ISDA, 2021. Overview of ESG-related derivatives products and transactions. ISDA - Res. Notes.
- Koekebakker, S., Lien, G., 2004. Volatility and price jumps in agricultural futures prices - evidence from wheat options. *Am. J. Agric. Econ.* 86 (4), 1018–1031.
- Kownatzki, C., Putnam, B., Yu, A., 2022. Case study of event risk management with options strangles and straddles. *Rev. Financ. Econ.* 40 (2), 150–167.
- Kramer, C., 1994. Macroeconomic seasonality and the January effect. *J. Finance* 49 (5), 1883–1891.
- Kruskal, W.H., Wallis, W.A., 1952. Use of ranks in one-criterion variance analysis. *J. Amer. Statist. Assoc.* 47 (260), 583–621.
- Lucia, J.J., Schwartz, E., 2002. Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Rev. Deriv. Res.* 5 (1), 5–50.
- Mirantes, A.G., Población, J., Serna, G., 2012. The stochastic seasonal behaviour of natural gas prices. *Eur. Financial Manag.* 18 (3), 410–443.
- Mirantes, A.G., Población, J., Serna, G., 2013. The stochastic seasonal behavior of energy commodity convenience yields. *Energy Econ.* 40, 155–166.
- Nikitopoulos, C.S., Squires, M., Thorp, S., Yeung, D., 2017. Determinants of the crude oil futures curve: Inventory, consumption and volatility. *J. Bank. Financ.* 84, 53–67.
- Pan, J., Poteshman, A.M., 2006. The information in option volume for future stock prices. *Rev. Financ. Stud.* 19 (3), 871–908.
- Pavlova, I., Whitworth, J., de Boyrie, M.E., 2022. The Sell-in-May effect in ESG indices. *Manag. Finance* (ahead-of-print).
- Putnam, B., 2020. *An Introduction to the Market Sentiment Meter: Reimagining Probability Distributions to Calibrate Event Risk*. CME Group, <https://www.cmegroup.com/education/files/an-introduction-to-the-market-sentiment-meter.pdf>.
- Schwartz, E., 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *J. Finance* 52 (3), 923–973.
- Schwartz, E., Smith, J., 2000. Short-term variations and long-term dynamics in commodity prices. *Manage. Sci.* 46 (7), 893–911.
- Shao, C., Bhar, R., Colwell, D.B., 2015. A multi-factor model with time-varying and seasonal risk premiums for the natural gas market. *Energy Econ.* 50, 207–214.
- Sørensen, C., 2002. Modeling seasonality in agricultural commodity futures. *J. Futures Mark.* 22 (5), 393–426.
- Suenaga, H., Smith, A., 2011. Volatility dynamics and seasonality in energy prices: implications for crack-spread price risk. *Energy J.* 32 (3), 27–58.