ARTICLE

Cloud cover and expected oil returns

Xianfeng Hao[1] & Yudong Wang [2✉]

Satellites can "sense" oil inventory, but cloud cover prevents observation, which reduces the flow of information into the oil market and creates uncertainty about information availability. The effects of the availability of such information on oil prices need to be thoroughly explored. Therefore, using time-series prediction, this paper examines the effects of the availability of satellite-based information on oil returns. The cloud cover above the floating roof tanks in major oil storage areas is measured to predict oil returns through regression approaches. The empirical results indicate that higher cloudiness in a week leads to lower oil returns in the following week. The predictive ability of cloudiness is significant from both in-sample and out-of-sample perspectives. The ability of cloudiness measures to predict oil returns can be explained by information uncertainty and information flow channels. The findings have important implications for asset pricing and risk management using big data.

[1] Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai, China. [2] School of Economics and Management, Nanjing University of Science and Technology, Nanjing, China. ✉email: wangyudongnj@126.com

Introduction

With the development of computer science, investors are beginning to utilize satellite images to improve their investment decisions. For example, satellite images of floating roof tanks (FRTs) are used to estimate oil inventory. Orbital Insight, a major alternative data vendor, provides satellite-based oil inventory data that help sophisticated investors formulate their trading strategies. This study investigates how the availability of satellite-based oil information affects future oil returns through time-series return prediction.

Cloud cover above FRTs prevents their observation and reduces the availability of satellite data on oil inventory. This paper uses cloud cover data over FRTs in eight major oil storage areas in the U.S. The FRTs are located using high-resolution historical maps and the You Only Look Once version 5 (YOLOv5) object detection model to generate cloud cover data. These data are collated using the Moderate Resolution Imaging Spectroradiometer (MODIS) cloud detection tool aboard NASA's Terra spacecraft[1]. The data pertaining to different storage areas are combined to obtain an aggregate cloudiness measure. Thus, eight area-specific cloudiness measures and one overall cloudiness measure are generated.

A simple predictive regression is employed to examine the ability of cloudiness to predict oil returns. During the sample period from January 2014 to December 2021, a significant and negative predictive relationship between cloudiness and oil returns is found for seven of the eight cloudiness measures. The aggregate cloudiness measure displays significant predictive ability that cannot be explained by existing predictors related to information uncertainty or oil market activity, such as oil inventory, volatility indices, or economic policy uncertainty. Out-of-sample evidence supports the ability of cloudiness to predict oil returns.

Three types of contrast tests are conducted. The first test is executed using data from February 2007 to December 2013, when no satellite-based information was available to market participants. This test examines whether such a cloudiness proxy is appropriate to measure the availability of satellite information. If the data can be used to predict returns for this period, the predictive ability may be linked to weather-related factors rather than to the availability of satellite data alone. The remaining two tests are executed using data from January 2014 to December 2021, the same period used for the main test. The second contrast test uses cloud cover above FRTs that do not store crude oil. The third test uses nighttime cloudiness above FRTs to predict oil returns. Satellite-based estimates of nighttime oil inventory are unavailable. The results of these tests imply that oil return predictability is related to satellite-based estimates of oil inventory. The results of three placebo tests show a much weaker or non-significant predictive ability of cloudiness measures for oil returns, indicating that the oil market indeed responds to satellite-based information.

A portfolio allocation framework is used to investigate the economic value of the ability to predict oil returns from satellite data. Specifically, an investor with mean–variance preferences allocates wealth between risky oil futures and risk-free Treasury bills. The weight of the risky asset in the portfolio is determined ex ante by the return forecasts. The results support the economic significance of forecasting portfolio gains from cloudiness measures. For example, under the appropriate constraints on oil futures, a portfolio constructed on the basis of forecasts using the aggregate cloudiness measure achieves annualized gains of 261.9 basis points in certainty-equivalent returns.

Cloudiness predicts oil returns through two potential channels. One is information uncertainty. The empirical results show that aggregate cloudiness has a significant and positive effect on future oil inventory. Specifically, a 1-point increase in aggregate cloudiness leads to a 0.398% increase in U.S. oil inventory in the following week. When cloud cover above the FRTs is heavy, market participants are uncertain about oil inventory. Such uncertainty can be resolved only after an Energy Information Administration (EIA) announcement. The theoretical model proposed by Gao et al. (2022) suggests that when oil-related uncertainty increases, firms rationally maintain higher oil inventories as a cushion against potentially large and adverse supply shocks. Because of the precautionary savings effect, the availability of oil as a production input decreases when aggregate demand is low and leads to low oil returns.

The second channel is information flow. A simple predictive regression of oil volatility on lagged volatility, lagged cloudiness, and their interaction term reveals that high cloudiness is significantly related to lower future oil volatility than is low cloudiness. Information releases increase volatility because investors are more likely to trade when they receive more information, and trading volume is positively related to volatility (Ederington and Lee, 1993; Kalev et al., 2004; Rangel, 2011; Bollerslev et al., 2018; Engle et al., 2020). High cloudiness means that less satellite-based information flows into the oil market, leading to low volatility. Uninformed investors are at risk of adverse selection and require a lower risk premium as compensation in such a situation (Wang, 1993). In addition, the coefficient of the interaction of lagged volatility and cloudiness indicates that the likelihood of persistently high oil volatility is significantly higher after cloudy days than it is after clear days. A possible explanation is that market participants tend to rely heavily on past information when cloudiness renders satellite information less available.

Background literature and contribution

Literature on the relationship between weather and asset prices

This paper investigates the effects of cloudiness on oil price dynamics. Thus, it is related to the literature on the influence of cloudiness on other assets' returns. Saunders (1993) finds that local weather systematically affects New York Stock Exchange prices. Hirshleifer and Shumway (2003) find that morning sunshine in the city of a country's leading stock exchange is significantly correlated with stock returns and that the mechanism is based on investors' mood. In contrast, Loughran and Schultz (2004) find little evidence that cloudy weather in the city where a firm is based affects its returns, illustrating that cloudy conditions near a firm's headquarters do not provide profitable trading opportunities. Novy-Marx (2014) uses predictive regression and finds that the weather in Manhattan, global warming, and El Niño can significantly predict portfolio returns based on anomalies. Lanfear, Lioui, and Siebert (2019) document that hurricanes making landfall in the U.S. have a strong abnormal effect on stock returns and lead to illiquidity across portfolios based on anomalies. The investor mood mechanism advocated by Hirshleifer and Shumway (2003) cannot explain empirical findings, because the hubs over which cloud cover is measured are not major oil-trading locations. New channels linking cloudiness to asset returns, namely, information uncertainty and information flow, are uncovered in this paper.

A stream of literature discusses the impacts of weather and climate change on food commodity prices because a warm climate directly reduces crop yield (e.g., Bradbear and Friel, 2013; Agnello et al., 2020). This paper is different from previous studies because the cloudiness above oil tanks has almost no influence on crude oil supply or demand fundamentals. Alternatively, it finds the channels of information uncertainty and information flow to explain the effects of cloudiness on oil returns. Some studies

investigate climate change as a factor affecting asset prices (e.g., Hong et al., 2019; Choi, Gao, and Jiang, 2020; Engle et al., 2020; Krueger, Sautner, and Starks, 2020). This paper is a departure from such research because it considers cloudiness to be the key factor preventing the flow of satellite information. Moreover, although other studies focus on cross-sectional stock prices, this study focuses on time-series prediction.

Literature on the relationship between information uncertainty and asset returns. Cloud cover creates uncertainty about oil inventory. According to information uncertainty theory, investors become overconfident when uncertainty is high (Zhang, 2006a; 2006b). The behavioral finance literature argues that over-confident investors overreact to private news, causing future price reversals (e.g., Daniel, Hirshleifer, and Subrahmanyam, 1998). Under a short-sale constraint, this situation is more likely when investors receive good news than it is when they receive bad news. Therefore, information uncertainty is negatively related to future returns. High uncertainty leads to more diverse investor opinions than does low uncertainty, in turn leading to low stock returns, as pessimistic valuations cannot be easily expressed under short-sale constraints (Deither, Malloy, and Scherbina, 2002; Yu, 2011). Cross-sectional research validates the negative relationship between information uncertainty (or disagreement) via a mechanism predicated on the key assumption of a short-sale constraint (Miller, 1977)[2]. However, such an assumption does not apply to commodity futures markets because short sales in futures transactions are not constrained.

This study is related to information uncertainty, but it considers different explanations for the effects of uncertainty on the crude oil and stock markets. The effect of satellite information uncertainty on oil returns is explained by inventory. For a storable commodity, inventory is an important tool for hedging against information uncertainty. This study contributes to the literature by revealing the statistical and economic significance of predicting asset returns from cloudiness measures through the use of a time-series forecasting framework.

Literature on the application of satellite images to asset pricing. Satellite data are increasingly being used by research in economics and finance to, for example, monitor aggregate economic activity (e.g., Henderson, Storeygard, and Weil, 2012) and predict asset prices. For example, using parking lot vehicle counts extracted from satellite imagery, Katona et al. (2018) discover that unequal access to satellite data increases information asymmetry among market participants. Using satellite-based estimates of the normalized number of cars in retail parking lots, Zhu (2019) finds that aggregate signals from such datasets can predict unannounced earnings. The introduction of alternative data improves price efficiency by reducing the cost of information acquisition. Mukherjee, Panayotov, and Shon (2021) show that introducing satellite estimates leads to a smoother response of oil prices to government inventory announcements, implying that satellite-based estimates can serve as an alternative to government data. The findings of Mukherjee, Panayotov, and Shon (2021) are corroborated by this paper, which reveals that cloud cover above oil tanks affects future oil returns. As an extension, a cloudiness measure is constructed to predict oil returns, broadening the application of satellite data in asset pricing studies.

Data
Satellite imagery, cloudiness, and oil inventory
Satellite images for oil market analysis. Satellite-based oil inventory data are increasingly used for oil market analysis. Optical satellites can "sense" oil inventory by observing the shadows of FRTs, which are used for storing crude oil. The floating roofs rise and fall with the level of liquid in the tanks to reduce the risk of product loss and increase safety. A higher floating roof indicates more oil in a tank.

On sunny days, each FRT has two shadows: an outer shadow projected by the tank wall onto the ground and an inner shadow projected by the tank wall onto the floating roof. From the solar altitude, one can compute the height of oil in an FRT from the difference between the areas of these two shadows and estimate the oil inventory accordingly. However, this methodology is not suitable for fixed roof tanks (FXRTs) because they do not have an inner shadow.

When satellite images of FRTs are available, market investors can estimate oil inventory without waiting for the EIA's announcements. However, cloud cover above FRTs prevents satellite observation. Therefore, the response of oil prices to government announcements is stronger on cloudier days because of the lack of satellite images and the heavy reliance of investors on such announcements (Mukherjee, Panayotov, and Shon, 2021).

Cloudiness measures. FRTs must be located before the cloudiness above them can be measured. Oil stored in the U.S. is distributed across five Petroleum Administration for Defense Districts (PADDs)[3]. Of these five PADDs, the Cushing (PADD2) and Gulf Coast (PADD3) PADDs account for more than 70% of U.S. oil inventory because of the strong demand for storage in their locations. This paper only considers the oil tanks in these two PADDs to balance the costs of requesting satellite imagery and the value of the information therein.

All oil tanks in the Cushing and the Gulf Coast PADDs are identified from map imagery by a YOLOv5 model, a mainstream deep learning method. The identified tanks are located in eight storage areas: Cushing, Oklahoma (COK); Houston, Texas (HTX); Texas City, Texas (TTX); Beaumont, Texas (BTX); Corpus Christi, Texas (CTX); Lake Charles, Louisiana (LLA); Baton Rouge, Louisiana (BLA); and the Mississippi River in Louisiana (MRL).

The distribution of FRTs within a storage area can change over time as new tanks are built, and old tanks are abandoned. Hence, yearly high spatial resolution maps from ArcGIS Online from 2014 to 2021 are used to update the geographic coordinates of the oil tanks for each year.

Cloud cover is measured using a cloud mask product from MODIS, a NASA tool utilizing two satellites, the first of which was launched in 1999. MODIS provides daily high-precision images of up to 250 m spatial resolution. This cloud mask product is widely used for cloud cover observation. The MODIS cloud mask data are collected from NASA's Level 1 and Atmosphere Archive and Distribution System. MODIS is installed on two satellites––Terra and Aqua. Following Zhang and Wang (2010), Hu (2011), and Purbantoro et al. (2019), Terra data are used because they are less affected by the sun's glint than Aqua data.

The satellite-based cloud mask data are analyzed as confidence levels calculated from a combination of clear-sky confidence levels from all tests performed and saved as two binary digits indicating longitude and latitude. The symbol "00" denotes a confident cloudy day (confidence ≤0.66), "01" denotes a probably cloudy day (0.95≥ confidence > 0.66), "10" denotes a probably clear day (0.99≥ confidence > 0.65), and "11" denotes a confident clear day (confidence > 0.99). Cloud mask values are extracted from the raw data using the script provided by NASA in the *MODIS Cloud Mask User's Guide*. Cloudiness symbols are arbitrarily transformed into numerals from 0 to 3. Specifically, 0, 1, 2, and 3 represent the confident clear, probably clear, probably cloudy, and confident cloudy days, respectively. The

cloudiness in each storage area is analyzed as the weighted average over all FRTs, where the weight is the footprint of each tank[4].

Oil market data

Oil price data. The price data analyzed are those of West Texas Intermediate (WTI) crude oil. Daily data are collected from the EIA website[5]. Weekly returns from the difference between the log price on the current Tuesday and the Tuesday of the previous week are used to isolate the effect of inventory announcements in the predictive models. This weekly return is equivalent to the cumulative sum of daily returns from Wednesday of the previous week to the current Tuesday. The sample period used in the main empirical analysis is January 2014 to December 2021 considering the availability of satellite-based oil information[6].

Control variables. Six economic variables are considered as competing predictors of oil returns. Their lagged terms are incorporated into the predictive regression as control variables to investigate the in-sample predictive ability of the proposed cloudiness measures. These six variables are the economic policy uncertainty (EPU) from Baker et al. (2016), the CBOE stock volatility index (VIX), the macroeconomic uncertainty (MU) from Bekaert et al. (2022), the unexpected oil inventory change (OI) from Mukherjee et al. (2021), WTI oil returns (OR), and asymmetric oil returns (AR) proxied by the positive component of OR.

The descriptive statistics of the cloudiness and economic variables and their pairwise Pearson correlations are presented in the online appendix. The magnitudes of all of the correlations are less than 0.1, suggesting weak correlations. These findings suggest that the proposed cloudiness measures and existing oil predictors provide markedly different information.

Methodology

In-sample test. A simple predictive regression is used to examine the ability of cloudiness above FRTs to predict oil returns:

$$r_{t+1} = \alpha^{(i)} + \beta_K^{(i)} K_t^{(i)} + \varepsilon_{t+1}^{(i)}, \tag{1}$$

where r_t denotes oil returns, $K_t^{(i)}$ is the measure of cloudiness above the ith oil storage area, $K_t^{(A)}$ denotes the average cloudiness across the eight considered areas $\left(\text{i.e., } K_t^{(A)} = \frac{1}{8} \sum_{i=1}^{8} K_t^{(i)}\right)$, and $\varepsilon_t^{(i)}$ is the error term. The parameter estimates of Model (1) are obtained using the ordinary least squares (OLS) method. Because the relationship between cloudiness and oil returns is expected to be negative, a one-sided test is conducted to test the null hypothesis of $H_0 : \beta_K^{(i)} \geq 0$ against the one-sided alternative of $H_1 : \beta_K^{(i)} < 0$. The Newey and West (1987) adjusted t statistic, which is robust to heteroskedasticity and autocorrelation, is employed to test for statistical significance.

It is necessary to investigate whether the predictive information offered by cloudiness measures overlaps that offered by existing oil return and uncertainty variables. Therefore, the control variables are incorporated to form a predictive regression as follows:

$$r_{t+1} = \alpha^{(i)} + \beta_K^{(i)} K_t^{(i)} + \sum_{j=1}^{N} \beta_{control}^{(j)} x_{t,control}^{(j)} + \varepsilon_{t+1}^{(i)}, \tag{2}$$

where $x_{t,control}^{(j)}$ is the jth control variable and N is the total number of control variables.

Out-of-sample test. As shown in Welch, Goyal (2008), it is difficult to find a predictive model that significantly outperforms the

benchmark of a historical average model assuming no predictability in an out-of-sample exercise. Therefore, to analyze out-of-sample performance, a univariate predictive regression employing a variable of interest other than the intercept is used:

$$r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}, \tag{3}$$

where r_{t+1} denotes weekly returns, which are the average daily returns from Wednesday of week t (EIA announcement day) to Tuesday of week $t + 1$. The return forecasts from the univariate predictive regression are computed as follows:

$$\hat{r}_{t+1} = \hat{\alpha}_t + \hat{\beta}_t x_t, \tag{4}$$

where $\hat{\alpha}_t$ and $\hat{\beta}_t$ denote the parameter estimates of α and β based on the information available before week t, respectively. Specifically, $\hat{\alpha}_t$ and $\hat{\beta}_t$ can be obtained by regressing $\{r_i\}_2^t$ on the constant and corresponding predictor $\{x_i\}_1^{t-1}$ using OLS. Then, the expected returns in week $t + 1$ can be obtained from Eq. (4) using $\hat{\alpha}_t$, $\hat{\beta}_t$, and x_t.

Forecast accuracy is evaluated using the out-of-sample R^2 value $\left(R_{OoS}^2\right)$ following Campbell and Thompson (2008). R_{OoS}^2 indicates the proportional reduction in the mean squared prediction error of a model relative to that of a historical average benchmark and is calculated as follows:

$$R_{OoS}^2 = 1 - \frac{\sum_{t=M+1}^{T} \left(r_t - \hat{r}_t\right)^2}{\sum_{t=M+1}^{T} \left(r_t - \bar{r}_t\right)^2}, \tag{5}$$

where $\hat{r}_t$ and r_t are the return forecast from the given model and realized return at time t, respectively; $\bar{r}_t$ denotes the historical average forecast; T denotes the total number of return observations; and M denotes the initial sample length for parameter estimation. A positive R_{OoS}^2 value implies that a model has greater predictive ability than the historical average benchmark. The Clark and West (CW, 2007) statistic is used to test whether R_{OoS}^2 is significantly different from 0.

The forecast encompassing test of Rapach, Ringgenberg, and Zhou (2016) is executed to investigate the uniqueness of the information provided by cloudiness measures. The test is based on the optimal combination of a predictive regression forecast and a benchmark forecast:

$$\hat{r}_{t+1}^* = \lambda \hat{r}_{t+1}^i + (1 - \lambda)\hat{r}_{t+1}^{bench}, \tag{6}$$

where $\hat{r}_{t+1}^i$ denotes the return forecast produced by the ith predictive regression model of interest and $\hat{r}_t^{bench}$ is the return forecast generated by the benchmark model. The parameter λ is the weight assigned to the forecast in the model of interest, where $0 \leq \lambda \leq 1$. If $\lambda = 1$, the optimal combination includes the return forecasts from only the model of interest, indicating that the information contained therein encompasses the information of the forecasts from the benchmark model. If $\lambda > 0$, the optimal combination considers the forecasts of the given model and the benchmark model. In this case, the forecast of the benchmark model does not encompass the forecast of the model of interest, which thus provides useful information for forecasting excess returns beyond that provided by the benchmark model. To test the significance of the weight estimate $\hat{\lambda}$, the Harvey, Leybourne, and Newbold (HLN, 1998) approach with the test of the null hypothesis of $\lambda = 0$ against the one-sided alternative $\lambda > 0$ is used.

Empirical results

In-sample predictive ability. Table 1 reports the results of Model (1), including the parameter estimates and in-sample R^2, interpreted as percentages. The period of weekly sample used in the in-sample analysis is from January 7, 2014, to December 21, 2021.

Table 1 In-sample predictive regression results, univariate models.

	$K^{(BLA)}$	$K^{(BTX)}$	$K^{(CTX)}$	$K^{(COK)}$	$K^{(HTX)}$	$K^{(MRL)}$	$K^{(LLA)}$	$K^{(TTX)}$	$K^{(A)}$
α	0.011	0.059***	0.059*	0.028*	0.104***	0.059***	0.081**	0.059*	0.113***
	(0.561)	(2.840)	(1.363)	(1.558)	(2.386)	(2.402)	(2.087)	(1.435)	(2.581)
β_K	−0.004	−0.021***	−0.021*	−0.010*	−0.036***	−0.021***	−0.028**	−0.021*	−0.040***
	(−0.553)	(2.763)	(−1.409)	(−1.573)	(−2.447)	(−2.500)	(−2.168)	(−1.503)	(−2.611)
R^2(%)	0.031	0.633	0.522	0.457	1.120	0.866	1.190	0.714	1.275

This table reports the in-sample estimation results for the predictive regression $r_{t+1} = \alpha^{(i)} + \beta_K^{(i)} K_t^{(i)} + \varepsilon_{t+1}^{(i)}$, where r_t denotes oil price returns, $K_t^{(i)}$ is a cloudiness indicator above the i-th oil storage district, $K_t^{(A)}$ denotes the aggregate cloudiness measure, and $\varepsilon^{(i)}$ represents the error term. The parameters are estimated via the ordinary least squares. The one-sided test is used for the null hypothesis $H_0 : \beta_K^{(i)} \geq 0$ against one-sided alternative $H_1 : \beta_K^{(i)} < 0$. The numbers in the parentheses are the t-statistics based on the Newey-West robust errors. The in-sample R^2 values are multiplied by 100 to interpret as percentages. The asterisks *, **, and *** denote statistical significance at 10%, 5%, and 1% levels, respectively.

Table 2 In-sample predictive regression results, control variables.

	$K^{(BLA)}$	$K^{(BTX)}$	$K^{(CTX)}$	$K^{(COK)}$	$K^{(HTX)}$	$K^{(MRL)}$	$K^{(LLA)}$	$K^{(TTX)}$	$K^{(A)}$
α	0.013	0.063***	0.043	0.025*	0.090***	0.057***	0.080***	0.061*	0.108***
	(0.681)	(2.935)	(1.102)	(1.463)	(2.478)	(2.590)	(2.656)	(1.570)	(2.731)
β_{EPU}	0.017	0.017	0.017	0.017	0.017	0.018	0.017	0.017	0.017
	(1.285)	(1.316)	(1.284)	(1.274)	(1.293)	(1.336)	(1.329)	(1.320)	(1.324)
β_{VIX}	−0.013	−0.014	−0.014	−0.013	−0.014	−0.014	−0.015	−0.014	−0.014
	(−1.033)	(−1.112)	(−1.053)	(−1.036)	(−1.069)	(−1.079)	(−1.135)	(−1.062)	(−1.106)
β_{MU}	−0.009	−0.009	−0.009	−0.009	−0.010	−0.009	−0.009	−0.010	−0.009
	(−0.652)	(−0.649)	(−0.636)	(−0.649)	(−0.666)	(−0.641)	(−0.641)	(−0.664)	(−0.647)
β_{OI}	−0.006	−0.005	−0.005	−0.005	−0.005	−0.006	−0.006	−0.006	−0.005
	(−1.031)	(−0.981)	(−0.975)	(−0.990)	(−0.959)	(−1.005)	(−1.022)	(−0.990)	(−0.961)
β_{OR}	−0.012	−0.013	−0.012	−0.012	−0.012	−0.012	−0.012	−0.012	−0.012
	(−0.743)	(−0.784)	(−0.737)	(−0.749)	(−0.774)	(−0.736)	(−0.765)	(−0.759)	(−0.777)
β_{AR}	0.022	0.022	0.022	0.022	0.022	0.021	0.022	0.022	0.022
	(1.144)	(1.165)	(1.133)	(1.153)	(1.143)	(1.115)	(1.134)	(1.155)	(1.148)
β_K	−0.005	−0.022***	−0.015	−0.010*	−0.031***	−0.020***	−0.028***	−0.021*	−0.038***
	(−0.683)	(−2.857)	(−1.143)	(−1.489)	(−2.500)	(−2.699)	(−2.689)	(−1.630)	(−2.748)
R^2(%)	7.598	8.276	7.820	7.929	8.391	8.355	8.690	8.314	8.700

This table reports the in-sample estimation results for the predictive regression, $r_{t+1} = \alpha^{(i)} + \beta_K^{(i)} K_t^{(i)} + \sum_{j=1}^{N} \beta_{control}^{(j)} x_{t,control}^{(j)} + \varepsilon_{t+1}^{(i)}$, where where r_t denotes oil price returns, $K_t^{(i)}$ is a cloudiness indicator above the i-th oil storage district, $K_t^{(A)}$ denotes the aggregate cloudiness measure, and $\varepsilon^{(i)}$ represents the error term, $x_{t,control}^{(j)}$ is the j-th economic variable, and N is the total number of such control variables. The parameters are estimated via the ordinary least squares. The one-sided test is used for the null hypothesis $H_0 : \beta_K^{(i)} \geq 0$ against one-sided alternative $H_1 : \beta_K^{(i)} < 0$. The numbers in the parentheses are the t-statistics based on the Newey-West robust errors. The in-sample R^2 values are multiplied by 100 to interpret as percentages. The asterisks *, **, and *** denote statistical significance at 10%, 5%, and 1% levels, respectively.

The coefficients are significant and negative for seven of the eight cloudiness measures related to oil storage areas. Notably, the aggregate cloudiness measure $K_t^{(A)}$ exhibits strong predictive ability in the negative direction for oil returns, as evidenced by its t statistic of −2.611, which is statistically significant at the 1% level. The in-sample R^2 of the univariate $K_t^{(A)}$ model is 1.275%.

Table 2 shows the results of Model (2). A significant and negative predictive relationship is evident for six of the eight cloudiness measures. That is, the predictive ability of the individual cloudiness measures remains almost unchanged after the incorporation of the control variables. The coefficient of $K_t^{(A)}$ is −0.038, with a t statistic of −2.748, which is statistically significant at the 1% level. Thus, for predicting oil returns, these cloudiness measures provide information markedly distinct from that provided by the control variables.

When investigating in-sample predictability, it is important to determine the statistical significance of the predictive regression coefficient of interest. The heteroskedasticity and autocorrelation affect the statistical inference about the significance of parameter estimates. The OLS residuals of the univariate $K_t^{(A)}$ model have a significant autocorrelation of 0.105 and a significant autoregressive conditional heteroskedasticity (ARCH) effect, according to the Engle (1982) test. The Newey and West adjusted t statistic used in the previous analysis accounts for the effects of autocorrelation and heteroskedasticity of oil returns on statistical

inferences of significance. To check the robustness of the statistical inferences, four tests are employed. The first uses p values based on the wild bootstrap procedure of Rapach, Ringgenberg, and Zhou (2016) to account for the nonnormal distribution. From 2,000 simulations, the p value of the $K_t^{(A)}$ model is 0.003. The second test uses the approach of Stambaugh (1999) for addressing autocorrelation. The adjusted slope coefficient of the $K_t^{(A)}$ series is −0.040, with a p value of 0.011. The third test is the Wald test of Kostakis, Magdalinos, and Stamatogiannis (2015), which accommodates the regressor's degree of persistence (unit root, local-to-unit root, near stationary, or stationary). The resulting slope coefficient for the IVX-Wald statistic is −0.039, with a p value of 0.024. The final test employs the weighted least squares (WLS) following Johnson (2019), in which the weights are the standard deviations of oil returns computed on the basis of a 50-week rolling window. WLS estimators are more efficient than OLS estimators for heteroskedastic data. The t statistic corresponding to $K_t^{(A)}$ is −2.581, suggesting the significance of predictive ability. The results of these tests indicate the robustness of the main findings regarding the ability of cloudiness measures to predict oil returns.

Placebo tests are conducted to further understand the predictive ability of cloudiness. These tests are executed on the aggregate cloudiness measure $K_t^{(A)}$ because it reflects the overall availability of satellite information on oil inventory. The first test

Table 3 The placebo test results.

	Earlier period results (2007M2-2013M12)	Results based on fixed roof tanks	Results based on Nighttime cloudiness
α	−0.012	0.137**	0.015
	(−0.517)	(1.653)	(0.560)
$\beta_K^{(A)}$	0.005	−0.047**	−0.006
	(0.527)	(−1.694)	(−0.543)
$R^2(\%)$	0.066	1.910	0.046

This table reports the in-sample estimation results for the predictive regression $r_{t+1} = \alpha^{(i)} + \beta_K^{(A)} K_t^{(A)} + \varepsilon_{t+1}$, where r_t denotes oil price returns, $K_t^{(A)}$ denotes the aggregate cloudiness measure, and $\varepsilon_t^{(i)}$ represents the error term. The second column reports the estimation results during the earlier period when the application of satellite data in investing is not prevalent (2007M2-2013M12). The third column shows the results using $K_t^{(A)}$ based on fixed roof tanks. The last column shows the results using $K_t^{(A)}$ based on nighttime cloudiness. The parameters are estimated via the ordinary least squares. The one-sided test is used for the null hypothesis $H_0 : \beta_K^{(A)} \geq 0$ against one-sided alternative $H_1 : \beta_K^{(A)} < 0$. The numbers in the parentheses are the t-statistics based on the Newey-West robust errors. The in-sample R^2 values are multiplied by 100 to interpret as percentages. The asterisks *, **, and *** denote statistical significance at 10%, 5%, and 1% levels, respectively.

Table 4 Out-of-sample results.

	Out-of-sample R^2		Forecast encompassing test Benchmark: historical average		Forecast encompassing test Benchmark: $K^{(A)}$ model	
	R_{OoS}^2	CW	λ	MHLN	λ	MHLN
EPU	−8.964	−0.796	0	−0.795	0	−0.787
VIX	−3.711	−0.097	0	−0.097	0	−0.105
MU	−6.054	−1.045	0	−1.044	0	−1.021
OI	−0.586	−0.089	0	−0.088	0	−0.077
OR	−6.304	−0.587	0	−0.586	0	−0.607
AR	−5.103	0.903	0.118	0.902	0.089	0.761
$K^{(A)}$	0.889	2.185**	1.000	2.182**	NA	NA

This table reports the out-of-sample forecasting results of the univariate predictive regressions, $r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}$, for, where x_t is a weekly variable of interest, r_{t+1} represents the weekly returns on WTI spot oil, and ε_{t+1} is the error term assumed to be independent and identically distributed. $K^{(A)}$ is the aggregate cloudiness measure, and the names of the economic variables are listed in Table 1. The forecast quality is evaluated by the out-of-sample R^2 (R_{OoS}^2) defined as the percent reduction of MSPE of the given model relative to the MSPE of the historical average benchmark, $R_{OoS}^2 = 1 - MSPE_{model}/MSPE_{bench}$, where $MSPE_{model} = \sum_{t=M+1}^{T}(r_t - \hat{r}_t)^2/(T - M)$ and $MSPE_{bench} = \sum_{t=M+1}^{T}(r_t - \bar{r}_t)^2/(T - M)$, where r_t, $\hat{r}_t$ and $\bar{r}_t$ are the realized returns and return forecasts recursively generated from the given model and the benchmark model, respectively. T is the total number of samples, and M denotes the initial in-sample length for parameter estimation. The full sample period is from January 2014 to December 2021 and the out-of-sample evaluation period starts from January 2015. The R_{OoS}^2 figure is multiplied by 100 to denote percentage values. The statistical significance relative to the prevailing mean model is examined using the Clark and West (2007) test statistic reported in the third column. The forth column reports the estimated weight λ on the predictive regression of interest in a convex combination forecast, composed of the forecasts from the corresponding univariate model and historical average forecasts. The seventh column reports the estimated weight on weight λ on the predictive regression of interest in a convex combination forecast, composed of the forecasts from the given univariate model and the $K^{(A)}$ model forecasts. The statistical significance is reported based on the Harvey, Leybourne, and Newbold (1998) statistics (MHLN). The null hypothesis of the HLN test is that the weight on the given model forecasts is zero against the alternative hypothesis that the weight is larger than zero. The asterisks *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively.

considers the sample period from February 2007 to December 2013, during which no satellite-based data were available to investors. The results in Table 3 suggest that $K_t^{(A)}$ has no predictive ability for this period. The second test investigates the predictive ability of $K_t^{(A)}$ based on the cloudiness above FXRTs. In this test, the sample period is the same as that used for our main test, namely, January 2014 to December 2021. Observations of FXRTs do not provide any oil storage–related information. Thus, the differences in the cloud cover of FXRTs provide compelling evidence for the necessity for the cloudiness above FRTs in oil return prediction. Cloudiness above FXRTs displays significant predictive ability, but its predictive power is weaker than that of cloudiness above FRTs, as evidenced by the weak significance of the slope coefficients. This is expected because FRTs and FXRTs are stored adjacently, and the cloudiness above them is similar. The third test uses the cloudiness above FRTs at night, when tank shadows are not visible, during the period from January 2014 to December 2021. The results again indicate that the predictive ability of cloudiness disappears, suggesting the importance of daytime observations. In summary, these three contrast test results confirm that daily cloud cover over oil storage facilities provides more predictive information than does cloud cover over other locations.

Out-of-sample predictive ability

Forecasting results. The data for the first 50 weeks (i.e., from January 7, 2014 to December 16, 2014) are used for the initial parameter estimation to generate first forecast. Thus, weekly data from December 23, 2014 through December 21, 2021 are employed for forecast evaluation. Table 4 shows the out-of-sample performance of $K_t^{(A)}$ and that of competing predictors. The $K_t^{(A)}$ model yields a R_{OoS}^2 of 0.889%, with a CW statistic of 2.185, which is statistically significant at the 5% level. In contrast, little evidence of predictive ability is observed among the competing predictors.

Table 4 also reports the results of the forecast encompassing test. When using the historical average benchmark, the weight of the $K_t^{(A)}$ model is 1, whereas the weights of the five competing models are 0, and the weight of the AR model is 0.118. Furthermore, the HLN statistics indicate that the estimate of λ for the $K_t^{(A)}$ model is significantly different from 0, whereas the λ coefficients for the remaining models are nonsignificant. When using the $K_t^{(A)}$ model as the benchmark, the weights of all of the competing models are 0 or close to 0.

Therefore, the forecasts from the $K_t^{(A)}$ model encompass the forecasts from the benchmark model and the model based on existing predictors.

Alternative multivariate analysis methods. The proposed $K_t^{(A)}$ indicator is a simple aggregate of the eight cloudiness measures. However, two alternative methods for dealing with multivariate information are considered: forecast combinations and the diffusion index method. This analysis is helpful for further understanding the predictive ability of satellite image information and serves as a robustness test of our findings on return predictability. Forecast combinations use the weighted average of forecasts from individual models, with the key inputs being the model weights. Following Rapach, Strauss, and Zhou (2010), a simple equally weighted combination is employed, which uses weights of $1/N$, where N is the number of models. The diffusion index method uses several factors drawn from a pool of predictors. For simplicity, the first principal component of the eight cloudiness variables is used as the predictor.

The out-of-sample results obtained using these two multivariate methods are significant and positive, with R_{OoS}^2 values of 0.821% and 0.751%, respectively. The forecast encompassing test results suggest that the combination and diffusion index forecasts provide content that is markedly different from that of the benchmark forecast. The quality of our finding on the return predictability of cloudiness information is thus not sensitive to the multivariate method employed.

Long-horizon prediction. The above analyses demonstrate the ability of $K_t^{(i)}$ to predict oil returns over a week. The predictive ability for longer horizons is investigated using the following

Table 5 Forecasting oil futures returns.

	$K^{(BLA)}$	$K^{(BTX)}$	$K^{(CTX)}$	$K^{(COK)}$	$K^{(HTX)}$	$K^{(MRL)}$	$K^{(LLA)}$	$K^{(TTX)}$	$K^{(A)}$
Panel A: In-sample predictive regression results									
α	0.016	0.063***	0.062	0.024	0.102**	0.066***	0.089**	0.058	0.117***
	(0.838)	(2.938)	(1.411)	(1.473)	(2.385)	(2.618)	(2.233)	(1.411)	(2.683)
β_K	−0.006	−0.022***	−0.022*	−0.009	−0.035***	−0.023***	−0.031**	−0.021*	−0.041***
	(−0.829)	(−2.897)	(−1.464)	(−1.468)	(−2.454)	(−2.740)	(−2.323)	(−1.483)	(−2.727)
R^2(%)	0.077	0.818	0.634	0.377	1.226	1.239	1.604	0.797	1.540
Panel B: Out-of-sample forecasting results									
R^2_{OoS}	−0.244	0.457	−0.008	−0.152	0.836	0.900	1.084	0.010	1.108
CW	0.060	2.005**	0.743	0.431	2.095**	2.205**	2.120**	1.049	2.326**
λ	0.038	0.896	0.493	0.294	0.941	0.848	1	0.507	0.959
MHLN	0.060	2.002**	0.742	0.431	2.092**	2.202**	2.117**	1.047	2.323**

This table reports the in-sample estimation results for the predictive regression $r_{t+1} = \alpha^{(i)} + \beta_K^{(i)} K_t^{(i)} + \varepsilon_{t+1}^{(i)}$, where r_t denotes oil price returns, $K_t^{(i)}$ is a cloudiness indicator above the i-th oil storage district, $K_t^{(A)}$ denotes the aggregate cloudiness measure, and $\varepsilon_t^{(i)}$ represents the error term. The parameters are estimated via the ordinary least squares. The one-sided test is used for the null hypothesis $H_0 : \beta_K^{(i)} \geq 0$ against one-sided alternative $H_1 : \beta_K^{(i)} < 0$. The numbers in the parentheses are the t-statistics based on the Newey-West robust errors. The in-sample R^2 values are multiplied by 100 to interpret as percentages. The out-of-sample forecast quality is evaluated by the out-of-sample R^2 (R^2_{OoS}) defined as the percent reduction of MSPE of the given model relative to the MSPE of the historical average benchmark, $R^2_{OoS} = 1 - MSPE_{model}/MSPE_{bench}$, where $MSPE_{model} = \sum_{t=M+1}^{T} (r_t - \hat{r}_t)^2 /(T-M)$ and $MSPE_{bench} = \sum_{t=M+1}^{T} (r_t - \bar{r}_t)^2 /(T-M)$, where r_t, $\hat{r}_t$ and $\bar{r}_t$ are the realized returns and return forecasts recursively generated from the given model and the benchmark model, respectively. T is the total number of samples, and M denotes the initial in-sample length for parameter estimation. The full sample period is from January 2014 to December 2021 and the out-of-sample evaluation period starts from January 2015. The R^2_{OoS} figure is multiplied by 100 to denote percentage values. The statistical significance relative to the prevailing mean model is examined using the Clark and West (2007) test statistic reported in the third column. The λ is the weight on the predictive regression of interest in a convex combination forecast, composed of the forecasts from the corresponding univariate model and historical average forecasts. The statistical significance is reported based on the Harvey, Leybourne, and Newbold (1998) statistics (MHLN). The null hypothesis of the MHLN test is that the weight on the given model forecasts is zero against the alternative hypothesis that the weight is larger than zero. The asterisks *, **, and *** denote significance at 10%, 5%, and 1% levels, respectively.

regression:

$$r_{t+1:t+h} = \alpha + \beta_K^{(A)} K_t^{(A)} + \varepsilon_{t+1}, \qquad (7)$$

where $r_{t+1:t+h}$ denotes the average returns from week $t+1$ to week $t+h$. Forecast horizons of 2 to 6 weeks are considered. The results indicate that the slope coefficient $\beta_K^{(A)}$ is significant and negative for 2- and 3-week horizons. For horizons longer than 3 weeks, the predictive ability of $K_t^{(A)}$ disappears[7]. The reason for this may be that the uncertainty of information due to cloud cover resolves within a short period, especially given that satellite images from any clear day allow for the quick release of inventory information.

Alternative benchmark model. Historical average forecasts generated through predictive regression based only on a constant are the most popular benchmark for testing forecasts of asset returns. Such historical average models assume constant excess returns and therefore no predictability. This benchmark is also consistent with the weak-form efficient market hypothesis of Fama (1970). The autocorrelation of oil spot returns used in the main empirical analysis is 0.110, which has a small magnitude but is statistically significant at the 5% level. This significant autocorrelation does not imply that the incorporation of a lagged dependent variable can improve out-of-sample forecasting performance, as the inclusion of autoregressive terms accounts for more past information but introduces parameter estimation error, harming out-of-sample performance. Nevertheless, a lagged dependent variable is useful for inferring the predictive ability of cloudiness measures relative to a benchmark model. This section compares the forecasting performance of a VAR(1,1) model with lagged returns and the aggregate cloudiness measure $K_t^{(A)}$ with that of an AR(1) model with only lagged returns. The in-sample estimation results of the VAR model indicate that the coefficient of $K_t^{(A)}$ is −0.038, with a t statistic of −2.605 significant at the 1% level. Using out-of-sample data, the mean squared prediction error of the VAR model is 0.837% lower than that of the AR model (i.e., $R^2_{OoS} = 0.837\%$), with a CW statistic of 2.294% significant at the 5% level. Thus, changing the benchmark model does not change the main finding of the significant predictive ability of the cloudiness measure, considering the R^2_{OoS} of 0.889% of the univariate $K_t^{(A)}$ model.

Forecasting futures returns for portfolio allocation. Oil investors are concerned about futures returns because of their high liquidity and low associated transaction costs. Therefore, this section investigates the ability of the proposed cloudiness measures to forecast oil futures returns and uncover the implications for portfolio allocation. The data on futures Contract 1 are available on the EIA website.

Table 5 shows the results of in-sample and out-of-sample forecasting performance. The in-sample predictive ability of six of the eight measures is significant, consistent with the results for oil spot returns. $K_t^{(A)}$ has a slope coefficient of −0.041, with a t statistic of −2.727, which suggests significant predictive ability. The out-of-sample R^2_{OoS} value indicates that four of the eight measures significantly outperform the historical average benchmark model. Notably, the aggregate cloudiness model has a significant and positive R^2_{OoS} of 1.108%. In summary, the results using futures return data confirm the robustness of the results based on spot returns.

Following the related literature (e.g., Rapach, Strauss, and Zhou, 2010), a portfolio comprising a risky asset (i.e., oil futures) and a risk-free asset (i.e., Treasury bills) is considered to evaluate the economic significance of return predictability. In this framework, the portfolio return is computed as follows:

$$r_{t,p} = \omega_t r_{t,c} + r_f, \qquad (8)$$

where $r_{t,c}$ and r_f denote crude oil futures returns and the 3-month Treasury bill rate as a proxy for the risk-free rate, respectively. The investor is assumed to use mean–variance utility when allocating wealth and computing the optimal weights for the futures contract as follows:

$$U_t(\omega_t) = E(r_{t,p}) - \frac{\gamma}{2} var(r_{t,p}), \qquad (9)$$

where $E(\cdot)$ and $var(\cdot)$ are expectation and variance operators and γ is a risk aversion coefficient. The ex ante weight of oil futures is obtained by maximizing utility as follows:

$$\omega_{t|t-1}^* = \frac{\hat{r}_{t,K}}{\gamma \hat{\sigma}_t^2}, \qquad (10)$$

where $\hat{r}_{t,K}$ and $\hat{\sigma}_t^2$ are the return and volatility forecasts, respectively. For simplicity, a 50-week rolling window is used to

forecast volatility. Thus, differences in portfolio performance depend on the return model. Three risk aversion coefficients are considered: $\gamma = 1$, 3, and 6. Using a limited interval for the weights is reasonable even though futures transactions theoretically allow for unlimited long and short positions. Three weight constraints are considered and imposed on the futures contract: (a) $-1 < \omega^*_{t|t-1} < 1$, indicating no financial leverage and at most a 100% short sale; (b) $-1 < \omega^*_{t|t-1} < 0$, precluding a long position and indicating at most a 100% short sale; and (c) $-2 < \omega^*_{t|t-1} < 2$, indicating at most 100% financial leverage and a 200% short sale.

Table 6 presents portfolio performance evaluated on the basis of two popular measures: the certainty equivalent return (CER) and the shape ratio (SR). The values reported are CER and SR gains, which are respectively defined as the difference between the CER and SR of portfolios based on the model of interest and the benchmark model. The CER differences are multiplied by 1,200 to obtain an annualized percentage value. The SR differences are multiplied by $\sqrt{12}$ to obtain an annualized value. The results indicate that portfolio performance depends on the risk aversion coefficient and weight constraints. Nevertheless, the CER and SR gains are prominent in all considered cases. For example, under the reasonable weight restriction of $-1 < \omega^*_{t|t-1} < 1$, the annualized CER gains are 422, 261.9, and 111.9 basis points for the risk aversion coefficients of 1, 3, and 6, respectively; the respective annualized SR gains are 0.194, 0.251, and 0.231. In summary, the portfolio performance results support the economic significance of predicting oil returns from the aggregate cloudiness measure.

Discussion

Information uncertainty channel. Inventory provides a useful instrument for hedging against oil risk. Accordingly, oil inventory increases during periods of high information uncertainty. Therefore, high information uncertainty negatively affects future oil prices. The following predictive regression is used to test this hypothesis:

$$y_{oi,t+1} = \alpha^{(i)} + \left(\beta^{(i)}_{oi} + \beta^{(i)}_{trend}K_t^{(i)}\right)y_{oi,t} + \delta_K^{(i)}K_t^{(i)} + \varepsilon_{t+1}^{(i)}, \quad (11)$$

where $K^{(i)}$ denotes the cloudiness measure for the ith storage area, $K^{(A)}$ denotes the aggregate cloudiness measure, and $y_{oi,\ t}$ is the change in oil inventory. Lagged oil inventory is incorporated to account for autocorrelation. The null hypothesis of $\delta_K^{(i)} \leq 0$ is tested against the one-sided alternative of $\delta_K^{(i)} > 0$. The results are reported in Table 7. The estimate of $\delta_K^{(A)}$ is 0.0034, with a marginally significant t statistic of 1.645 This evidence supports oil inventory as a channel explaining how $K_t^{(A)}$ can predict oil returns.

Information flow channel. The other plausible explanation for the predictive ability of cloud cover is based on volatility. When the sky above FRTs is cloudy, less satellite-based information flows into the oil market. Consequently, investors may hesitate to trade, reducing trading volume and, by extension, volatility. In addition, on very cloudy days, investors may tend to heavily rely on past volatility information, resulting in high volatility persistence. This hypothesis is tested using the following predictive regression:

$$logRV_{t+1} = \alpha^{(i)} + \left(\beta^{(i)}_{RV} + \beta^{(i)}_{trend}K_t^{(i)}\right) \times logRV_t + \delta_K^{(i)}K_t^{(i)} + \varepsilon_{t+1}^{(i)}, \quad (12)$$

where RV_t denotes the realized volatility (RV) computed as the sum of the squared daily returns from Wednesday of the previous week to the current Tuesday. Following Paye (2012), log RV is used instead of the original series, because the RV series exhibits strong leptokurtosis, which can lead to biased inferences about parameter significance in OLS estimation; in contrast, the distribution of log RV is nearly normal (Andersen et al., (2001); Andersen et al., (2001)). The null hypothesis of is tested against the one-sided alternative of The OLS results are reported in Table 8. The estimate of β_{trend} is 0.161. The corresponding Newey and West (1987) adjusted t statistic (robust to heteroskedasticity and autocorrelation) is 1.725, indicating significance at the 5% level. The coefficient of $K_t^{(A)}$ is -1.396, with a t statistic of -1.699, which is also significant at the 5% level. Therefore, $K_t^{(A)}$ negatively predicts oil volatility, providing evidence of a volatility channel from cloudiness to oil returns.

Table 6 Portfolio results.

Constraint	CER gains			Sharpe ratio gains		
	$\gamma = 1$	$\gamma = 3$	$\gamma = 6$	$\gamma = 1$	$\gamma = 3$	$\gamma = 6$
$-1 < \omega < 1$	4.220	2.619	1.119	0.194	0.251	0.231
$-1 < \omega < 0$	2.459	1.986	0.706	0.118	0.204	0.170
$-2 < \omega < 2$	7.490	2.238	0.539	0.241	0.231	0.202

This table reports the out-of-sample forecasting performance for the predictive regression $r_{t+1} = \alpha^{(A)} + \beta_K^{(A)}K_t^{(A)} + \varepsilon_{t+1}^{(A)}$, where r_t denotes oil price returns, $K_t^{(A)}$ denotes the aggregate cloudiness measure, and $\varepsilon_t^{(i)}$ represents the error term. At each time t, an investor with mean-variance preference assigns the ex-ante weight $\omega^*_{t+1|t} = \frac{\hat{r}_{t+1,K}}{\gamma\hat{\sigma}_{t+1}^2}$ to crude oil futures, where γ denotes risk aversion coefficients; $\hat{r}_{t,K}$ and $\hat{\sigma}_t^2$ are the return and volatility forecasts, respectively. A 50-week rolling window method is employed to compute the volatility forecasts. Two popular evaluation measures: certainty equivalent return (CER) and Sharpe ratio (SR) are applied to evaluate portfolio performance. The numbers reported in the table are annualized CER (SR) gains, defined as the differences between the CER (SR) of given model relative to that of the benchmark model. The CER differences are further multiplied by 1,200 to denote annualized percentage value. The SR differences are multiplied by $\sqrt{12}$ to interpret as the annualized value.

Table 7 information uncertainty and oil inventory.

	$K^{(BLA)}$	$K^{(BTX)}$	$K^{(CTX)}$	$K^{(COK)}$	$K^{(HTX)}$	$K^{(MRL)}$	$K^{(LLA)}$	$K^{(TTX)}$	$K^{(A)}$
α	-0.002	-0.008	-0.009^*	-0.002	-0.007	-0.006	-0.004	-0.004	-0.011
	(-0.389)	(-1.582)	(-1.820)	(-0.634)	(-1.179)	(-1.401)	(-0.949)	(-0.923)	(-1.804)
β_{oi}	0.107	0.671	0.532	0.488^{**}	0.362	0.484	0.739^{**}	0.807^{**}	0.766^*
	(0.296)	(1.551)	(1.058)	(2.204)	(0.698)	(1.464)	(2.152)	(2.321)	(1.712)
β_{trend}	0.091	-0.104	-0.057	-0.044	0.003	-0.040	-0.128	-0.152	-0.138
	(0.714)	(-0.698)	(-0.331)	(-0.532)	(0.017)	(-0.338)	(-1.050)	(-1.231)	(0.862)
δ_K	0.001	0.003^{**}	0.003^{**}	0.001	0.003	0.002^*	0.002	0.002	0.004^{**}
	(0.497)	(1.670)	(1.901)	(0.801)	(1.249)	(1.509)	(1.035)	(1.012)	(1.861)
$R^2(\%)$	13.68	14.07	14.11	13.66	13.80	13.91	13.78	13.81	14.03

This table reports the in-sample estimation results for the predictive regression, $y_{oi,t+1} = \alpha^{(i)} + \left(\beta^{(i)}_{oi} + \beta^{(i)}_{trend}K_t^{(i)}\right)y_{oi,t} + \delta_K^{(i)}K_t^{(i)} + \varepsilon_{t+1}^{(i)}$, where $y_{oi,t}$ denotes oil inventory change at time t, where $K^{(i)}$ is the cloudiness measure for the i-th district, $K^{(A)}$ denotes the aggregate cloudiness, and ε_t represents the error term. The parameters are estimated via the ordinary least squares. The one-sided test is used for null hypothesis $\delta_K^{(i)} \leq 0$ against the alternative $\delta_K^{(i)} > 0$. The numbers in the parentheses are the t-statistics based on the Newey-West robust errors. The in-sample R^2 values are multiplied by 100 to interpret as percentages. The asterisks *, **, and *** denote statistical significance at 10%, 5%, and 1% levels, respectively.

Table 8 Cloudiness and oil volatility.

	$K^{(BLA)}$	$K^{(BTX)}$	$K^{(CTX)}$	$K^{(COK)}$	$K^{(HTX)}$	$K^{(MRL)}$	$K^{(LLA)}$	$K^{(TTX)}$	$K^{(A)}$
α	0.327	1.752	4.148**	1.488*	3.754	2.367	4.611***	3.954**	4.452*
	(0.288)	(1.256)	(2.279)	(1.696)	(1.592)	(1.348)	(2.614)	(2.472)	(1.892)
β_{RV}	0.957***	0.790***	0.525**	0.826***	0.555**	0.733***	0.480**	0.555***	0.487*
	(7.302)	(5.071)	(2.568)	(8.261)	(2.095)	(3.642)	(2.418)	(3.048)	(1.829)
β_{trend}	−0.003	0.055	0.148**	0.045	0.136	0.075	0.163**	0.137**	0.161**
	(0.053)	(1.016)	(2.064)	(1.142)	(1.500)	(1.087)	(2.339)	(2.151)	(1.725)
δ_K	0.040	−0.454	−1.295**	−0.377	−1.144*	−0.669	−1.443***	−1.218**	−1.396**
	(0.097)	(−0.933)	(−2.033)	(−1.109)	(−1.415)	(−1.107)	(−2.343)	(−2.182)	(−1.699)
R^2(%)	90.56	90.60	90.74	90.60	90.70	90.61	90.76	90.72	90.69

Notes: This table reports the in-sample estimation results for the predictive regression, $logRV_{t+1} = \alpha^{(i)} + \left(\beta_{RV}^{(i)} + \beta_{trend}^{(i)} K_t^{(i)}\right) \times logRV_t + \delta_K^{(i)} K_t^{(i)} + \varepsilon_{t+1}^{(i)}$, where RV_t denotes realized volatility at time t, where $K^{(i)}$ is the cloudiness measure for the i-th district, $K^{(A)}$ denotes the aggregate cloudiness, and ε_t represents the error term. The parameters are estimated via the ordinary least squares. The one-sided test is used for against the one-sided alternative hypothesis ($H_1 : \delta_K^{(i)} < 0$). The numbers in the parentheses are the t-statistics based on the Newey-West robust errors. The in-sample R^2 values are multiplied by 100 to interpret as percentages. The asterisks *, **, and *** denote statistical significance at 10%, 5%, and 1% levels, respectively.

Satellite data, cloudiness, and information asymmetry. The availability of satellite-based oil inventory information depends on the cloud cover over FRTs. This paper demonstrates that such cloud cover affects oil returns via two channels. In practice, satellite images are widely applied to aid in investment decisions. For example, Orbital Insight provides parking lot traffic data based on satellite images to help clients estimate the sales of retail firms before earnings announcements. Because of high acquisition and processing costs, market participants have unequal access to satellite images. Generally, high-resolution satellite images are priced by square kilometer, and even a small image covering just $25\,km^2$ costs more than US$600. Moreover, a profitable investment strategy requires daily observations of multiple locations and thus thousands of images per year, leading to staggering costs. The introduction of satellite information creates new information asymmetry between sophisticated and unsophisticated investors. Whether such information asymmetry is priced into security and commodity markets is not investigated in the literature. This issue has important implications for asset pricing and requires further quantitative analysis, as such information asymmetry is always analyzed in cross-sectional analyses.

Changes in predictive ability. Predictive ability may change over time. Mclean, Pontiff (2016) find that the predictive ability of variables weakens over time after they are taken as return predictors in academic publications. A potential explanation for such weakening predictive ability is investor learning. To address this issue, an interaction term of the cloudiness measure and time t $(t \times K_t^{(A)})$ is added to Model (1) as an additional explanatory variable. Because cloudiness negatively predicts oil returns, a significantly positive coefficient of the interaction term suggests that the predictive ability weakens over time (Jacobsen et al., (2019)). The coefficient is positive but nonsignificant (t statistic = 0.635), indicating no effect of investor learning.

Conclusions
Market participants are applying big data techniques to improve investment decisions. However, the literature does not pay adequate attention to the asset pricing implications of the availability of alternative data. Satellite imagery data are typical alternative data, and cloudiness is a natural proxy for the availability of such data and thus the flow of such data into the market. Therefore, this paper investigates the effects of cloudiness over FRTs on oil returns through time-series prediction.

This study constructs eight measures of the cloudiness above FRTs in eight oil storage areas in the U.S. The aggregate of these measures exhibits significant in-sample and out-of-sample

predictive ability for oil returns, and this predictive ability is statistically and economically significant. Moreover, this predictive ability cannot be explained by existing predictors, such as uncertainty or oil market variables. The channels from cloudiness to oil returns include the effects of both oil inventory and oil volatility.

Notes
1 This tool provides daily cloud observations at spatial resolutions of 1 km and 250 m (nadir).
2 Deither, Malloy, and Scherbina, 2002 relax the short-sale constraint and show that friction, which prevents the revelation of negative opinions, produces a negative relationship between forecast dispersion and future returns. The literature documents that oil price increases significantly and negatively affect real economic activity, but oil price decreases have nonsignificant effects (e.g., Mork, 1989; Hamilton, 1996). The "rockets and feathers" literature shows that retail gasoline prices respond quickly to oil price increases but slowly to oil price decreases (Borenstein, Cameron, and Gilbert, 1997). However, no research investigates the effect of revelations of oil price decreases on futures markets.
3 "PADD" is a term created by the U.S. government and is the regional definition used by U.S. government agencies when collecting and publishing industry data.
4 Individual cloudiness measures and their aggregate measure are presented in the online appendix.
5 www.eia.gov.
6 The graphical representations of oil returns are illustrated in the online appendix.
7 See the results in the online appendix.

References
Agnello L, Castro V, Hammoudeh S, Sousa RM (2020) Global factors, uncertainty, weather conditions and energy prices: On the drivers of the duration of commodity price cycle phases. Energy Econ. 90
Andersen TG, Bollerslev T, Diebold FX, Ebens H (2001) The distribution of realized stock return volatility. J. Financ. Econ 61(1):43–76
Andersen TG, Bollerslev T, Diebold FX, Labys P (2001) The distribution of realized exchange rate volatility. J. Am. Stat. Assoc 96(453):42–55
Baker SR, Bloom N, Davis SJ (2016) Measuring economic policy uncertainty. Q. J. Econ 131(4):1593–1636
Bekaert G, Engstrom EC, Xu NR (2022) The Time Variation in Risk Appetite and Uncertainty. Manage. Sci. 68(6):3975–4004
Bollerslev T, Hood B, Huss J, Pedersen LH (2018) Risk Everywhere: Modeling and managing volatility. Rev. Financ. Stud 31(7):2730–2773
Borenstein S, Colin Cameron A, Gilbert R (1997) Do gasoline prices respond asymmetrically to crude oil price changes? Q. J. Econ. 112(1)

Bradbear C, Friel S (2013) Integrating climate change, food prices and population health. Food Policy 43:56–66

Campbell JY, Thompson SB (2008) Predicting excess stock returns out of sample: Can anything beat the historical average? Rev. Financ. Stud. 21(4):1509–1531

Choi D, Gao Z, Jiang W (2020) Attention to global warming. Rev. Financ. Stud 33(3):1112–1145

Clark TE, West KD (2007) Approximately normal tests for equal predictive accuracy in nested models. J. Econom. 138(1):291–311

Daniel K, Hirshleifer D, Subrahmanyam A (1998) Investor psychology and security market under- and overreactions. J. Finance 53(6):1839–1885

Diether KB, Malloy CJ, Scherbina A (2002) Differences of opinion and the cross section of stock returns. J. Finance 57(5):2113–2141

Ederington LH, Lee JH (1993) How Markets Process Information: News Releases and Volatility. J. Finance 48(4):1161

Engle RF (1982) Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. Econometrica 50(4):987

Engle RF, Giglio S, Kelly B, Lee H, Stroebel J (2020) Hedging climate change news. Rev. Financ. Stud 33(3):1184–1216

Fama EF (1970) Efficient Capital Markets: A Review of Theory and Empirical Work. J. Finance 25(2):383

Gao L, Hitzemann S, Shaliastovich I, Xu L (2022) Oil volatility risk. J. Financ. Econ. 144(2):456–491

Hamilton JD (1996) This is what happened to the oil price - Macroeconomy relationship. J. Monet. Econ. 38(2):215–220

Harvey DS, Leybourne SJ, Newbold P (1998) Tests for forecast encompassing. J. Bus. Econ. Stat. 16(2):254–259

Henderson JV, Storeygard A, Weil DN (2012) Measuring economic growth from outer space. Am. Econ. Rev. 102(2):994–1028

Hirshleifer D, Shumway T (2003) Good Day Sunshine: Stock Returns and the Weather. J. Finance 58(3):1009–1032

Hong H, Li FW, Xu J (2019) Climate risks and market efficiency. J. Econom 208(1):265–281

Hu C (2011) An empirical approach to derive MODIS ocean color patterns under severe sun glint. Geophys. Res. Lett. 38(1)

Jacobsen B, Marshall BR, Visaltanachoti N (2019) Stock market predictability and industrial metal returns. Manage. Sci. 65(7):3026–3042

Johnson TL (2019) A fresh look at return predictability using a more efficient estimator. Rev. Asset Pricing Stud 9(1):1–46

Kalev PS, Liu WM, Pham PK, Jarnecic E (2004) Public information arrival and volatility of intraday stock returns. J. Bank. Financ 28(6):1441–1467

Katona Z, Painter MO, Patatoukas PN, Zeng J (2018) On the Capital Market Consequences of Alternative Data: Evidence from Outer Space. 9th Miami Behav. Financ. Conf. 23529(2):1–45

Kostakis A, Magdalinos T, Stamatogiannis MP (2015) Robust econometric Inference for stock Return predictability. Rev. Financ. Stud. 1506–1553

Krueger P, Sautner Z, Starks LT (2020) The importance of climate risks for institutional investors. Rev. Financ. Stud 33(3):1067–1111

Lanfear MG, Lioui A, Siebert MG (2019) Market anomalies and disaster risk: Evidence from extreme weather events. J. Financ. Mark. 46

Loughran T, Schultz P (2004) Weather, stock returns, and the impact of localized trading behavior. J. Financ. Quant. Anal 39(2):343–364

Mclean RD, Pontiff J (2016) Does Academic Research Destroy Stock Return Predictability? J. Finance 71(1):5–32

Miller EM (1977) Risk, Uncertainty, and Divergence of Opinion. J. Finance 32(4):1151–1168

Mork KA (1989) Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results. J. Polit. Econ. 97(3)

Mukherjee A, Panayotov G, Shon J (2021) Eye in the sky: Private satellites and government macro data. J. Financ. Econ 141(1):234–254

Newey WK, West KD (1987) A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55(3):703–708

Novy-Marx R (2014) Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars. J. Financ. Econ 112(2):137–146

Paye BS (2012) "Déjà vol": Predictive regressions for aggregate stock market volatility using macroeconomic variables. J. Financ. Econ 106(3):527–546

Purbantoro B, Aminuddin J, Manago N, Toyoshima K, Lagrosas N, Sumantyo JTS, Kuze H (2019) Comparison of aqua/terra MODIS and Himawari-8 satellite data on cloud mask and cloud type classification using split window algorithm. Remote Sens. 11(24)

Rangel JG (2011) Macroeconomic news, announcements, and stock market jump intensity dynamics. J. Bank. Financ 35(5):1263–1276

Rapach DE, Ringgenberg MC, Zhou G (2016) Short interest and aggregate stock returns. J. Financ. Econ 121(1):46–65

Rapach DE, Strauss JK, Zhou G (2010) Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. Rev. Financ. Stud. 23(2):821–862

Saunders EM (1993) Stock prices and Wall Street weather. Am. Econ. Rev. 83(5):1337–1345

Stambaugh RF (1999) Predictive regressions. J. Financ. Econ. 54(3):375–421

Wang J (1993) A model of intertemporal asset prices under asymmetric information. Rev. Econ. Stud. 60(2):249–282

Welch I, Goyal A (2008) A comprehensive look at the empirical performance of equity premium prediction. Rev. Financ. Stud. 21(4):1455–1508

Yu J (2011) Disagreement and return predictability of stock portfolios. J. Financ. Econ 99(1):162–183

Zhang H, Wang M (2010) Evaluation of sun glint models using MODIS measurements. J. Quant. Spectrosc. Radiat. Transf. 111(3):492–506

Zhang XF (2006a) Information uncertainty and analyst forecast behavior. Contemp. Account. Res 23(2):565–590

Zhang XF (2006b) Information uncertainty and stock returns. J. Finance 61(1):105–137

Zhu C (2019) Big Data as a Governance Mechanism. Rev. Financ. Stud 32(5):2021–2061

Acknowledgements

Author contributions

Y.W., and X.H. conceived of the study. X.H. collected, analyzed satellite imagery, designed data preprocessing pipelines, trained deep learning models, and analyzed their output. X.H. and Y.W. finished the forecasting process. Y.W. wrote the paper, managed the project and provided funding support. Both authors contributed equally to this work.

Competing interests

The authors declare no competing interests.

Ethical approval

All experimental data are obtained from publicly available data sources and do not include any personal information.

Informed consent

This article does not contain any studies with human participants performed by any of the authors.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1057/s41599-023-02128-5.

Correspondence and requests for materials should be addressed to Yudong Wang.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.