

Empirical Tests of the Green Paradox for Climate Legislation

Maya Norman* Wolfram Schlenker†

December 6, 2024

Abstract

The Green Paradox posits that fossil fuel markets respond to evolving expectations about climate legislation, intended to limit future consumption, by shifting consumption to the present. This shift results in lower present-day prices. We demonstrate that oil futures reacted negatively to daily changes in prediction market expectations regarding the Waxman-Markey bill—the U.S. climate bill discussed in 2009–2010. This effect is observed across various maturities, with an even larger impact on more distant maturities, as the proposed legislation would have reset the entire price and consumption path, unlike temporary supply or demand shocks, which tend to phase out over time. The decrease in oil prices is accompanied by a decrease in storage levels as consumption accelerates while it takes time for new production to come online. Furthermore, two unexpected court rulings that implied limited future fossil fuel use were also associated with negative abnormal oil future returns. Taken together, these findings confirm that restricting future fossil fuel use accelerates current-day consumption. The passage of the Waxman-Markey bill would have increased current global oil consumption by 2%.

Key words: Q41, Q54, G18

JEL Codes: Green Paradox, oil consumption, climate change.

Acknowledgments: We would like to thank Kyle Meng and Derek Lemoine for sharing the prediction market data and for helpful feedback, as well as participants of the Virtual Seminar on Climate Economics by the Federal Reserve Bank of San Francisco, the Harvard Seminar in Environmental Economics and Policy, and seminar participants at the University of Illinois.

*School of International and Public Affairs, Columbia University

†Harvard Kennedy School of Government and National Bureau of Economic Research (NBER)

Climate legislation often establishes goals for the future to give companies and consumers time to adapt and plan for a transition away from fossil fuels. For example, the European Union enacted the goal to be climate neutral (net zero emissions) by 2050, while China established the same goal for 2060. Fossil fuels are exhaustible resources, and their finite availability dictates their use and price path (Hotelling 1931). This scarcity leads to a price that exceeds the marginal extraction cost, resulting in resource rents that ensure less is consumed today and fossil fuels are saved for the future.¹ The literature on the “Green Paradox” highlights that climate legislation, which limit future fossil fuel use, give resource owners an incentive to extract more in the present and medium term before the regulation binds, leading to lower prices, accelerated resource depletion, and higher consumption today. This theoretical literature is based on Hotelling’s seminal model (Hoel 2010a, Sinn 2008a, Sinn 2008b, Van der Ploeg and Withagen 2012, Van der Ploeg and Withagen 2015). By the same logic, a global carbon tax on fossil fuels with scarcity rents will not be passed on to consumers. If producers did so, then demand for fossil fuels would fall, cumulative consumption would decrease and not all fossil fuels would be used, incentivizing resource owners to lower fossil fuel prices to sell all units. In the end, much of a carbon tax would be absorbed by producer rents with limited effects on fossil fuel use or consumers (Dasgupta, Heal and Stiglitz 1980, Heal and Schlenker 2019).

What is common to both a carbon tax and future carbon quantity regulation is the concern that such legislation might not lead to the desired reductions in cumulative fossil fuel use and may even accelerate consumption today through lower prices, limiting the regulation’s effectiveness. On the other hand, the limited pass-through alleviates widely held apprehensions that such policy would have distributional consequences and high personal costs. If the dynamics of Hotelling’s rule shield consumers from a global carbon tax, it would significantly improve public opinion of such a policy (Dechezleprêtre, Fabre, Kruse, Planterose, Sanchez Chico and Stantcheva 2022). The “Green Paradox” hence has important implications for both the effectiveness of climate legislation in limiting fossil fuel use as well as the pass-through of carbon taxes, yet most of the literature to date has been theoretical.

Our paper adds to the emerging empirical literature on the “Green Paradox.” Specifically, we test the predictions of the “Green Paradox” using several data sources on different timescales. We consistently find evidence of the mechanism underlying the “Green Paradox:” additional restrictions on oil use, or an increased likelihood of future restrictions, reduce both the oil spot price and oil futures prices in the 24 months for which futures data are traded, i.e., the following two years. This occurs as supply is reallocated from the future to the present. The stickiness of the price change is a key prediction of the Green Paradox that sets apart from temporary shocks whose effect phase out over time, and a key piece of evidence in our paper. We find that futures contracts with more distant maturities show a larger sensitivity than those with shorter maturities. Consistent with a larger price sensitivity of more distant maturities, we observe a reduction in storage levels in conjunction with a reduction in oil prices.² Other temporary shocks that lower prices generally lead to an *increase* in storage levels.

Empirically testing the “Green Paradox” is difficult because the analysis requires information about agent’s expectations of climate policy stringency for which there is very little data. We address this challenge through the use of prediction market prices, a measure that captures the market’s expectations, in our

¹For example, Saudi Arabia’s extraction cost are less than 10 dollars per barrel for most of its oil fields, while the oil price trades in the 60-80 dollar per barrel range.

²Arbitrage in the storage market requires that the expected price increase covers the storage cost. If future prices decrease by more than current prices, storage becomes less profitable.

baseline specification. We confirm the sign of the discovered relationship using unexpected news regarding court cases that mandated changing climate policy and a news based index proxying for information shocks related to climate policy. Specifically, our analysis provides four piece of evidence.

We begin by documenting that oil price shocks, i.e., changes in the daily oil price, have become stickier over the last three decades, implying that shocks have become more permanent rather than transitory. Under the “Green Paradox,” uncertainty from climate legislation deliberations leads to persistent (sticky) price changes, as changes in expectations around future climate legislation reset the entire future oil price path and hence impact all maturities equally, while the effect of temporary supply and demand shocks phase out over time. Consistent with this prediction, we find the persistence of spot price shocks among maturities ranging from 1 to 24 months greatly increased during periods when climate bills were under consideration (the 2000s and 2010s). Daily changes in the oil spot price translate into roughly the same change in oil futures prices with a maturity one month into the future throughout the sample period. However, the story differs for longer-term maturities: in the 1990s, spot price shocks phased out for oil futures with longer maturities. Specifically, only about a third of the spot price change was reflected in the oil futures with a 24-month maturity. Around 2010, the fraction doubled to two thirds, i.e., daily shocks phased out slower with longer-term maturity futures. This finding only reverses in the 2020s, when COVID-related temporary supply disruptions lead to a decoupling of future and spot price movements. The time profile of how shocks phase out provides an important baseline for subsequent analysis.

In a second step, what we find to be our strongest empirical evidence, we construct a direct test of the “Green Paradox” by pairing daily oil price data with daily estimates of the market’s expectations that Waxman-Markey would pass. Waxman-Markey was a climate bill that intended to limit economy wide greenhouse gas emissions in the U.S. through cap and trade. We retrieve the market’s expectations using prediction market contract prices in 2009 and 2010. Theory predicts that increases in the probability of a cap and trade bill passing should reduce contemporaneous oil prices and vice versa. Consistent with this prediction, we find a significant negative coefficient; prices of oil futures decline whenever the expected likelihood that the bill will pass increases. This effect is persistent across all futures contracts, even increasing for longer-term maturities, suggesting that the relationship reflects long-term adjustments in the expected oil price path rather than temporary shocks. Through our analysis we find (i) the passage of the Waxman-Markey bill would have increased global oil consumption 2% and (ii) Waxman-Markey deliberations increased oil consumption by 8 million metric tons equivalent to roughly the daily global oil consumption. We present several pieces of evidence to rule out the possibility that reverse causality could explain our finding, i.e., that lower oil prices associated with negative returns increased the likelihood that a climate bill would pass, or stated differently, opposition to the bill was higher when oil prices were higher. We find that the effect is even larger and more significant when we limit the sample to days with major changes in prediction market prices – these major changes were driven by political negotiations at committee meetings that were scheduled in advance and should not have been influenced by day-to-day oil price movements, ruling out reverse causality. In a falsification check, we find no significant effect if we use a one-period lead in prediction market price changes,³ confirming that prediction market movements on a particular day and the implied news on that day lead to changes in oil future prices.

³We look at returns between consecutive closing prices, which generally implies a period of one day, except over weekends and holidays, when a period covers 2-3 days.

In a third step, we construct the abnormal oil price returns on the days two historic climate court cases were resolved. Specifically, we study the effect of the surprise *Urgenda v. Netherlands* rendering, when a Dutch court sided with an environmental group and ordered the Dutch government to have stricter limits on future fossil fuel emissions. When the ruling was announced, people predicted it would set a precedent for all countries subject to the European Convention. Additionally, we study the effect on the day news coverage suggested that Justice Kennedy, the swing vote in the *Massachusetts v. EPA* Supreme Court case, would support the states suing the EPA to regulate automobile carbon dioxide emissions. At the time that the case was deliberated, numerous additional climate court cases awaited the Supreme Court’s verdict, including a case challenging the EPA’s refusal to regulate power plant carbon dioxide emissions. In both instances, we find significant negative coefficients, i.e., oil futures prices declined when new information increased the expected likelihood of limitations on future fossil fuel use.

In a fourth step, we revert from daily prices over a year and a half to a longer time series of monthly estimates of U.S. renewable policy and international climate negotiation salience over several decades and pair it with monthly oil price data. We measure policy salience using Noailly, Nowzohour and Van Den Heuvel (2021)’s news-based indices generated by text-mining articles from ten leading US newspapers published between 1981 and 2019. The indices reflect the monthly number of articles covering US renewable policy and international climate negotiations, respectively, relative to the total number of articles published. While the “Green Paradox” makes no direct predictions of the effect of climate policy salience on oil prices, the measure of climate policy salience used in this paper generally tracks events that strengthened future climate policy, i.e., the renewable policy index peaks after the passage of renewable policy. Hence, pairing oil prices with the news based indices can provide a suggestive yet compelling test of the paradox’s prevalence over the last four decades. The “Green Paradox” predicts that the indices should be negatively correlated with oil prices. For example, increases in the international negotiations index indicate international cooperation around climate likely strengthened, elevating the expected stringency of future climate policy, causing oil producers to supply more today, and consequently reducing prices. Consistent with the paradox’s predictions, we find increases in the salience of international climate negotiations significantly reduce oil prices.

On the other hand, we find increases in the salience of renewable energy policy has no significant effect on oil prices. Renewable energy programs have two countervailing effects: strengthening renewable energy policy could reduce the backstop price causing oil producers to increase supply today, reducing oil prices. Alternatively, strengthening of renewable policy has often occurred in place of climate policy, thereby easing the concern that there might be future restriction on fossil fuel use, resulting in higher oil prices as future supply is no longer threatened.

Taken together, these findings show that the oil market is sensitive to climate laws and that expected restrictions on future fossil fuel use will lead to increased consumption today. The economics of exhaustible resources predicts that discoveries of an exhaustible resource influence the scarcity of a resource and its price (Ekeland, Schlenker, Tankov and Wright 2022). If total availability of an exhaustible resource goes up through a new discovery, the expected future price path resets and is lowered. The effect of climate legislation is analogous: by limiting resource use in the long-term, available resources are shifted towards the short and medium-term. By the same token, scarcity rents will absorb carbon taxes, shielding consumers.

Our paper contributes to the emerging empirical literature on the “Green Paradox,” by documenting how environmental laws can increase present-day oil demand through lowering prices. Grafton, Kompas, Long

and To (2014) show that increases in biofuel production, a substitute to fossil fuels, increase oil production. Di Maria, Lange and van der Werf (2014) show that the passage of the acid rain program decreased the price of high-sulphur coal. Merrill (2018) finds the out-of-committee introduction of climate related bills in congress accelerate oil and gas firm wellhead investments. Lemoine (2017) observes an abnormal return in coal futures on the day Senator Graham abandoned the Waxman-Markey bill on Monday, April 26, 2010. Barnett (2023) observe abnormal returns in the oil spot price in response to climate transition events – major events in the energy industry, election results, and other major shifts in climate policy.

A challenge of previous papers is to determine when the market updated its beliefs about the likelihood of a policy change. The “Green Paradox” is derived from expectations of future prices, and markets might see and react to an impending regulation before it is officially implemented and ratified (Dube, Kaplan and Naidu 2011, McDermott, Meng, McDonald and Costello 2019, Langer and Lemoine 2020). Lemoine (2017), Merrill (2018), and Di Maria et al. (2014) test the paradox’s predictions using policy or information shocks occurring at a specific moment, relying on comparisons between the period before and after a single shock. Grafton et al. (2014) uses annual variation in biofuel production as a proxy for annual variation in biofuel subsidies to test the paradox. Barnett (2023) interacts climate transition events with a portfolio tracking the response of firms highly exposed to climate transition risk to account for and recover dynamics in the magnitude and timing of transition events.

We see the novelty of our paper as follows: our approach’s reliance on *repeated* daily (and monthly in a later check) information shocks allow us to employ more convincing variation in the market’s beliefs regarding the stringency of future climate policy rather than focusing on single events or pre-post comparisons. These repeated shocks capture the policy making process that goes through announcements, deliberation, redrafting, and upheavals. Previous literature on the “Green Paradox” that relied on a single pre-post comparison might be susceptible to reverse causality. For political economy reasons, a bill might be more easily passed when a resource was declining in economic importance and the price was falling. Or, a subsidy might be more easily implemented when there is generally more demand for fuel. A key innovation of our paper is that our approach is more defensible against the concerns of reverse causality as we employ a panel data set encompassing numerous daily changes in expectations that occurred in response to previously scheduled committee meetings. These meetings, scheduled weeks or months in advance, are unlikely to coincide with daily oil spot price shocks.

Our paper builds most closely on Lemoine (2017) and reinforces its findings. We make several additional contributions: first, our analysis is less likely to be confounded by other events as our analysis relies on various time scales including high-frequency daily data. For example, during the week of Senator Graham’s abandonment, the focus of Lemoine (2017), a catastrophic oil spill in the Gulf of Mexico, known as Deep Water Horizon, unfolded. This spill makes focusing on the effect of Graham’s abandonment on fossil fuel prices difficult, as uncoupling the effect of the spill from the bill’s abandonment is empirically challenging. It is also not entirely clear when the market learned of Graham’s announcement.⁴ Instead, we rely on variation in prediction market prices over the course of one-and-a-half years, during which Waxman-Markey

⁴The immigration bill that led to the demise of Waxman-Markey was announced Thursday April 22nd (Lemoine 2017). The market’s expectation that Waxman-Markey would pass did not change until Saturday April 24th when it decreased by 5 probability points. The market’s expectation again decreased by 5 probability points on Sunday April 25th before increasing by just under 5 probability points on Monday the 26th. Additionally, the Waxman-Markey Google Trends index exceeded 5% of the most popular day on Wednesday April 21st and 10% on Saturday April 24th. The index was zero on the 22nd, 23rd, 25th and 26th indicating these days had too little search volume to determine relative popularity.

bill deliberations took place, to identify the effect of changes in the expected stringency of future climate policy on oil prices. Second, coal supply, the major focus of Lemoine’s paper may be less sensitive to changes in future US climate policy than oil because air quality regulation protecting local environments already heavily restricts US coal consumption, limiting resource rents today, and consequently limiting the “Green Paradox” (Hoel 2010b).

Our paper also contributes more generally to an empirical literature demonstrating that the anticipation of or uncertainty around a new environmental policy can induce behaviors counteracting or enhancing the policy’s intended benefits (Polasky and Doremus 1998, McDermott et al. 2019, Gowrisankaran, Langer and Zhang 2022, Dorsey 2019, Mukanjari and Sterner 2018, Bruno and Hagerty 2023). Even more broadly, policy anticipation is a well considered and documented topic in public economics; policy anticipation can impact consumption, investment or job selection prior to policy implementation, affecting a policy’s effectiveness (Judd 1987, Gründler and Sauerhammer 2018, Kleven, Landais and Saez 2013). The literature on anticipation effects consistently struggles to determine when actors update beliefs and results are sensitive to these researcher choices. Unique to this literature, our paper focuses on oil prices’ reaction to the policy making process, testing “Green Paradox” predictions through the use of high-frequency variation in expectations around the stringency of future regulation to identify how markets react to policy making.

1 Model

We will briefly review the theory behind the Green Paradox to motivate our empirical analysis. Starting with the simplest case where there are no extraction costs and extraction can be started and stopped instantaneously, the fossil-fuel resource endowment is S_0 and demand is iso-elastic, i.e., quantity consumed $q = \frac{\alpha}{p^\eta}$ where α is a constant and the price-elasticity is η . The inverse demand curve is $p = \sqrt[\eta]{\frac{\alpha}{q}}$.

The arbitrage condition implies that price has to rise at the rate of interest δ over time t , or $p(t) = p_0 e^{\delta t}$. The extraction quantity is hence $q(p_t) = \frac{\alpha}{p_t^\eta} = \frac{\alpha}{p_0^\eta} e^{-\delta \eta t}$. Additionally, all reserves are eventually used up. There is no incentive for a competitive firm not to extract everything, as it has no influence on the price and extraction costs are zero.

$$S_0 = \int_0^\infty q(t) dt = \int_0^\infty \frac{\alpha}{p_0^\eta} e^{-\delta \eta t} dt = \left[-\frac{1}{\delta \eta} \frac{\alpha}{p_0^\eta} e^{-\delta \eta t} \right]_0^\infty = \frac{1}{\delta \eta} \frac{\alpha}{p_0^\eta}$$

Hence the initial endowment combined with the arbitrage condition pins down the optimal extraction path in a competitive equilibrium (as $\frac{\alpha}{p_0^\eta} = \delta \eta S_0$):

$$q(t) = \delta \eta S_0 e^{-\delta \eta t}$$

The initial quantity consumed is $q_0 = \delta \eta S_0$, and $q(t)$ is decreasing at an exponential rate, approaching zero, but never equaling zero. For any quantity $\bar{q} < \delta \eta S_0$ there is a time \bar{t} when $q_{\bar{t}} = q_0 e^{-\delta \eta \bar{t}} = \bar{q}$ or $\bar{t} = \frac{1}{\delta \eta} [\ln(\delta \eta S_0) - \ln(\bar{q})]$.

The Waxman-Markey bill limited the amount that was allowed to be consumed from some future point onward i.e., $q(t) \leq \bar{q} \quad \forall t \geq \hat{t}$. There are two possible cases: first, if $\bar{t} \leq \hat{t}$ the regulation is not binding as the quantity extracted from \hat{t} onward is already below the mandate and hence nothing changes.

The second, and more interesting case, is when the regulation is binding as $\bar{t} > \hat{t}$ and the competitive equilibrium would have implied a consumption $q(t) > \bar{q}$ for the time interval $[\hat{t}, \bar{t}]$. A sample case is shown in the top graph of Figure 1. It uses a discount rate of 3 percent, and a demand elasticity of $\eta = 0.59$, which is consistent with empirical estimates (Hamilton 2009). While the Markey-Waxman bill had several features that evolved over time, we display the mandated 83% reduction by 2050. Please note that the figure is used as a motivating example, not a projection of actual oil use and prices. Our example does not include expected outward shifts in demand (global population and income growth) and new discoveries, as we assume a stationary demand and a fixed known endowment, S_0 . In our motivating example, we have $[\hat{t}, \bar{t}] = [2050, 2110]$. Under the competitive equilibrium, consumption on this time interval would have been

$$\int_{\hat{t}}^{\bar{t}} q(t) dt = \int_{\hat{t}}^{\bar{t}} \delta \eta S_0 e^{-\delta \eta t} dt = S_0 e^{-\delta \eta \hat{t}} - \frac{\bar{q}}{\delta \eta}$$

The regulation only allows for $\bar{q}[\bar{t} - \hat{t}]$ in consumption, $\Delta S = S_0 e^{-\delta \eta \hat{t}} - \bar{q} \left[\frac{1}{\delta \eta} + \bar{t} - \hat{t} \right]$ less than the consumption in competitive equilibrium on this interval, shown as area A in Figure 1. If the proposed climate bill had passed, the optimal extraction path reallocates extraction that would have otherwise occurred during the time interval $[\hat{t}, \bar{t}]$ to before $\hat{t} = 2050$ and after $\bar{t} = 2110$.

We denote the revised quantity and price paths as $q'(t)$ and $p'(t)$, respectively. The revised optimal quantity path starts with a larger $q'_0 > q_0 = \delta \eta S_0$, which leads to higher initial consumption on $[0, \hat{t}]$, shown as area B_1 . For the revised optimal extraction path to be optimal, it must be the case that $p'_{t_1} = e^{-\delta[t_1 - \hat{t}]} p'_{\hat{t}}$ for any $t_1 > \hat{t}$ when the quantity regulation is non-binding, otherwise arbitrageurs would benefit from shifting production to utilize the fact that prices are not rising at the rate of interest. The exponential rise in prices implies an exponential decline in quantity consumed $q(t) = q'_0 e^{-\delta \eta t_1}$ when the regulation is non-binding. The unrestricted continuation path from the re-optimized consumption path is shown as a grey dashed line in the bottom panel of Figure 1. When the regulation is non-binding, extraction occurs at a faster rate than the extraction path in the absence of regulation.⁵

The time period when regulation stops binding is hence no longer at $\bar{t} = 2110$ in the bottom panel of Figure 1, but when the re-optimized path $q'_0 e^{-\delta \eta t} = \bar{q}$. Since $q'_0 > q_0$, we know that $\bar{t}' > \bar{t}$. The time when regulation is binding is extended to $\bar{t}' = 2122$ under the re-optimized path in Figure 1. Consumption is higher under the re-optimized path from $\bar{t} = 2110$ onward, as shown by the area B_2 . The reallocation of consumption across time implies that the sum of area B_1 and B_2 equals A .

As highlighted in Figure 1, an increase in the stringency of future climate policy theoretically increases the optimal rate of extraction before the policy enters into force and hence reduces the resource price in the present. In our primary analysis, we never observe any actual change in the stringency of future climate policy. We do observe changes in market expectations around the stringency of future climate policy. In our prediction market analysis, we exploit variation in these expectations to identify how the expected price path of oil would differ if these policies went into effect. Since the prediction market reflects the probability that US climate law will pass, the effect on global oil prices is also influenced by the market's belief that other countries will follow suit or free ride. We assume that these follow-on effects are positive, i.e., a US regulation makes it weakly *more* likely that other countries will follow suit. Any expected free riding does not compensate for the expected demand reductions caused by US regulation as well as other

⁵This is equivalent to assuming a higher initial endowment where part of the time quantity consumed is restricted.

countries' subsequent policies. This assumption seems plausible as political pressure by the United States likely would attenuate free riding.⁶ Under this assumption, the Hotelling model predicts that an increase in the probability of the US law passing will lower oil prices.

This motivating example makes several crucial simplifying assumptions, both on the demand side (constant demand) and supply side (instantaneous start and stop of production). Regarding the latter, Anderson, Kellogg and Salant (2018) stress that once an oil well is tapped, it becomes costly to close it and it is best to just continue extraction, which declines exponentially over time. Regarding former, (Kellogg 2024) models a case where demand initially continues to grow until 2030, but then lets the intercept of the linear demand curve decline monotonically towards zero over the next 75 years. Combining decreasing demand with long-lived investment decisions implies that oil producers might no longer find it profitable to invest in new wells, implying that not all oil will be used. The crucial assumption is that demand indeed vanishes towards zero over 75 years. In a sensitivity check where 15% of demand remains after 75 years as some sectors are difficult to decarbonize (e.g., aviation fuels) or there is a demand for petrochemicals, the oil reserves will be almost fully exploited again, as is the case for a slower (100 year) demand decline.

2 Data

Fossil Fuel Futures Prices

Oil futures are obtained from NYMEX, specifically futures on the West Texas Intermediate (WTI) crude price. Oil futures contracts are the market's assessment of future oil prices. If traders are risk neutral and the market efficiently aggregates information, then contract prices reflect society's best guess of the future oil price (Kellogg 2014). We obtained oil futures prices for futures with liquid contracts. These range between 1 and 24 months into the future, reflecting the market's best guess of oil prices over the upcoming two years on a monthly basis (24 different contracts). Additionally, in some sensitivity checks we use coal futures, also obtained from NYMEX, specifically the Central Appalachian Contract. Lastly, we use daily WTI crude oil prices from Cushing, Oklahoma, recorded by the US Energy Information Administration as a measure of the oil spot price. Oil spot price data allow us to document the evolution of the stickiness of shocks to the oil spot price during the sample period and test for reverse causality.

Market Controls

Oil prices respond to macro-economic shocks, so throughout our analysis we generally control for daily changes in the S&P 500 index, a stock market index that tracks the performance of the 500 largest publicly traded companies in the US. In a sensitivity check, we use S&P 500 index futures to control for expected economic growth. Unlike oil futures, where actively traded maturities are available up to 24 months into the future, the furthest maturity for S&P 500 contracts is 8 months.⁷

Oil Production and Stores

To support our claim that the observed oil price shifts are caused by supply rather than demand shocks, we expand our analysis to include crude oil production and storage. We obtain weekly estimates from the EIA

⁶For example, leading into the COP meetings in Paris where the Paris agreement was signed, the United States and China announced a willingness to work together to solve this problem.

⁷We downloaded the oil futures, coal futures, and S&P 500 index from a Bloomberg Terminal in March 2023.

of U.S. oil ending stocks excluding the strategic petroleum reserve as well as of U.S. oil field production.

Prediction Market on Probability of Climate Law Passing

Market beliefs on the probability that the US government would enact a cap-and-trade system for emissions by the end of 2010 are obtained from prediction market contract prices from Intrade. A prediction market contract is a bet on the realization of a particular event by a given date. If the event is realized by the specified date in the contract, holders receive a dollar. If the event is not realized by the specified date, holders do not receive anything. In an efficient market, the price of the betting market in cents should equal the probability of the law passing. We use prediction market prices from Intrade from May 1, 2009 to December 31, 2010. The end date is given by the contract, which was on the passage of a US climate bill by the end of 2010. Specifically, the prediction market contract was for “A cap-and-trade system for emissions trading to be established before midnight ET on 31 Dec 2010.” While the market did not explicitly cite the Waxman-Markey bill, prediction market participants were likely reacting to the Waxman-Markey bill, legislation to limit economy wide greenhouse gas emissions in the U.S. (Meng 2017).

Internet Searches about Waxman-Markey and Court Cases

We pull daily Google Trends data for the search term “Waxman-Markey” during the period in which the prediction market operated as well as the search terms “Urgenda v. Netherlands” and “Massachusetts v. EPA” in the six months preceding each court cases’ verdict announcement.

For the prediction market analysis, Google Trends data provides an independent measure whether Waxman-Markey was of general interest, which we use in a sensitivity analysis by restricting our sample to days with high search volumes, when it is more likely that belief updates occurred. Google Trends data can only be downloaded at the daily level for a maximum period of nine months. Each Google Trends data download is normalized based on the observations in the download. The day when the search term was the most popular has an index of 100. A day during the period when the search term was half as popular as the most popular day has an index of 50. Days with not enough search volume to determine the relative popularity of the search term have an index of zero. We construct daily Google Trends data during the period in which the Waxman-Markey prediction market operated, between May, 2009 and June, 2010, by making three data queries. The first query covers days between May, 2009 and December, 2009, the second covers October, 2009 to June 2010, and the third covers March to November, 2010. Using the period of overlap between queries, we re-scale the indices such that the most popular day during the whole sample has an index of 100.

Additionally, we use the Google Trends data regarding the court cases to isolate the days with the most substantial information shocks regarding both court cases’ verdicts. Specifically, we focus our event study of each court case on the day in which the court case peaked in popularity on Google in the months leading up to and including verdict announcement. This approach for isolating information shocks is consistent with the literature documenting the anticipation or impacts of environmental policy (McDermott et al. 2019, Carattini and Sen 2019).

Measures of Policy Saliency

We measure the saliency of US renewable policy and international cooperation around climate using news based indices from Noailly et al. (2021). Noailly et al. (2021) develop their indices by text-mining 15 million articles published between 1981 and 2019 by the *New York Times*, *Washington Post*, *Wall Street Journal*, *Houston Chronicle*, *Dallas Morning News*, *San Francisco Chronicle*, *Boston Herald*, *Tampa Bay Times*, *San Jose Mercury News*, and *San Diego Union Tribune*. The authors identify articles pertaining to environmental policy generally using a support vector machine algorithm trained on 2,464 labeled articles. To classify environmental policy articles into sub-topics, Noailly et al. (2021) use topic modeling, an unsupervised learning algorithm. The approach recognizes recurring patterns in the set of environmental policy articles, creating groupings of articles that cover similar topics. The authors use this approach to construct twenty five different groupings, two of which they made publicly available, renewable energy policy and international climate negotiations. Articles classified as covering renewable energy policy include words/phrases such as renewable energy, wind, solar, energy, turbine, energy, power, electricity, renewable, wind power, farm, solar energy, turbine, etc. Articles classified as covering international climate negotiations include agreement, united, international, government, country, state, world, trade, president, European, Mexico, China, etc. They construct news indices for environmental policy generally, renewable policy and international climate negotiations by counting the number of articles in a given category each month and scaling the count by the total monthly volume.

Noailly et al. (2021) document that both the international climate negotiations and renewable policy index are predominantly associated with events that make US policy in these respective areas more stringent. For example, major changes in the international climate negotiations occurred during the Rio de Janeiro Earth Summit, the Kyoto Protocol signing, the Bonn Climate Change Conference, the Copenhagen Climate Change Conference, Paris Agreement, Trump’s withdrawal from the Paris Agreement, and the Katowice Climate Change Conference. Most of these events mark moments when international cooperation around climate change strengthened, with the one exemption being Trump’s withdrawal from the Paris Climate Accord. The renewable policy index is elevated during the Obama era, a period of strong support for renewables, and peaks during deliberations around Bush’s National Energy Policy, the decision to build the first US offshore wind farm, the announcement of the Green New Deal, Al Gore’s call for a move towards ending dependence on carbon, First Solar and China’s agreement to build the largest photovoltaic power plant, investigations into solar panel dumping by China, and the passage of Obama’s Clean Power Plan. Moreover, Noailly et al. (2021) finds their index generally measuring environmental policy saliency in the news is strongly correlated with the OECD’s Environmental Policy Stringency Index for the US.⁸ Throughout our paper we generally interpret increases in the international climate negotiations and renewable policy index as moments when expected US policy became more stringent.

3 Empirical Strategy

We test theoretical predictions of the “Green Paradox” by linking daily changes in oil futures to daily changes in market expectations around the probability that a cap-and-trade bill in the US will be enacted. If financial

⁸The OECD Environmental Policy Stringency Index measures the extent to which a country prices environmentally harmful behavior explicitly or implicitly.

markets expect climate change regulation to be more stringent in the future, resource owners will shift some of the supply they can no longer sell in the future towards the present, thereby lowering both the spot price as well as all futures prices with maturities between 1 and 24 months from today. Specifically, we are able to test the following two predictions: an increase in the probability that a cap-and-trade bill will be enacted (i) decreases prices and (ii) decreases prices by similar amounts for *all* futures contracts that are actively traded, i.e., for maturities ranging from 1 to 24 months into the future. We support findings from this analysis using broad trends in oil price shock persistence over the last three decades, examining the association between oil price returns and climate policy salience over the last three decades with monthly level data, and estimating oil price returns on days when news coverage peaked regarding historic climate court cases. These supporting analyses speak to the external validity of our findings by testing for (i) the pervasiveness of the “Green Paradox” throughout the last three decades and (ii) the “Green Paradox” against actual changes in climate policy.

The stickiness of price shocks induced by climate change legislation is a key feature of the “Green Paradox” that sets it apart from other temporary shocks whose effect phases out over time. We therefore begin our empirical analysis by documenting correlations between the oil spot price and oil futures prices at different points in time. This analysis serves two purposes. First, it provides descriptive evidence that the persistence of oil price shocks has changed remarkably over time, with the highest persistence during years when climate change legislation was debated. This is consistent with climate change legislation being a key factor of uncertainty that induces persistent price shocks. Second, this analysis constructs an important baseline for interpreting our causal results and motivation for robustness analyses.

Second, in our main analysis, we focus on US cap and trade policy deliberations in 2009 and 2010, a particularly uncertain period in terms of future climate policy, to provide causal evidence of the “Green Paradox.” Using high-frequency data, we link daily oil future returns to daily prediction market prices for a contract tied to the passage of a US cap and trade policy. As in Meng (2017), we argue that price changes in the prediction market approximate market expectations of the probability that climate regulation will be passed.⁹ We would expect increases in prediction market prices to be linked with negative oil future returns on average.

Third, we test if the surprising *Urgenda v. Netherlands* verdict and the prevailing news that Justice Kennedy, the key swing vote, would support the states in *Massachusetts v. EPA* were associated with anomalous negative oil future returns using an event study framework. The *Urgenda v. Netherlands* ruling indicated future climate regulation in the Netherlands would become more stringent, and people speculated that the verdict laid the foundation for more stringent climate regulation in Europe more broadly. The *Massachusetts v. EPA* ruling stipulated that the EPA has the power to regulate greenhouse gas automobile emissions. The ruling did not force the agency to regulate said emissions; however, people speculated that if the EPA failed to regulate greenhouse gas emissions, the agency would face additional legal battles. Thus, both the *Urgenda v. Netherlands* verdict and the prevailing news that Justice Kennedy would back the states created unexpected shocks in market expectations around future climate regulation stringency, giving us the opportunity to test for evidence of the paradox in response to information shocks. In both cases, we would expect the consequential new information revealed by the verdict announcement and the media to be

⁹Meng (2017) pioneered the use of prediction markets for climate change legislation in a different context to derive the abatement costs of climate legislation.

linked with negative oil future returns.

Finally, we continue our empirical analysis by documenting associations between oil price returns and monthly measures of renewable policy and international climate negotiation salience in the news. The stringency of future climate policy has been in continual flux in recent history. The “Green Paradox” predicts that information shocks informing expectations around the stringency of future climate policy should impact oil prices, i.e., information shocks indicative of more stringent future climate policy should reduce prices today. As discussed in the data section, the news based indices predominantly track events that strengthened future climate policy. Hence, the indices proxy for information shocks indicative of more stringent future climate policy, and our analysis can provide suggestive yet compelling evidence of a relationship between increases in the expected stringency of future climate policy and oil prices. More generally, this analysis tests if the “Green Paradox” has been a pervasive phenomenon throughout the last three decades.

Oil Price Shock Permanence In a first step, we link daily changes in oil futures to daily changes in oil spot prices. As described in the data section, the analysis includes oil futures with maturities ranging from 1 to 24 months into the future, i.e., for each day we have 24 different future returns. The regression equation is:

$$\Delta y_{ft} = \alpha_{fm(t)} + \beta_f \Delta p_t + \gamma_f \Delta z_t + \epsilon_{ft} \quad (1)$$

We regress Δy_{ft} , the percent change in the future price¹⁰ on day t for the oil future with a maturity of $f = 1 \dots 24$ months into the future on the percent change in the oil spot price Δp_t on day t . The main coefficients of interest are β_f , the association between the spot price and oil future f . We allow the coefficients to vary by maturity f to show how spot price shocks impact expected prices at various points in the future. We control for changes in the overall economy by including Δz_t , the percent change in the S&P 500 index on day t . The effect of the S&P 500 is also allowed to vary by future (maturity) f . Finally, we include future-by-month fixed effects $\alpha_{fm(t)}$, thereby focusing the identification on changes in a particular oil future f that occur within a given month $m(t)$. Errors ϵ_{ft} are clustered by day allowing the returns of the 24 different maturities to be correlated as they might be influenced by the same market events. In a sensitivity check in the appendix we allow for different clustering options and fixed effects.

We run the analysis for different subsets of days. The degree to which climate change legislation was at the forefront of political agendas varied greatly over the last three decades. As discussed previously, we expect oil price shocks to be more permanent during periods when climate regulation is under consideration by major governing bodies, as the passage of such legislation would permanently alter future price paths. Thus, we replicate the analysis described by equation (1) for periods that begin every 5-years, starting in 1990 to see if oil price shock stickiness varied with uncertainty around future climate legislation.

Cap & Trade Prediction Market We link daily changes in oil futures to daily changes in prediction market prices, which capture the market’s assessment of the probability that a US climate bill would pass. During the prediction market’s lifespan, the contract price ranged from 0 to 57 cents, implying the market predicted between a 0 and 57% chance that the US government would enact a cap-and-trade bill by the end of 2010. The price peaked at 57 cents when the House passed the Waxman-Markey bill. The price increased by 10 cents when a cap-and-trade bill in the Senate garnered support from some Republicans, including

¹⁰We use the percent change in the closing price relative to the previous closing price.

Lindsay Graham of South Carolina (Meng 2017). While the main bill that was discussed in 2009-2010 was the Waxman-Markey bill, the prediction market contract was for the event that *any* cap-and-trade bill passed by the end of 2010. We include returns between May 1, 2009 and Dec 31, 2010, as these are the dates with prediction market prices. The regression equation for the pooled effect is:

$$\Delta y_{ft} = \alpha_{fm(t)} + \beta \Delta x_t + \gamma_f \Delta z_t + \epsilon_{ft} \quad (2)$$

As in equation (1), we use Δy_{ft} , the percent change in the future price on day t for the oil future with a maturity of $f = 1 \dots 24$ months into the future. We now regress the percent change in the future price on the change in the prediction market probability Δx_t of a cap-and-trade bill passing¹¹ on day t . The main coefficient of interest is β , the effect of changes in the probability of the bill passing on oil future returns. All other controls and clustering are identical to equation (1).

We run the analysis for different subsets of days. Oil futures might be especially responsive on days when there are major changes in the probability of a bill passing and hence major changes in the price of the prediction market. We conduct the analysis for all days as well as for subsets of days when the absolute change in the probability of a climate bill passing exceeds various cutoffs ranging from 1 to 5 cents. We do this in two ways, first by limiting the subset of observations:

$$\Delta y_{ft} = \alpha_{fm(t)} + \beta \Delta x_t + \gamma_f \Delta z_t + \epsilon_{ft} \quad \text{for days where } |\Delta x_t| \geq c. \quad (2a)$$

and second by relying on all observations but including an interaction term with a dummy for days below the cutoff $\mathbb{1}_{[|\Delta x_t| < c]}$ to test whether the effect is different for small changes:

$$\Delta y_{ft} = \alpha_{fm(t)} + \beta_c \Delta x_t + \beta_2 \Delta x_t \mathbb{1}_{[|\Delta x_t| < c]} + \gamma_f \Delta z_t + \epsilon_{ft} \quad (2b)$$

The coefficient β in equation (2b) measures the sensitivity on days where the prediction market moved by at least c cents, while days below the threshold have a sensitivity equal to $\beta + \beta_c$.

In an additional auxiliary analysis, we instead classify days based on Google Trends data for the search term “waxman markey.” We thereby restrict the analysis to days where there was broad interest in the topic by the general population. The Google Trends filter is applied to the day of and day after a price shock as the prediction market may internalize new information faster than the general population. Finally, we present a sensitivity-check where we further limit the analysis to days when the trade volume exceeds the median, ensuring the prediction market was liquid enough to incorporate updates in the probability of the bill passing.

We further relax the linearity and symmetry assumption of oil prices’ response to prediction market shocks, modelling it using restricted cubic splines $g(\Delta x_t)$ with knots at -5, -2, 0, 2, and 5 percent. Such an approach forces the response to be linear below the minimum knot (-5) and above the maximum knot (+5), but uses third-order polynomials in between. We expect the effect of a positive prediction market shock on oil prices to be negative and a negative shock to be positive; however, this approach allows the response to be asymmetric. A positive shock could have a negative effect while a negative shock can have no effect. The

¹¹The prediction market values range from 0 to 100 cents, which in an efficient market should reflect the probability of the bill passing. We take the difference in the closing price relative to the previous closing price, thereby obtaining the change in the probability of the bill passing.

equation becomes:

$$\Delta y_{ft} = \alpha_{fm(t)} + g(\Delta x_t) + \gamma_f \Delta z_t + \epsilon_{ft} \quad (2c)$$

Returning to the linear model, as in our analysis of oil shock persistence, we allow for heterogeneity of the effect β by oil future maturity, i.e., β_f , while all other controls remain identical to equation (2). The modified regression equation is:

$$\Delta y_{ft} = \alpha_{fm(t)} + \beta_f \Delta x_t + \gamma_f \Delta z_t + \epsilon_{ft} \quad (2d)$$

In the appendix, we include the yield curve for S&P 500 futures, specifically, the change in the difference between a future with 8 months maturity and the current index. This test controls for confounding variation in expected economic growth, e.g., if the climate bill was part of a larger set of regulation that induce short to medium-term slowdown in growth that would reduce the price of oil through a reduction in demand.

Historic Climate Court Cases We examine oil future returns on the days when major news regarding historic climate court cases broke. Specifically, we analyze the abnormal return on the day *Urgenda v. Netherlands* was rendered (June 24, 2015) and the day when the news broke that Justice Kennedy would likely back the states suing the EPA (December 6, 2006).

Urgenda v. Netherlands was the first successful climate liability suit brought under human rights and tort law. In the ruling released on June 24, 2015, the judge acknowledged that climate change’s threat was severe and stated that under Dutch law a threat of damage suffices for injunctive relief. The verdict stated that by the end of 2020, the Dutch state had to reduce greenhouse gas emission by at least twenty five percent relative to 1990 levels. The *Urgenda v. Netherlands* ruling was unexpected, notable and historic. In the six months leading up to the verdict announcement, the court case’s Google Trends index rarely exceeded zero indicating there was not enough search volume to determine the court case’s relative popularity. On only six days, the Google search index exceeded zero reaching at most one third the level of popularity reached on June 24th. On the 24th, the New York Times article “Ruling Says Netherlands Must Reduce Greenhouse Gas Emissions” quoted Marjan Minnesma, the director of *Urgenda*, saying “Everybody in the legal scene said, ‘This will never happen — this is just a P.R. stunt.’ This is not a P.R. stunt.” The same article quoted Michael Gerrard, director of the Sabin Center for Climate Change Law at Columbia University, saying “I think this will encourage lawyers in several other countries to see if they have opportunities in their domestic courts to pursue similar litigation.” People predicted the verdict set a precedent that other countries subject to the European Convention would follow. Given the Dutch share of global fossil fuel consumption is minimal, the estimated effect is driven by an update in the probability that other countries would follow suit. There have subsequently indeed been similar suits, some of them successful, e.g., a suit brought against the state of Montana.

Massachusetts v. EPA was a landmark case in which states and environmental groups argued that the EPA was obligated under the Clean Air Act to regulate carbon dioxide in automobile emissions. The case was argued on November 29, 2006 and decided on April 2, 2007. Similar to the *Urgenda v. Netherlands* ruling, people predicted the *Massachusetts v. EPA* verdict would be followed by numerous additional climate cases, including a case challenging the EPA’s refusal to regulate power plant carbon dioxide emissions.¹² Unlike the *Urgenda v. Netherlands* ruling, the likely outcome in the *Massachusetts v. EPA* case was actively discussed

¹²<https://www.nytimes.com/2007/04/03/washington/03scotus.html>

before the ruling, and the eventual ruling was arguably expected beforehand. In the six months leading up to the verdict announcement, the court case peaked in popularity on Google Trends on December 6, 2006, when headlines read “Key Justice Appears to Back States’ Standing to Sue in CO₂ Case.” Numerous news articles published on the 6th, found using a Lexis-Nexis search on the terms “Massachusetts v. EPA”, documented how Justice Kennedy, the key swing vote for the case, appeared to be backing the states suing the EPA. This is the date when the market should have incorporated new information, not the date of the eventual ruling, highlighting again the importance of accurately identifying when market beliefs update.

In both instances, we do not have a direct estimate on the market update in the probability that either court case would be settled in favor of restricting future fossil fuel use. Moreover, in the Urgenda v. Netherlands ruling, we do not have an estimate of the market’s belief that other countries would follow suit. Nor do we have an estimate of the market’s belief that the Massachusetts v. EPA ruling would be followed on by additional court case verdicts enabling additional fossil fuel use restrictions in the US. Hence, our coefficient estimates for the event studies in response to both court cases are not directly comparable to our prediction market analysis, in which we are able to identify the effect of a 100% change in the probability of Waxman-Markey by scaling by the change in prediction market probability. Our event study estimates should be scaled (divided) by the change in the probability to be comparable, a statistic we do not observe.

Our analysis of court cases provide another source of identification using discrete events documented by Google Trends data rather than a change in probability documented by a prediction market. Some researchers have levied concern that prediction markets do not reliably reflect market beliefs (Manski 2006, Fountain and Harrison 2011, Lemoine 2017). Our analysis offers a cross-check using major news updates without having to rely on prediction markets: we are interested in the sign and significance of the coefficient, but the magnitude of the coefficients requires further assumptions to be comparable to our primary analysis.

As in the previous equations, the analysis includes oil futures with maturities ranging from 1 to 24 months. We construct the abnormal returns on the Urgenda v. Netherlands verdict announcement date, June 24, 2015, and the day when the prevailing news regarding the Massachusetts v. EPA case suggested Kennedy, the swing vote, would support the states suing the EPA, December 6, 2006. As documented in Figure A1, the permanence of oil price shocks varied significantly between 1990 and 2022. We compare returns on June 24th, 2015 and December 6, 2006 to five different subsets of days to demonstrate robustness to these heterogeneous comparison groups. The regression equation for the pooled effect is:

$$\Delta y_{ft} = \alpha_{f_{m(t)}} + \beta \mathbb{1}_v + \gamma_f \Delta z_t + \epsilon_{ft} \quad (3)$$

We regress Δy_{ft} , the percent change in the future price on day t for the oil future with a maturity of $f = 1 \dots 24$ months into the future on a dummy $\mathbb{1}_v$ for June 24, 2015, the day the Urgenda v. Netherlands verdict was rendered. The main coefficient of interest is β , the effect of the verdict announcement on oil futures. Other controls and clustering are identical to the previous equations. We conduct an identical analysis for the Massachusetts v. EPA court case replacing the dummy $\mathbb{1}_v$ with $\mathbb{1}_k$ for December 6, 2006, the day the news that Justice Kennedy would support the states’ position.

We subsequently allow again for heterogeneity in the effect β by oil future maturity, i.e., β_f , while all other controls remain identical to equation (3). The modified regression equation is:

$$\Delta y_{ft} = \alpha_{f_{m(t)}} + \beta_f \mathbb{1}_v + \gamma_f \Delta z_t + \epsilon_{ft} \quad (3a)$$

Climate Policy Saliency Unlike the futures data, which is available daily, the climate policy index is available monthly. We hence switch the analysis to a monthly level in an analogous fashion. We link monthly oil futures returns to monthly measures of renewable policy and international climate negotiation saliency, a proxy for information shocks about future climate policy. For easier interpretation, the monthly news indices are standardized to be mean zero and to have unit standard deviation. The news indices report monthly saliency measures between 1981 and 2019. The international climate negotiation index peaks during the Copenhagen Climate Change Conference at 7.5 standard deviations above the average and reaches its second highest peak during the adoption of the Paris Climate agreement at 6 standard deviations above average. The renewable policy index peaks when a group of US based companies first accused China of dumping solar panels in the US at 2.9 standard deviations above the average. The regression equation is:

$$\Delta y_{fm} = \alpha_{fq(m)} + \beta I_m + \theta R_m + \lambda E_m + \gamma_f \Delta z_m + \epsilon_{fm} \quad (4)$$

Δy_{fm} reflects the percent change in the future price at the end of month m relative to the end of month $m - 1$ for the oil future with a maturity of $f = 1 \dots 24$ months into the future. We regress the percent change in the future price on standardized measures of three news-based indices: I_m , R_m , and E_m .¹³ The main coefficients of interest are β and θ , the effect of a one standard deviation in the saliency of international climate negotiations (I_m) and renewable policy (R_m), respectively, on oil future returns. We include environmental policy saliency generally (E_m) to control for general environmental policy saliency and as a placebo.¹⁴ As in equation (1), we control for changes in the overall economy by including Δz_m , the percent change in the S&P 500 index at the end of month m relative to the end of month $m - 1$. The effect is allowed to vary by future (maturity) f . Finally, we include future-by-quarter fixed effects $\alpha_{fq(m)}$, thereby focusing the identification on changes in a particular oil future f that occur within a given quarter $q(m)$. Since we have monthly data, there are three observations per quarter for each maturity.

4 Empirical Results

Oil Price Shock Permanence Figure 2 highlights the persistence of oil spot price shocks on returns of oil futures across maturities. We regress daily changes in oil futures on the daily change in the spot price while accounting for daily changes in the S&P 500 as outlined in equation (1). The x-axis indicates the maturity ranging from 1 to 24 months into the future, while the y-axis gives the point estimate as well as the 90% confidence band. The displayed results reflect the average shock persistence, the combination of all shocks, both transitory shocks as well as permanent shocks. Estimates using all days between 1990 and 2022 are shown in dark blue. Contracts with a maturity in one month have a coefficient of 0.95, i.e., on average 95% of the change in the daily oil spot price is reflected in the futures price with a one-month maturity. On the other

¹³Unlike our other analyses, we do not first-difference these indices. While prices are first-difference stationary, the environmental policy index is trend stationary. In other words, we would expect prices from one day to the next to be roughly equivalent, so our analysis relies on variation in prices from one day to the next. We do not have a similar expectation for the environmental policy index; news coverage of environmental policy in one month has little bearing on the amount of coverage in the preceding or proceeding months. Recall that we are including maturity-by-quarter fixed effects and hence limit the analysis to the three monthly observations within a quarter, but unlike the case of first-differencing, we also use the comparison how much the third month differs from the first.

¹⁴We do not expect environmental policy saliency on average to impact oil prices because on average most environmental policy news is not about climate policy (Noailly et al. 2021).

extreme, only 13% of the change in the daily oil spot price is reflected in the future price with a maturity 24 months (2 years) away. Maturities in between decay roughly exponentially, suggesting that on average spot price shocks during the sample period were temporary and phased out over two years. The light blue (cyan) line replicates the same analysis using only days between May 2009 and December 2010, days when the cap-and-trade prediction market was active. During this period the House passed Waxman-Markey and the Senate deliberated a very similar bill. For this period, the coefficient is 0.90 for a one-month maturity and 0.58 for a 24-month maturity. In other words, when climate legislation was under consideration, shocks to the spot price were on average much more persistent, phasing out much slower relative to the whole sample. This is consistent with the theory of the “Green Paradox”, as climate change legislation should reset the *entire* expected price path and lead to permanent rather than transitory changes. Uncertainty around future legislation should yield stickier price shocks rather than transient shocks.

Appendix Figure A1 further splits the sample period into five year intervals ranging from 1990-1994 to 2015-2019, as well as a 3-year end period 2020-2022. Shock persistence increases from 1990-1994 to the 2010-2014 period, when climate legislation was most actively discussed. The persistence slightly declines again in 2015-2019, before collapsing in 2020-2022, as COVID-related supply disruptions lead to short-term price fluctuations – the oil price was briefly even negative when storage levels reached capacity. Notably, shocks were most persistent during the period when cap and trade policy was heavily considered by the US Federal Government and the European Union established an emission trading system. The average decay in the translation of spot price shocks to changes in futures with various maturities sets an important baseline for following analyses when we again present results and how they vary across maturities.

Cap & Trade Prediction Market Table 1 highlights the average effect of changes in the probability of a cap and trade policy passing on expected oil prices, oil futures contracts that would have expired before the policy entered into force. Specifically, we regress daily changes in oil futures prices on daily changes in prediction market prices, a market-based measure of the probability that a climate bill passes, as outlined in equation (2a) and (2b), respectively. We display coefficient estimate β (and β_c in Panel B), the effect of changes in the prediction market on future prices, while we suppress other coefficients (future-by-month fixed effects and future-specific controls for the movement of the S&P 500).¹⁵ Columns differ by the cutoff c to classify “major” news updates. The first column in panels A as well as all panels in panel B use all days between May 2009 and December 2010, the period when the prediction market was active. A coefficient of -3.44 suggests that changing the probability of the bill passing by December 2010 from 0% (certainly not passing) to 100% (definitely passing) decreases oil futures prices by 3.44%. Recall that the prediction market contract is on a climate bill passing by the end of 2010, so market participants might still expect a bill to pass at a later point. The estimated coefficient is hence for a climate bill net of later subsequent expected climate bills. Nonetheless we find a sizable effect of 3.44 percent.

Columns (2)-(6) of Panel A in Table 1 restrict the sample to days when the absolute change was at least 1, 2, 3, 4 or 5%, respectively, i.e., when the prediction market saw increasingly major updates. Accordingly, the number of days in the analysis successively decreases, but the coefficient estimate increases in magnitude from -3.44 in column (1) when we include all days to -7.08 in column (6) when we use only days that had at least a 5% change in the prediction market price. We would expect the oil market to be especially responsive

¹⁵Figure A2 displays the evolution of the sensitivity to the S&P 500 over time.

to highly visible events, e.g., when Republican Senator Lindsay Graham withdrew his support, and we would expect the oil market to be less responsive to smaller day-to-day changes in the probability of a bill passing as small fluctuations might be a result of round number bias and the bid-ask spread of a market with limited liquidity. Column (7) restricts the data to when the internet search volume for “Waxman Markey” on Google Trends was at least 5% relative to the day with the highest search volume. Table A2 documents the sensitivity of the results in Table 1, Panel A to various clustering options. Our baseline approach generally results in the most conservative significance levels, i.e., widest standard errors.

Panel B uses all days across all columns, but adds an interaction term for days when the prediction market movement (or Google Trend index) is below the corresponding cutoff used in panel A. Both panels control for maturity-specific sensitivity to the S&P500, as oil prices are sensitive to overall economic conditions. By restricting the data to a few days in column (6) of panel A, these controls are identified using very few days and hence estimated with lots of noise. The advantage of panel B is that we always use the entire period to construct these controls consistently and derive the “abnormal” oil futures return net of overall market movements. While the coefficients slightly increase from column (1) to column (6), the increase is less pronounced.

As outlined in Meng (2017), major prediction market movements were associated with politicians joining or abandoning the bill after negotiation rounds that were scheduled in advance. Hence, their timing was unlikely to have aligned with daily oil price changes, ruling out reverse causality. Moreover, public statements about why politicians joined or abandoned the bill do not mention contemporaneous oil prices. The New Yorker had a background story¹⁶ that outlined the key events that led to the unraveling of the coalition supporting the climate bill. Oil prices as well as oil price changes are never mentioned. Rather political events that are not related to oil prices caused most fluctuations in the probability of the bill passing. Table A1 lists news stories for the 26 days where prediction market prices changed by at least 5%, i.e., the days used in column (6) of Panel A in Table 1. While we do not know what exact news the prediction market responded to, positive changes are usually associated with news stories in which the bill’s sponsors speak up in its support, while negative changes are associated with opponents voicing their dissent. These events sometimes occur on consecutive days with opposite signs. Importantly, the timing of these events is unlikely to be related to day-to-day oil price fluctuations, but rather the result of committee meetings that were scheduled days in advance.

One alternative explanation of the observed relationship is that climate regulation limits future economic growth, thereby lowering fossil fuel demand and resulting in more fossil fuel being consumed in the present. We distinguish between this channel and future supply restrictions in Table 2. Panel A of Table 2 replicates Panel B of Table 1 with one modification: it only allows for a common control of how the S&P 500 affects oil futures of various maturities.¹⁷ We find a highly significant relationship between the S&P 500 and oil futures prices, with a coefficient of 0.84. If the S&P 500 increases by one percent, oil futures increase by 0.84 percent. Panels B1-B3 replicate the analysis for three metals: copper, aluminum, and steel. Copper futures have the same sensitivity (0.84) to fluctuations in the S&P 500 as oil, while aluminum’s sensitivity is slightly lower at 0.67, yet steel shows no significant relationship at all. This is not surprising: while iron ore is an exhaustible

¹⁶“As the World Burns: How the Senate and the White House missed their best chance to deal with climate change,” October 11, 2010 issue.

¹⁷Recall that Table 1 allowed for maturity-specific controls, resulting in 24 coefficients.

resource, it is much more abundant than oil, copper, and aluminum.¹⁸ Today’s consumption of steel is less influenced by projections of economic growth given its relative abundance. On the other hand, a reduction in future growth will also influence demand for copper and aluminum given its comparable dependence on the S&P500; however, there is no significant response as observed in the case of oil. This suggests that the relationship observed for oil is due to the limitation of future oil use rather than a deterioration of economic growth. It also highlights that our results are unlikely to be driven by overall events that influence both the S&P 500 as well as the oil price, as these events should also have affected copper and aluminum prices.

Table A5 presents another possible test using more immediate change in expectations of economic growth/ For instance, a slowdown in economic growth, e.g., a recession, leads to a reduction in the oil price and an increase in the probability that the climate bill would pass. This would be a demand-driven rather than a supply-driven response. Table A5 not only includes changes in today’s S&P 500 as a control but also changes in the yield-curve, which we define by the difference between the S&P 500 future with a maturity in eight months and the present day index value.¹⁹ Controlling for daily changes in this yield curve has no discernible effect on the estimated coefficients, making a story that changes in expected future economic growth drive our results unlikely. Robustness to controlling for the yield curve also confirms our earlier argument that the events with major prediction market movements shown in Table A1 were mostly driven by political events. The change in the S&P 500 index on the 26 days, in which the prediction market changed by at least 5%, has a positive but insignificant correlation with changes in prediction market prices. This is inconsistent with a story where changes in economic growth lead to changes in oil demand and prediction market prices – such a story implies a negative correlation.

Figure 3 relaxes the linearity and symmetry assumptions by estimating a flexible response function using restricted cubic splines from equation (2c) using all observations, i.e., the data from column 1 of Table 1. The green histogram displays the density of the prediction market movements – there is a clear mass point at zero, i.e., days where the prediction market had no price change. The spline implementation forces the response of oil prices to prediction market movements to be linear below the lowest knot (-5) and above the highest knot (5). However, within $[-5, 5]$, the relationship is allowed to evolve flexibly using third-order polynomials, yet the linearity assumption still holds reasonably well between $[-5, -1]$ and $[1, 5]$. For small changes in probability $[-1, 1]$ the linearity assumption does not hold as well, which might not be surprising as the market has limited liquidity, there is a bid-ask spread, and prices are predominantly discrete integer values. In this sense, changes smaller than one might not reflect belief updates and may lead to attenuation bias. These results support our choice to use a linear model, assuming symmetric effects of positive and negative shocks, and to focus on estimates that restrict our sample to price changes larger than 1.

Table 3 conducts a falsification check where we purposefully offset the prediction market prices by one period into the future, e.g., regress today’s oil future return on tomorrow’s change in prediction market prices.²⁰ None of the coefficients are statistically significant and the sign switches from negative to positive for some of the estimates. We see this as additional evidence that the daily changes in prediction market

¹⁸Current global oil reserves are roughly 50 times as large as annual global oil consumption. The ratio of reserves to current consumption is more than 70,000 for iron ore <https://pubs.usgs.gov/periodicals/mcs2022/mcs2022-iron-ore.pdf>. The sensitivity of oil and copper prices to the S&P 500, a proxy for economic growth, showcases the underlying rationale behind the “Green Paradox,” i.e., that the optimal allocation of a scarce resource depends on its demand across time.

¹⁹While oil futures are traded up to 24 months into the future, the longest maturity for S&P 500 futures is 8 months into the future and we hence rely on it.

²⁰Recall that we construct daily returns using consecutive closing prices. In most cases the period between consecutive closing prices is one day, except for holidays or weekends where the period covers two or three days between consecutive closing prices.

prices convey significant information, at least in the eyes of oil futures’ traders that respond to them. Our results are not driven by periods of high versus low prediction market prices, as overall prediction market price levels are preserved when offsetting the information by one day. Instead, our results are driven by day-to-day changes in prediction market prices. We also present results of a distributed lag model in Table A6 where we include prediction markets in the current as well as next day’s period. The prediction market (Intrade) is based in Ireland, which predates the timezone of the US.²¹ We find the effect for the concurrent day, not the next day.

Table A4 limits the analysis to days with above-median trade volume, restricting the sample to days when the contract is actively traded, i.e., the market is liquid enough to accurately reflect the market’s assessment.²² This sensitivity check also ensures that a price jump is not the result of one “rogue” trade (sticky-finger) but rather occurred due to active trading. When we impose the above-median trade volume restriction, the estimated coefficients slightly increase in magnitude in Panel B relative to Table 1. As in our primary analysis, columns vary in the sample of days we use to estimate the effect of prediction market price changes on oil futures. Column (1) does not restrict the sample based on the absolute change in the prediction market price, while column (2) restricts the sample to days when the absolute change was at least 1%. In addition, we restrict the sample used in all columns to days with above-median trade volume. Imposing this restriction greatly reduces the number of days (clusters) in column (1), as days with zero price change often have zero trade volume. While the point estimate remains rather robust, the significance level drops somewhat. As we restrict the sample to days with absolute price changes above a certain threshold in columns (2)-(7), requiring an above-median trade volume does not significantly alter the set of included days: larger price jumps go together with above-median trade volume as new information gets priced in.

The “Green Paradox” is a supply shift, which implies that prices and storage levels should move in the same direction. If future climate legislation shifts supply to the present, both the current price and current storage levels should decline as the optimal allocation over time is adjusted. The reduced price increases current-period consumption, which is achieved through depleting storage levels as production cannot be increased in the very short-term. On the other hand, demand shocks generally lead to an inverse relationship between oil prices and storage levels. An inward shift in demand reduces prices but leads to a buildup in storage levels, as production cannot be adjusted in the short-term. Table 4 further aims to distinguish between demand and supply shocks. While oil price data is available on a daily level, storage and production numbers are only available at the weekly level. Table 4 replicates the same analysis for oil storage levels. The regression coefficients are again negative with one exception (column 6 in Panel A), but that specification rests on very few observations and has wide standard errors. Coefficients in Panel B are statistically significant for the first five of the six columns. The reduction in storage levels while prices decline points towards a supply rather than a demand response.

While consumption and storage respond immediately, increasing oil production takes time as new wells need to be drilled. Figure 4 therefore presents the response of Table 1 but allows the effect to vary by maturity. Panels A and B use the same models as in the Table, i.e., one using only days where market movements or internet search volume exceed a threshold (Panel A) or where a separate interaction term

²¹One technical issue we noticed in the Intrade data that Derek Lemoine and Kyle Meng shared with us is that they are offset by one day. We rely on the day definition of Kyle Meng in this paper.

²²We thank Derek Lemoine for sharing the trade volume data with us, which is highly rightward skewed: the highest trade volume is 2901, while the median for days when trades occurred (dropping days without any trade) is 12.

is included (Panel B). The different cutoffs are now shown in different colors. In both panels, the x-axis displays the results by maturity, while the y-axis gives the point estimates as well as the 90% confidence band. Similar to the pooled analysis, the coefficients become larger in magnitude when the data is limited to days when the prediction market price changed by at least 4% (orange and red color in Figure 4 or column (5) and (6) in Table 1). Remarkably, estimated coefficients *weakly increase* in magnitude for maturities that are further into the future. The effect of deviations in prediction market prices is smallest on the oil spot price (not shown) relative to oil price futures with maturities further into the future. This pattern is in sharp contrast to the average relationship between temporary shocks that phase out over time. Price adjustments associated with prediction market fluctuations are not phasing out but rather phasing in. This is consistent with Anderson et al. (2018)’s finding that oil production from *existing* wells do not respond to oil prices in the short-term due to physical constraints, i.e., closing an existing producing well is costly, explaining why oil prices briefly turned negative during COVID shutdowns. Letting the oil flow, once a well is tapped, is usually the most economical irrespective of short-term price dynamics. Kellogg (2014) shows that drilling decisions react quickly to changes in market expectations, but it again takes a while for the oil to flow. Notably, development of infill wells respond to oil future contracts with 18 months to maturity. While the standard Hotelling model suggests that the price path will reset right away in response to a climate bill passing, physical constraints (both in closing existing wells and in opening new ones) imply that it can take a few months for adjustments to be realized, explaining why the coefficients get larger in magnitude for maturities that are further into the future. The overall effect of the climate legislation shock needs time to materialize as new oil production cannot be turned on or off overnight. That said, short-run adjustments can be made through altering the amount of oil in storage, which may reconcile why we see expected oil price changes in as little as one month.

Table A7 replicates the analysis for coal futures prices. We again regress daily changes in futures prices, this time from coal contracts, on daily changes in prediction market prices, a market-based measure of the probability that a climate bill passes. The relationship is not statistically significant when we use the full sample to estimate the effect in column 1. We find that changing the probability of the bill passing by December 2010 from 0% (certainly not passing) to 100% (definitely passing) decreases coal futures prices by 1.02%. As we restrict the number of days in the analysis, to days when the absolute change in prediction market prices exceeded increasingly restrictive cutoffs, the coefficient estimate increases in magnitude from -1.02 in column (1) when we include all days to -5.50 in Panel A, column (6) when we include only days that had at least a 5% change in the prediction market price. On the other hand, none of the coefficients in Panel B are statistically significant, albeit all are negative.

One possible reason why coal futures prices might show a lower responsiveness to prediction market prices is that coal use in the U.S. already faces binding future constraints due to air quality regulation. In nations such as the U.S. and Germany, coal has been phased out to improve local air quality. If air quality regulation limit future coal use by more than the discussed domestic climate policy, the latter will be non-binding and changes to its probability will not impact coal prices. Consistent with this line of argument, the repressiveness of coal futures to changes is much lower for coal (0.22) relative to oil (0.84) reported in Table 2. That said, demand for coal in China and India continues to grow. Assuming sufficiently low transport and trade costs, the paradox’s predictions should still apply to coal.

In summary, our analysis of the cap & trade prediction market provides empirical evidence that oil

markets respond to changing expectations about climate legislation that limit future fossil fuel consumption. What are the policy implications of our findings? First, we can derive the effects if the law had passed on global oil consumption. Using the average long-term demand elasticity of -0.6 from Hamilton (2009, Table 3), the price coefficients in Table 1 imply that the passage of the cap and trade policy considered in the US in 2009-2010 would have increased *global* oil consumption in that year, i.e., before its binding constraints go into effect, by 2.0%²³, accelerating the depletion of the resource. The projection reflects the market’s best guess of how US regulation and follow-on regulation in other countries would have shifted supply towards the present. Second, even though the legislation never passed, its discussion temporarily altered the oil price path. We derive the additional oil consumption induced by climate policy deliberation with a back-of-the-envelope calculation. To begin, we derive the difference in the daily oil price caused by climate policy deliberation relative to a counterfactual in which the bill was never discussed by multiplying the price coefficient associated with the one month maturity in Figure 4 by the prediction market price. We then compute the effect on oil consumption by multiplying the daily price differences with the average short-term demand elasticity of -.26 from Hamilton (2009, Table 3). As the probability of the law passing increases, the price is suppressed, leading to additional oil consumption, which in our back-of-the-envelope calculation is simply taken to be the coefficient estimate times the prediction market prices times the demand elasticity. As the probability of the law passing falls back to zero, the price path returns to the trajectory of the initial undisturbed price path and our calculation attenuates. The total effect of climate policy deliberations on consumption is then the sum of the temporary increase in oil consumption due to the temporary price reduction for the time period for which the prediction market data is available, i.e., May 1, 2009 and Dec 31, 2010. The combined *additional* oil consumption is roughly 10 million metric tons, equivalent to the daily global oil consumption, highlighting how uncertainty about legislation, even if never passed, can still influence fossil fuel consumption.

Historic Climate Court Cases We derive the oil market’s response to surprising news regarding two historic climate court cases. Specifically, we estimate the effect of the surprise *Urgenda v. Netherlands* ruling on June 24, 2015 as well as the effect of prevailing news coverage on December 6, 2006 that Justice Kennedy, the swing vote, would support the states suing the EPA. The results for the event study examining the abnormal return on June 24, 2015 as specified in equation (3) are shown in Panel A of Table 5. The results for December 6, 2006 are shown in Panel B. As always, in both panels, we control for overall market movements by controlling for the daily returns of the S&P 500, which are allowed to vary by maturity of the future contract. Different columns use different time spans around the event day. While the estimated coefficient is always for one day, June 24, 2015 in Panel A, and December 6, 2006 in Panel B, the inclusion of further days around the event itself influences the coefficient estimates γ_f in equation (3) and hence the prediction of the “normal” return on that day, which forms the basis for constructing the abnormal return.²⁴ In Panel A (Panel B), the first column uses the smallest time period, year 2015 (2006), while the second column adds two years before and after 2015 (2006), and the third column adds an additional 2 years on either side. The fourth column uses all days between 1990 and the end of 2019 and hence stops before the COVID-related disruptions. Finally, the fifth column also adds the COVID years. The estimated effect of the

²³We obtain this number by multiplying the coefficients from Table 1 by the demand elasticity.

²⁴The inclusion of further days around the event itself in equation (3) would have no impact on the estimated coefficient in the absence of controls for changes in the overall economy.

Urgenda v. Netherlands ruling is always negative and significant ranging from -0.55 to -0.9. The estimated effect of the news that Justice Kennedy would support the states suing the EPA is always negative and significant ranging from -0.47 to -0.55. Figure A4, Panel A (Panel B) allows the pooled effect of Table 5, Panel A (Panel B) to vary by maturity. Coefficients increase with maturity until about 4 months and then start to decrease again.

Coefficients in Table 5 are not directly comparable to the point estimates in Table 1, which were scaled to reflect the impact of a change in the probability of a US cap and trade bill from certainly not passing (0% probability) by December 2010 to it passing with certainty (100% probability). As outlined above, the Netherlands account for a small fraction of global emissions, and the bigger issue of the ruling was whether courts in other countries would follow suit. Similarly, the Massachusetts v. EPA ruling stipulated that the EPA had the authority to regulate tail pipe carbon dioxide emissions, and the bigger issue of the ruling was whether EPA would regulate greenhouse gas emissions broadly. Moreover, unlike the Urgenda v. Netherlands case where a ruling was rendered, the Kennedy comments offered an early indication on how the justice was leaning, but did not offer certainty on what the ruling eventually would be.

To make the coefficients between the court cases and prediction market tables comparable, the estimates of Table 5, Panel A would need to be divided by the change in probability that enough other countries adopt similar measures to add up to the same oil use restrictions as Waxman-Markey. Panel B estimates would need to be divided by both (i) the change in probability that the court sided with the suing parties as well as (ii) the probability that EPA's regulatory action of greenhouse gas emissions under the Clean Air Act would add up to the same oil use restrictions as Waxman-Markey. We unfortunately do not know these probabilities, but instead note the inverse: if Urgenda v. Netherlands caused a 12-20% increase in market beliefs that other countries would reduce emissions to a similar extent as Waxman-Markey and potential follow on policies, then the coefficient estimates from the two tables would be consistent.²⁵ Similarly, if news that Justice Kennedy supported the states caused a 12% increase in market beliefs that future EPA regulation of greenhouse gas emission would restrict future oil use to a similar extent as Waxman-Markey, then the coefficient estimates from the two tables are consistent.

More importantly, both daily changes in the prediction market, the abnormal return on the day of the surprise Urgenda v. Netherlands ruling, as well as the news that Justice Kennedy would support the states provide strong evidence that the financial market quickly updates its beliefs about possible future restrictions to oil use, with implications for the optimal extraction and price path. Notably, restriction in the future leads to lower prices and more consumption today, offsetting some of the savings in the future through higher consumption today.

Climate Policy Salience In a final step we examine a time series spanning several decades that provides additional suggestive evidence of the “Green Paradox.” Table 6 gives the average effect of climate policy salience in the news, generally a strong proxy for moments when future climate policy strengthened, on oil future returns. Specifically we regress monthly changes in oil futures prices on three monthly indices measuring the prevalence of articles covering environmental policy generally, international climate negotiations, and renewable policy, as outlined in equation (4). Indices are standardized to be mean zero and have unit standard deviation for interpretability. We display coefficient estimates β , θ and λ that are jointly estimated.

²⁵Dividing the coefficients of Table 5, Panel A by 0.12-0.2 gives similar estimates to Table 1.

The first row displays β , the effect of a one standard deviation increase in the international climate negotiations index on future prices. The second row gives θ , which reflects the effect of a one standard deviation increase in the renewable policy index on future prices. Finally, λ displayed in the last row, reflects the effect of a one standard deviation increase in the environmental policy index on future prices. Columns differ in the included controls. The first column includes only quarter by year fixed effects. The final column includes quarter by year by maturity fixed effects and controls for monthly changes in the S&P 500 index.

We find a one standard deviation increase in the international climate negotiations index is associated with a 0.76-0.81% decrease in expected oil prices for various maturities over the next two years as shown in the first row of Table 6. For reference, the Paris Climate agreement is associated with an index six standard deviations above the mean. While only suggestive, this finding is consistent with “Green Paradox” predictions. Moments demonstrating international cooperation around climate change enforce expectations that climate policy will be more stringent moving forward, increasing production in the near term and decreasing prices.

We find a one standard deviation increase in the renewable index has no significant effect on oil prices across maturities as shown in the second row of Table 6. There are two effects at work that move in opposite directions. On the one hand, renewable policies that lead to technological progress, ultimately lowering backstop prices, should decrease oil prices in the short-run as the increased competition with renewable technology in the future shifts oil supply to the present. On the other hand, if the passage of renewable legislation is seen as a substitute for climate legislation, thereby lowering the probability of future fossil fuel restrictions, it should increase the oil price today as market participants anticipate less restrictions in the future. T

Finally, we find that our placebo measure, the general environmental policy index, is not significantly correlated with oil prices as shown in the third row of Table 6. The environmental policy index tracks the salience of environmental policy generally. From 1981-2000, the most prevalent topics covered by the index include water and air pollution as well as court cases and clean ups (Noailly et al. 2021). Once controlling for the salience of international negotiations and renewable policy, the index predominantly reflects the salience of non-climate environmental policy in the news proxying for moments when non-climate related environmental policy became more stringent. Increases in the stringency of non-climate related environmental policy should have no effect on oil prices, except through the indirect effect if they are to limit economic growth and hence change the demand for oil. The resulting null effect is reassuring.

Conclusions

We provide novel evidence of the “Green Paradox” for climate change legislation using various panel data sets. We consistently find that proposed climate bills that limit future oil use, shift oil consumption from the future towards the present, thereby lowering oil prices in the present. Previous papers on the “Green Paradox” have conducted pre-post comparisons around the passage or discussion of environmental regulation to construct evidence of legislation’s effect on fossil fuel prices. Forward-looking futures prices do not respond when laws are enacted or fail to be enacted, but rather to the release of new information on whether or not a law will pass. Whether or not a bill will pass is oftentimes clear long before it is enacted, making bill passage not a surprise. Moreover, discerning exactly when major market belief updates occur can be difficult. One

of our paper’s contributions is our reliance on daily oil price data and market estimates of the probability that a bill will pass, insulating us from the difficulty in determining when the market participants update their beliefs. Additionally, our reliance on daily variation and use of direct estimates of market expectations allow us to overcome reverse causality challenges common in pre-post comparisons. For example, downward trends in prices can decrease resource owners’ resistance towards a law increasing the likelihood that a bill passes, which would be a concern in a pre-post analysis but not in our analysis using day-to-day changes. We get around this challenge by exploiting market estimates of the probability that the US climate law would pass from prediction markets with price variation driven by political processes unrelated to daily changes in oil prices.

We provide three pieces of evidence that are consistent with the “Green Paradox” for climate change legislation. First, we document that daily shocks to the oil spot price (changes relative to the previous day) historically phase out quickly over time, i.e., maturities further into the future show less responsiveness to changes in present oil prices. This is consistent with temporary spot price shocks, e.g., due to temporary demand spikes (cold winters) or temporary supply disruptions. However, during the time period when a US climate bill was deliberated, the average daily shock in the spot price became much more persistent, indicating that the underlying uncertainty was less transitory and more fundamental. Under the “Green Paradox,” market belief updates around future climate policy should yield persistent shocks, as each belief update the entire price path resets.

Second, in what we see as our most defensible evidence of a causal relationship, we link daily changes in oil futures to changes in the probability that the US will pass a climate bill and find a highly significant negative relationship consistent with a story that legislation limiting future oil use increases current consumption. This relationship persists when we limit the data to a few dates of key political events, ruling out reverse causality, as the timing of these political events was not driven by oil prices. Furthermore, it is unlikely that there are other persistent demand shocks (e.g., higher than expected GDP growth) or persistent supply shocks (e.g., new discoveries) that coincide on exactly the 26 days that major updates on the bill occurred. Moreover, a falsification test where we offset the prediction market prices by one day yields null results, highlighting that daily variation drives our results rather than price level variation across months or quarters. The relationship is symmetric: increases (decreases) in the probability of the law passing lead to decreases (increases) in oil future price of similar magnitude.

By the same token, the surprise ruling of a Dutch Court that ordered the government to limit fossil fuel use was associated with a significant negative abnormal oil price return. Moreover, the verdict release date was pre-determined and not affected by daily oil price movements, again ruling out reverse causality. Similarly, the news that Justice Kennedy, the swing vote, would support the states suing the EPA to regulate automobile carbon dioxide emissions was associated with a significant negative abnormal oil price return.

Third, the maturity profile of the negative coefficients from the prediction market, the Dutch court verdict, and news regarding Justice Kennedy show very high persistence. Effects continuously increase in magnitude for all 24 months for which oil futures are available in the case of prediction market changes. This is consistent with price path and consumption adjustments that are constrained by short-term supply constraints, implying that the full effect will only be felt later on. These findings are in sharp contrast to temporary spot price shocks that tend to phase out rather than in.

One limitation of our study is that liquid oil futures are only available for the next two years, and we


hence only identify the effect over that time span. If demand for oil starts to decrease from 2030 onward as suggested in Kellogg (2024), it might no longer be profitable to drill new oil wells by mid-century. We don't see our paper as giving conflicting results. That paper also finds a short-term price drop as consumption shifts toward the present, but the results reverse for time periods further into the future. That latter finding is largely driven by how fast demand for fossil fuels collapses toward zero. As that paper shows, if only 15% of the demand remains, virtually all of the oil would eventually be extracted. The most significant effect on cumulative extraction that we don't know is how quickly renewable energy development will lead to a decline in demand for fossil fuels or if some demand (e.g., aviation, petrochemicals) will remain.

References

- Anderson, Soren T, Ryan Kellogg, and Stephen W Salant**, “Hotelling under pressure,” *Journal of Political Economy*, 2018, 126 (3), 984–1026.
- Barnett, Michael**, “A Run on Fossil Fuel? Climate Change and Transition Risk,” SSRN Working Paper 436 2023.
- Bruno, Ellen Marie and Nick Hagerty**, “Anticipatory Effects of Regulation in Open Access,” Working Paper, UC Berkeley: Department of Agricultural and Resource Economics August 2023.
- Carattini, Stefano and Suphi Sen**, “Carbon Taxes and Stranded Assets: Evidence from Washington State,” SSRN Working Paper 7785 2019.
- Dasgupta, Partha, Geoffrey Heal, and Joseph E Stiglitz**, “The Taxation of Exhaustible Resources,” Working Paper 436, National Bureau of Economic Research 1980.
- Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva**, “Fighting Climate Change: International Attitudes Toward Climate Policies,” Working Paper 30265, National Bureau of Economic Research July 2022.
- der Ploeg, Frederick Van and Cees Withagen**, “Is there really a green paradox?,” *Journal of Environmental Economics and Management*, 2012, 64 (3), 342–363.
- and — , “Global warming and the green paradox: A review of adverse effects of climate policies,” *Review of Environmental Economics and Policy*, 2015.
- Di Maria, Corrado, Ian Lange, and Edwin van der Werf**, “Should we be worried about the green paradox? Announcement effects of the Acid Rain Program,” *European Economic Review*, 2014, 69, 143–162. Sustainability and Climate Change: From Theory to Pragmatic Policy.
- Dorsey, Jackson**, “Waiting for the Courts: Effects of Policy Uncertainty on Pollution and Investment,” *Environmental and Resource Economics*, 2019, 74 (4), 1453–1496.
- Dube, Arindrajit, Ethan Kaplan, and Suresh Naidu**, “Coups, Corporations, and Classified Information,” *The Quarterly Journal of Economics*, 08 2011, 126 (3), 1375–1409.
- Ekeland, Ivar, Wolfram Schlenker, Peter Tankov, and Brian Wright**, “Optimal Exploration and Price Paths of a Non-renewable Commodity with Stochastic Discoveries,” *NBER Working Paper 29934*, 2022.
- Fountain, John and Glenn W. Harrison**, “What do prediction markets predict?,” *Applied Economics Letters*, 2011, 18 (3), 267–272.
- Gowrisankaran, Gautam, Ashley Langer, and Wendan Zhang**, “Policy Uncertainty in the Market for Coal Electricity: The Case of Air Toxics Standards,” Working Paper 30297, National Bureau of Economic Research July 2022.
- Grafton, R. Quentin, Tom Kompas, Ngo Van Long, and Hang To**, “US biofuels subsidies and CO2 emissions: An empirical test for a weak and a strong green paradox,” *Energy Policy*, 2014, 68, 550–555.
- Gründler, Klaus and Sarah Sauerhammer**, “Do expectations matter? Reassessing the effects of government spending on key macroeconomic variables in Germany,” *Applied Economics Letters*, 2018, 25 (15), 1045–1050.
- Hamilton, James D.**, “Understanding Crude Oil Prices,” *Energy Journal*, 2009, 30 (2), 179–206.

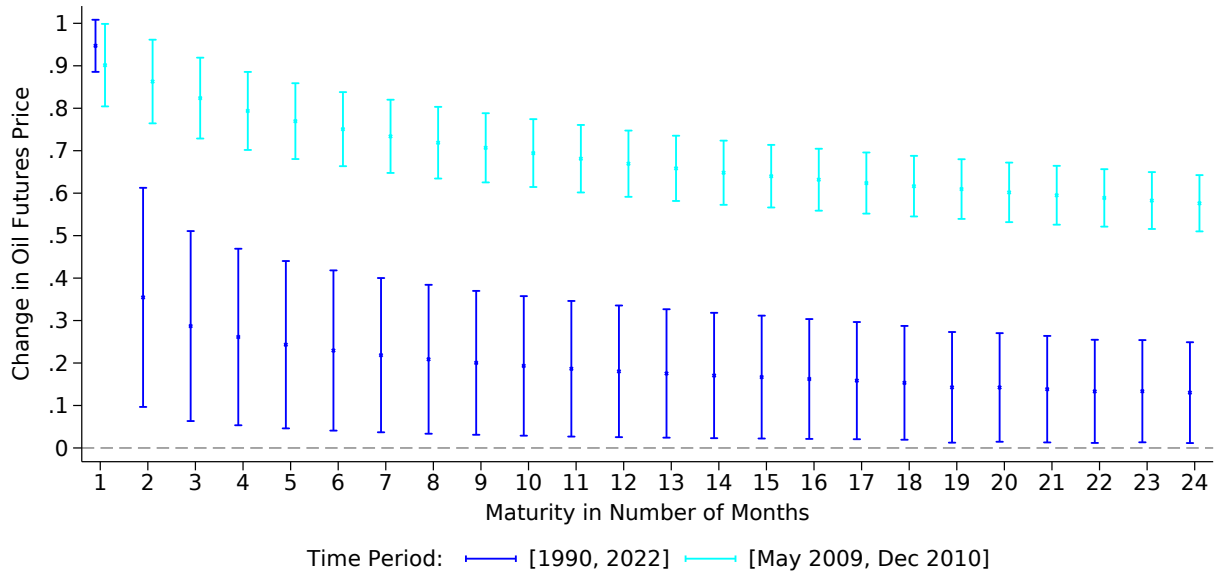
- Heal, Geoffrey and Wolfram Schlenker**, “Coase, Hotelling and Pigou: The Incidence of a Carbon Tax and CO Emissions,” Technical Report, National Bureau of Economic Research 2019.
- Hoel, Michael**, “Is there a green paradox?,” 2010.
- , “Is there a green paradox?,” *CESifo Working Paper 3168*, 2010.
- Hotelling, Harold**, “The Economics of Exhaustible Resources,” *Journal of Political Economy*, 1931, 39 (2), 137–175.
- Judd, Kenneth L.**, “The Welfare Cost of Factor Taxation in a Perfect-Foresight Model,” *Journal of Political Economy*, 1987, 95 (4), 675–709.
- Kellogg, Ryan**, “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling,” *American Economic Review*, June 2014, 104 (6), 1698–1734.
- , “The End of Oil,” *NBER Working Paper 33207*, 2024.
- Kleven, Henrik Jacobsen, Camille Landais, and Emmanuel Saez**, “Taxation and International Migration of Superstars: Evidence from the European Football Market,” *American Economic Review*, August 2013, 103 (5), 1892–1924.
- Langer, Ashley and Derek Lemoine**, “What Were the Odds? Estimating the Market’s Probability of Uncertain Events,” Working Paper 28265, National Bureau of Economic Research December 2020.
- Lemoine, Derek**, “Green Expectations: Current Effects of Anticipated Carbon Pricing,” *The Review of Economics and Statistics*, 07 2017, 99 (3), 499–513.
- Manski, Charles F.**, “Interpreting the predictions of prediction markets,” *Economics Letters*, 2006, 91 (3), 425–429.
- McDermott, Grant R., Kyle C. Meng, Gavin G. McDonald, and Christopher J. Costello**, “The blue paradox: Preemptive overfishing in marine reserves,” *Proceedings of the National Academy of Sciences*, 2019, 116 (12), 5319–5325.
- Meng, Kyle C.**, “Using a free permit rule to forecast the marginal abatement cost of proposed climate policy,” *American Economic Review*, 2017, 107 (3), 748–784.
- Merrill, Ryan Knowles**, “Carbon Risk, the Green Paradox, and Strategic Behavior in the US Petroleum Sector,” *Academy of Management Proceedings*, 2018, 2018 (1), 15273.
- Mukanjari, Samson and Thomas Sterner**, “Do Markets Trump Politics? Evidence from Fossil Market Reactions to the Paris Agreement and the U.S. Election,” Working Paper 728, Department of Economics, University of Gothenburg March 2018.
- Noailly, Joelle, Laura Nowzohour, and Matthias Van Den Heuvel**, “Heard the news? Environmental policy and clean investments,” *CIES Research Paper 299407*, 2021.
- Polasky, Stephen and Holly Doremus**, “When the Truth Hurts: Endangered Species Policy on Private Land with Imperfect Information,” *Journal of Environmental Economics and Management*, 1998, 35 (1), 22–47.
- Sinn, Hans-Werner**, “Das grüne Paradoxon; Plädoyer für eine illusionsfreie Klimapolitik,” *Monographs in Economics*, 2008.
- , “Public policies against global warming: a supply side approach,” *International tax and public finance*, 2008, 15, 360–394.

Figure 1: Motivation: Green Paradox



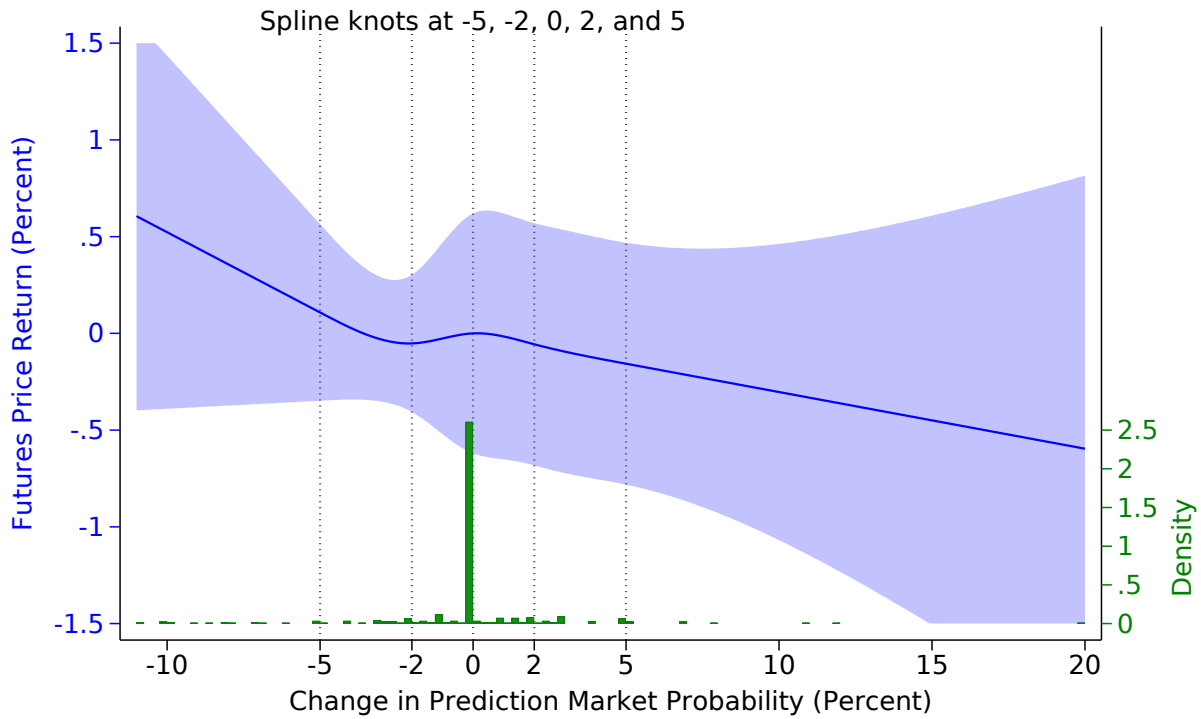
Notes: Figure provides a motivating example of the “Green Paradox.” The black line in the top panel shows a hypothetical optimal fossil fuel consumption profile with iso-elastic demand, a price-elasticity of $\eta = 0.59$ and an interest rate of $\delta = 0.03$. The Waxman-Markey bill, among other potential policies, would have limited consumption to 83% from 2050 onward. Under the assumptions of our model, the bill would have been binding from 2050-2110, reducing consumption by the area A. The bottom panel shows the corresponding re-optimized consumption path, again as black line. The reduction in consumption (area A) is redistributed across time: consumption is higher in period 2010-2050 as shown by area B1. Note that the path does not revert to the old unconstrained path in 2110 (dashed black line) as the arbitrage condition between 2050 and 2110 would be violated. When the regulation is non-binding, the consumption profile is equivalent to assuming a higher initial endowment (dashed grey line), which drops below the bill’s threshold in 2122 rather than 2110. Areas B1 and B2 together are of the same size as area A as consumption is reallocated in time.

Figure 2: Oil Spot Price and Oil Futures for Various Maturities



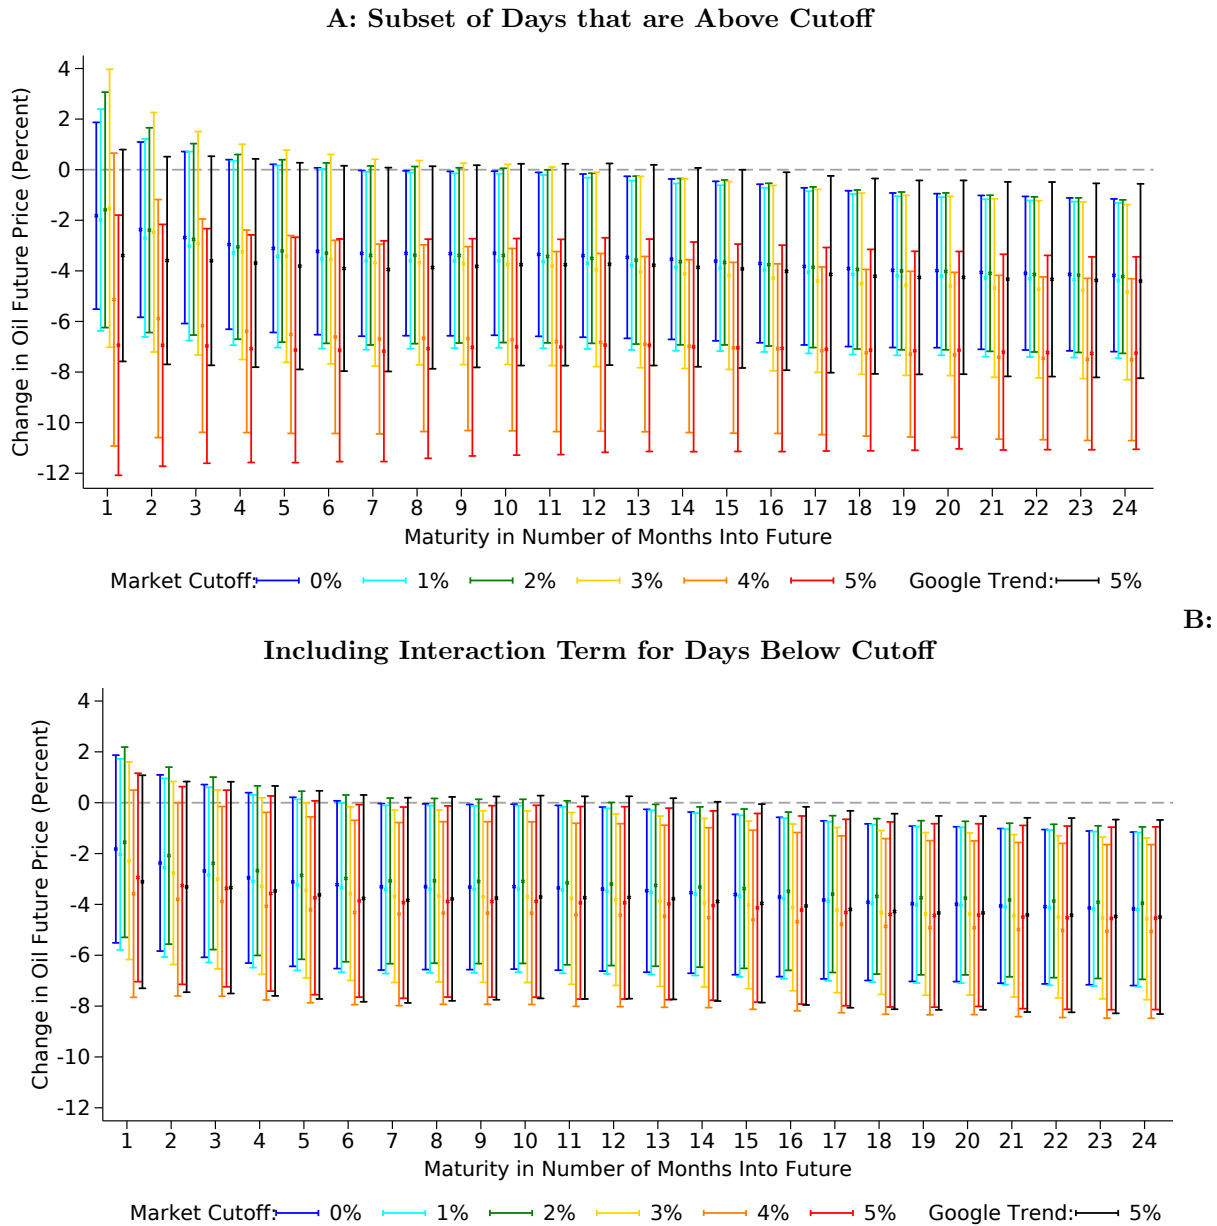
Notes: Figure plots the results when we regress the change in daily oil futures prices on the corresponding change in the oil spot price. The coefficients and 90% confidence intervals are allowed to vary by maturity ranging from 1 to 24 months. Point estimates (marked as x) give the change in oil futures price for a given change in the spot price, e.g., a coefficient of 0.95 implies that on average 95% of the change in the daily oil spot price is reflected in the futures price. The two colors represent various temporal subsets of the data. Regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Figure 3: Allowing Nonlinear Relationship between Prediction Market and Oil Futures Returns



Notes: Figure plots the effect of a change in the probability of a US cap-and-trade bill passing on oil futures prices by pooling across all maturities and allowing for a non-linear relationship. Specifically, we use restricted cubic splines with 5 knots (indicated by dashed lines). The point estimates (blue line) as well as the 90% confidence band are shown on the left y-axis. The density in observed price changes is shown in green on the right axis - there is a mass point at zero as the price does not change for the majority of days.

Figure 4: Changes in the Probability of the Climate Bill Passing and Oil Futures by Maturity



Notes: Panel A and B plots the effect of a change in prediction market probability of a US cap-and-trade bill passing on oil futures prices. Coefficients in Panel A and B as well as 90% confidence intervals are analogous to the coefficients in Table 1, except that the effect is allowed to vary by maturity. Point estimates (marked as x) give the change in oil price in percent for a 100% change in the probability of the bill passing (i.e., from certainty it won't pass to that it will pass). Panel A uses subsets of days included in the analysis, with the seven colors representing the different cutoffs of the seven columns in Table 1. For example, the blue lines include all days, while the red lines include only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and the black lines use days where the google trend data is least 5% of the maximum search index. Panel B always includes all days but introduces an interaction term for days that are below the cutoff - only coefficient for days above the cutoff are shown. Regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table 1: Prediction Market for Climate Bill and Price of Oil Futures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Using Subset of Days Above Cutoff							
Prediction Market	-3.44*	-3.71*	-3.49*	-3.89*	-6.83***	-7.08***	-3.95*
	(1.87)	(1.97)	(1.97)	(2.26)	(2.07)	(2.39)	(2.32)
Observations	10080	2832	1920	1344	912	624	6912
Days	420	118	80	56	38	26	288
Fixed Effects	480	432	384	360	264	240	480
B: All Observations incl. Interaction Below Cutoff							
Prediction Market	-3.44*	-3.53*	-3.21*	-3.83**	-4.50**	-4.01*	-3.91*
	(1.87)	(1.88)	(1.85)	(1.94)	(2.07)	(2.17)	(2.31)
x below cutoff		11.35	-4.34	3.96	6.59	1.64	1.38
		(34.23)	(12.23)	(8.00)	(5.70)	(4.79)	(4.02)
Observations	10080	10080	10080	10080	10080	10080	10080
Days	420	420	420	420	420	420	420
Fixed Effects	480	480	480	480	480	480	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability that a US cap-and-trade bill will pass. Coefficients give the percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty it won't pass to that it will pass). Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. Panel B uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table 2: Prediction Market for Climate Bill and Price of Other Minerals

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Oil							
Prediction Market	-3.44*	-3.53*	-3.21*	-3.83**	-4.50**	-4.01*	-3.91*
	(1.87)	(1.88)	(1.85)	(1.94)	(2.07)	(2.17)	(2.30)
x below cutoff		11.35	-4.34	3.96	6.59	1.64	1.38
		(34.19)	(12.21)	(7.99)	(5.70)	(4.79)	(4.02)
S&P500	0.84***	0.84***	0.84***	0.84***	0.84***	0.84***	0.84***
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Observations Above	10080	2832	1920	1344	912	624	6912
Days Above	420	118	80	56	38	26	288
Panel B1: Copper							
Prediction Market	-1.59	-1.53	-1.04	-2.44	-1.83	-2.67	-0.19
	(2.35)	(2.35)	(2.42)	(2.40)	(2.47)	(2.42)	(3.00)
x below cutoff		-10.03	-12.21	6.98	-2.72	3.30	-4.09
		(37.73)	(15.12)	(9.80)	(7.68)	(6.69)	(4.52)
S&P500	0.84***	0.84***	0.84***	0.84***	0.84***	0.84***	0.84***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
Observations Above	5194	1280	883	602	408	260	3462
Days Above	420	119	80	56	38	26	288
Panel B2: Aluminum							
Prediction Market	1.32	1.47	0.95	0.26	-0.01	-1.29	2.86
	(2.20)	(2.17)	(2.21)	(2.31)	(2.44)	(2.65)	(2.61)
x below cutoff		6.85	18.50	12.10	7.64	10.97*	-4.40
		(41.19)	(15.81)	(10.16)	(7.64)	(6.03)	(4.53)
S&P500	0.67***	0.67***	0.67***	0.67***	0.67***	0.67***	0.67***
	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
Observations Above	9672	2688	1872	1344	912	624	6648
Days Above	403	112	78	56	38	26	277
Panel B3: Steel							
Prediction Market	-0.65	-0.72	-0.82	-0.85	-0.85	-0.55	-0.59
	(1.09)	(1.11)	(1.15)	(1.12)	(1.19)	(1.26)	(1.45)
x below cutoff		4.33	4.94	1.78	1.07	-0.56	-0.18
		(6.63)	(4.59)	(5.67)	(3.81)	(2.79)	(2.14)
S&P500	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Observations Above	6996	1958	1301	931	636	422	4668
Days Above	416	117	81	58	40	26	286
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table replicates panel B of Table 1 for other metals with one modifications: regression no longer control for S&P500 by maturity but rather force the relationship to be the same across all maturities so it can be displayed in the Table. Panel A shows the results of Panel B Table 1 with a common sensitivity to fluctuation in the SP&500. Panels B1-B3 present placebo regressions for metals. All regressions still include maturity-by-month fixed effects and errors are clustered by day.

Table 3: Falsification: Next Day’s Prediction Market for Climate Bill and Price of Oil Futures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Using Subset of Days Above Cutoff							
Prediction Market	-0.88 (2.25)	-2.21 (2.24)	-2.58 (2.34)	-1.97 (2.60)	-0.69 (2.56)	-0.21 (3.17)	-0.41 (3.09)
Observations	10080	2832	1896	1320	888	600	6936
Days	420	118	79	55	37	25	289
Fixed Effects	480	432	360	336	240	216	480
B: All Observations incl. Interaction Below Cutoff							
Prediction Market	-0.88 (2.25)	-0.85 (2.25)	-0.55 (2.29)	-0.09 (2.28)	0.46 (2.33)	0.25 (2.49)	-0.14 (3.02)
x below cutoff		-6.73 (40.59)	-8.83 (10.71)	-9.04 (9.55)	-9.26 (6.88)	-4.69 (5.42)	-1.91 (4.73)
Observations	10080	10080	10080	10080	10080	10080	10080
Days	420	420	420	420	420	420	420
Fixed Effects	480	480	480	480	480	480	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table replicates Table 1 except that it uses *one lead* of prediction market prices (i.e., from the next period). It again regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability on the next day that a US cap-and-trade bill will pass. Coefficients give the percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty in won’t pass to that it will pass). Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. Panel B uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table 4: Prediction Market and Oil Storage

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Oil Storage						
A: Using Subset of Days Above Cutoff						
Prediction Market	-4.75*	-3.66	-3.48	-3.49	-3.78	0.80
	(2.54)	(2.38)	(2.40)	(2.56)	(2.88)	(2.71)
Observations	85	42	35	29	23	13
Fixed Effects	8	7	7	7	6	4
B: All Observations incl. Interaction Below Cutoff						
Prediction Market	-4.75*	-4.68*	-4.47*	-4.49*	-5.13*	-2.93
	(2.54)	(2.55)	(2.52)	(2.54)	(2.64)	(2.64)
x below cutoff		0.85	-16.59	-11.02	3.64	-10.78*
		(35.00)	(19.58)	(14.21)	(9.90)	(6.37)
Observations	85	85	85	85	85	85
Fixed Effects	8	8	8	8	8	8
Min market change	0	1	2	3	4	5

Notes: Table presents a sensitivity analysis of the main results in Table 1 using weekly storage data. Panel A and B give the percent change in oil stored for a 100% change in the probability of the bill passing. Columns differ by what weeks are used in the analysis, with the bottom row listing the cutoff value for the absolute change in the prediction market price required for a week to be included. For example, column (1) in panel A includes all weeks, while column (6) includes only the weeks when the prediction market price changed by at least 5 cents. Regressions control for changes in the S&P 500 index well as maturity-by-quarter fixed effects. Errors are clustered by week.

Table 5: News About Historic Climate Court Cases and Prices of Oil Futures

	(1)	(2)	(3)	(4)	(5)
Panel A: Urgenda v. Netherlands Court Ruling					
I _{June 24, 2015}	-0.549** (0.274)	-0.595** (0.257)	-0.628** (0.256)	-0.902*** (0.278)	-0.871*** (0.275)
Observations	6024	30192	54260	174150	192270
Fixed Effects	288	1440	2592	8459	9323
Clusters	251	1258	2261	7531	8286
Years	[15,15]	[13,17]	[11,19]	[90,19]	[90,22]
Panel B: Justice Kennedy Indicates Support					
I _{December 6, 2006}	-0.564** (0.229)	-0.501** (0.231)	-0.488** (0.232)	-0.500** (0.231)	-0.491** (0.232)
Observations	6000	30070	54140	174150	192270
Fixed Effects	288	1440	2592	8459	9323
Clusters	250	1254	2258	7531	8286
Years	[06,06]	[04,08]	[02,10]	[90,19]	[90,22]

Notes: Table regresses the change in oil futures (24 different maturities ranging from 1 to 24 months into the future) on dummies representing the days two historic climate court cases peaked in popularity on Google. Panel A regresses the change in oil futures on dummy for June 24, 2015, the day the Urgenda v. Netherlands verdict was rendered. Panel B regresses the change in oil futures on a dummy for December 6, 2006, the day prevailing news suggested that Justice Kennedy would back the states suing the EPA. Coefficients give the change in oil price in percent. Columns differ by what days are included in the analysis with the bottom row in each panel displaying the range of years that are used to derive the controls. Column (1) focuses only on days in the year the court cases peaked in popularity on Google, while column (5) includes all days between 1990-2022. Regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table 6: Media Coverage of Climate Policy and Prices of Oil Futures

	(1)	(2)	(3)	(4)	(5)	(6)
International Climate Negotiations	-0.76** (0.36)	-0.77** (0.37)	-0.79** (0.34)	-0.79** (0.35)	-0.81** (0.34)	-0.81** (0.35)
Renewable Policy	0.30 (1.29)	0.33 (1.35)	0.24 (1.22)	0.25 (1.27)	0.20 (1.21)	0.22 (1.25)
Environmental Policy	0.44 (0.93)	0.42 (0.97)	0.55 (0.91)	0.54 (0.94)	0.62 (0.89)	0.61 (0.93)
Observations	9216	9210	9216	9210	9216	9210
Fixed Effects	37	811	37	811	37	811
Year FE	Yes	No	Yes	No	Yes	No
Maturity x Year FEs	No	Yes	No	Yes	No	Yes
S&P 500 Control	No	No	Yes	Yes	No	No
S&P 500 x Maturity Control	No	No	No	No	Yes	Yes

Notes: Table regresses the change in the closing price of oil futures (24 different maturities ranging from 1 to 24 months into the future) on monthly indices measuring the share of news articles covering international climate negotiations, renewable energy policy, and environmental policy generally. Each column shows results from a joint regression. All indices are standardized to a mean of zero and unit standard deviation. Coefficients give the percent change in the closing oil price for a one standard deviation increase in each news index. As a reference, the Paris Climate Agreement is associated with an International Climate Negotiation index 6 standard deviations above the mean. Columns differ in the controls included. For example, column (1) includes year fixed effects and column (4) includes maturity by year fixed effects and controls for changes in the S&P 500 index by maturity (Columns 1 and 2 force the effect of the S&P 500 index to be the same across maturities). Errors are clustered by month.

A Appendix


List of Figures

A1	Oil Spot Price and Oil Futures for Various Maturities - Temporal Evolution	ii
A2	Sensitivity Check: Oil Futures Sensitivity to S&P 500 Over Time	iii
A3	Sensitivity Check: Prediction Market and Oil Futures by Maturity Controlling for Spot Price	iv
A4	Oil Future Returns in Response to Breaking News About Court Cases by Maturity	v

List of Tables

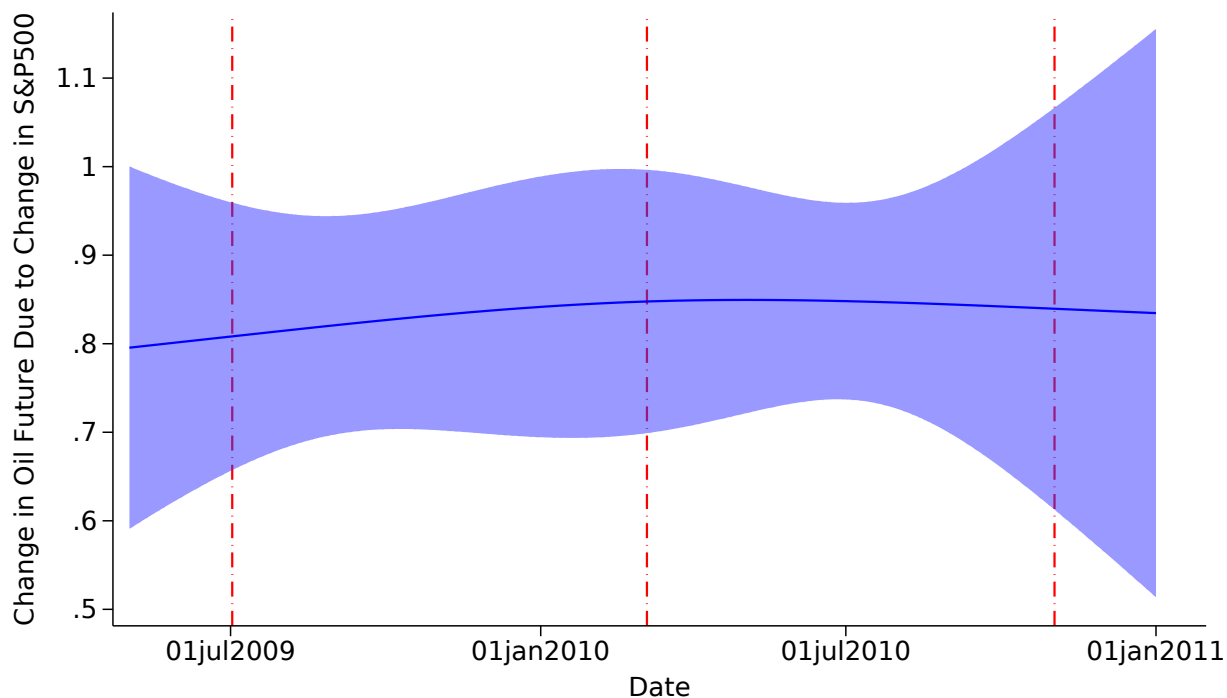
A1	News Coverage on Days when Prediction Market Price Changed by at Least 5 Cents	vi
A2	Sensitivity Check: Prediction Market and Oil Futures - Clustering	vii
A3	Sensitivity Check: Prediction Market and Oil Futures - Clustering	viii
A4	Sensitivity Check: Prediction Market and Oil Futures - High Trade Volume	ix
A5	Sensitivity Check: Prediction Market and Oil Futures - Controlling for S&P 500 Yield Curve	x
A6	Sensitivity Check: Distributed Lag Model of Prediction Market and Price of Oil Futures	xi
A7	Prediction Market for Climate Bill and Price of Coal Futures	xii

Figure A1: Oil Spot Price and Oil Futures for Various Maturities - Temporal Evolution



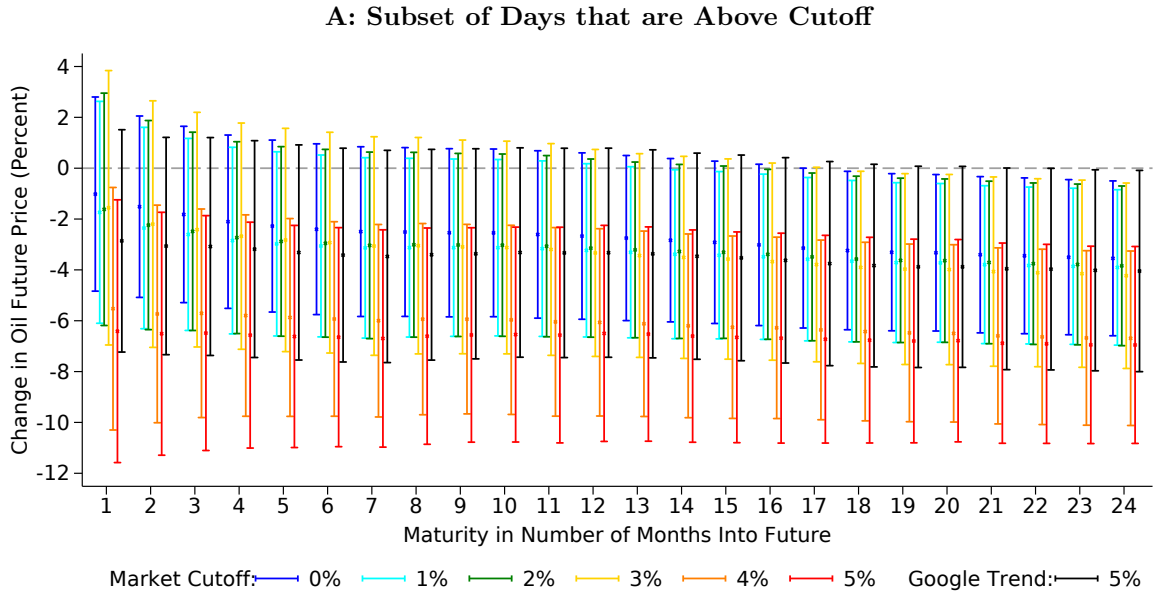
Notes: Figure replicates Figure 2 but breaks the overall period into sub-periods. It again plots the results when we regress the change in daily oil futures prices on corresponding change in the oil spot price. The coefficients and 90% confidence intervals are allowed to vary by maturity ranging from 1 to 24 months. Point estimates (marked as x) give the change in oil futures prices for a given change in the spot price. The seven colors represent various temporal subsets of the data. Regressions controls for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Figure A2: Sensitivity Check: Oil Futures Sensitivity to S&P 500 Over Time

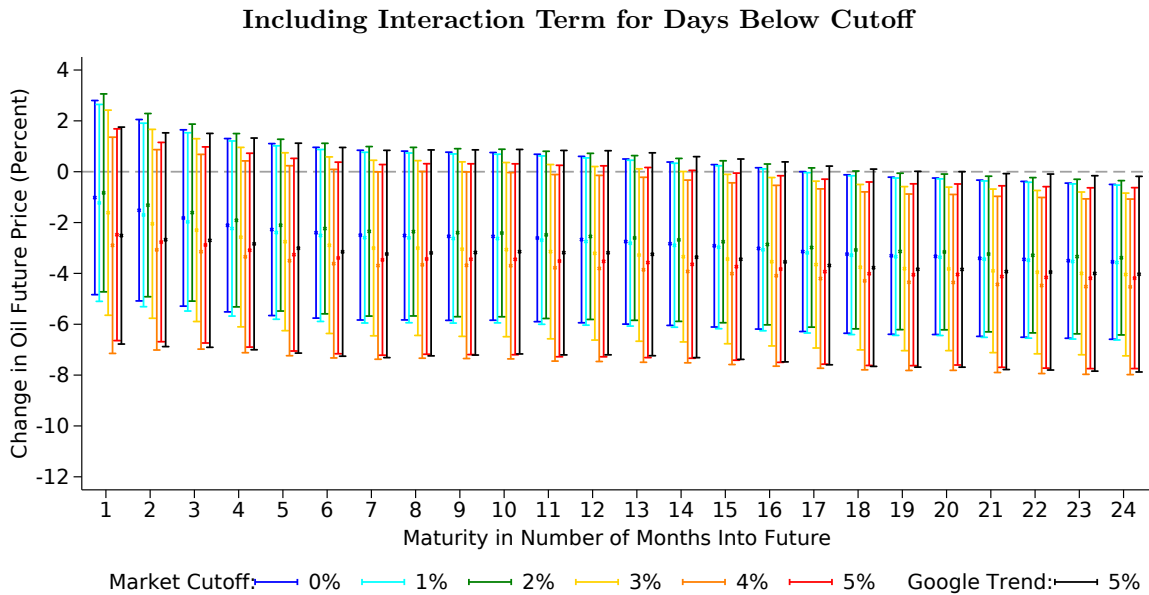


Notes: Figure presents the results when we regress the daily changes in oil futures prices on changes in the S&P 500 index during the time period the prediction market is active. The effect is forced to be the same across all 24 maturities, but allowed to vary over time as the world just recovered from the 2008 financial crisis. Specifically, we interact the change in the S&P 500 index with a restricted cubic spline with three knots (shown as green dashed lines) in time. They allow for a flexible response function (3rd-order polynomials between knots). The shaded area reflects the 90% confidence band. Errors are clustered by day. The sensitivity of oil futures to the S&P 500 did not vary over the sample period when prediction market was active.

Figure A3: Sensitivity Check: Prediction Market and Oil Futures by Maturity Controlling for Spot Price

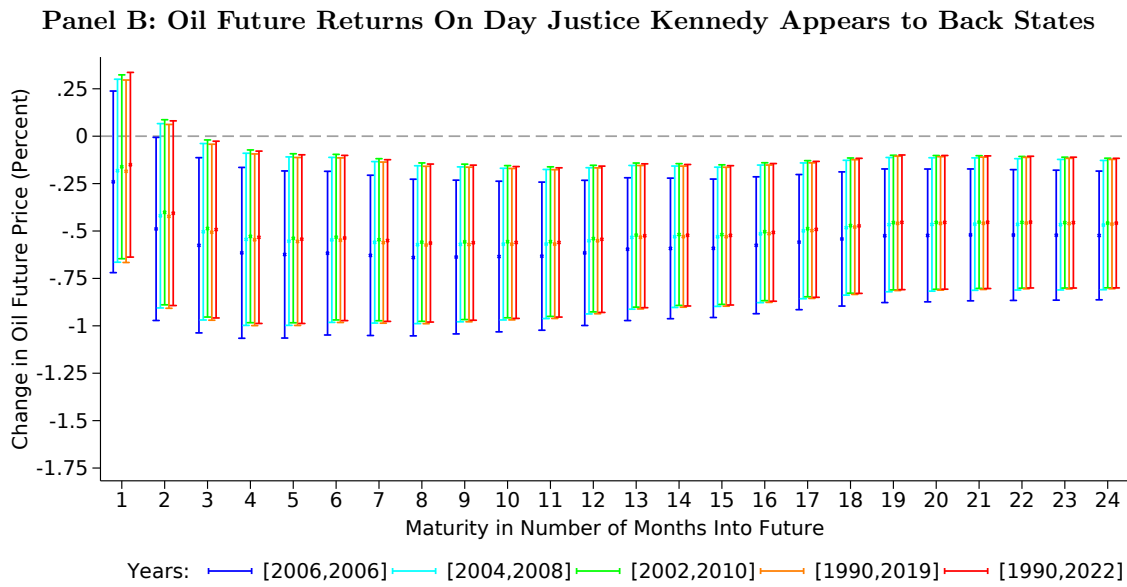



B:



Notes: Figure presents a sensitivity analysis to the main results in Figure 4, where we control for the WTI oil spot price. Both panels plot the effect of a change in prediction market probability for the passage of a US cap-and-trade bill on oil futures prices. Point estimates (marked as x) give percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty it won't pass to that it will pass). Colors differ by what days are included in the analysis. The six colors represent the different cutoffs of the six columns in Table 1. For example, in Panel A the blue lines include all days, while the red lines include only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing). In Panel B, like Panel A, the blue lines include all days, while the red lines include only the days when the prediction market price changed by at least 5 cents and the Google search volume was at least 5% of the day with maximum search activity. All regressions control for changes in the S&P 500 index by maturity, oil spot price, as well as maturity-by-month fixed effects. Errors are clustered by day.

Figure A4: Oil Future Returns in Response to Breaking News About Court Cases by Maturity



Notes: Figure plots the effect of significant information shocks regarding the verdicts of two historic climate court cases on oil futures prices. Panel A plots the effect of the Urgenda v. Netherlands verdict announcement, which was also the day when the interest peaked in search volume on Google Trends. Panel B plots the effect of the prevailing news that Justice Kennedy – a key swing vote – appeared to be backing the states suing the EPA; the day when the Massachusetts v. EPA court case peaked in search volume on Google Trends. The coefficients and 90% confidence intervals are allowed to vary by maturity ranging from 1 to 24 months. Point estimates (marked as x) give the percent change in oil futures on the day each court case peaked in popularity on Google. Colors differ by what days are included in the analysis with the five colors representing the range of years used in the five columns of Table 5. For example, the blue lines focus only on days in the year each court case peaked in popularity on Google, while the red lines include all days between 1990-2022. Regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table A1: News Coverage on Days when Prediction Market Price Changed by at Least 5 Cents

Date	Change	Lexis-Nexis Search for News Coverage
05/11/09	-10.0	“While Democrats met behind closed doors, Republicans held a public energy summit to consider alternative solutions to what they dub a ‘cap and tax’ program.” (SNL Energy Dataset, S&P Global Marketplace)
05/12/09	+20.0	“The U.S. House of Representatives will pass a sweeping climate change bill by the end of next week, House Energy Committee Chairman Henry Waxman said.” (Reuters)
05/13/09	-10.0	“The Economic Impact of Waxman-Markey.” (States News Service on Heritage Foundation Report)
06/01/09	+5.2	“The 30-page report, commissioned by the U.S. Department of Energy, focuses on draft bills introduced individually by Senate Energy and Natural Resources Committee Chairman Jeff Bingaman, D-N.M., and House Subcommittee on Energy and Environment Chairman Edward Markey, D-Mass., and a joint bill by Markey and House Energy and Commerce Committee Chairman Henry Waxman, D-Calif. Bingaman’s bill did not fare as well as the others in terms of raising renewable capacity and reducing emissions, according to the report.” (SNL Energy Dataset, S&P Global Marketplace)
06/08/09	+5.0	“A new analysis of the bill by the Congressional Budget Office (CBO) shows the legislation is a fiscally-responsible clean energy plan.” (States News Service on press release by Markey)
07/01/09	+5.0	“Duke CEO: New state-federal relationship needed to meet Waxman-Markey targets.” (SNL Power Daily Northeast, S&P Global Marketplace)
07/09/09	-6.9	“The Waxman-Markey bill passed by the U.S. House of Representatives last month would set strict new carbon dioxide emissions levels for new coal plants, requiring them to come close to current natural gas plants in CO2 emissions.” (SNL Power Daily Northeast, S&P Global Marketplace)
08/17/09	-9.8	“EPA denies senators’ request to redo Waxman-Markey analysis.” (SNL Electric Utility Report, S&P Global Marketplace)
08/27/09	-9.8	“The National Association of Manufacturers today launched a multi-state, multi-million-dollar comprehensive advertising campaign opposing the American Clean Energy and Security Act (H.R. 2454), also known as the Waxman-Markey climate change bill.” (States News Service)
11/04/09	-11.0	“Consulting firm Point Carbon notes that recent legislative proposals in the US Congress hold oil companies accountable for both refinery and tailpipe emissions, making them more vulnerable to carbon controls than the coal-dominated electric utility sector. And the Energy Policy Research Foundation (EPRINC) calculates that climate change legislation currently being debated in Congress could put as much as 8 million barrels per day of US refining capacity at risk of closure – an astounding 45% of total operable domestic capacity.” (Oil Daily)
11/19/09	-5.0	“The American Recovery and Reinvestment Act recommitted or country to science and technology. And the Waxman-Markey clean energy legislation that the House passed this past June will extend this commitment by investing \$200 billion through 2025 to unleash the clean energy revolution waiting to happen across America.” (US Fed News)
12/21/09	-7.0	“Sens. Maria Cantwell, D-Wash., and Susan Collins, R-Maine, unveiled a climate change bill in the Senate on Dec. 11 that would auction carbon permits to producers and importers of coal, natural gas and oil, which is an approach that differs dramatically from the Waxman-Markey cap-and-trade bill that the House of Representatives passed in June.” (SNL Electric Utility Report, S&P Global Marketplace)
12/28/09	-10.8	“The findings, contained in a new analysis from the environmental think tank Resources for the Future, bolster the rationale for a cap-and-dividend plan introduced earlier this month by Sens. Maria Cantwell (D-WA) and Susan Collins (R-ME), which calls for auctioning all allowances and returning 75 percent of the revenue raised to the public in the form of monthly rebates. The Cantwell-Collins bill is a competitor to the leading Senate cap-and-trade proposal authored by Sens. John Kerry (D-MA) and Barbara Boxer (D-CA), which mirrors the House bill.” (Carbon Control News)
03/15/10	+12.0	“A new report prepared for the environmental group Natural Resources Defense Council (NRDC) finds that requiring carbon capture and storage (CCS) technology to be installed on new power generation and industrial facilities would not cause severe damage to the U.S. economy, but could provide economic benefits by boosting domestic oil production 3 million to 3.6 million barrels a day by 2030 if the CO2 were injected underground for enhanced oil recovery.” (Carbon Control News)
03/17/10	-8.0	“Bingaman: Comprehensive energy legislation not likely in Senate in 2010.” (SNL FERC Power Report, S&P Global Marketplace)
03/23/10	+7.0	“In a March 22 letter addressed to Sen. Maria Cantwell, D-Wash., the International Emissions Trading Association said the Carbon Limits and Energy for America’s Renewal Act, or CLEAR Act, is fundamentally flawed as written, and the group expressed its concern the legislation will not achieve its stated emissions reduction objectives in the most cost-effective manner possible. On Dec. 11, 2009, Cantwell and Sen. Susan Collins, R-Maine, unveiled a climate change bill in the Senate that would auction carbon permits to producers and importers of coal, natural gas and oil, an approach that differs dramatically from the Waxman-Markey cap-and-trade bill, which the House of Representatives passed in June of 2009.” (SNL Power Daily, S&P Global Marketplace)
03/24/10	-9.0	“The American Petroleum Institute released the following statement today from its President and CEO Jack Gerard commenting on some media reports characterizing API’s position on the Kerry-Graham-Lieberman climate discussions: [...] Moving away from the House Waxman-Markey approach was imperative. The House bill would have eliminated millions more jobs than it created and unfairly burdened families, farmers, truckers and other regular users of gasoline, diesel and other petroleum products.” (States News Service)
03/26/10	+5.1	“Cap and Trade Loses Its Standing as Energy Policy of Choice.” (New York Times)
04/05/10	-10.0	“The House of Representatives-passed Waxman-Markey climate bill allows holders of RGGI allowances to exchange them for federal emission permits based on the average auction price paid for the allowances in a given year. However, the passage of similar climate legislation in the Senate has faced an uphill battle since the end of last year.” (SNL Power Daily, S&P Global Marketplace)
04/14/10	+7.0	“Chairman Markey: Climate Bill Has Multiple Benefits.” (Congressional Documents and Publications)
04/19/10	+5.1	“Congress weighs far-reaching global warming bill. [...] If Congress balks, the Obama administration has signaled a willingness to use decades-old clean air laws to impose tough new regulations for motor vehicles and many industrial plants to limit their release of climate-changing pollution.” (Associated Press International)
04/26/10	-6.0	“South Carolina Republican Sen. Lindsey Graham is getting an enormous amount of flack for subtracting his initial from the Kerry-Graham-Lieberman energy bill that was due to be revealed this morning. Graham’s decision delays debate and could possibly be fatal for the bill’s prospects.” (Atlantic Online)
05/14/10	-8.0	“Sens. John Kerry, D-Mass., and Joe Lieberman, I-Conn., released the details of their long-awaited Senate energy and climate change legislation.” (SNL Renewable Energy Weekly, S&P Global Marketplace)
05/19/10	+8.0	“Markey Statement on New National Academy of Science Reports.” (States News Service)
07/13/10	-7.8	“Maryland Republican Party Spokesman Ryan Mahoney issued the following statement today: [...] Cap And Trade Could Cause The Loss Of Up To 41,500 Jobs In Maryland.” (States News Service)
07/23/10	-8.5	“Democrats Call Off Climate Bill Effort.” (New York Times)

Notes: Table provides the results from a Lexis-Nexis search on the terms “Waxman Markey” on the 26 days when the prediction market price changed by at least 5%, i.e., the days used in column (6) of Table 1. The first column gives the date, the second the change in the prediction market probability the legislation will pass, and the third the news story.

Table A2: Sensitivity Check: Prediction Market and Oil Futures - Clustering

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Cluster by Day							
Prediction Market	-3.44*	-3.71*	-3.49*	-3.89*	-6.83***	-7.08***	-3.95*
	(1.87)	(1.97)	(1.97)	(2.26)	(2.07)	(2.39)	(2.32)
Clusters	420	118	80	56	38	26	288
Panel A1: Cluster by Future-Maturity and Month							
Prediction Market	-3.44**	-3.71***	-3.49**	-3.89*	-6.83***	-7.08***	-3.95*
	(1.50)	(1.27)	(1.55)	(1.91)	(1.44)	(2.07)	(2.14)
Clusters	44	42	40	39	35	34	44
Panel A2: Cluster by Future-Maturity-Month							
Prediction Market	-3.44***	-3.71***	-3.49***	-3.89***	-6.83***	-7.08***	-3.95***
	(0.31)	(0.27)	(0.32)	(0.40)	(0.31)	(0.42)	(0.44)
Clusters	480	432	384	360	264	240	480
Panel A3: Robust Standard Errors							
Prediction Market	-3.44***	-3.71***	-3.49***	-3.89***	-6.83***	-7.08***	-3.95***
	(0.39)	(0.41)	(0.41)	(0.47)	(0.43)	(0.49)	(0.48)
Observations Above	10080	2832	1920	1344	912	624	6912
Fixed Effects	480	432	384	360	264	240	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table reproduces panel A of Table 1, and presents sensitivity analysis for various clustering structures in panels A1-A3. Table regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability that a US cap-and-trade bill will pass. Coefficients give the percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty in won't pass to that it will pass). Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects.

Table A3: Sensitivity Check: Prediction Market and Oil Futures - Clustering

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel B: Cluster by Day							
Prediction Market	-3.44*	-3.53*	-3.21*	-3.83**	-4.50**	-4.01*	-3.91*
	(1.87)	(1.88)	(1.85)	(1.94)	(2.07)	(2.17)	(2.31)
x below cutoff		11.35	-4.34	3.96	6.59	1.64	1.38
		(34.23)	(12.23)	(8.00)	(5.70)	(4.79)	(4.02)
Clusters	420	420	420	420	420	420	420
Panel B1: Cluster by Future-Maturity and Month							
Prediction Market	-3.44**	-3.53**	-3.21**	-3.83**	-4.50**	-4.01**	-3.91*
	(1.50)	(1.46)	(1.51)	(1.76)	(1.74)	(1.87)	(2.11)
x below cutoff		11.35	-4.34	3.96	6.59	1.64	1.38
		(38.77)	(14.53)	(11.13)	(5.76)	(4.36)	(3.84)
Clusters	44	44	44	44	44	44	44
Panel B2: Cluster by Future-Maturity-Month							
Prediction Market	-3.44***	-3.53***	-3.21***	-3.83***	-4.50***	-4.01***	-3.91***
	(0.31)	(0.31)	(0.31)	(0.37)	(0.36)	(0.39)	(0.43)
x below cutoff		11.35	-4.34	3.96*	6.59***	1.64*	1.38*
		(7.99)	(2.99)	(2.28)	(1.19)	(0.89)	(0.79)
Clusters	480	480	480	480	480	480	480
Panel B3: Robust Standard Errors							
Prediction Market	-3.44***	-3.53***	-3.21***	-3.83***	-4.50***	-4.01***	-3.91***
	(0.39)	(0.39)	(0.38)	(0.40)	(0.43)	(0.45)	(0.48)
x below cutoff		11.35	-4.34*	3.96**	6.59***	1.64*	1.38*
		(7.10)	(2.52)	(1.65)	(1.18)	(0.99)	(0.83)
Observations Above	10080	10080	10080	10080	10080	10080	10080
Fixed Effects	480	480	480	480	480	480	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table reproduces panel B of Table 1, and presents sensitivity analysis for various clustering structures in panels B1-B3. Table regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability that a US cap-and-trade bill will pass. Coefficients give the percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty in won't pass to that it will pass). Columns differ by what cuto is used, where the coefficient is allowed to differ above and below the cutoff. The bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index). Panels uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects.

Table A4: Sensitivity Check: Prediction Market and Oil Futures - High Trade Volume

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Using Subset of Days Above Cutoff							
Prediction Market	-4.09 (3.09)	-4.08 (3.42)	-3.93 (3.34)	-3.80 (3.70)	-3.22 (3.86)	-3.74 (6.13)	-2.46 (3.91)
Observations	2568	1176	720	504	384	240	1656
Days	107	49	30	21	16	10	69
Fixed Effects	432	312	216	168	120	96	360
B: All Observations incl. Interaction Below Cutoff							
Prediction Market	-4.09 (3.09)	-4.03 (3.18)	-2.90 (3.21)	-6.51** (3.21)	-5.68* (3.21)	-6.51* (3.65)	-1.18 (4.05)
x below cutoff		5.41 (43.01)	-12.43 (16.06)	15.65 (13.99)	10.54 (9.57)	8.85 (7.96)	-7.27 (5.85)
Observations Above	2568	2568	2568	2568	2568	2568	2568
Days	107	107	107	107	107	107	107
Fixed Effects	432	432	432	432	432	432	432
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table presents a sensitivity analysis of the main results in Table 1, where we only include days when the trade volume in the prediction market was at least 12, the median trade volume for days with more than zero trades. Table regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability that a US cap-and-trade bill will pass. Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. Panel B uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table A5: Sensitivity Check: Prediction Market and Oil Futures - Controlling for S&P 500 Yield Curve

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Using Subset of Days Above Cutoff							
Prediction Market	-3.38*	-4.06**	-2.83	-2.97	-7.29***	-9.78***	-3.58
	(1.90)	(2.00)	(1.97)	(2.64)	(2.49)	(2.66)	(2.37)
Observations	10080	2832	1920	1344	912	624	6912
Days	420	118	80	56	38	26	288
Fixed Effects	480	432	384	360	264	240	480
B: All Observations incl. Interaction Below Cutoff							
Prediction Market	-3.38*	-3.46*	-3.14*	-3.77*	-4.44**	-3.94*	-3.82
	(1.90)	(1.91)	(1.88)	(1.97)	(2.12)	(2.22)	(2.33)
x below threshold		11.10	-4.44	3.96	6.53	1.62	1.35
		(33.91)	(12.21)	(8.02)	(5.74)	(4.81)	(4.02)
Observations	10080	10080	10080	10080	10080	10080	10080
Days	420	420	420	420	420	420	420
Fixed Effects	480	480	480	480	480	480	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table presents a sensitivity analysis of the main results in Table 1, where we account for market expectations on economic growth by additionally including the change in the difference between the S&P 500 Future 8 months into the future and the current S&P 500 index. Recall that all regressions already control for changes in the contemporaneous S&P 500 index, but this additional control is the market expectation on future growth. Table regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability that a US cap-and-trade bill will pass. Coefficients give the percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty in won't pass to that it will pass). C Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. Panel B uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table A6: Sensitivity Check: Distributed Lag Model of Prediction Market and Price of Oil Futures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: All Observations incl. Interaction Below Cutoff							
Prediction Market in t	-2.78 (1.95)	-2.75 (1.99)	-2.59 (1.96)	-3.19 (2.05)	-4.03* (2.19)	-3.25 (2.37)	-2.96 (2.43)
Prediction Market in t-1	3.09 (2.09)	3.20 (2.10)	3.19 (2.18)	3.30 (2.17)	2.95 (2.25)	3.72 (2.34)	2.57 (2.60)
x below cutoff in t		5.93 (34.85)	-2.90 (12.59)	4.28 (8.15)	7.64 (5.79)	1.67 (4.80)	0.98 (4.13)
x below cutoff in t-1		-24.91 (29.43)	-1.59 (10.94)	-2.98 (8.76)	1.20 (6.75)	-2.37 (5.53)	2.22 (4.21)
Observations	10056	10056	10056	10056	10056	10056	10056
Days	419	419	419	419	419	419	419
Fixed Effects	480	480	480	480	480	480	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table replicates Table 1 except that it uses both the contemporaneous as well as *one lag* of prediction market prices. It again regresses daily changes in oil futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability on the next day that a US cap-and-trade bill will pass. Coefficients give the percent change in oil futures for a 100% change in the probability of the bill passing (i.e., from certainty in won't pass to that it will pass). Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. Panel B uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.

Table A7: Prediction Market for Climate Bill and Price of Coal Futures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A: Using Subset of Days Above Cutoff							
Prediction Market	-1.02 (1.63)	-1.37 (1.66)	-2.60 (1.81)	-3.00 (1.89)	-3.75* (1.99)	-5.50* (2.91)	-3.06 (1.92)
Observations	10080	2832	1920	1344	912	624	6912
Days	420	118	80	56	38	26	288
Fixed Effects	480	432	384	360	264	240	480
B: All Observations incl. Interaction Below Cutoff							
Prediction Market	-1.02 (1.63)	-0.92 (1.66)	-1.34 (1.68)	-1.44 (1.78)	-1.16 (1.83)	-1.26 (2.03)	-2.93 (1.92)
x below threshold		-23.65 (19.26)	9.02 (9.50)	4.69 (5.95)	2.71 (4.73)	2.46 (3.87)	5.51 (3.46)
Observations	10080	10080	10080	10080	10080	10080	10080
Days	420	420	420	420	420	420	420
Fixed Effects	480	480	480	480	480	480	480
Min market change	0	1	2	3	4	5	0
Min google trend	0	0	0	0	0	0	5

Notes: Table regresses daily changes in coal futures (24 different maturities ranging from 1 to 24 months into the future) on the change in the prediction market probability that a US cap-and-trade bill will pass. Coefficients give the percent change in coal futures for a 100% change in the probability of the bill passing (i.e., from certainty it won't pass to that it will pass). Columns differ by what days are used in the analysis, with the bottom two rows listing the cutoff value for the absolute change in the prediction market price (and Google Trends index) required for a day to be included. For example, column (1) of panel A includes all days, while column (6) of panel A includes only the days when the prediction market price changed by at least 5 cents (equivalent to a 5% change in the probability of passing), and column (7) of panel A includes only days when the google trends search volume was at least 5% of the maximum day. Panel B uses all observations (10,080 for 420 days with 480 maturity-by-month fixed effects) but adds an interaction term for the prediction market with an indicator that the daily change / google search volume was below the cutoff. All regressions control for changes in the S&P 500 index by maturity as well as maturity-by-month fixed effects. Errors are clustered by day.