

GREEN EXPECTATIONS: CURRENT EFFECTS OF ANTICIPATED CARBON PRICING

Derek Lemoine*

Abstract—I report evidence that an anticipated strengthening of environmental policy increased emissions. I find that the breakdown of the U.S. Senate's 2010 climate effort generated positive excess returns in coal futures markets. This response appears to be driven by an increase in coal storage. The proposed legislation aimed to reduce U.S. greenhouse gas emissions after 2013, but the legislative process itself may have increased emissions by over 12 million tons of carbon dioxide leading up to April 2010.

I. Introduction

IN 2009, the U.S. House of Representatives passed a bill to cap carbon dioxide emissions. The Senate soon began drafting complementary legislation. The emission cap was to begin in 2013, four years after the House bill's passage, twelve years after Congress first voted on cap-and-trade legislation, and twenty-five years after attention-grabbing congressional hearings on global climate change. Energy markets had ample forewarning. The "green paradox" literature suggests that energy producers will increase near-term extraction when they anticipate a future emission policy (Sinn, 2008). In many models, the increase in emissions during the period of delayed implementation completely offsets the eventual environmental benefits of the emission policy. However, the existence and magnitude of green paradox effects are both open questions. For instance, van der Werf and Di Maria (2012) conclude their review by lamenting that "the most striking void in this literature is an empirical assessment of the green paradox, without which it is hard (if not impossible) to provide even order of magnitude estimates of green paradox effects." I use an event study of the sudden and unexpected collapse of the U.S. Senate's climate effort to establish the existence of a green paradox in U.S. energy markets, to bound its magnitude, and to evaluate the primary channel through which it operates.

I find that if coal market participants had anticipated that the proposed emission policy would actually pass, they would have increased coal consumption as if offered an emission subsidy of between \$1 and \$12 per ton of carbon dioxide. This green paradox effect is the opposite of the emission tax of between \$5 and \$65 per ton required to internalize the climate externality (Greenstone, Kopits,

& Wolverton, 2013). The unsuccessful legislative process distorted prices to the same degree as did railroads' market power (Busse & Keohane, 2007) and utility regulation (Cicala, 2015). Green paradox effects are economically significant, and, given the past decades' discussion of potential carbon regulation, they may have been distorting energy markets for years.

I also report suggestive evidence that storage markets are the primary channel for green paradox effects, in contrast to the resource extraction channel commonly analyzed in the theoretical literature. Coal markets are linked over time in three ways. First, when the cost of extracting coal increases in cumulative extraction (Heal, 1976), extraction firms must consider how today's extraction affects tomorrow's costs. Anticipating an emission tax tomorrow reduces the marginal benefit of delaying extraction until tomorrow, which increases extraction and lowers coal prices today (Gerlagh, 2011; Hoel, 2012). Second, large quantities of coal are stored above ground at power plants and other facilities. Anticipating an emission tax tomorrow reduces the marginal benefit of storage, which decreases today's storage, lowers today's coal spot prices, and increases today's consumption. Finally, coal-burning power plants are long-lived investments. Having anticipated an emission tax tomorrow reduces the number of coal-burning plants operating today, which decreases today's coal prices and also today's coal consumption.

Distinguishing these channels is important because they have different implications for the magnitude and prevalence of green paradox effects. For instance, Jones, Keen, and Strand (2013) question the relevance of the green paradox because Hotelling models of physical exhaustion do not describe actual prices very well and because the abundance of cheap coal deposits undercuts the relevance of models of economic exhaustion. Green paradox effects become more widespread if they require only a market for physical storage. However, the potential magnitude of a storage channel may be smaller than the potential magnitude of an extraction channel. And green paradox-like effects could actually increase welfare if energy demand responds to future policies more strongly than does supply.

All three channels predict that an anticipated emission cap reduces the price of coal on futures markets. I find that coal futures prices indeed responded to the collapse of the Senate's climate effort with a statistically significant jump in the theoretically predicted direction. This result provides strong evidence for intertemporal linkages that could generate green paradox effects. Intuitively, we might expect storage to be the primary channel over the time frame of the legislative process: large stocks of coal are held in storage, but extraction decisions and demand-side investments may respond only over longer time horizons. I report evidence

Received for publication March 4, 2014. Revision accepted for publication June 6, 2016. Editor: Bryan S. Graham.

*University of Arizona.

I am grateful for comments from the editor, three anonymous referees, and seminar participants at Arizona, HEC Montréal, the Occasional California Workshop in Environmental and Resource Economics, the Stanford Institute for Theoretical Economics, Toronto, UC Berkeley, UC Davis, and Yale FES. Hoa Nguyen provided research assistance. This work was supported by the University of Arizona's Renewable Energy Network, with special thanks to Ardeth Barnhart and Stan Reynolds. I also thank Keisuke Hirano and Ashley Langer for helpful discussions.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00627.

that storage increased substantially over the event month, whereas extraction did not change by much and consumption actually declined. These results are consistent with the predictions of the storage channel but not with the other two channels. They suggest that the breakdown of the Senate's climate effort reduced near-term coal consumption by shifting spot supply into storage.

The anticipation effect at the heart of the green paradox is a more general feature of public economics. Anticipating future changes in taxes on investment, income, or consumption can induce smoothing behavior or offsetting behavior (Hall, 1971; Branson, Fraga, & Johnson, 1986; Judd, 1985; Auerbach, 1989; Yang, 2005; Mertens & Ravn, 2011). These anticipation effects can strongly affect the excess burden of taxation (Judd, 1987). The Hall (1971) consumption tax has particularly similar effects to a proposed emission price: the future change in the tax causes a jolt to real flows that an anticipation effect offsets by shifting earlier real flows in the opposite direction. The empirical literature on tax anticipation has obtained mixed results that depend on how one constructs the timing of tax changes (Mertens & Ravn, 2012; Perotti, 2012). Identifying anticipation effects in commodity markets is more straightforward for two reasons: these markets' participants do not face the liquidity constraints that could mask the effect for individuals or households, and proposals for emission pricing are not themselves responding to daily commodity market movements as, for instance, personal taxation responds to quarterly aggregate output or consumption. My results indicate that hypothesized anticipation effects need not make overly strong demands of market efficiency.

I take advantage of a unique event that occurred among the parties developing climate legislation to isolate an exogenous shift in the probability of regulation.¹ Narrative evidence, contemporary news accounts, and prediction market trades all indicate that expectations of regulation shifted when Senator Lindsey Graham walked out of his climate collaboration over the weekend before its scheduled unveiling. I measure the strength of green paradox effects by estimating how futures markets responded to Senator Graham's decision. Event studies typically examine stock market returns. In particular, event studies of energy policy have used changes in equity prices to learn about the cost of regulation and the distributional consequences of cap-and-trade programs (Lange & Linn, 2008; Linn, 2010; Bushnell, Chang, & Mansur, 2013; Meng, 2017). Commodity futures markets instead

¹Di Maria, Lange, and van der Werf (2014) examine how the price of U.S. coal changed with the announcement of future sulfur dioxide regulations. They use the final passage of the 1990 Clean Air Act Amendments to separate a period with anticipation effects from a period without anticipation effects. However, the final passage of legislation is itself usually anticipated, which generally challenges the event study methodology (Binder, 1998; Lamdin, 2001). Indeed, President Bush had announced his intent to regulate sulfur dioxide when campaigning for office in 1988, and he had formally proposed a cap-and-trade scheme in June 1989. By taking a finer view of the policymaking process, I more precisely identify a shift in expectations, and do so in the standard green paradox setting of carbon dioxide regulation.

teach us about the consequences for energy consumption and emissions. In contrast to equity market outcomes, these consequences should be largely independent of how regulators choose to allocate emission permits in a cap-and-trade program.

I begin by theoretically analyzing how commodity prices, extraction, and consumption respond to new information about future emission policies. Section III combines several lines of evidence to argue that Senator Graham's week-end withdrawal from his climate bill conveyed unanticipated information that altered expectations. Section IV introduces the estimation framework by which I identify the event's effect on futures markets. Section V demonstrates the anticipation effect in coal futures. Section VI shows that other days with similar prediction market movements also saw sizable increases in coal prices. It also evaluates the channels for intertemporal leakage and the magnitude of the anticipation effect. Section VII concludes with the policy implications of finding that markets are distorted by the suggestion of regulation. The online appendix contains supporting evidence for the event narrative and further empirical results.

II. The Effects of an Anticipated Emission Policy

An anticipated emission policy affects current emissions when markets are linked across time. In coal markets, intertemporal linkages can arise in three ways. First, extracting lower-cost deposits of coal today raises the cost of tomorrow's extraction. Second, coal is stored above ground. Third, the primary source of coal demand is long-lived power plants, whose construction depends on expectations of future profitability. I show that each type of intertemporal linkage implies that anticipated emission pricing lowers the nearer-term price of coal.

A. Theoretical Analysis

Consider a two-period, perfectly competitive economy. All firms use a common discount factor $\beta \in (0, 1]$. In each period, extraction firms decide the quantity of resource to extract. A representative firm's period t cost of extraction $C^t(q^t, Q^t)$ increases in period t extraction q^t and in cumulative extraction Q^t from periods prior to t . The cost of extraction is strictly convex, and the marginal cost of current extraction increases in cumulative extraction ($C_{qQ}^t \geq 0$, where subscripts indicate partial derivatives). Normalize cumulative extraction prior to period 1 as 0.

Storage firms buy quantity $s^1 \geq 0$ from extractors in the first period and sell it to resource consumers in the second period. The cost of storage $k(\cdot)$ is an increasing, convex function of the quantity stored. The convenience yield $\psi(\cdot)$ is a common reduced-form means of capturing the additional benefit from owning the physical commodity, as opposed to owning a financial contract for its delivery. The convenience yield is an increasing, concave function of the quantity stored.

Resource consumers obtain benefit $\pi^t(z^t, K^t)$ in period t , where z^t is time t consumption and K^t is the time t capital stock. This benefit is strictly concave and increasing in both arguments, and π_z^t goes to infinity as consumption goes to 0. The marginal benefit of consumption increases in the capital stock ($\pi_{zK}^t \geq 0$). In the first period, resource consumers choose that capital stock, which persists for both periods. The cost of installing capital K^1 is an increasing and convex function $G(K^1)$.

We take as given a tax τ on second-period emissions, which we treat as identical to second-period consumption. This tax is equivalent to the emission price under a second-period cap-and-trade instrument. All firms anticipate this tax. We consider how marginal changes in the tax affect first-period decisions. Because the economy is perfectly competitive, the first welfare theorem holds: equilibrium outcomes maximize welfare conditional on the tax. We therefore solve for equilibrium outcomes via a social planner's problem.

The quantity of period 2 extraction solves the following maximization problem:

$$V(q^1, s^1, K^1, \tau) = \max_{q^2} \{ \pi^2(q^2 + s^1, K^1) - C^2(q^2, q^1) + \psi(s^1) - k(s^1) - \tau[q^2 + s^1] \},$$

where we take the first-period decisions q^1 , s^1 , and K^1 as given. The first-order condition implicitly defines the equilibrium outcome q^{2*} as a function of those first-period decisions and the tax τ :

$$\pi_z^2(q^{2*} + s^1, K^1) = C_q^2(q^{2*}, q^1) + \tau.$$

Using the implicit function theorem, it is easy to show that q^{2*} increases in K^1 and decreases in q^1 , s^1 , and τ .

The quantities of period 1 extraction, storage, and investment solve the following maximization problem:

$$\max_{q^1, s^1, K^1} \{ \pi^1(q^1 - s^1, K^1) - C^1(q^1, 0) - G(K^1) + \beta V(q^1, s^1, K^1, \tau) \}, \text{ s.t. } s^1 \geq 0.$$

Using the envelope theorem to obtain partial derivatives of V , we find the following first-order necessary conditions:

$$\begin{aligned} 0 &= \pi_z^1 - C_q^1 - \beta C_Q^2, \\ 0 &= -\pi_z^1 + \beta [C_q^2 + \psi_s - k_s] + \lambda, \\ 0 &= \pi_K^1 - G_K + \beta \pi_K^2, \\ s^1 &\geq 0, \lambda \geq 0, \lambda s^1 = 0, \end{aligned}$$

where λ is the Karush-Kuhn-Tucker multiplier on the storage nonnegativity constraint. The equilibrium first-period resource price p^1 is equal to π_z^1 .

We consider each type of intertemporal linkage in turn. Begin with the extraction cost channel. Fix storage s^1 and investment K^1 and consider how raising the second-period emission tax affects first-period extraction. Totally

differentiating the first-order conditions and substituting yields

$$\frac{dq^1}{d\tau} = \frac{\beta C_{qQ}^2}{[-\pi_{zz}^2 + C_{qq}^2][-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] - \beta C_{qQ}^2 C_{qQ}^2} \geq 0.$$

The convexity of $C^2(q^2, Q^2)$ ensures that the denominator is positive. Raising the second-period emission tax increases first-period extraction q^1 and thus reduces the first-period resource price. From the first-order conditions, we see that resource extractors choose q^1 so that $p^1 = C_q^1 + \beta C_Q^2$. They equate marginal revenue to a forward-looking measure of marginal cost that accounts for how period 1 extraction increases costs in period 2. Because greater first-period extraction imposes additional costs in the second period, firms have an incentive to conserve resources for the second period. When the second-period tax increases, the quantity of second-period extraction declines, which reduces the marginal value of conserved resources. Firms therefore increase extraction in the first period. This increase in extraction prior to the imposition of the tax is familiar from Heal (1976)-type models of green paradox effects (Gerlagh, 2011; Hoel, 2012). The effect is strongest when C_{qQ}^2 is large.

Now consider the storage channel. Set $C_{qQ}^2 = 0$ and fix K^1 so that we turn off the extraction cost and investment channels. When $s^1 > 0$, totally differentiating the first-order conditions and substituting yields

$$\begin{aligned} \frac{ds^1}{d\tau} &= \left\{ -\beta C_{qq}^2 [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] \right\} \\ &\quad \times \left\{ \beta [k_{ss} - \psi_{ss}] [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] \right. \\ &\quad \times [-\pi_{zz}^2 + C_{qq}^2] - \pi_{zz}^1 [C_{qq}^1 + \beta C_{QQ}^2] [-\pi_{zz}^2 + C_{qq}^2] \\ &\quad \left. - \beta \pi_{zz}^2 C_{qq}^2 [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] \right\}^{-1} \\ &< 0. \end{aligned}$$

The optimal quantity of storage equalizes the marginal value of the extracted resource across the two periods, adjusted for storage costs, the convenience yield, and the emission tax. Raising the second-period emission price reduces the value of the extracted resource in the second period, which reduces both second-period extraction and the incentive to store the resource between periods. Changes in storage account for more of the reduction in second-period consumption when the second period's supply of the extracted resource is inelastic (i.e., when C_{qq}^2 is large). First-period consumption z^1 changes as

$$\frac{dz^1}{d\tau} = \frac{dq^1}{d\tau} - \frac{ds^1}{d\tau} = - \left(1 - \frac{-\pi_{zz}^1}{-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2} \right) \frac{ds^1}{d\tau} > 0.$$

The reduction in storage increases first-period spot supply, which works to increase first-period consumption directly

but also to reduce first-period consumption by crowding out first-period extraction. We see that the decline in first-period extraction does not fully offset the reduction in storage, and it offsets an especially small fraction of the reduction in storage when first-period resource demand is elastic (π_{zz}^1 is small) and resource supply is inelastic (C_{qq}^1 is large). The storage channel is similar to the extraction cost channel in predicting that an anticipated emission price increases first-period consumption and thus reduces the first-period resource price, but it differs in predicting that an anticipated emission price reduces first-period extraction.

Now consider the investment channel. Set $C_{qQ}^2 = 0$ and fix s^1 , so that we turn off the extraction cost and storage channels. Totally differentiating the first-order conditions and substituting yields

$$\begin{aligned} \frac{dK^1}{d\tau} = & \left\{ -\beta \pi_{zK}^2 [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] \right\} \\ & \times \left\{ G_{KK} [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] [-\pi_{zz}^2 + C_{qq}^2] \right. \\ & - \pi_{KK}^1 [C_{qq}^1 + \beta C_{QQ}^2] [-\pi_{zz}^2 + C_{qq}^2] \\ & - \beta \pi_{KK}^2 [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] C_{qq}^2 \\ & + [\pi_{KK}^1 \pi_{zz}^1 - \pi_{zK}^1 \pi_{zK}^1] [-\pi_{zz}^2 + C_{qq}^2] \\ & \left. + \beta [\pi_{KK}^2 \pi_{zz}^2 - \pi_{zK}^2 \pi_{zK}^2] \right\}^{-1} \\ & \times [-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2] \leq 0. \end{aligned}$$

Concavity of $\pi'(z', K')$ ensures that the denominator is positive. The second-period emission tax reduces the marginal benefit of durable capital by reducing second-period consumption. First-period investment therefore declines in the anticipated emission price, and it declines by an especially large amount when the second period's marginal benefit of capital is sensitive to consumption (i.e., when π_{zK}^2 is large). The decline in first-period investment reduces first-period extraction:

$$\frac{dq^1}{d\tau} = \frac{\pi_{zK}^1}{-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2} \frac{dK^1}{d\tau} \leq 0.$$

The first-period resource price changes as

$$\begin{aligned} \frac{dp^1}{d\tau} &= \frac{d\pi_z^1}{d\tau} = \pi_{zz}^1 \frac{dq^1}{d\tau} + \pi_{zK}^1 \frac{dK^1}{d\tau} \\ &= \left(1 - \frac{-\pi_{zz}^1}{-\pi_{zz}^1 + C_{qq}^1 + \beta C_{QQ}^2} \right) \pi_{zK}^1 \frac{dK^1}{d\tau} \\ &\leq 0. \end{aligned}$$

The decline in investment works to decrease the first-period resource price by reducing the marginal benefit from consuming the resource, but the decline in extraction works to increase the first-period resource price by reducing supply. The first effect dominates, and it dominates most strongly when first-period resource demand is elastic (π_{zz}^1 is small)

TABLE 1.—THEORETICAL PREDICTIONS FOR HOW THE PROPOSED EMISSION LEGISLATION'S COLLAPSE SHOULD HAVE AFFECTED VARIABLES IN THE YEARS BEFORE THE POLICY WOULD HAVE BEEN IMPLEMENTED (I.E., FOR HOW A REDUCTION IN τ AFFECTS PERIOD 1 VARIABLES)

Channel	Preimplementation Variable				
	Extraction q^1	Consumption z^1	Price p^1	Storage s^1	Investment K^1
Extraction cost	↓	↓	↑	-	-
Storage	↑	↓	↑	↑	-
Demand-side investment	↑	↑	↑	-	↑

and resource supply is inelastic (C_{qq}^1 is large). The demand-side investment channel therefore generates the same type of price declines as do the other two channels. But in this case, the lower first-period resource price corresponds to reduced first-period extraction (in contrast to the extraction cost channel) and reduced first-period consumption (in contrast to both the extraction cost and storage channels).

B. Predictions for the Empirical Setting

In the empirical setting, a cap-and-trade proposal collapsed. This collapse reduced the anticipated emission price (τ fell). Table 1 summarizes how the proposal's collapse would affect first-period extraction, consumption, prices, storage, and investment through each of the three channels. We observe prices for coal futures contracts that expired before the cap was to have been implemented. All three theoretical channels for intertemporal linkages predict that the proposal's collapse increased the prices of these coal futures.² The extraction cost channel predicts that the proposal's collapse decreased resource extraction in the preimplementation period, but the storage and investment channels predict that the proposal's collapse increased extraction in the preimplementation period. Finally, the extraction cost and storage channels predict that the bill's collapse decreased coal consumption in the years leading up to the emission cap's planned implementation in 2013, whereas the demand-side investment channel predicts that the bill's collapse increased coal consumption in the preimplementation period.

III. Shifting Regulatory Expectations: Senator Graham Abandons His Bill

Having analyzed how evolving expectations of future carbon pricing should affect current energy futures, I now argue that Senator Lindsey Graham's abandonment of the American Power Act isolates a shift in expectations that can test

²Futures prices are the risk-adjusted expectation of spot prices. Our comparative statics for spot prices directly map into comparative statics for futures prices under the plausible assumptions that, after the change in the anticipated policy, the new distribution for future spot prices first-order stochastically dominates (or is dominated by) the original one and that the change in the anticipated policy does not change the economy-wide pricing kernel prior to time T (e.g., through general equilibrium effects).

the theory. The main identification challenge is establishing that even insiders did not see the event coming. I establish surprise using contemporary news reports, retrospective accounts, prediction market outcomes, and Internet search patterns.

On June 26, 2009, the U.S. House of Representatives passed the American Clean Energy and Security Act, also known as Waxman-Markey after its sponsors. That bill aimed to reduce U.S. greenhouse gas emissions to 17% below the 2005 level by 2020 and to 83% below the 2005 level by 2050 through a cap-and-trade scheme of tradable emission permits. The vote was tight (219–212) and largely followed party lines, with 44 Democrats voting against it and only 8 Republicans supporting it. The Senate now had a year and a half to pass a similar bill, but it faced a stiffer requirement: it needed 60 votes (out of 100) to overcome a Republican filibuster. By spring 2010, there were 57 Democratic senators and 2 independent senators (who both typically voted with the Democrats). The 60-vote hurdle meant that any bill would need to attract the support of at least one Republican in order to become law.

Senator Graham was that Republican. In a joint October 2009 op-ed in the *New York Times*, the senator from South Carolina announced a partnership with Democratic senator John Kerry to pass climate legislation. Senator Joseph Lieberman, an independent, soon joined their efforts to craft what became known as the Kerry-Graham-Lieberman bill. Their effort was the focal point for Senate climate legislation. Their emission limits would begin with a 4.75% reduction from the 2005 level in 2013, starting one year later than in Waxman-Markey but otherwise matching the House bill's emission trajectory. As the primary source of near-term emission reductions, the electricity sector would be subject to a cap-and-trade program with a price floor and a form of price ceiling.³ To address petroleum refiners' concerns about permit price fluctuations, the transportation sector would be regulated by a novel "linked fee" that provided price stability. The senators lined up support from a range of interest groups by offering perks such as free emission permits, expanded guarantees for nuclear power, and expanded oil drilling.

On Thursday, April 22, 2010, after months of meetings and negotiations, the three senators completed the bill with a final compromise to bring the Edison Electric Institute on board (Lizza, 2010). They scheduled the bill's unveiling for Monday, April 26. The press conference would include top business, environmental, and religious figures to demonstrate the breadth of support for their comprehensive package.

On Friday, April 23, Arizona governor Jan Brewer signed SB 1070 into law. This controversial bill provoked Latino

³The electricity sector cap-and-trade program was expected to be the dominant policy component over the first decade or more. Home natural gas use and industrial facilities were not to be regulated until 2016, three years after the electricity and transportation sectors.

leaders by its measures to deter illegal immigration. The governor had not been closely involved in the legislative effort, and her support was uncertain as the window for her to sign the bill began to close. Meanwhile, Senate majority leader Harry Reid was trailing his challenger in the polls leading up to an election in neighboring, Latino-heavy Nevada. The day before the signing (April 22), he announced that the Senate would begin advancing a federal immigration bill.⁴ Yet the Senate calendar could not handle immigration reform at the same time as the climate bill. Moreover, there was already a completed climate bill and no immigration bill near ready for advancement. As the Republican involved in both the climate and immigration efforts, Senator Graham knew this well. According to Lizza (2010), Senator Graham considered the announcement cheap political point scoring that demonstrated a lack of commitment from the Senate leadership to take care of climate change.

On Saturday morning, the *Washington Post* gave high billing to the upcoming unveiling of Kerry-Graham-Lieberman. It detailed the bill's array of industry support and quoted Senator Graham saying the bill would not raise the price of gasoline. Monday's *Platts Coal Trader* featured a story on the bipartisan bill's expected unveiling.⁵ But this story must have been filed before the weekend. According to Lizza (2010), at around 10:00 p.m. on Friday night, one of Senator Graham's aides e-mailed his counterpart on Senator Lieberman's side to say, "Sorry buddy." Lieberman's aide later described the note as "soul-crushing." On Saturday morning, Senator Graham abandoned his bill. In his formal statement, he refused to delay a climate bill for immigration.

Neither industry insiders nor the bill's other sponsors seemed to know until Saturday that the bill was about to be held up. Senator Kerry promptly flew back to Washington, DC from Boston to meet with Senator Graham. Senator Lieberman broke his Sabbath restrictions to call Senator Graham. Senator Lieberman's aide received the confirmation by text message while describing the bill's final provisions to key natural gas lobbyists. At that point, says Lizza, the Kerry-Graham-Lieberman bill, "perhaps the last best chance

⁴Internet search data (see the appendix) show that interest in "Arizona immigration" began to register on April 20 after the state senate passed the bill and spiked dramatically upon Governor Brewer's signature on April 23. Arizona's immigration policy grabbed the spotlight rather suddenly, which created the political opening for majority leader Reid.

⁵The *Platts Coal Trader's* daily "Emissions Roundup" regularly featured the Senate bill's progress. Its daily "Market Commentary" ascribed coal's price increase on Friday, April 23, to increases in the price of natural gas and ascribed the increase on Monday, April 26, to "strong perceived demand from the German power sector." The *Platts Coal Trader International* had previously ascribed a Friday, April 23, increase in European coal swaps to Germany's "brighter economic sentiment." Our specifications with covariates will aim to maintain identification even if there were a further change in German economic prospects over the weekend. Further, because natural gas is both subject to a future carbon price and also the primary substitute for coal, our theoretical analysis differs from contemporary market analysis in expecting only a small—and potentially opposite—comovement in the natural gas market. This is in fact what the appendix reports.

to deal with global warming in the Obama era, was officially dead.”⁶

This narrative strongly suggests that the prospect of the Senate passing climate legislation took a sharp hit between markets’ closing on Friday, April 23 and their opening on Monday, April 26. The appendix shows that Senator Graham’s Saturday withdrawal made headline news and altered prediction markets. Hints of potential problems due to immigration were present by Friday, but the climate bill’s press conference was still scheduled as of Friday’s closing. Indeed, key players were still planning to hold the press conference as late as Saturday morning. Events after Friday afternoon appear to have surprised the bill’s authors, their aides, and pundits. We now examine how they affected futures markets.

IV. Estimation Framework

I use the exogenous shift in carbon price expectations over the April 24 weekend to identify the effect of an anticipated emission policy on current commodity prices.⁷ Two primary choices in an event study are the event window and the estimation window (MacKinlay, 1997). The event window is the period in which expectations shifted, and the estimation window is the period on either side of the event window used to establish counterfactual returns. The event happened on Saturday, April 24 while markets were closed. The event day is therefore Monday, April 26. Following common practice, the event window also includes the day before the event (Friday, April 23) and the day after the event (Tuesday, April 27).⁸

We must navigate a trade-off when choosing the estimation window: longer periods provide more observations, but shorter periods define counterfactual returns under market conditions more similar to the date of interest and have more usable futures contracts. I employ two estimation windows to assess robustness. The shorter window uses the 60 trading days centered around the event window, and the longer window uses the 200 trading days centered around the event window.⁹ As described below, these two window lengths

⁶ On May 12, Senators Kerry and Lieberman finally unveiled their bill as the American Power Act. They did so without Senator Graham. On June 8, Senator Graham confirmed that he would not vote for his own bill. Nobody would get the chance. In late July, Senate Majority Leader Reid announced that he would pursue smaller, more narrowly targeted energy bills.

⁷ I focus on commodity prices instead of equity prices for two reasons. First, the theoretical predictions are for equilibrium quantities, prices, and emissions, not for firm profits. Second, profits depend on the expected permit allocation rule, whereas quantity outcomes should depend only on the expected permit price.

⁸ I use a different dummy variable for each day in the event window. Extending the event window beyond Monday, April 26, merely removes those additional days from the estimation window.

⁹ The exact dates of the estimation windows vary slightly depending on the covariates because of differences between U.S. and British banking holidays. For specifications without covariates, the 60-day estimation window extends from March 11, 2010, to June 9, 2010, and the 200-day estimation window extends from November 27, 2009, to September 17, 2010. For specifications with covariates, the 60-day estimation window extends from March 10, 2010 to June 10, 2010, and the 200-day estimation window extends from November 24, 2009 to September 21, 2010.

allow us to detect effects at the conventional 10%, 5%, and, for the 200-day window, 1% significance levels when using the sample quantile test.

My estimation strategy compares the unexplained variation on the event day to the distribution of the estimation window’s residuals. Two factors make it challenging to detect an effect. First, when futures returns contain a lot of unexplained variation, it will be difficult to pick up the event’s signal. Second, expectations of emission policies surely shifted at other times in the estimation window, and the presence of these pseudo-events in the control group will reduce the degree to which the treatment group looks extreme (Hakala, 2010). To mitigate the first problem, some specifications include covariates and more refined error structures in order to explain more of the nonevent variation in futures returns. The second problem is mitigated by removing Friday, April 23 from the estimation window, but it still biases the results toward no statistically significant effect.

In order to assess how the event affects a given contract, I regress daily futures returns on the event window dummies and, in some specifications, on covariates:

$$F_{it} = X_{it}\beta_i + D_t^{Fri}\gamma_i^{Fri} + D_t^{Mon}\gamma_i^{Mon} + D_t^{Tues}\gamma_i^{Tues} + \epsilon_{it}, \quad (1)$$

where $i \in \{1, \dots, N\}$ indexes the futures contract (by month of expiration) and t indexes the trading day. Each contract is run in its own regression.¹⁰ There are $T + 3$ trading days in the combined estimation and event windows. F_{it} is the log return. X_{it} is the $1 \times k$ vector of covariates, with β_i the $k \times 1$ vector of coefficients. Each D_t^d is a dummy variable indicating a day in the event window, with d corresponding to Friday, April 23 (*Fri*); Monday, April 26 (*Mon*); or Tuesday, April 27 (*Tues*). Each D_t^d equals 1 if day t corresponds to day d and 0 otherwise. Excess (or unpredicted) returns for days in the event window are given by the dummy variables’ coefficients γ_i^d , and excess returns for other days are given by the residuals ϵ_{it} . For each contract, the coefficient of interest is γ_i^{Mon} . The GARCH(1,1) specification represents the error term as a generalized autoregressive conditional heteroskedasticity model (Bollerslev, 1986). This formulation, previously found to work well in futures returns (McKenzie, Thomsen, & Dixon, 2004), allows the variance of the error term to vary over time in a fashion that possibly exhibits clustering. To the extent that it better captures the true error process, the GARCH specification should increase the power of tests for an event effect.

Exogeneity of the event requires that some type of unusual nonevent news did not move markets on April 26. Specifications without covariates include only the constant in X_{it} . In this case, the dummy coefficients are simply the recentered futures returns on those particular days. Including covariates helps ensure identification by absorbing other news around

¹⁰ Durbin-Watson tests indicate positive autocorrelation in many specifications. Reported standard errors use the Newey-West automatic bandwidth selection procedure to make them robust to arbitrary autocorrelation.

the event day,¹¹ but including covariates can also bias the estimates toward 0 by absorbing some of the event's effect into the explained portion of returns.¹² The specifications with covariates include a constant and variables (in log return form) to capture general economic conditions and broader commodity market movements:¹³ the S&P 500 stock price index, the ten-year U.S. Treasury rate (constant maturity), the three-month London Interbank Offered Rate (LIBOR), the benchmark West Texas Intermediate crude oil contract for the month nearest to contract i ,¹⁴ the Baltic Dry Index to represent the price of transporting raw materials by sea, and the Thomson Reuters/Jefferies CRB Index to represent broader commodity market movements.¹⁵ Figure 1 shows that three of the covariates that are potentially most sensitive to the event do not behave abnormally in the event window.

Coal futures prices use the NYMEX Central Appalachian contract.¹⁶ This is the reference contract for coal in the North American market. The price of the nearest-dated contract is conventionally used as the spot price in this market. The longest-dated liquid contracts expire before the proposed cap-and-trade program would have begun constraining emissions, which means that the theoretical model predicts the same type of price change for each contract. The top panel of figure 2 plots futures prices in the sixty-day estimation window. Longer-dated contracts tend to have

¹¹ Specifications with covariates act like a difference-in-differences model with a synthetic control group.

¹² The Senate bill's primary near-term policy was a cap-and-trade program for the electric power sector. Anticipating this policy probably did not strongly affect broad equity market indices, in part because the free permit allocations were designed to dampen the effect on retail electricity prices. The event also probably did not strongly affect crude oil prices: the refinery regulations would have acted like a gasoline tax of around 15 cents per gallon in only one portion of a global oil market.

¹³ S&P 500, Treasury rates, and LIBOR are available from the Federal Reserve Economic Data (FRED) website. I purchased the Baltic Dry Index series from EODData and the CRB index and crude oil series from the Commodity Research Bureau.

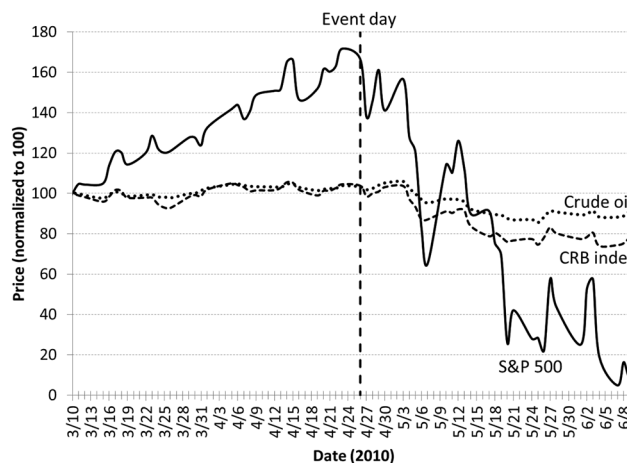
¹⁴ This is the only covariate that varies by contract. Because the covariates are basically the same across contracts, there is little efficiency gain from using a seemingly unrelated regression framework to estimate equation (1).

¹⁵ I construct a rolling contract using the nearest month's price. Note that natural gas forms 6% of the index. Including this index increases power insofar as it picks up nonevent noise, but it reduces power insofar as event-related changes in the price of natural gas might move the index. Geman (2005) recommends this index in part because no single commodity tends to move it.

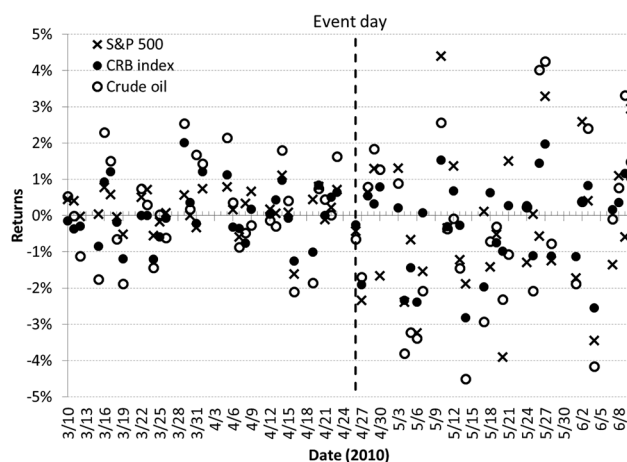
¹⁶ Purchased from the Commodity Research Bureau (NYMEX: QL). The coal contract is the benchmark "Big Sandy" contract for delivery in a specified section of river in Central Appalachia. The coal could have any origin, provided it meets the quality specifications. In particular, the energy content must be at least 12,000 Btu/lb and the sulfur content must be no greater than 1%. In contrast to the Central Appalachian contract, the ICE futures contract for Powder River Basin coal often shows zero open interest and zero volume. The coal contract is not subject to a limit move restriction, but the crude oil contract pauses trading after a large move in the nearest contract months. Limit moves do not appear in the event or estimation windows.

Each specification uses only those contracts that are listed throughout the estimation window and have nonzero volume nearly every day throughout the estimation window. These restrictions yield eighteen contracts in the 60-day estimation window and fifteen contracts in the 200-day estimation window.

FIGURE 1.—PRICES (TOP) AND RETURNS (BOTTOM) FOR THREE OF THE COVARIATES IN THE SIXTY-DAY ESTIMATION WINDOW



(a) Covariates' prices



(b) Covariates' returns

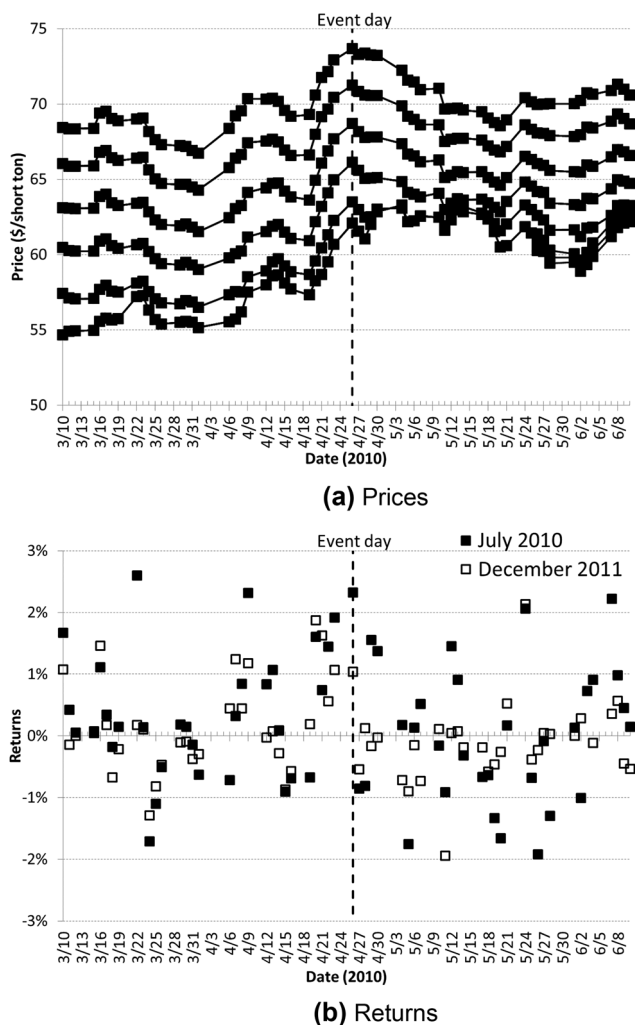
The crude oil plot uses the July 2010 contract.

higher prices. Each series remains within a relatively narrow band throughout the window. The price increases on the event day by \$1.50 per short ton for the nearest-dated contracts and by \$0.75 per short ton for the longest-dated contracts.

The bottom panel of figure 2 plots returns for the nearest- and longest-dated contracts in the sixty-day estimation window. The event day's coal returns are unusually large: the July 2010 contract provided greater returns on only one other day in the estimation window, and the December 2011 contract provided greater returns on only eight other days. This visual evidence suggests a positive event effect for coal at the edge of the "usual" variation.

Figure 3 plots the change in price (top) and the return (bottom) for each futures contract on the event day, on the trading day before the event, and on the trading day after the event. The event appears to have increased prices and

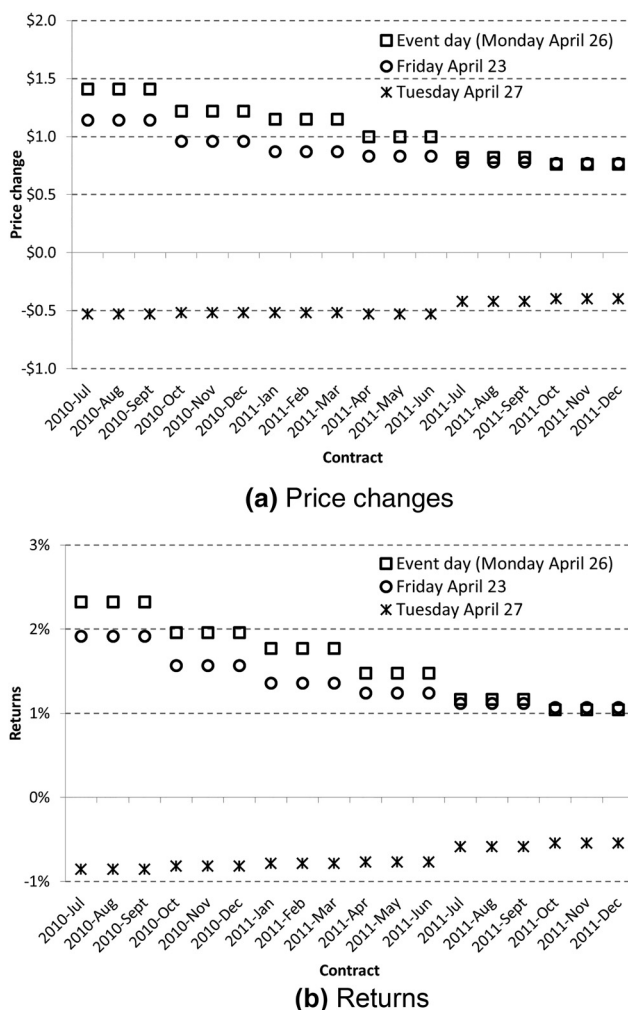
FIGURE 2.—PRICES FOR ALL COAL CONTRACTS (TOP) AND RETURNS FOR THE NEAREST- AND LONGEST-DATED COAL CONTRACTS (BOTTOM) IN THE SIXTY-DAY ESTIMATION WINDOW



returns across all maturities.¹⁷ As discussed in section III, the first phase of the event may have begun unfolding on Friday, April 23. There, too, we see positive returns across all maturities. In contrast, we have no evidence of legislative news on the day after the event. Indeed, that day's returns were less than half as large as the event day's returns and of the opposite sign.

¹⁷ This plot also shows that the event's effect appears to decline in time to maturity. For now, this term structure remains a bit of a puzzle. Continuous-time theoretical analysis (not reported) with quadratic functional forms suggests that this term structure is unlikely to be consistent with the extraction cost and demand-side investment channels but is consistent with the storage channel if the event directly increased the marginal convenience yield. It is also possible that this term structure reflects the potential for different types of regulations at a later date (credit to Brian Wright for the suggestion). Finally, this term structure could merely reflect differences in factors such as liquidity across futures contracts. The online appendix shows that nearer-dated contracts tend to have greater volume, and both volume and prices seem to regularly move together within groups of three contracts.

FIGURE 3.—TERM STRUCTURE OF FUTURES PRICE CHANGES (TOP) AND RETURNS (BOTTOM) ON THE EVENT DAY AND THE DAYS IMMEDIATELY BEFORE AND AFTER



A complication arises when testing whether the coefficient of primary interest (γ_i^{Mon}) is significantly different from 0 (Conley & Taber, 2011; Gelbach, Helland, & Klick, 2013). Because the coefficient is identified from a single observation (the excess return on the event day), central limit theorem arguments do not apply to the distribution of the estimator. The assumptions behind a standard t -test therefore hold only if the event day residual is drawn from a population of excess returns that is itself normal.¹⁸ Gelbach et al. (2013) propose the sample quantile (SQ) test for this case where the central limit theorem does not apply. The test assumes that the process generating excess returns is

¹⁸ It has been recognized since at least Brown and Warner (1985) that individual securities' returns are often nonnormal. Indeed, Jarque-Bera and Shapiro-Wilk tests for normality produce mixed results. For specifications with covariates, neither test rejects normality at a reasonable significance level for contracts in the 200-day window or for the six nearest-dated contracts in the 60-day window. On the other hand, both tests reject normality at the 5% level for the longer-dated contracts in the 60-day window.

stationary, in which case the event day's counterfactual distribution can be approximated by the realized distribution of excess returns.¹⁹ The quantile corresponding to the event day's excess return then determines its likelihood of being generated by that same process. In order to avoid distorting the size of the test, I choose the length of the estimation windows so that the 90th, 95th, and, for the longer estimation window, 99th percentiles are directly observed in the data.

A final test achieves greater power by using the theoretical prediction that the event effect has the same sign across contract months. This approach tests alternate hypotheses about the sum of the unexplained returns for a commodity's contracts. The key step is to construct a summary index for each day's excess returns (Anderson, 2008). I then test the event's effect on the summary index using t -tests and SQ tests.²⁰ I consider two summary indices in specifications with covariates. Both average a day's standardized excess returns across contract lengths. The first approach uses an unweighted average, and the second uses an efficient generalized least squares (GLS) weighting procedure.²¹ Contracts that are less correlated with the others provide more novel information and so receive higher weights; contracts that are more correlated provide more redundant information and so receive lower weights. The summary index approach increases the power of tests for the event's effect if standard noise is not strongly correlated across contract months.

V. Results: The Event's Effect on Coal Futures

We are first interested in whether coal futures' event-day coefficients γ_i^{Mon} are positive, as predicted by theory. Table 2 reports the coefficients on the event-day dummy for each coal contract. The nearest-dated contracts have event-day excess returns of around 2% that are significant at the 5% level.²² The longest-dated contracts have event-day excess returns of around 1% and still sometimes achieve statistical significance at the 10% level. The central estimates are stable across all six specifications. We see the greater power of the GARCH specification: under either estimation window, its

¹⁹ Autocorrelation in the error term does not pose a special problem as long as the estimation window is sufficiently large and there is only a single event observation (Andrews, 2003).

²⁰ Gelbach et al. (2013) establish the validity of the SQ test for a real-valued function of the outcome vector. The Jarque-Bera and Shapiro-Wilk tests reject normality of the summary index at the 1% and 5% levels, respectively, for the 60-day estimation window. They fail to reject normality at standard significance levels for the 200-day window.

²¹ The weights are as described in Anderson (2008). His "outcomes" are our contract lengths, his "areas" are our commodities, and his "individuals" are our trading days.

²² When including covariates, the coal contracts have an R^2 of around 0.21 to 0.32 in the 60-day estimation window and around 0.22 to 0.24 in the 200-day estimation window. The commodity index and S&P 500 equity index often have significantly positive coefficients, while the ten-year Treasury rate often has a significantly negative coefficient. Other covariates' coefficients are not typically significant. The estimated GARCH and ARCH terms are not statistically significant.

estimated positive event effect is statistically significant at the 1% level for nearly all contracts because it produces smaller standard errors.

I have argued that Friday could have been the first phase of the event, in which case we would expect a positive coefficient γ_i^{Fri} . Indeed, the central estimates of the coefficient for the day-before dummy (see the online appendix) are similar to those of the event day's coefficient, with the nearest-dated contracts having excess returns of around 1.7% and the longest-dated contracts having excess returns of around 1.3%. Some contracts' day-before coefficients are significant at the 10% level with conventional error terms, and most are significant at the 1% level with GARCH error terms. In contrast, we have no reason to expect any particular sign for the day-after coefficient. In fact, none of the estimates for γ_i^{Tues} (see the online appendix) are statistically significant at even the 10% level, and the central estimates suggest small negative returns of less than 1%.

The SQ test replaces the t -test's assumption that the distribution of excess returns is normal with the assumption of stationarity. Figure 4 plots the distribution of the nearest-dated contract's excess returns for specifications (2) and (5). The event day's excess returns are far in the positive tail of this distribution. Table 3 confirms this visual evidence. The SQ test judges the event-day coefficient to be significant at the 10% level (and often at the 5% level) across nearly all contracts and specifications. Excess returns of the event day's magnitude are observed in fewer than 10% of the trading days in either estimation window, and often in fewer than 5%.

The summary index tests consider hypotheses about the sum of the contracts' event-day coefficients (table 4). The event-day index is positive and significantly different from 0 at the 5% level. This result holds up under standard t -tests and also under the SQ test, and it is robust to the choice of weighting scheme. The summary index tests support the stronger hypothesis that coal's contracts jointly demonstrate a consistently positive event effect.

VI. Discussion

We have seen that coal futures jumped by a statistically significant amount in the theoretically predicted direction on the event day. I next discuss evidence of anticipation effects from other days' price movements. I then assess the dominant channel for intertemporal linkages and evaluate the magnitude of the Senate proposal's preimplementation effect on coal prices.

A. Evidence from Other Days' Price Movements

I have focused on a particular event because it clearly contains surprising, significant news. Coal futures responded in the predicted direction on the main event day (Monday, April 26) and even on the day when the event's first

TABLE 2.—COEFFICIENTS (γ_i^{Mon}) ON THE DUMMY VARIABLE FOR THE EVENT DAY, AS IN EQUATION (1)
Dependent Variable: Log Futures Returns for the Indicated Contract

Contract (<i>i</i>)	(1)	(2)	(3)	(4)	(5)	(6)
July 2010	0.021** (0.011)	0.021** (0.011)	0.021*** (0.0043)			
August 2010	0.021** (0.010)	0.021** (0.011)	0.021*** (0.0041)			
September 2010	0.021** (0.010)	0.021** (0.011)	0.021*** (0.0040)			
October 2010	0.018** (0.0085)	0.017** (0.0085)	0.017*** (0.0041)	0.020** (0.010)	0.018** (0.0094)	0.018*** (0.0050)
November 2010	0.018** (0.0085)	0.017** (0.0085)	0.017*** (0.0042)	0.019** (0.010)	0.018** (0.0093)	0.018*** (0.0025)
December 2010	0.018** (0.0085)	0.017** (0.0085)	0.017*** (0.0042)	0.019** (0.0099)	0.018** (0.0092)	0.018*** (0.0024)
January 2011	0.017** (0.0078)	0.016** (0.0081)	0.016*** (0.0028)	0.018** (0.0090)	0.016** (0.0084)	0.016*** (0.0019)
February 2011	0.017** (0.0078)	0.016** (0.0081)	0.016*** (0.0028)	0.018** (0.0090)	0.016** (0.0084)	0.016*** (0.0019)
March 2011	0.017** (0.0078)	0.016** (0.0081)	0.016*** (0.0028)	0.018** (0.0090)	0.016** (0.0084)	0.016*** (0.0019)
April 2011	0.014** (0.0077)	0.013* (0.0082)	0.013*** (0.0028)	0.015** (0.0085)	0.014** (0.0080)	0.014*** (0.0019)
May 2011	0.014** (0.0077)	0.013* (0.0082)	0.013*** (0.0029)	0.015** (0.0085)	0.014** (0.0080)	0.014*** (0.0019)
June 2011	0.014** (0.0077)	0.013* (0.0082)	0.013*** (0.0029)	0.015** (0.0085)	0.014** (0.0080)	0.014*** (0.0019)
July 2011	0.011* (0.0075)	0.0096 (0.0078)	0.0096*** (0.0027)	0.012* (0.0080)	0.011* (0.0075)	0.011*** (0.0017)
August 2011	0.011* (0.0075)	0.0096 (0.0078)	0.0096*** (0.0027)	0.012* (0.0080)	0.011* (0.0075)	0.011*** (0.0017)
September 2011	0.011* (0.0075)	0.0096 (0.0078)	0.0096*** (0.0027)	0.012* (0.0080)	0.011* (0.0075)	0.011*** (0.0017)
October 2011	0.010* (0.0077)	0.0082 (0.0079)	0.0082 (0.029)	0.011* (0.0077)	0.0095* (0.0072)	0.0095*** (0.0017)
November 2011	0.010* (0.0077)	0.0082 (0.0079)	0.0082 (0.039)	0.011* (0.0077)	0.0095* (0.0072)	0.0095*** (0.0017)
December 2011	0.010* (0.0077)	0.0082 (0.0079)	0.0082 (0.063)	0.011* (0.0077)	0.0095* (0.0072)	0.0095*** (0.0017)
Estimation window	60 days	60 days	60 days	200 days	200 days	200 days
Covariates	No	Yes	Yes	No	Yes	Yes
GARCH error	No	No	Yes	No	No	Yes

Standard errors, in parentheses, use the Newey-West automatic bandwidth selection procedure to make them robust to arbitrary autocorrelation or, in the GARCH specifications, use the Huber-White estimator to make them robust to heteroskedasticity. Significant at *10%, **5%, ***1% for a *t*-test of the null hypothesis of a weakly negative effect. Covariates are the S&P 500, the ten-year U.S. Treasury rate, three-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract. Each entry in this table corresponds to a separate regression. R^2 values and other coefficients are described in note 22.

stages began unfolding (Friday, April 23).²³ As detailed in the appendix, the price of the Intrade prediction market contract betting on congressional passage of climate legislation declined through the event window and experienced two of the five greatest volumes in its 610-day lifetime. The appendix also shows that this contract is highly illiquid: a typical trading day sees no activity. It would be a mistake to

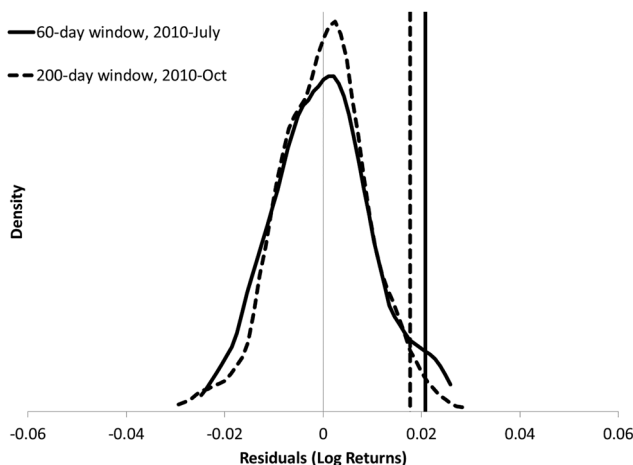
²³ Some might be concerned that the event-day coefficient is merely picking up a “Monday” or “weekend” effect, despite controlling for covariates. For instance, if the stochastic process generating returns operates in calendar time rather than trading time, then we expect greater variance and greater average returns when markets reopen after weekend closures. However, evidence from equity markets supports neither hypothesis (Fama, 1965; French, 1980). Indeed, the Monday or weekend effect has instead come to denote the frequent finding of average negative returns upon reopening (French, 1980; de Jong, Kemna, & Kloek, 1992; Alt, Fortin, & Weinberger, 2011). Such an effect would work against our finding significantly positive returns on the event day. Regardless, of the twelve other postweekend trading days in figure 2(b), all but three have returns of less than 1% in either direction for the July 2010 contract, and all but one have returns of less than 0.5% in either direction for the December 2011 contract. This section discusses supporting evidence from other days of the week.

place much weight on the precise prediction market prices.²⁴ However, we can use notable movements in this contract to identify other potential events that can test the robustness of our results. In particular, the sixty-day estimation window contains two other days on which the Intrade prediction market price changed by a large amount with exceptional volume: the declines in the Intrade price on Friday, April 30, and Thursday, May 13, were greater than on any single day around the event window and were accompanied by the sixth- and fifteenth-greatest volumes in the contract’s lifetime.²⁵ The April 30 price decline on Intrade probably reflects the heightening of the immigration-climate standoff: Senator Graham presented an uncompromising posture in an

²⁴ Even if this prediction market contract were liquid, its price is endogenously determined with futures prices if prediction market participants look to futures or equity markets when placing bets.

²⁵ Unfortunately for the potential of a further robustness check, this prediction market contract never combines a large price increase with high volume.

FIGURE 4.—KERNEL DENSITY PLOT OF THE RESIDUALS IN THE ESTIMATION WINDOW, WITH THE COEFFICIENT (γ^{Mon}) ON THE MONDAY, APRIL 26 DUMMY INDICATED BY A VERTICAL LINE



This plot uses the nearest-dated contract under specifications (2) and (5).

interview published in the middle of the previous day, and Senate Democrats nonetheless unveiled their immigration framework after markets closed that evening. The May 13 price decline on Intrade is probably related to the May 12 press conference at which Senators Kerry and Lieberman finally unveiled the American Power Act, but did so without any Republican support.

The combination of highly unusual prediction market movements and plausibly new information about legislative prospects suggests the potential for anticipatory price jumps in coal markets. Indeed, the nearest-dated coal contracts increased by 1.6% on April 29, 1.4% on April 30, 1.5% on May 12, and 0.9% on May 13. All of these movements are smaller than on the event day but are nonetheless relatively

large movements for these contracts (see figure 2). These price changes would not be diagnostic regardless of their directions, but they do provide further suggestive evidence that current coal markets respond to news about potential emission regulations in ways that match green paradox predictions.

B. Channels for Intertemporal Leakage

We have seen that the unexpected collapse of the legislative effort increased coal prices in accord with the theoretical analysis of section II. Of the three channels discussed there, the storage channel is a priori the most plausible. At any given point between 2009 and 2013, the U.S. electric power sector stored 140 to 200 million tons of coal, which would have allowed the average plant 50 to 100 days of typical operation (Energy Information Association, 2013). The relative ease of physical storage should pass anticipated price changes into current prices. Meanwhile, the other two channels should be weak over the timescale of this analysis: electric power plants require years to plan and build (K^1 is approximately fixed), and coal resources are sufficiently large that their quality is unlikely to be sensitive to near-term extraction (C_{qQ}^2 is small).

The appendix reports suggestive evidence in favor of the storage channel in monthly data from January 1973 through August 2013 for U.S. coal production, consumption, and electric power sector inventories. First, the extraction cost channel predicts that coal extraction decreased over the event month. However, the statistical analysis suggests that, if anything, seasonally adjusted extraction increased slightly over the event month. This pattern is consistent with the storage and investment channels. Second, the demand-side investment channel predicts that coal consumption increased over the event month. However, the statistical analysis shows that

TABLE 3.—PERCENTILE RANK OF THE EVENT-DAY RESIDUAL

Contract (<i>i</i>)	(1)	(2)	(3)	(4)	(5)	(6)
July 2010	97**	96**	96**			
August 2010	97**	97**	97**			
September 2010	97**	100***	100***			
October 2010	98**	97**	97**	97**	97**	97**
November 2010	98**	97**	97**	97**	97**	97**
December 2010	98**	97**	97**	97**	97**	97**
January 2011	96**	97**	97**	97**	97**	97**
February 2011	96**	97**	97**	97**	97**	97**
March 2011	96**	97**	97**	97**	97**	97**
April 2011	94*	94*	94*	95*	96**	96**
May 2011	94*	94*	94*	95*	96**	96**
June 2011	94*	94*	94*	95*	96**	96**
July 2011	90	92*	92*	92*	93*	93*
August 2011	90	92*	92*	92*	93*	93*
September 2011	90	92*	92*	92*	93*	93*
October 2011	88	88	88	91*	91*	91*
November 2011	88	88	88	91*	91*	91*
December 2011	88	88	88	91*	91*	91*
Estimation window	60 days	60 days	60 days	200 days	200 days	200 days
Covariates	No	Yes	Yes	No	Yes	Yes
GARCH error	No	No	Yes	No	No	Yes

Significant at *10%, **5%, ***1% for a null hypothesis of a weakly negative effect. The sixty-day estimation window cannot distinguish between significance levels smaller than 1.67%. Covariates are the S&P 500, the ten-year U.S. Treasury rate, three-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract. Each entry corresponds to a separate regression.

TABLE 4.—SUMMARY INDEX TESTS ON THE EVENT-DAY DUMMY IN SPECIFICATIONS WITH COVARIATES AND STANDARD (NON-GARCH) ERROR TERMS

	60-Day Estimation Window		200-Day Estimation Window	
	Unweighted	Weighted	Unweighted	Weighted
<i>Coefficients, standard errors, and t-tests</i>				
	1.8**	1.7**	1.7**	1.7**
	(0.92)	(0.79)	(0.97)	(0.92)
<i>SQ tests (percentile rank)</i>				
	96**	97**	96**	97**

The unweighted columns take the average of the standardized residuals across each day's contracts, while the weighted columns use the efficient weighting scheme. Significant at *10%, **5%, ***1% for a null hypothesis of a weakly negative effect.


seasonally adjusted consumption decreased over the event month, with borderline statistical significance. This pattern is consistent with the storage and extraction cost channels. Third, the storage channel predicts that coal inventories increased over the event month. Visual inspection suggests that coal inventories increased by an unusual amount over the event month, and the statistical analysis suggests that the seasonally adjusted increase is economically significant, though not statistically significant. Taken together, these results match the reasonable expectation that physical storage is the primary link between the 2010 coal market and the 2013 coal market.

C. Magnitude of the Event's Effect on Coal Markets


The reported coefficients likely underestimate the total effect of the Senate's legislation because the pre-event probability of passage was less than 1 and the post-event probability of passage was greater than 0. For instance, there is some evidence that the probability of passage began dropping on Friday, April 23. Including this coefficient estimate in the event's effect would nearly double all reported results while still underestimating the total effect of the bill. To develop an upper bound for the bill's total effect, note that prediction market contracts fell by around \$5 over the event weekend. If, despite concerns about liquidity and theoretical questions about whose beliefs prediction market prices represent (Manski, 2006; Fountain & Harrison, 2011), we take these changes as suggesting that the Senate bill's probability of passage fell by 5 percentage points over the event weekend, then the effect of the full Senate bill would be twenty times our estimates. Finally, the collapse of the Senate's 2010 effort may not have strongly reduced the probability of eventual greenhouse gas regulations, whether through later legislation or executive branch actions. Not only do our reported numbers underestimate the effect of the Senate bill by a factor of 2 to 20, but they are also a lower bound for the effect of anticipating that some type of greenhouse gas regulations will be implemented at some future time.

The top panel of figure 5 translates the event's estimated effect on coal futures' returns into its effect on futures prices. Coal prices around the event day were between \$60 and

FIGURE 5.—ESTIMATED EFFECT OF THE SENATE EVENT ON COAL FUTURES PRICES UNDER SPECIFICATION (2) AND THE CARBON PRICE THAT WOULD EQUAL THE CHANGE IN THE FUTURES PRICE



(a) Effect on coal prices



(b) Equivalent change in the carbon price

Error bars give the 90% confidence interval.

\$70 per short ton. The bill's collapse increased the nearest-dated contracts' prices by around \$1.30 per short ton, with the 90% confidence interval extending from around \$0.10 to \$2.50 per short ton.²⁶ The effects are smaller for longer-dated contracts, with several having a central estimate of around \$1.00 per short ton and the longest-dated contract having a central estimate just above \$0.60 per short ton. Combining the rough probability bounds with the central estimates, the bill's actual passage would have reduced the nearest-dated contracts' prices by between \$2.60 and \$26.00 per short ton in the years prior to the policy's implementation. These price changes are nontrivial when we consider that the U.S. electric power sector consumed 975 million short

²⁶ Confidence interval calculations assume that the noise introduced through β and ϵ is trivial in comparison to the noise in $\sqrt{v_i^{Mon}}$ and that the estimator is normally distributed.

tons of coal in 2010 (Energy Information Administration, 2012a).

For comparison, Busse and Keohane (2007) estimate that the introduction of the Clean Air Act's sulfur dioxide emission trading program enabled railroad companies to practice price discrimination that raised delivered prices for low-sulfur coal by over \$2.00 per short ton for plants with less attractive outside options and by as little as \$0.09 per short ton for plants with better outside options. The confidence interval for our event's effect on the nearest-dated contracts covers a similar range. Because the event's effect might underestimate the full effect of the bill by a factor of 2 to 20, anticipation of the Senate bill probably distorted coal markets more strongly than did railroads' market power in the wake of the Clean Air Act. Further, Cicala (2015) estimates that electricity sector deregulation lowered the price paid by coal-fired generators by \$0.25 per million Btu. The energy content per unit of coal consumption varies by state, but our central estimates translate into an effect of \$0.05 to \$0.10 per million Btu.²⁷ The estimated effect of divestiture is therefore about twice as great as our central estimates for the event's effect. It is just above our event's 90% confidence interval and around the effect obtained from combining the estimate for Friday with the event day's effect. If the weekend event shifted the bill's probability of passage by fewer than 50 percentage points, then the Senate bill's overall effect would have been greater than that of deregulation-induced divestiture.

Another way to consider the size of the change in coal prices is to consider the change in the carbon price that would affect the coal price as much as the event did, assuming the carbon price were fully passed through to coal consumers. The bottom panel of figure 5 shows that the weekend collapse of the Senate bill acted like taxing coal's carbon dioxide (CO₂) emissions by around \$0.60 per metric ton, with the 90% confidence interval for the nearest-dated contracts extending from \$0.05 to \$1.10 per metric ton of CO₂.²⁸ Combining the central coefficient estimate with the probability bounds described above, the anticipation-induced implicit subsidy from the bill's actual passage would have been between \$1.20 and \$12.00 per ton of CO₂.

The projected 2013 allowance prices under the Senate's bill were \$14 to \$25 per ton of CO₂ (Congressional Budget Office, 2010; Energy Information Administration, 2010; Environmental Protection Agency, 2010), and the draft legislation set a hard price floor at \$12 per ton. Anticipating the bill's passage could have implicitly subsidized near-term coal emissions by as much as the bill would have begun implicitly taxing them. As another point of comparison, the

U.S. government's central value for the social cost of carbon in 2010 was \$21 per ton of CO₂, with the recommended range stretching from \$5 to \$65 per ton (Greenstone et al., 2013). The subsidy-like effect of anticipating the passage of some climate legislation at some future time could easily be of comparable magnitude to the optimal carbon tax but acting in the opposite direction. The bill's passage would have amplified the climate externality from coal plants before its implementation eventually began to internalize it.

What do the estimated changes in the coal price mean for electricity generation? Demand for electricity is highly price inelastic. The set of generators that could potentially supply electricity is fixed in the short run, but the set of generators that actually supply electricity depends on the relative costs of the existing generators. Futures prices for Central Appalachian coal affect the cost of coal-fired generation by affecting negotiations for long-term contracts and, because the nearest-month futures contract defines the spot price for coal, by affecting mark-to-market valuations of coal inputs to electricity generation (Energy Information Administration, 2012b). The nearest-dated coal futures contract thus directly affects the variable cost that coal-fired power plants use to determine their hourly bids to supply electricity. A change in the coal price of \$2 to \$27 per short ton would swap the ordering of the supply bids from some natural gas and coal plants. Especially at a time of low gas prices, that magnitude of a reduction in the coal price should increase coal-fired generation at the expense of gas-fired generation, thereby increasing power sector carbon emissions.

Finally, we can use the estimated effect on event-month storage to approximate the event's implications for emissions. Using the specification with month and year fixed effects, the estimated April 2010 increase in coal inventories (see the appendix) translates into 12 million tons of carbon dioxide. This is about 0.1% of average monthly emissions from U.S. coal-fired power plants. As described above, this estimate serves as a lower bound for the full effect of passing this particular piece of cap-and-trade legislation, which in turn serves as a lower bound for the full effect of anticipating some type of climate legislation at some future time. While intertemporal emission leakage does not appear to be a major fraction of coal emissions, it is nonetheless detectable, and it could plausibly be many times greater than estimated here.

VII. Conclusion

The relevance of green paradox effects has been disputed, but I have found evidence of their presence in the market for the largest contributor to climate change. I first showed that anticipated climate policies can reduce futures prices via an extraction cost channel, a storage channel, and a demand-side investment channel. The first two channels generate positive intertemporal emission leakage. The third channel generates negative intertemporal emission leakage. I then empirically identified the anticipation effect in coal futures

²⁷ See table C1 in the U.S. Energy Information Administration's July 2010 "Electric Power Monthly."

²⁸ Calculations use 2.29 metric tons CO₂ per short ton of coal, which follows from the Central Appalachian futures contract's minimum requirement of 12,000 Btu per pound and from the U.S. electric sector's value of 94.31 kg CO₂ per million Btu of coal (see table 1.C.3 in the 2007 "Technical Guidelines: Voluntary Reporting of Greenhouse Gases (1605(b)) Program," from the U.S. Department of Energy).

markets: as predicted by theory, coal futures spiked with the surprising collapse of the U.S. Senate's 2010 climate effort. Consistent with intuition, the storage market appears to have been the primary channel for intertemporal leakage. The proposed policy to restrict carbon emissions seems to have increased carbon emissions prior to the proposal's collapse.

Green paradox effects are a form of intertemporal leakage in which the future policy regime displaces emissions toward earlier periods. Leakage between the present unregulated economy and the future regulated economy is an analog of leakage between regulated economies and unregulated economies via trade. In both cases, this leakage should affect the optimal policy's stringency and broaden its scope. In the trade context, broadening the policy means including more countries; in the anticipation context, broadening the policy means including earlier years. Delaying a policy's implementation might be valuable for giving expectations, institutions, and investment a chance to catch up, but these benefits must be balanced against the costs imposed by higher emissions during the interval of delay. Cost-benefit analyses and policy design discussions have typically ignored these costs, which we have seen acted like subsidizing coal-fired plants' carbon dioxide emissions by between \$1.20 and \$12.00 per ton.

More generally, we have found that markets are distorted by the suggestion of regulation. Major legislative proposals take time to enact. And once enacted, policies like health care reform, financial regulation, and spending cuts then intentionally delay implementation by years. The suggestion of regulation is likely to have particularly acute consequences in energy markets for three reasons. First, the primary actors in these markets tend to have strong intertemporal linkages, and they do not suffer the liquidity constraints that might hamper individuals from smoothing consumption. Second, these markets generate a suite of externalities that make them particularly prone to regulation. Third, demand for regulating these externalities tends to increase over time as populations become richer, scientific knowledge of their effects advances, and damages increase. This is especially true of climate change: major actors plausibly expect to be regulated at some point in the not-too-distant future.

This potential for future climate regulation must have affected energy markets for many years. President Lyndon Johnson warned of fossil fuel-driven climate change in a special message to Congress in 1965, Congress has held hearings since at least 1987, prominent international negotiations began in 1992, and Congress began voting on cap-and-trade legislation in 2001. The last decades' trends in energy prices should be reevaluated for the possibility that prices were systematically lowered by the anticipation of future policy. Not only have energy markets been distorted by failing to price their externalities correctly, but the expectation of future regulations has actually acted like subsidizing their externalities. The last years' policy of not pricing carbon for now may have generated more emissions than simply

not pricing carbon at all. Which policy ultimately generates more emissions depends on what that carbon price looks like, when it occurs, and how policymakers shape expectations beforehand.

REFERENCES

- Alt, Raimund, Ines Fortin, and Simon Weinberger, "The Monday Effect Revisited: An Alternative Testing Approach," *Journal of Empirical Finance* 18 (2011), 447–460.
- Anderson, Michael L., "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association* 103 (2008), 1481–1495.
- Andrews, D. W. K., "End-of-Sample Instability Tests," *Econometrica* 71 (2003), 1661–1694.
- Auerbach, Alan J., "Tax Reform and Adjustment Costs: The Impact on Investment and Market Value," *International Economic Review* 30 (1989), 939–962.
- Binder, John J., "The Event Study Methodology since 1969," *Review of Quantitative Finance and Accounting* 11 (1998), 111–137.
- Bollerslev, Tim, "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics* 31 (1986), 307–327.
- Branson, William, Arminio Fraga, and Robert Johnson, "Expected Fiscal Policy and the Recession of 1982" (pp. 109–128), in M. H. Peston and R. E. Quandt, eds., *Price, Competition and Equilibrium* (Totowa, NJ: Barnes & Noble Books, 1986).
- Brown, Stephen J., and Jerold B. Warner, "Using Daily Stock Returns: The Case of Event Studies," *Journal of Financial Economics* 14 (1985), 3–31.
- Bushnell, James B., Howard Chong, and Erin T. Mansur, "Profiting from Regulation: Evidence from the European Carbon Market," *American Economic Journal: Economic Policy* 5 (2013), 78–106.
- Busse, Meghan R., and Nathaniel O. Keohane, "Market Effects of Environmental Regulation: Coal, Railroads, and the 1990 Clean Air Act," *RAND Journal of Economics* 38 (2007), 1159–1179.
- Cicala, Steve, "When Does Regulation Distort Costs? Lessons from Fuel Procurement in US Electricity Generation," *American Economic Review* 105 (2015), 411–444.
- Congressional Budget Office, "Discussion Draft of the American Power Act," Congressional Budget Office technical report (2010).
- Conley, Timothy G., and Christopher R. Taber, "Inference with 'Difference in Differences' with a Small Number of Policy Changes," this REVIEW 93 (2011), 113–125.
- de Jong, Frank, Angelien Kemna, and Teun Kloek, "A Contribution to Event Study Methodology with an Application to the Dutch Stock Market," *Journal of Banking and Finance* 16 (1992), 11–36.
- 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* 69 (2014), 143–162.
- Energy Information Administration, "Energy Market and Economic Impacts of the American Power Act of 2010," U.S. Energy Information Administration report SR/OIAF/2010-01 (2010).
- "Annual Coal Report 2011," U.S. Energy Information Administration technical report (2012a).
- "Central Appalachian (CAPP) Coal Spot Prices Affect Markets for Coal and Electric Power" (2012b).
- "Electricity Monthly Update: With Data for September 2013," U.S. Energy Information Administration technical report (2013).
- Environmental Protection Agency, "EPA Analysis of the American Power Act of 2010," Office of Atmospheric Programs, U.S. Environmental Protection Agency technical report (2010).
- Fama, Eugene F., "The Behavior of Stock-Market Prices," *Journal of Business* 38 (1965), 34–105.
- Fountain, John, and Glenn W. Harrison, "What Do Prediction Markets Predict?" *Applied Economics Letters* 18 (2011), 267–272.
- French, Kenneth R., "Stock Returns and the Weekend Effect," *Journal of Financial Economics* 8 (1980), 55–69.
- Gelbach, Jonah B., Eric Helland, and Jonathan Klick, "Valid Inference in Single-Firm, Single-Event Studies," *American Law and Economics Review* 15 (2013), 495–541.

- Geman, Hélyette, *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy* (Chichester, West Sussex: Wiley, 2005).
- Gerlagh, Reyer, "Too Much Oil," *CESifo Economic Studies* 57 (2011), 79–102.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton, "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation," *Review of Environmental Economics and Policy* 7 (2013), 23–46.
- Hakala, Scott D., "Lessons from Single-Company Event Studies: The Importance of Controlling for Company-Specific Events," unpublished manuscript (2010), https://works.bepress.com/scott_hakala/4/.
- Hall, R. E., "The Dynamic Effects of Fiscal Policy in an Economy with Foresight," *Review of Economic Studies* 38 (1971), 229–244.
- Heal, Geoffrey, "The Relationship between Price and Extraction Cost for a Resource with a Backstop Technology," *Bell Journal of Economics* 7 (1976), 371–378.
- Hoel, Michael, "Carbon Taxes and the Green Paradox" (pp. 203–224), in Robert W. Hahn and Alistair Ulph, eds., *Climate Change and Common Sense: Essays in Honour of Tom Schelling* (New York: Oxford University Press, 2012).
- Jones, Benjamin, Michael Keen, and Jon Strand, "Fiscal Implications of Climate Change," *International Tax and Public Finance* 20 (2013), 29–70.
- Judd, Kenneth L., "Short-Run Analysis of Fiscal Policy in a Simple Perfect Foresight Model," *Journal of Political Economy* 93 (1985), 298–319.
- , "The Welfare Cost of Factor Taxation in a Perfect-Foresight Model," *Journal of Political Economy* 95 (1987), 675–709.
- Lamdin, Douglas J., "Implementing and Interpreting Event Studies of Regulatory Changes," *Journal of Economics and Business* 53 (2001), 171–183.
- Lange, Ian, and Joshua Linn, "Bush v. Gore and the Effect of New Source Review on Power Plant Emissions," *Environmental and Resource Economics* 40 (2008), 571–591.
- Linn, Joshua, "The Effect of Cap-and-Trade Programs on Firms' Profits: Evidence from the Nitrogen Oxides Budget Trading Program," *Journal of Environmental Economics and Management* 59 (2010), 1–14.
- Lizza, Ryan, "As the World Burns," *New Yorker*, October 11, 2010.
- MacKinlay, A. Craig, "Event Studies in Economics and Finance," *Journal of Economic Literature* 35 (1997), 13–39.
- Manski, Charles F., "Interpreting the Predictions of Prediction Markets," *Economics Letters* 91 (2006), 425–429.
- McKenzie, Andrew M., Michael R. Thomsen, and Bruce L. Dixon, "The Performance of Event Study Approaches Using Daily Commodity Futures Returns," *Journal of Futures Markets* 24 (2004), 533–555.
- Meng, Kyle C., "Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy," *American Economic Review* 107:3 (2017), 748–784. doi:10.1257/aer.20150781.
- Mertens, Karel, and Morten O. Ravn, "Understanding the Aggregate Effects of Anticipated and Unanticipated Tax Policy Shocks," *Review of Economic Dynamics* 14:1 (2011), 27–54.
- , "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks," *American Economic Journal: Economic Policy* 4 (2012), 145–181.
- Perotti, Roberto, "The Effects of Tax Shocks on Output: Not So Large, But Not Small Either," *American Economic Journal: Economic Policy* 4 (2012), 214–237.
- Sinn, Hans-Werner, "Public Policies against Global Warming: A Supply Side Approach," *International Tax and Public Finance* 15 (2008), 360–394.
- van der Werf, Edwin, and Corrado Di Maria, "Imperfect Environmental Policy and Polluting Emissions: The Green Paradox and Beyond," *International Review of Environmental and Resource Economics* 6 (2012), 153–194.
- Yang, Shu-Chun Susan, "Quantifying Tax Effects under Policy Foresight," *Journal of Monetary Economics* 52 (2005), 1557–1568.