

Green Expectations: Current Effects of Anticipated Carbon Pricing*

Derek Lemoine

Department of Economics, University of Arizona
McClelland Hall 401, 1130 E Helen St, Tucson, AZ, 85721-0108, USA
dlemoine@email.arizona.edu

University of Arizona Working Paper 13-06

April 9, 2013

Markets respond not just to current emission policies but also to expectations of future policies. We theoretically demonstrate a “green paradox” in a multiple-commodity market and empirically identify it in the U.S. energy market. We show that a future carbon pricing policy has an anticipation effect which should decrease the near-term price of coal. If the future policy also decreases the near-term price of natural gas, then the policy increases emissions prior to its implementation while nonetheless decreasing cumulative emissions. An event study confirms the theoretical prediction: the price of coal increased by \$1.30 per short ton upon the unexpected collapse of the U.S. Senate’s 2010 climate effort. The response of natural gas futures suggests that the legislative process was increasing U.S. emissions prior to 2013, but if the bill had passed, it would have eventually decreased cumulative emissions.

JEL: G13, H23, H30, Q41, Q58

*I am grateful for conversations with Keisuke Hirano. 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.

In March of 2009, U.S. Representatives Henry Waxman and Edward Markey unveiled draft climate legislation. Their bill is the closest the U.S. has come to a formal federal price on carbon: the House of Representatives passed their American Clean Energy and Security Act three months later, and the Senate spent the next year working on a close cousin. A carbon cap was to begin in 2013, four years after the unveiling and twelve years after the Senate began drafting climate legislation. Cap-and-trade programs in Europe, California, and elsewhere have experienced comparable delays in implementation. What happens to energy markets in the interim? Does energy producers' anticipation of a carbon price undercut the policy's intended emission reductions?

The "green paradox" literature suggests that anticipating a future climate policy increases emissions today: the future emission price reduces the return to future extraction of energy resources, which reduces today's incentive to delay extraction. The extraction profile thereby tilts further towards the present. This intertemporal leakage undercuts the benefits of a future climate policy and argues for implementing the policy sooner rather than later. However, the existence and magnitude of green paradox effects are both open questions. To this end, van der Werf and Di Maria (2012) conclude their review by noting 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."

We use the collapse of the U.S. Senate's 2010 climate effort to establish the existence of a green paradox in U.S. energy markets and to estimate a lower bound for its magnitude. We find that anticipating the proposed carbon policy acted like subsidizing coal's current emissions by at least \$0.60 per ton of carbon dioxide. This subsidy would have grown as the bill became more likely to pass, suggesting that the overall effect of the bill could have been several times as large. Not pricing carbon emissions "for now" distorts current coal markets more than not pricing carbon at all. The identified green paradox distortion is of the same magnitude as distortions from railroads' market power (Busse and Keohane, 2007) and from utility regulation (Cicala, 2012). Green paradox effects are economically significant, and they may have been distorting energy markets for years.

But are green paradox effects climatically significant? Even if an anticipated carbon price increases emissions today, it should still decrease emissions in the future and might decrease emissions overall. As long as the additional near-term emissions do not tip us over a threshold, we would then still come out ahead relative to never having adopted the policy. Here the green paradox literature diverges. In Hotelling (1931) models of physical exhaustion, producers have a finite endowment of energy resources to allocate over time (e.g., Sinn, 2008, 2012; Di Maria et al., 2012). In these settings, an anticipated policy increases current emissions but has no effect on cumulative emissions, which are completely determined by the physical endowment. By tilting the emission profile towards the present, the future policy tends to increase the present value of climate damages. However, physical exhaustion is not a prime

concern for coal or natural gas, the two commodities most closely linked to climate policy.¹ In contrast, Heal (1976) models endogenize total resource use by making marginal cost increase in cumulative extraction (e.g., Gerlagh, 2011; Hoel, 2012).² Commodity producers no longer decide merely how to allocate a given resource over time but also decide how much to extract in total. The future carbon price reduces planned future extraction, which also reduces the incentive to conserve low-cost resources for the future. The anticipated carbon policy now decreases cumulative emissions because it decreases cumulative extraction.

Previous models with endogenous cumulative extraction consider only a single polluting commodity. Yet the U.S. Senate’s bill would have primarily regulated electricity markets, in which a carbon price taxes both coal and natural gas. We therefore extend the green paradox literature to consider the case of multiple commodities with endogenous cumulative extraction. The second commodity complicates the standard green paradox story by introducing substitution effects. The future policy taxes both commodities, but it taxes the higher-emission commodity (coal) more than the lower-emission commodity (natural gas). The future policy could therefore increase future natural gas consumption. As in the one-commodity setting, current extraction of the higher-emission resource generally increases in anticipation of a future emission price. But now this increase reduces demand for the lower-emission resource in the pre-tax period. This pre-tax substitution effect and anticipation of the post-tax substitution effect both decrease pre-tax extraction of the lower-emission resource. We demonstrate that the possible outcomes obey a clear ordering: as the lower-emission resource becomes relatively cleaner, we switch from a world in which it follows the standard green paradox story to a world in which its pre-tax extraction begins decreasing in the anticipated tax, to a world in which its post-tax extraction begins increasing in the anticipated tax, to a world in which its cumulative extraction even begins increasing in the anticipated tax. Crucially, while the anticipated emission price now has an ambiguous effect on emissions, we show that a necessary condition for the policy to either decrease pre-tax emissions or increase cumulative emissions is that it decreases pre-tax consumption of the lower-emission resource.

¹In their review of the fiscal implications of climate change, Jones et al. (2013) question the relevance of the green paradox on the basis of the common assumption of exhaustibility: “Perhaps most fundamentally, however, whether fossil fuels are best modeled as exhaustible is questionable: empirically, the evolution of resource prices is not well-described by simple Hotelling-type models; and stocks—especially of coal—are so large that the relevance of exhaustion is moot.” According to the June 2012 BP Statistical Review of World Energy, global proved coal reserves would last 112 years at current production rates, in contrast to 54 years for oil. Coal is the most carbon-intensive commodity, and its primary substitute is natural gas. Global proved reserves of natural gas would last 64 years at current production rates. The advent of shale gas will substantially increase these estimates of “proved” reserves.

²Smulders et al. (2012) also generate a green paradox in a setting without exhaustibility. Their intertemporal linkages arise not from resource owners but from households’ consumption-savings decisions in a general equilibrium setting. When households anticipate a future carbon price, they increase their savings in order to smooth the shock of its implementation. These savings increase the capital stock in the pre-tax period, and because energy and capital are complements, these savings also increase pollution in the pre-tax period.

We use this necessary condition to learn about the Senate bill's effect on emissions. The estimated response of natural gas futures to the climate effort's collapse is not consistent with the anticipated policy reducing pre-tax emissions or increasing cumulative emissions. Therefore, while the proposed cap-and-trade bill was increasing emissions prior to its implementation in 2013, it would have decreased emissions after 2013 and in total. While not discussing a climate bill at all would have generated fewer contemporary emissions, the climate would nonetheless have eventually benefited from the bill's passage. However, the bill did not pass. The additional emissions from its discussion in 2009–2010 were therefore never offset by a carbon price. While the legislation would have decreased cumulative emissions, it is likely that the unsuccessful legislative process actually increased cumulative emissions.

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 et al., 1986; Judd, 1985; Auerbach, 1989; Yang, 2005; Mertens and 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 tax: the future change in the tax causes a jolt to real flows which an anticipation effect offsets by shifting earlier real flows in the opposite direction.³ The empirical literature on tax anticipation has obtained mixed results that depend on how one constructs the timing of tax changes (Mertens and Ravn, 2012; Perotti, 2012). Identifying anticipation effects in commodity markets is more straightforward for two reasons: these markets probably lack liquidity constraints that can mask the effect, and proposals for emission pricing are not themselves responding to commodity market movements as, for instance, personal taxation responds to aggregate output or consumption. Our results indicate that hypothesized anticipation effects need not make overly strong demands of market efficiency.

Our identification strategy uses an event study in futures markets. Most event studies examine stock market returns. In energy policy, event studies have used changes in equity prices to learn about the cost of regulation and about the winners and losers from specific permit allocations (Lange and Linn, 2008; Linn, 2010; Bushnell et al., 2013). Using futures markets allows us to learn about the consequences for real variables such as consumption and emissions. They also tell us about the shadow price of emissions. All of these consequences are largely independent of the permit allocation scheme.

The main challenge in using regulatory changes as events is that these changes are often anticipated (Binder, 1998; Lamdin, 2001). A regulatory announcement or a policy's passage

³In models of rational addiction, present consumption increases the marginal utility from future consumption. Anticipated taxes therefore reduce present and future consumption. These effects have been identified in, for instance, cigarette markets (e.g., Becker et al., 1994; Gruber and Köszegi, 2001). In our setting, present production increases the marginal cost of future production. Anticipated taxes therefore increase present production and decrease future production, though we show that substitution between commodities can reverse these effects.

should only affect markets insofar as it represents new news, but information often leaks ahead of time. Further, the results of many votes and rulemakings are clearly anticipated well beforehand. Meng (2012) considers the same regulatory setting as this paper, using 2009–2010 prediction market contracts to estimate the effect of proposed cap-and-trade regulations on firm profits. We avoid using prediction markets to identify shifting regulatory probabilities because these markets are highly illiquid and because they are vulnerable to endogeneity concerns if participants look to other markets’ data when placing bets. In contrast, we use an event that occurred among the parties developing legislation to isolate an exogenous shift in the probability of regulation. Narrative evidence, contemporary news accounts, prediction market trades, and Internet search patterns support a shift in regulatory expectations when Senator Graham walked out of his climate collaboration over the weekend before its scheduled unveiling.

We begin by analyzing how commodity spot prices, emissions, and futures markets respond to new information about future emission policies. We focus on the interplay between higher- and lower-emission commodities and on theoretical predictions relevant to proposed climate legislation. Section 2 then combines several lines of evidence to argue that Senator Graham’s weekend withdrawal from his climate bill represented “new” news that altered expectations. Section 3 introduces the estimation framework by which we identify the event’s effect on futures markets. Section 4 demonstrates the anticipation effect in coal futures across a range of specifications. Section 5 discusses the potential for confounding events. It then compares the magnitude of the effect on coal markets to recent studies of market power and deregulation and to recent estimates of the social cost of carbon. It also uses the effect on natural gas markets to learn whether a green paradox in fact increased contemporary and/or cumulative emissions. Section 6 concludes with policy implications. The appendix contains proofs.

1 The current effects of anticipated taxation

We develop a two-period model with two commodities, indexed by H and L . A representative consumer obtains utility $U(q_t^H, q_t^L)$ from consuming quantities q_t^H and q_t^L at time t , where utility is quadratic, increasing, and strictly concave. The two commodities are at least partially substitutable: $U_{HL} < 0$, where subscripts indicate partial derivatives. Consuming commodity H generates emissions at rate e^H , and consuming commodity L generates emissions at rate e^L , with $e^H > e^L > 0$. A two-commodity market captures the main features of our empirical setting: the two major sources of dispatchable generation in U.S. electricity markets are coal (the higher-emission product) and natural gas (the lower-emission product).

The emissions are externalities, with the usual first-best policy pricing them in each period at their marginal damage. We consider the case where emissions are not priced in the first period, but the regulator imposes a linear emission tax τ in the second period. There are

two equivalent ways of interpreting this setting: the government does not succeed in adopting the pollution policy until the second period, or the government adopts the policy in the first period but delays its implementation to allow regulators and firms time to prepare. Firms anticipate the second-period tax when selecting their first-period production schedules.

Firms are identical price-takers, allowing us to analyze them via a representative firm. Each period's production cost $C(Q_t^H, Q_t^L)$ is a quadratic, increasing, and strictly convex function of cumulative production Q . It is additively separable in the type of product: $C(\cdot) \equiv C^H(Q^H) + C^L(Q^L)$. Neither coal nor natural gas markets face a clear threat of exhaustion, but both have faced the cost of moving towards less accessible reserves: coal mining has progressed from surface seams to mountaintop removal, while gas companies developed techniques to access ever deeper reservoirs and even "unconventional" reserves. The representative firm maximizes the present value of profits, with discount factor $\beta \in (0, 1]$:

$$\max_{\{q_1^H, q_1^L, q_2^H, q_2^L\}} \sum_{k \in \{H, L\}} \{p_1^k q_1^k - C^k(q_1^k) + \beta [p_2^k q_2^k - C^k(q_1^k + q_2^k) - \tau e^k q_2^k]\}, \quad (1)$$

where subscripts index time and superscripts index commodities.

The firm's optimal quantities solve the following four first-order conditions:

$$p_1^k = C_k(q_1^k) + \beta C_k(q_1^k + q_2^k), \quad \text{for } k \in \{H, L\}, \quad (2)$$

$$p_2^k = C_k(q_1^k + q_2^k) + \tau e^k, \quad \text{for } k \in \{H, L\}, \quad (3)$$

where subscripts on the cost function indicate partial derivatives. The firm equates marginal revenue to a comprehensive measure of marginal cost. In both periods, marginal cost includes the marginal extraction cost. In the first period, it also includes the (discounted) additional cost imposed by extracting from more costly reserves in the second period. In the second period, it also includes emission tax payments. Quantity decisions are linked over time by the dependence of marginal extraction cost on cumulative extraction.

Substituting among the four equations and recognizing that equilibrium prices are equal to marginal utility, we obtain:

$$\begin{aligned} U_H(q_1^H, q_1^L) + \beta \tau e^H &= C_H(q_1^H) + \beta U_H(q_2^H, q_2^L), \text{ and} \\ U_L(q_1^H, q_1^L) + \beta \tau e^L &= C_L(q_1^L) + \beta U_L(q_2^H, q_2^L). \end{aligned}$$

These equations characterize the firm's intertemporal decisions. The left-hand side measures the marginal benefit of extracting in period 1 instead of period 2: it includes the revenue from selling the commodity in period 1 and the gain from not paying the tax in period 2. The right-hand side measures the marginal cost of extracting in period 1 instead of period 2: it includes the immediate cost of extraction and the revenue forgone from not selling in period 2. At an optimum, the firm is indifferent between extracting in these two periods. The anticipated tax provides an incentive to extract more of the resource earlier, and that incentive is stronger for the more emission-intensive commodity.

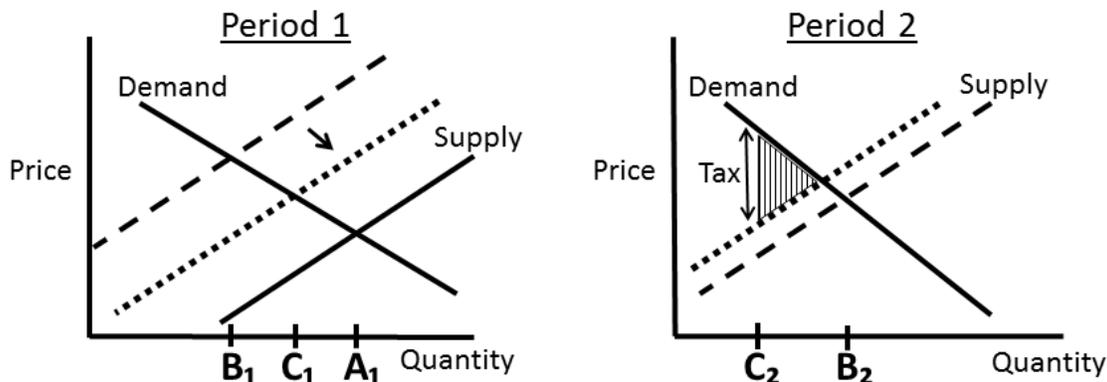


Figure 1: When costs depend on cumulative extraction, first-period marginal cost (dashed line) includes the effect on second-period costs, which reduces equilibrium quantity to B_1 from the myopic equilibrium A_1 . In a one-commodity model, imposing a second-period tax reduces the second-period market-clearing quantity from B_2 to C_2 (direct effect) and increases the first-period quantity from B_1 to C_1 (anticipation effect).

In a one-commodity system, introducing intertemporal linkages has two effects. The solid lines in the left plot of Figure 1 show demand and supply in a model in which second-period extraction costs are independent of first-period extraction or, equivalently, in which the firm is myopic ($\beta = 0$). The firm's profit-maximizing quantity (point A_1) occurs at the intersection of these curves. When extraction costs depend on cumulative extraction and the firm is forward-looking, the first period's marginal cost curve is higher (dashed line) because additional extraction increases costs in the second period. These intertemporal linkages therefore reduce first-period extraction (to point B_1).

1.1 The tax's effect on market-clearing quantities

Now consider the effect of introducing an emission tax in the second period. The right plot in Figure 1 shows second-period supply and demand. Introducing the tax reduces the second period's equilibrium quantity from B_2 . This is the tax's *direct effect*. Because the last unit of second-period consumption now comes from cheaper deposits, additional first-period extraction has a smaller effect on second-period costs. The first-period marginal cost curve therefore shifts down (dotted line) part of the way towards the myopic curve. Raising the second-period tax thereby increases first-period extraction from point B_1 to C_1 , but extraction is still lower than it would have been in a world without a second period. Greater first-period extraction shifts the cost curve up in the second period (dotted line), which further reduces the equilibrium quantity all the way to C_2 . The increase in first-period extraction due to the second-period tax is an *anticipation effect*.

Introducing a second, partially substitutable commodity complicates the analysis by creating substitution effects. The first proposition establishes how a second-period tax affects market-clearing quantities, indicated with an asterisk.

Proposition 1 (Comparative Statics of Equilibrium Quantities). *Consider a marginal increase in the second-period emission tax. There exist k_1^i , k_2^i , and k_3^i such that $k_3^i > k_2^i > k_1^i > 0$ and:*

(i) $\partial q_1^{i*} / \partial \tau > 0$ if and only if $e^j / e^i < k_1^i$,

(ii) $\partial q_2^{i*} / \partial \tau < 0$ if and only if $e^j / e^i < k_2^i$, and

(iii) $\partial [q_1^{i*} + q_2^{i*}] / \partial \tau < 0$ if and only if $e^j / e^i < k_3^i$.

Holding the other parameters fixed at some level, there exists $x < 0$ such that $k_1^i > 1$ for all $U_{ij} > x$. Holding the other parameters fixed at some level, there exists $x < 0$ such that $k_1^i < 1$ for all $U_{ii} < x$, and there exist $y < 0$ and $z > 0$ such that $k_1^i < 1$ for all pairs (U_{jj}, C_{jj}) with $U_{jj} > y$, $C_{jj} < z$, and $U_{ij} < U_{jj}$.

Proof. See appendix. □

In the two-commodity setting, an emission tax changes the relative prices of the two commodities. The resulting *substitution effect* pulls second-period consumption away from the high-emission product and towards the low-emission product. The second-period tax affects each commodity's first-period consumption via its direct effect on that commodity in the second period and also through its effects on the other commodity in both periods. The strength of substitution effects depends on the products' relative emission intensities.

Figure 2 separates the space of emission factors into regions that describe how quantity i responds to the tax. It plots that product's own emission intensity e^i as the abscissa and the other product's emission intensity e^j as the ordinate. Above the diagonal, product i is the low-emission product, and below the diagonal, product i is the high-emission product. Proposition 1 says that the possible marginal effects are divided by rays from the origin. Region A includes all cases in which product i is the low-emission product and substitution effects are weak. It also includes all cases in which product i is the high-emission product and particular combinations of emission intensities and elasticities do not hold. In this region, the intuition from the single-product setting carries into the two-product setting: the first-period quantity of product i increases in the tax, whereas the second-period quantity and cumulative quantity both decrease in the tax.

Now consider the other regions above the diagonal. For a relative emission intensity e^j / e^i above k_1^i (Regions B, C, and D), substitution effects are strong enough to overturn the anticipation effect for the low-emission product. The quantity of low-emission product decreases in the first period because the increase in high-emission product sufficiently reduces the value of additional low-emission product, while substitution towards the low-emission product in

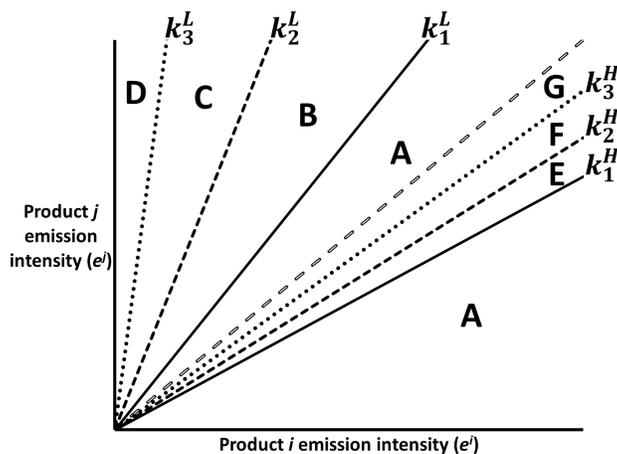


Figure 2: The diagonal divides areas where product i is the high-emission product (lower) from areas where product i is the low-emission product (higher). In region A, raising the tax increases the first-period quantity of product i and decreases both the second-period and cumulative quantity of product i . In all other regions, product i 's first-period quantity decreases in the tax. In regions C, D, F, and G, product i 's second-period quantity increases in the tax. In regions D and G, the cumulative quantity of product i increases in the tax.

the second period also reduces the anticipation effect's pull on first-period consumption. For e^j/e^i above k_2^L (Regions C and D), substitution effects are now strong enough to overturn the second period's direct effect for the low-emission product. The tax raises the marginal cost of the low-emission product, but its second-period consumption increases via substitution for the more heavily taxed high-emission product. Now the anticipation effect decreases first-period consumption of the low-emission product, which works in the same direction as the first-period substitution effect.⁴ For e^j/e^i above k_3^L (Region D), the substitution effect is so strong that the second-period increase in product i overwhelms its first-period decrease. Cumulative extraction of the low-emission product increases in the tax.⁵

Next consider the space below the diagonal. Region A fills this whole area unless demand and supply for product H are sufficiently inelastic or demand for product L is sufficiently elastic. When these special cases combine with high values of e^L/e^H , the effects on the first-period (regions E, F, and G), second-period (regions E and F), and cumulative quantities (region G) of product H can be reversed as already described for product L . Here the intuition is that the quantity of product L responds strongly to small changes in price while product H does not respond strongly. When the two products' emission intensities are similar, the per-unit taxes on each product are similar and product L responds with a strong

⁴An increase in second-period consumption is therefore never compatible with an increase in first-period consumption, as indicated by $k_2^L > k_1^L$.

⁵If product i does not generate emissions ($e^i = 0$), then we are always in region D.

decrease. This change increases the value of additional product H , and that higher value (induced by $U_{HL} < 0$) can outweigh the direct cost of the tax. The difference between these special cases and those described for product L is that there are always emission intensities which can reverse the standard (single-product setting) effect on product L because each unit of product L is taxed less than each unit of product H . However, because the tax differential tends to make substitution effects reinforce the standard (single-product setting) effect on product H , only under special combinations of elasticities do substitution effects work to overcome the tax's greater effect on the cost of H .

Finally, these rays are drawn for fixed U_{HL} . As U_{HL} rises to 0, the substitution effect weakens and region A comes to fill the whole positive quadrant. In contrast, increasing the magnitude of U_{HL} (i.e., increasing the effectiveness of each product in substituting for the other) increases the strength of the substitution effect and decreases the slope of each ray, which shrinks region A.⁶

Di Maria et al. (2012) also consider a case with two resources, but do so in a setting of physical exhaustion. They describe two effects from announcing a future emission constraint. First, an abundance effect occurs when the emission constraint requires a reduction in future resource use. Yet resource owners want to exhaust their stocks. They therefore shift more extraction into the time before the emission constraint so that they will still ultimately achieve physical exhaustion. Second, an ordering effect describes how extraction of low-emission and high-emission resources is reorganized around the policy's introduction. The low-emission resource becomes relatively more valuable once the emission constraint binds, which leads resource owners to conserve it by tilting earlier extraction towards the high-emission resource. This effect increases the emission intensity of consumption in the time leading up to the announced emission constraint. Because we allow stock-dependent extraction costs, our results are not driven by the desire to extract an exogenous endowment of each resource. We do find that first-period extraction of the high-emission commodity usually increases in the emission tax (as with the abundance effect), but in our setting its cumulative extraction usually decreases. We also describe cases in which extraction of the low-emission resource declines in the first period (similar to the ordering effect), but in our setting this change need not be fully offset by increased extraction in the second period and is even compatible with declining second-period extraction. In addition, we describe cases in which second-period extraction of the high-emission product actually increases in the tax while its first-period extraction actually decreases in the tax.

1.2 The tax's effect on emissions

Emission reductions are likely to be the primary motivation for an emission tax. The next proposition describes how time t aggregate emissions E_t change with the tax.

⁶If U_{LL} is sufficiently large in magnitude or U_{HH} and C_{HH} are sufficiently small in magnitude, then k_1^L is below the diagonal and there is no region A above the diagonal.

Proposition 2 (Comparative Statics of Equilibrium Emissions). *If U_{LL} and C_{LL} are of sufficiently small magnitude relative to U_{HH} and C_{HH} , then there exist $h_1, h_2, h_3, h_4, h_5,$ and h_6 such that $h_i < h_j$ if and only if $i < j$, $h_i > k_i^L$ for $i \in \{1, 2, 3\}$, and:*

(i) $\partial E_1 / \partial \tau < 0$ if and only if $e^H / e^L \in (h_1, h_6)$,

(ii) $\partial E_2 / \partial \tau > 0$ if and only if $e^H / e^L \in (h_2, h_5)$, and

(iii) $\partial[E_1 + E_2] / \partial \tau > 0$ if and only if $e^H / e^L \in (h_3, h_4)$.

If U_{LL} and C_{LL} are of sufficiently large magnitude relative to U_{HH} and C_{HH} , then the signs are as if $e^H / e^L < h_1$.

Proof. See appendix. □

When both products respond in the same direction, aggregate emissions clearly respond in the same direction. In particular, if the anticipated tax increases the first-period quantity of each product, then it increases first-period emissions, decreases second-period emissions, and decreases cumulative emissions. When the low-emission product is relatively responsive while the second-period product is relatively unresponsive, then the low-emission product's response can reverse these emission outcomes, but only if its emission factor has an intermediate value. If the low-emission product is too clean, then its consumption cannot drive aggregate emissions, and if it is too dirty, then substitution effects will not be strong enough to overwhelm the high-emission product's change in emissions.

Figure 3 separates the space of emission factors into regions that describe how aggregate emissions respond to the second-period tax. As in the analysis of quantity responses, all possible effects are divided by rays from the origin. In the unshaded region, raising the second-period tax increases first-period emissions and decreases both second-period and cumulative emissions. If the utility and cost functions are not much more strongly curved in the direction of the high-emission product, then the entire space is filled by this region. When the utility and cost functions do curve more strongly towards the high-emission product, then this region is limited to the space where the low-emission product is very clean or is almost as dirty as the high-emission product.

For relative emission intensities e^j / e^i between the outermost rays (the shaded regions), raising the tax decreases first-period emissions. In fact, in the lightly shaded region, raising the tax decreases emissions in both periods. In the two darkest regions, raising the tax actually increases second-period emissions because of the strong second-period substitution effect. This increase in emissions could not happen in a single-period model. However, in a two-period model with intertemporally linked production costs, greater second-period tax payments can be offset by changes in first-period production patterns. Finally, in the darkest (middle) region, we have the most perverse outcome: increasing the anticipated tax actually increases cumulative emissions. Here the increase in second-period emissions dominates the

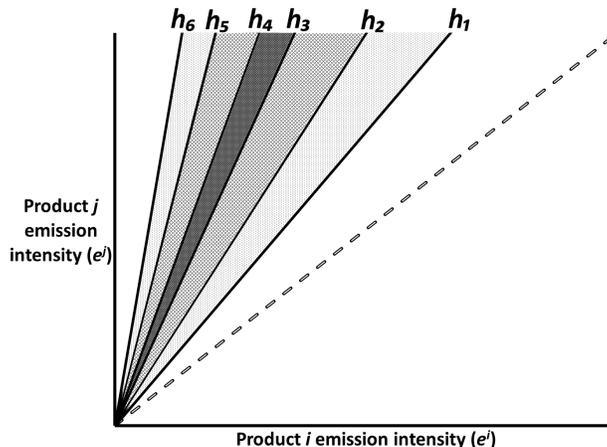
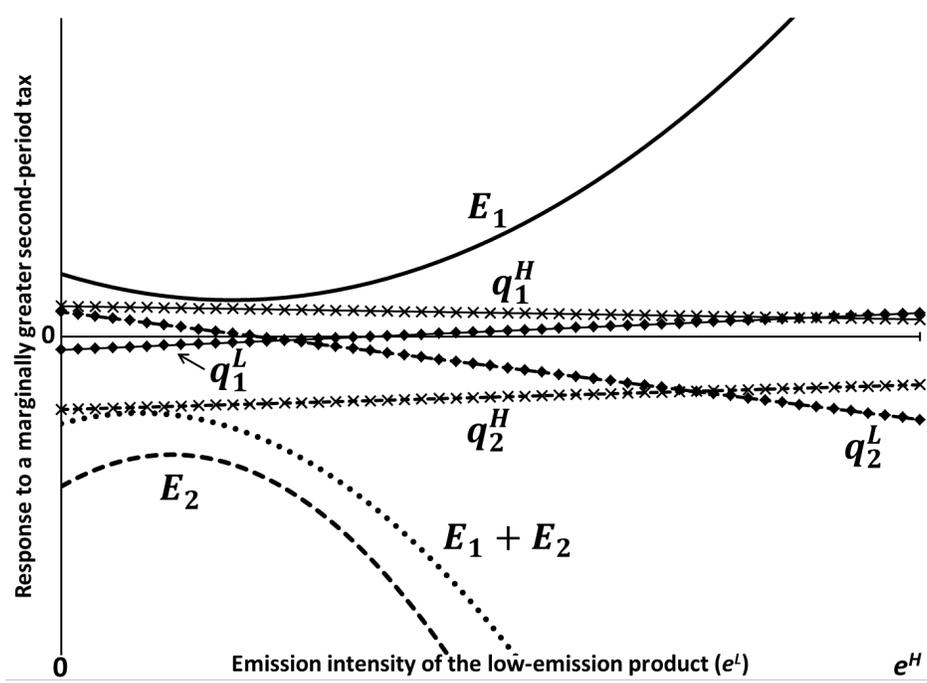


Figure 3: In the unshaded region, first-period emissions increase in the second-period tax while second-period emissions and cumulative emissions both decrease in the tax. In the shaded regions, first-period emissions decrease in the tax. In the two darkest regions, second-period emissions increase in the tax. In the darkest (middle) region, cumulative emissions increase in the tax.

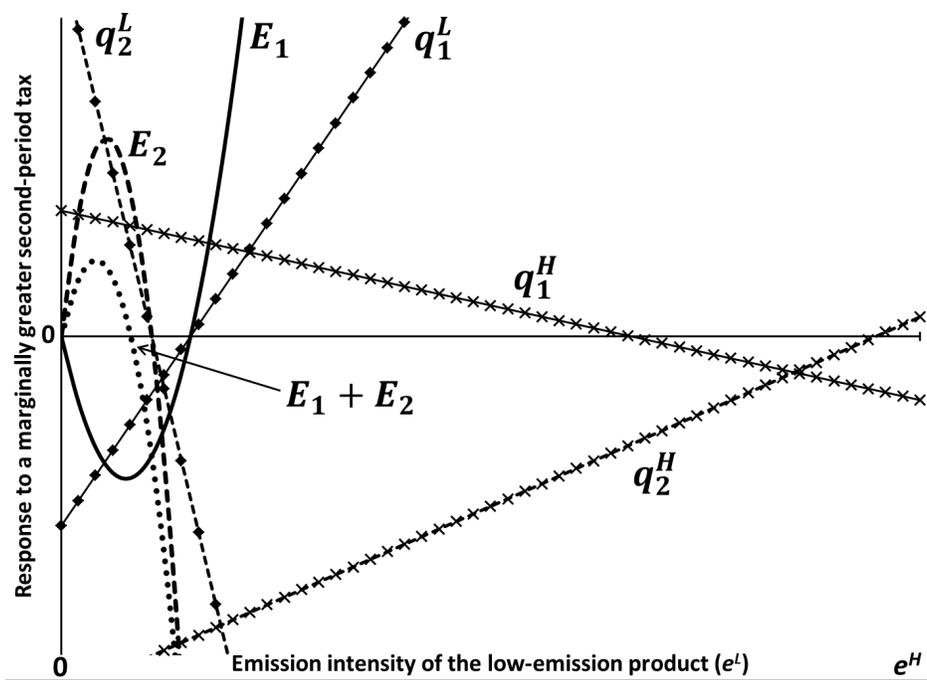
decrease in first-period emissions. Importantly, though, this darkest region is the smallest of all, and it can disappear completely depending on the curvature of the utility and cost functions. In the vast majority of potential parameterizations, raising the tax increases emissions in some period but decreases cumulative emissions.

To provide further intuition, Figure 4 plots how emissions and quantities respond to a marginally greater second-period tax. The left-hand plot is a case in which $U_{HH} = U_{LL}$ and $C_{HH} = C_{LL}$, and the right-hand plot is a case in which $U_{HH} \ll U_{LL}$ and $C_{HH} \gg C_{LL}$. In each case we see that as we make the low-emission product relatively cleaner (moving to the left along the abscissa), its first-period quantity response is the first to change signs, its second-period quantity response is the next to change signs, and its cumulative quantity response is the last to change signs. In the right-hand plot, we see that when the low-emission product is sufficiently dirty, then the high-emission product's quantity responses can be the ones to change signs relative to the single-product setting. This is the case in which there are rays k_i^H below the diagonal in Figure 2.

Whereas the quantity responses are linear functions of the relative emission intensities, the emission responses are quadratic functions. When the low-emission product is very clean, aggregate emissions are determined by the high-emission product, and when the low-emission product is nearly as dirty as the high-emission product, the change in aggregate emissions is nearly proportional to the change in aggregate quantity. In the left-hand plot, there are no combinations of emission intensities that reverse the effects seen in the single-product setting. However, in the right-hand plot, these combinations of elasticities are such that the



(a) $U_{HH} = U_{LL}$ and $C_{HH} = C_{LL}$



(b) $U_{HH} \ll U_{LL}$ and $C_{HH} \gg C_{LL}$

Figure 4: The effect of a marginal increase in the second-period tax on equilibrium quantities and emissions, as a function of the low-emission product's relative emission intensity.

quadratics have real roots. The larger root occurs just to the left of where the corresponding quantity of low-emission product goes to 0. The smaller root occurs just to the right of where the low-emission product is perfectly clean. First-period aggregate emissions demonstrate “non-standard” effects across the broadest range of e^L (as illustrated by the full set of shaded regions in Figure 3); a greater emission tax increases cumulative emissions only over a smaller set of e^L (as illustrated by the darkest region in Figure 3).

How do these emission outcomes translate into a green paradox? Gerlagh (2011) distinguishes a weak and a strong form of the paradox. In our setting, the weak form holds when imposing an emission tax in the second period increases emissions in the first period, and the strong form holds when imposing an emission tax in the second period increases the net present value of pollution damages, evaluated with per-period discount rate r .

Corollary 1 (Green Paradox). *A weak green paradox occurs if and only if $e^H/e^L \notin [h_1, h_6]$. Assume time t damages are a linear, increasing function of time t emissions. For $e^H/e^L \in [h_1, h_3] \cup [h_4, h_6]$, a strong green paradox never occurs for $r \geq 0$. For $e^H/e^L \in (h_3, h_4)$, there exists $\gamma > 0$ such that a strong green paradox occurs if and only if $r < \gamma$, and for $e^H/e^L \notin [h_1, h_6]$, there exists $\lambda > 0$ such that a strong green paradox occurs if and only if $r > \lambda$.*

Proof. See appendix. □

The review by van der Werf and Di Maria (2012) shows that while Hotelling (1931) models of physical resource exhaustion typically demonstrate a weak green paradox, it is a more fragile result in Heal (1976) models where stock-dependent extraction costs make cumulative extraction endogenous. Our setting also endogenizes cumulative extraction and has a weak green paradox usually emerge even when there are multiple types of polluting commodities. However, there are special cases in which strong substitution effects prevent a weak green paradox. There are also special cases in which raising the tax decreases emissions in both periods. In these cases, a strong green paradox never occurs. In general, though, first-period emissions and second-period emissions change in opposite ways, with the change in second-period emissions dominating the change in cumulative emissions. When first-period emissions decrease, a strong green paradox requires a sufficiently low discount rate for the second-period increase to dominate the present damage calculation. In the more common case in which first-period emissions increase, a strong green paradox requires a sufficiently large discount rate for the near-term change in emissions to dominate present damages.

1.3 Implications for coal and natural gas futures markets

In our empirical application, the high-emission commodity is coal and the low-emission commodity is natural gas. Both are inputs to electricity generation and have some degree of substitutability in this application. Coal is generally thought to have a relatively flat

production cost function, while natural gas probably has a more curved cost function due to the smaller size of each reserve. If both products were only used for electricity generation, they might have a similar price elasticity of demand, but natural gas also has major markets in relatively price-inelastic residential and industrial applications. In sum, the cost and utility functions probably curve less in the direction of coal than in the direction of natural gas.

Under this combination of curvatures, proposals to implement a carbon price at some future date should increase near-term coal consumption, which would manifest as a decrease in near-term coal spot prices. The effect on near-term natural gas consumption and spot prices is complicated by the presence of substitution effects that mitigate the quantity increase and potentially even lead to a net decrease in its quantity. For low degrees of substitutability, near-term natural gas consumption should increase and spot prices should fall, but for high degrees of substitutability, near-term natural gas consumption could decrease and spot prices could rise. We know by Proposition 2 that a near-term increase in natural gas consumption implies a weak green paradox. Indeed, the expected combination of curvatures favors a world more like that in the left-hand plot of Figure 4, in which an anticipated carbon price should increase near-term emissions while decreasing emissions in the future and in total.

In order to test our predictions in commodity markets, we now complete our theoretical analysis by connecting changes in future spot prices to changes in futures markets. We are interested in how the time t futures price for delivery at time $T > t$ changes in response to a first-order stochastic dominant shift in the probability distribution for the time s emission tax, where $s > T$ and the support of each distribution is the set of weakly positive numbers. Assume that the time s tax τ_s will be known at time T . In that case, the spot price at time T is $g(S_T, \tau_s)$, where S_T is the spot price in the absence of a time s tax (implying $g(S_T, 0) = S_T$). The time T spot price increases in S_T : $\partial g(S_T, \tau_s) / \partial S_T > 0$. In region A of Figure 2, the time T spot price decreases in τ_s ($\partial g(S_T, \tau_s) / \partial \tau_s < 0$), with the opposite effect in all other regions.

Let $F_{t,T}^k$ be the futures price at time t for delivery (or cash termination) at time T under distribution k for the tax.⁷ Define \mathbb{Q}_k as the risk-adjusted measure under tax distribution k .⁸ The proposition describes how the futures price changes when current events shift the time s tax distribution from distribution i to distribution j , where we assume that current events shift tax probabilities such that distribution j first-order stochastically dominates (FSD) distribution i (so that the expected tax increases).

Proposition 3 (Futures Prices). *Assume that the commodity price follows an Ito process in the absence of taxes. The following conditions are jointly sufficient for the sign of $F_{t,T}^j - F_{t,T}^i$ to match the sign of $\partial g(S, \tau) / \partial \tau$:*

⁷We assume deterministic interest rates. The futures price therefore equals the forward price (Cox et al., 1981).

⁸The time t futures price is then $F_{t,T}^k = \mathbb{E}_t^{\mathbb{Q}_k} [g(S_T, \tau_s)]$. The risk-adjusted measure determines how the parties to a futures contract are compensated for bearing risk.

- (i) either the market price of risk does not change or its change in each instant has the opposite sign as $\partial g(S, \tau)/\partial \tau$,
- (ii) the expected time T stochastic discount factor weakly increases, and
- (iii) the covariance between the time s tax and the time T stochastic discount factor weakly increases.

Proof. See appendix. □

The proof shows that we can approximate the change in the futures price as:

$$F_{t,T}^j - F_{t,T}^i \approx \underbrace{\left(\mathbb{E}_t^{\mathbb{Q}_j} [S_T] - \mathbb{E}_t^{\mathbb{Q}_i} [S_T] \right)}_{\text{Commodity expectation channel}} + \frac{\partial g(S, \tau)}{\partial \tau} \Big|_{\left(\mathbb{E}_t^{\mathbb{Q}_0} [S_T], 0 \right)} \underbrace{\left(\mathbb{E}_t^{\mathbb{Q}_j} [\tau_s] - \mathbb{E}_t^{\mathbb{Q}_i} [\tau_s] \right)}_{\text{Tax expectation channel}}, \quad (4)$$

where $k = 0$ indicates the degenerate distribution with the tax certainly equal to 0. Beyond the effect on the future spot price $g(\cdot)$ analyzed previously, there are two effects due to changing the risk-adjusted measure from \mathbb{Q}_i to \mathbb{Q}_j . We call their net effect the *risk effect*. First, we have an altered expectation of the no-tax commodity spot price (the “commodity expectation channel”). The proof shows that this change is equal to the integrated difference between the market price of risk under the old and new pricing kernels. The market price of risk reflects the degree to which the commodity’s volatility factor is diversifiable. For instance, the new tax distribution might decouple energy prices and economic growth. If high energy prices are driven by a booming economy, then this decoupling makes the payoff from holding the commodity less negatively correlated with marginal utility and so lowers the market price of risk. On the other hand, if high energy prices cause recessions, then decoupling makes the commodity less countercyclical and so raises the market price of risk. A lower (higher) market price of risk increases (decreases) the futures price in order to account for the contract holder’s increased (reduced) diversification. The first (second) case reinforces the effects described in previous sections if and only if $\partial g(\cdot)/\partial \tau > (<) 0$. We expect the commodity expectation channel to be small for small changes in the distribution of emission prices.

The second component of the risk effect arises from the difference between the tax’s expectation under the two risk-adjusted measures (the “tax expectation channel”). If both expectations were taken under the physical measure, we know this term would be positive by the assumption that the new tax distribution first-order stochastically dominates the old one. However, in analyzing this term we have to account for the change of risk-adjusted measure in addition to the change in the actual distribution. The proof derives an expression for this channel. Several stories are consistent with the expression being positive and so reinforcing the effects analyzed previously. First, the tax expectation channel is unambiguously positive if neither the expectation of the time T stochastic discount factor nor the covariance between

the time s tax and the time T stochastic discount factor change appreciably. Second, if the FSD shift in the tax raises the expected time T stochastic discount factor by making investors feel poorer, then the whole term is still unambiguously positive if the covariance does not decrease.⁹ Third, if the FSD shift in the time s tax is accompanied by a rule reducing the tax in bad times and raising it in good times, then the covariance decreases and the tax expectation channel is positive as long as the stochastic discount factor does not change by too much. Finally, if a higher tax directly hurts economic output, then the covariance increases and the tax expectation channel is positive as long as the stochastic discount factor does not decrease too strongly. While these stories all support a positive channel, a negative channel would depend on the change in the stochastic discount factor being sufficiently large and of opposite sign to the change in the covariance. By raising the expected value under the physical measure, the FSD shift stacks the deck against a negative tax expectation channel that would reverse the effects analyzed previously.

In sum, the risk effect is itself composed of a commodity expectation channel and a tax expectation channel. The first is probably small and the second is probably positive. In that case, the sign of the net effect of a future tax change on current futures prices is determined by the direct, substitution, and anticipation effects described previously. In our empirical application, the longest-dated liquid futures contracts expire before emission pricing was to have taken effect. We predict that an FSD change in the distribution of emission prices (e.g., establishing a cap-and-trade program) would decrease futures prices for coal, with an ambiguous (but probably smaller) effect for natural gas. In our event, a cap-and-trade proposal collapses. We therefore predict that this collapse increased coal futures prices and expect a smaller—and potentially negative—effect on natural gas futures.

2 Shifting regulatory expectations: Graham walks out

Having analyzed how evolving expectations of future carbon pricing should affect current energy futures, we now argue that Senator Lindsey Graham’s abandonment of the American Power Act isolates a shift in expectations that can test the theory. The main identification challenge is establishing that even insiders did not see the event coming. For instance, neither a vote on legislation nor a press conference provides the crucial element of surprise. We 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

⁹If we think of the FSD shift in the tax as taking the time s tax closer to its optimal level and think of the stochastic discount factor as being the one a social planner would use, then the expected stochastic discount factor decreases when we change the tax’s distribution from i to j . Assuming the covariance is still approximately the same, the tax expectation channel then becomes negative only if the expected stochastic discount factor decreases by a sufficiently large amount.

reduce U.S. greenhouse gas emissions to 17% below 2005 by 2020 and 83% below 2005 by 2050 using a “cap-and-trade” scheme of tradable emission permits. The vote was tight (219–212) and broke down along 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 the spring of 2010, there were 57 Democratic senators (including several from oil- and coal-producing areas) and 2 independent senators. This 60-vote hurdle therefore meant that any bill would need to attract the support of at least one Republican in order to become law.

Senator Lindsey Graham was that Republican. In a joint October 2009 op-ed in the *New York Times*, he and Democratic Senator John Kerry announced a partnership to pass climate legislation. Senator Joseph Lieberman, an independent who caucused with the Democrats, 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 targets matched those in the Waxman-Markey bill, except beginning a year later with a 4.75% reduction in 2013. They 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. 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.¹⁰ To address petroleum refiners’ concerns about permit price fluctuations, the transportation sector would be regulated via a novel “linked fee” that provided price stability.

This last bargain proved their undoing, as documented by Lizza (2010). On April 15, 2010, with the bill nearly complete and its opening on the legislative calendar fast approaching, Fox News broke a story about the linked fee: “WH Opposes Higher Gas Taxes Floated by S.C. GOP Sen. Graham in Emerging Senate Energy Bill” (Garrett, 2010). The bill’s authors had worked hard to avoid any impression of a politically poisonous gasoline tax, and now not only had the label cropped up but it had done so via unnamed Democratic White House sources pinning the idea on Senator Graham. He was enraged: not only had this “linked fee” since been modified specifically to avoid having the Congressional Budget Office describe it as a tax, but Graham felt scapegoated for the bill’s least popular provision. And this after already taking heat for working with Democrats on climate legislation. In the story’s wake, Senate Majority Leader Harry Reid declined to provide Graham the political cover he sought.

Still the three senators pushed ahead. On Thursday April 22, after months of meetings and negotiations, they completed the bill with a final compromise to bring the Edison Electric Institute on board. They scheduled the bill’s unveiling for Monday April 26. The joint press conference would include top business, environmental, and religious figures to demonstrate

¹⁰The 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.

the breadth of support for their comprehensive package.

On Friday April 23, Arizona Governor Jan Brewer signed SB 1070 into law. This controversial bill stirred up Latino leaders by its intrusive measures to deter illegal immigration. Meanwhile, Senate Majority Leader Reid faced a tough reelection fight in neighboring, Latino-heavy Nevada. The day before the signing (April 22), he announced that the Senate would begin advancing a federal immigration bill. 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 a point man in both the climate and immigration efforts, Graham knew this well. He considered the announcement cheap political point-scoring that demonstrated a lack of commitment from the Senate leadership to take care of climate change.

The Saturday morning Washington Post reported on the upcoming unveiling of Kerry-Graham-Lieberman (Eilperin, 2010a). It detailed the bill's array of industry support and quoted 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 (Cash, 2010). But this story must have been filed before the weekend. According to Lizza (2010), at 10 PM on Friday night one of Graham's aides e-mailed his counterpart on Lieberman's side to say "Sorry buddy." Lieberman's aide later described the note as "soul-crushing." On Saturday morning, Graham walked out on his bill. His formal statement refused to delay a climate bill for immigration, but he also sent Majority Leader Reid a note about the lack of cover for the modified linked fee. Neither industry insiders nor the bill's other sponsors seemed to know until Saturday that the bill was about to be held up. Kerry promptly flew back to Washington, DC from Boston to meet with Graham. Lieberman broke his Jewish Sabbath restrictions to call Graham. Lieberman's aide got 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 to deal with global warming in the Obama era, was officially dead."

Reaction was swift. On Saturday evening, the Washington Post's climate blog quoted a lawyer for electric utilities as saying that Graham's departure "diminishes significantly" the chance of passing "a large climate package" (Eilperin, 2010b). The New York Times said the move cast "its already cloudy prospects deeper into doubt" (Broder, 2010). A prominent climate policy blogger posted that losing Graham "would certainly kill any chances of a climate bill this year" (Romm, 2010). Reuters led its breaking update with "severe setback," saying the effort to pass climate legislation "could be doomed" without Graham on board to gather Republican support (Cowan and Ferraro, 2010). On Sunday evening, Politico reported that Reid was holding firm on moving immigration reform first (Thrush and Cogan, 2010). On Monday, the Oil & Gas Journal quoted an energy lobbyist congratulating Graham "for apparently backing out of this job-killing legislation" (Snow, 2010). Reuters cautioned that the bill was "stalled—but not yet officially dead" (Cowan, 2010). The Los Angeles Times referred to the bill as "on life support" (Tankersley, 2010).

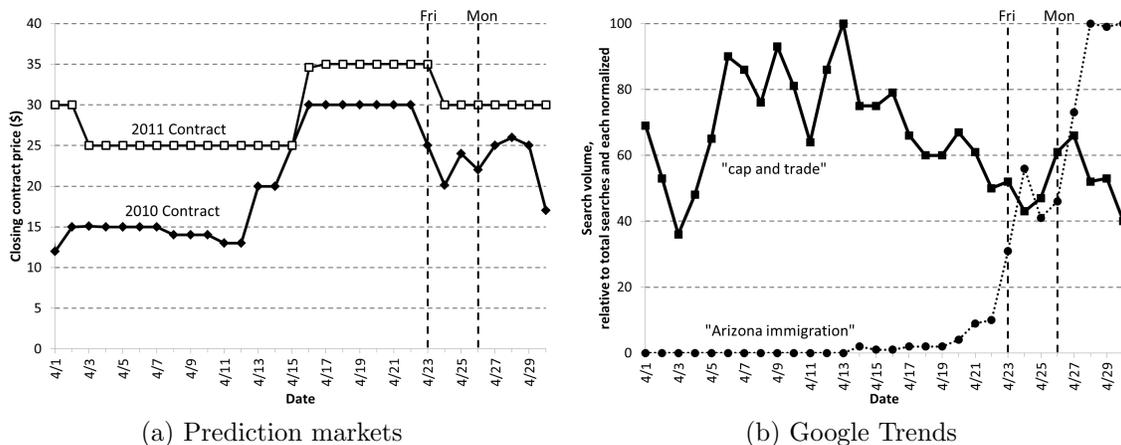
Over the next weeks, the ongoing Deepwater Horizon oil spill made the bill's drilling compromise increasingly untenable even were Graham to come back on board. On May 13, Kerry and Lieberman finally unveiled their bill as the American Power Act. They did so without Graham. On June 8, Graham announced that he would not vote for his own bill. In late July, Senator Reid acknowledged that he could not find 60 votes for Kerry-Lieberman and would instead pursue smaller bills more narrowly targeted to energy and the oil spill.

This narrative strongly suggests that the prospect of the Senate passing climate legislation took a sharp hit between the markets' close on Friday April 23 and their close on Monday April 26. Hints of potential problems due to immigration were present by Thursday and Friday, but the press conference was still scheduled as of Friday's closing. Indeed, key players were not planning to cancel the press conference as late as Saturday morning. Events after Friday afternoon appear to have surprised the bills' authors, their aides, energy lobbyists, and pundits.

Events after Friday afternoon also appear to have surprised prediction markets. Intrade offered contracts that would pay \$100 in the case that a cap-and-trade system was "established" by the end of 2010 or 2011. The price of each contract is conventionally read as its probability of paying off.¹¹ Figure 5a shows the evolution of these prices over the course of April 2010. The price of the 2011 contract was consistently higher because enacting a cap-and-trade system by the end of 2010 would also trigger a payoff under the 2011 contract. Both contracts were stable in the week leading up to Friday April 23: the 2010 contract traded around \$30, and the 2011 contract traded around \$35. On that Friday, the 2010 contract dropped to \$25. On Saturday, the proposed event day, the 2010 contract dropped further to \$20 while the 2011 contract also dropped to \$30. The 2011 contract remained at this new level through the rest of the month. The 2010 contract bounced back to \$24 on Sunday before dropping again to close at \$22 on Monday. In line with our narrative evidence and with contemporary news accounts, prediction market data suggest that expectations of climate legislation shifted on both Friday and Saturday.

Google Trends data on web searches in April 2010 provide further evidence for the "newness" of the April 24 weekend's information (Figure 5b). Saturdays and, especially, Sundays are generally times of relatively low search volume. Yet we see interest in "cap and trade" rise from Saturday to Sunday and spike into Monday April 26. This evidence is not conclusive, however, as we may have expected a spike even if the press conference had gone forward. What more clearly supports the above narrative is the interest in "Arizona immigration". Here we see trivial search volume up until April 20. At that point interest in the topic increases, and we see a dramatic spike begin on Friday April 23. Governor Brewer's

¹¹While it is common practice to interpret prediction market prices as probabilities, it is not theoretically clear what beliefs these prices represent (Manski, 2006; Fountain and Harrison, 2011). Further, it is not clear that market participants will correctly assess small changes in probabilities (Wolfers and Zitzewitz, 2006). Most importantly, these particular contracts are highly illiquid. We therefore emphasize the general pattern and the event day's change in value rather than placing too much weight on the precise contract prices.



(a) Prediction markets

(b) Google Trends

Figure 5: Intrade prediction market contracts for “a cap and trade system for emissions trading to be established before midnight ET on 31 Dec” of 2010 and 2011. Also, normalized search volume (from Google Trends) for “cap and trade” and “Arizona immigration.”

signature appears to have kicked the political controversy up several levels. Graham’s formal statement cited Majority Leader Reid’s decision to move immigration as the prime reason for withdrawing from the climate effort. Internet search data strongly suggest that the immigration concern was not an issue that had been festering as a likely back-breaker but instead bubbled up relatively quickly the day before the proposed event day.

All of this evidence indicates that the probability of passing a climate bill in 2010 dropped on April 24. We use this shift in expectations to identify the effect of anticipated climate legislation in contemporary futures markets. Importantly, this bill would have begun pricing the electricity sector’s carbon emissions in 2013, many months beyond the longest-dated futures contracts for natural gas and coal. Our theoretical analysis predicts that coal futures prices jumped upwards between markets’ close on Friday April 23 and their close on Monday April 24. The theoretical prediction for natural gas is ambiguous, so we cannot use its response to test the theory. Instead, if coal futures support the theory, then we can use natural gas futures to learn about the effect on emissions: from Proposition 2, a weak green paradox failed to hold only if natural gas prices responded in the opposite direction as coal prices.

3 Estimation framework

We use the exogenous shift in carbon price expectations to identify the effect of anticipated climate policy on current commodity prices.¹² We are interested in two questions. First, does anticipated policy lower coal futures prices as predicted by the theoretical framework? Second, if the theory does hold for coal futures, then does it imply that the bill's anticipation was increasing or decreasing emissions prior to its implementation? The response of natural gas futures identifies which of the theoretically possible emission outcomes was consistent with actual market dynamics.

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 normal returns. Our event happened on Saturday April 24, while markets were closed. The event day is therefore Monday April 26. As is common, the event window also includes the day before the event (Friday April 23) and the day after (Tuesday April 27). Indeed, this makes special sense in our case because of the evidence that expectations began shifting on Friday.¹³

The choice of estimation window must navigate a trade-off: longer periods provide more observations, but shorter periods define normal returns under market conditions more similar to the date of interest and have more usable futures contracts. We employ two estimation windows to assess robustness. The shorter estimation window uses the 60 trading days centered around the event window, and the longer estimation window uses the 200 trading days centered around the event window.¹⁴ As described below, these two window lengths allow us to detect effects at the conventional 10%, 5%, and, for the 200-day window, 1% significance levels when using the SQ test.

Our 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 climate policy surely shifted at other

¹²We focus on commodity prices instead of equity prices for three 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. Third, forward-looking equity markets integrate over the changes in profits that occur before and after the policy's implementation. Even if we did know the allocation rule and could predict profit outcomes in each period, equity market outcomes would still confound pre- and post-implementation effects.

¹³We 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.

¹⁴The 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.

times in our estimation window (for instance, we see the prediction market move on other days in Figure 5a), 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 non-event variation in futures returns. The second problem is mitigated by removing Friday April 23 from the estimation window, but it does still bias the results towards no statistically significant effect.

We regress daily futures returns on the event window dummies and, in some specifications, on covariates:

$$F_{cit} = X_{it}\beta_{ci} + D_t^{Fri}\gamma_{ci}^{Fri} + D_t^{Mon}\gamma_{ci}^{Mon} + D_t^{Tues}\gamma_{ci}^{Tues} + \epsilon_{cit}, \quad (5)$$

where c indicates the commodity of interest (coal or natural gas), i indexes the futures contract (by month of expiration), and t indexes the trading day. There are $T + 3$ trading days in the combined estimation and event windows. F_{cit} is the log return. X_{it} is the $1 \times k$ vector of covariates, with β_{ci} 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 equals 0 otherwise. For each combination of a commodity and contract, the coefficient of interest is γ_{ci}^{Mon} . The error term ϵ_{cit} gives the excess (or unpredicted) returns on day t .¹⁵ 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 et al., 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 our estimation framework.

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. These specifications impose minimal structure. As discussed in Section 5.1, including covariates helps ensure identification by removing any potentially confounding non-event news that also affected the covariates, but including covariates can also bias our estimates towards zero by absorbing some of the event's effect into the explained portion of returns. The specifications with covariates include a constant and the following variables in log return form. First, the S&P 500 stock price index, the 10-year U.S. Treasury rate (constant maturity), and the 3-month London Interbank Offered Rate (LIBOR) capture general economic conditions.¹⁶ Additional covariates reflect broader commodity markets: the Baltic Dry Index (shipping),¹⁷

¹⁵Durbin-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.

¹⁶All of these are available from the Federal Reserve Economic Data (FRED) web site.

¹⁷Purchased from EODData.

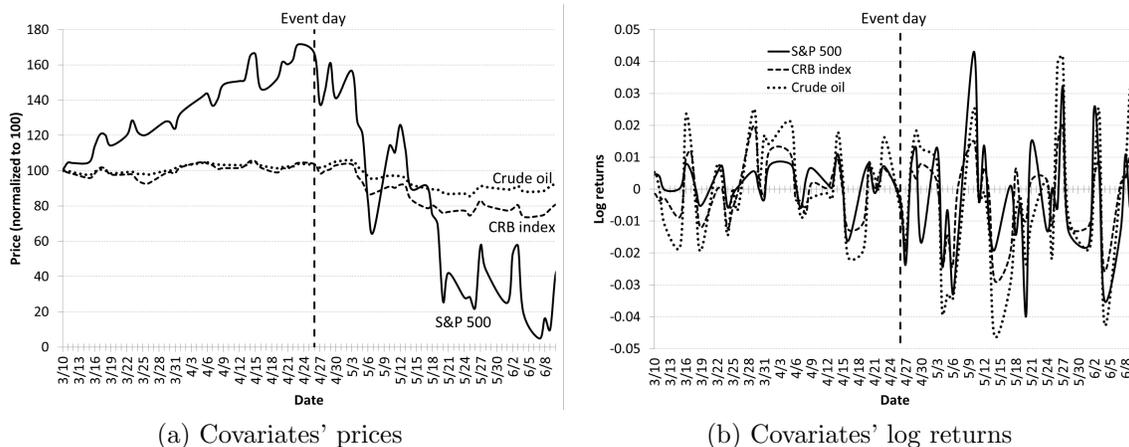


Figure 6: Prices and log returns for three of the covariates in the 60-day estimation window. The “crude oil” plot uses the July 2010 contract.

the Thomson Reuters/Jefferies CRB Index,¹⁸ and the benchmark West Texas Intermediate crude oil contract for the month nearest to contract i (NYMEX symbol: CL).¹⁹ Figure 6 shows that three of the covariates that are potentially most sensitive to the event do not appear to behave abnormally in the event window.

Coal futures prices are from the NYMEX Central Appalachian contract (symbol: QL), and natural gas futures prices are from the NYMEX Henry Hub contract (symbol: NG).²⁰ These are the reference contracts for coal and gas in the North American market.²¹ Each specification uses only those contracts that are listed throughout the estimation window

¹⁸Purchased from the Commodity Research Bureau (CRB). We construct a rolling contract using the nearest month’s price. Note that natural gas forms 6% of the index. Including this index increases our power insofar as it picks up non-event noise, but it reduces our power insofar as changes in the price of natural gas might move the index. Geman (2005) recommends this index in part because no one commodity tends to move it.

¹⁹Purchased from the Commodity Research Bureau. 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.

²⁰Futures data come from the Commodity Research Bureau. The coal contract is not subject to a limit move restriction. The natural gas contract pauses trading after a large move in the nearest contract months. Among the covariates, the crude oil contract is subject to a similar restriction. However, limit moves are not a problem in our event or estimation windows.

²¹The 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. Uranium futures contracts (NYMEX: UX) also showed zero volume.

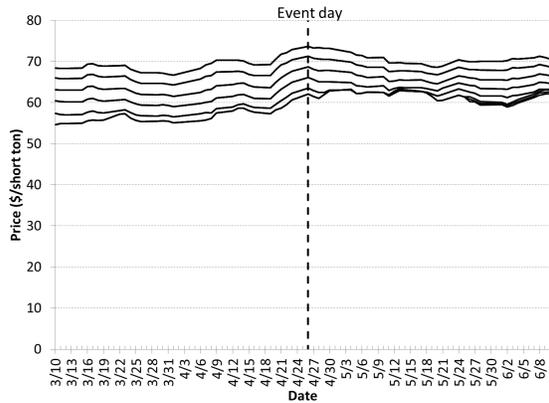
and that have nonzero volume nearly every day throughout the estimation window. These restrictions yield 18 coal contracts and 22 natural gas contracts in the 60-day estimation window, and they yield 15 coal contracts and 19 natural gas contracts in the 200-day estimation window.

The top panel of Figure 7 plots futures prices in the 60-day estimation window. The longer-dated contracts tend to have higher prices. Each series remains within a relatively narrow band throughout the window, without particularly dramatic changes. The event day does not appear noticeably abnormal, though it does occur near a local maximum for the coal contracts. The middle panel plots log returns. Natural gas returns exhibit greater variance and cluster tightly around zero on the event day. In contrast, the event day's coal returns are unusually large. This visual evidence suggests a positive event effect for coal at the edge of the "usual" variation. Finally, the bottom panel plots log volume. The natural gas contracts usually trade at greater volume than do the coal contracts, and nearer-dated contracts usually trade at greater volume than do longer-dated contracts.²² For both commodities, the event day's volume is typical of volumes in the 60-day estimation window.

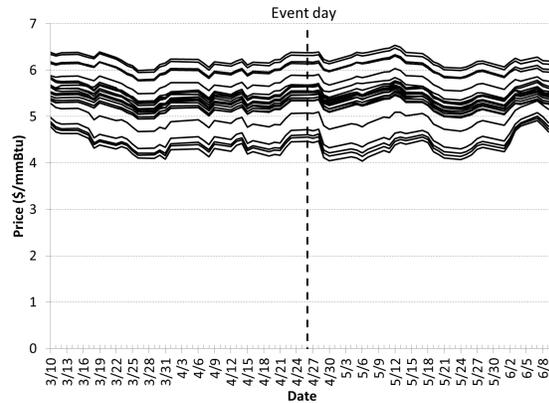
Figure 8 isolates returns in the days around the event day, again using the 60-day estimation window's liquid contracts. The coal contracts have positive returns on Thursday, Friday, and Monday and negative returns on Tuesday. In line with the theoretical predictions, coal's highest returns are on Monday and its second-highest returns are on Friday. In contrast, the returns from natural gas contracts are positive on Thursday and Friday and near zero on Monday and Tuesday. These plots suggest a positive event effect for coal and no significant event effect for gas.

A complication arises when testing whether the coefficient of interest (γ_{ci}^{Mon}) is significantly different from zero (Conley and Taber, 2011; Gelbach et al., forthcoming). 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. Yet it has been recognized since at least Brown and Warner (1985) that individual securities' returns are often non-normal. We test for normality in two ways. First, a Jarque-Bera test uses the skewness and kurtosis of the realized excess returns to test the null hypothesis that the distribution is normal. Second, a Shapiro-Wilk test uses the observed order statistics to test the null hypothesis of normality. We obtain mixed results. For specifications with covariates, neither test rejects normality at a reasonable significance level for natural gas in the 60-day window, coal in the

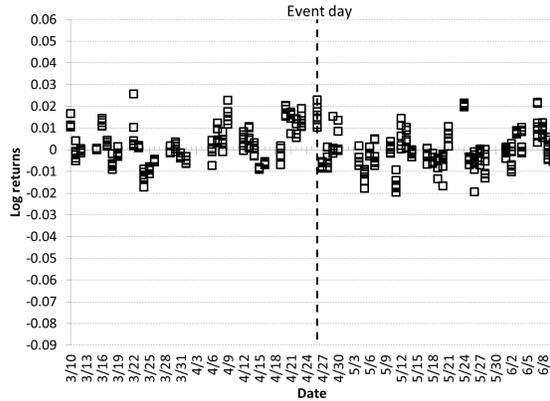
²²Two of the included natural gas contracts have a single day with zero volume, and 110 instances of included coal contracts have zero volume. However, many of these zero-volume observations occur when multiple longer-dated contracts have zero volume on the same day, so that instances of coal contracts having zero volume actually occur on only 23 distinct trading days. No instances of zero volume occur within a week of the event day.



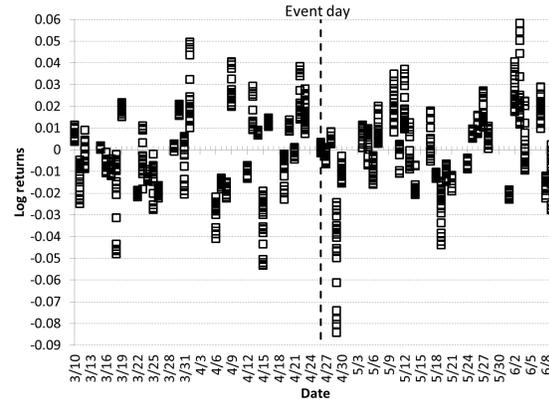
(a) Coal prices



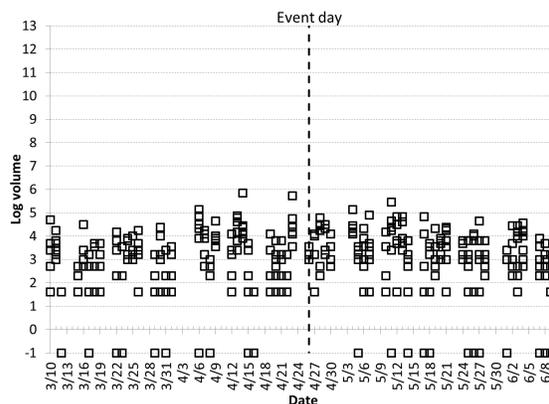
(b) Natural gas prices



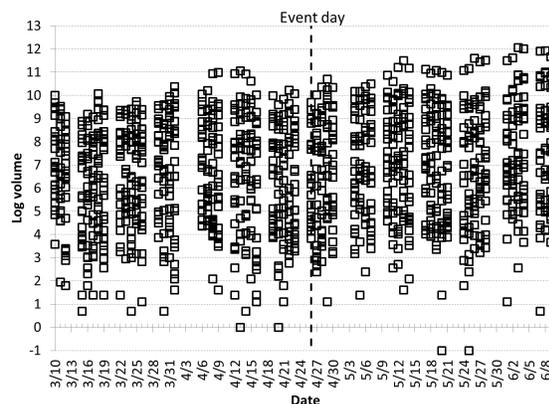
(c) Coal log returns



(d) Natural gas log returns

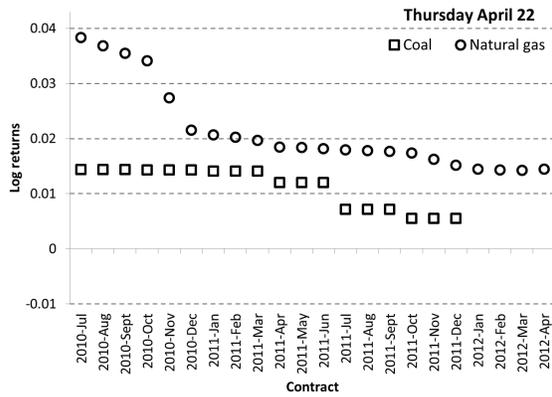


(e) Coal log volume

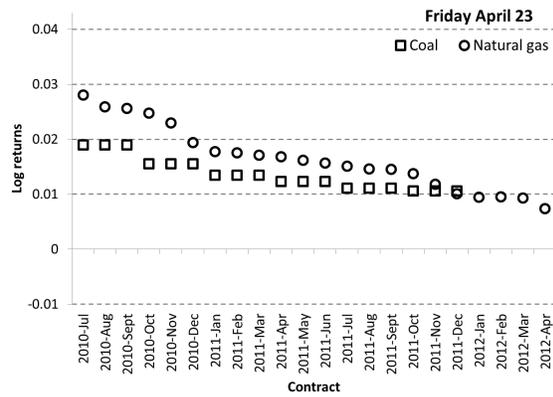


(f) Natural gas log volume

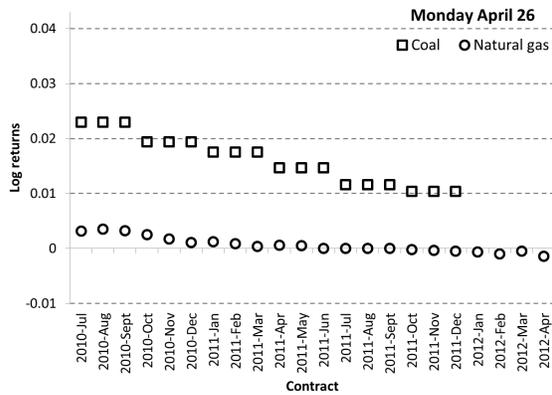
Figure 7: Futures prices (top), log returns (middle), and log volume (bottom) in the 60-day estimation window. Observations with zero volume are plotted as -1 on the log scale.



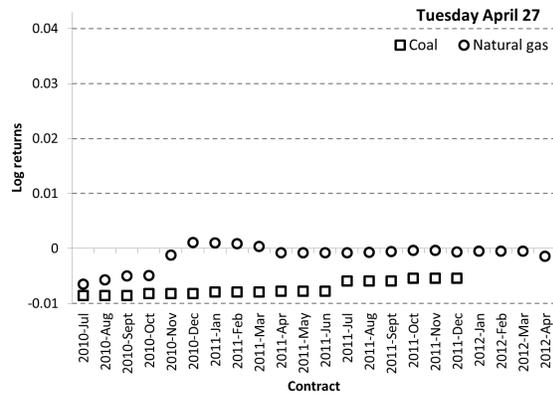
(a) Thursday, April 22



(b) Friday, April 23



(c) Monday, April 26 (event day)



(d) Tuesday, April 27

Figure 8: Log returns around the event window.

200-day window, or the six nearest coal contracts in the 60-day window. On the other hand, both tests reject normality at the 5% level for most natural gas contracts in the 200-day window and for the longer-dated coal contracts in the 60-day window. Standard t-tests may be appropriate for some contracts in some estimation windows but not for others.

Gelbach et al. (forthcoming) 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 stationary, in which case it can be approximated by the realized distribution of excess returns. We can then assess the probability that the event day’s excess return came from that same population by considering its quantile in the excess return distribution. Consider the alternate hypothesis of a positive event effect. We reject the null hypothesis of a weakly negative event-specific effect at the $(\alpha \times 100)\%$ level if and only if the event day’s excess return is above the $(1-\alpha)$ -quantile of the estimation window’s distribution of excess returns. The $(1-\alpha)$ -quantile is itself the $(1 - \alpha) \times T$ -largest realized excess return. Assuming stationarity, this estimator of the distribution of excess returns improves as the estimation window increases (i.e., as T grows). Autocorrelation in the error term does not pose a special problem as long as T is sufficiently large and there is only a single event observation (Andrews, 2003).

The size of the SQ test is distorted when $\alpha \times T$ is not an integer (Gelbach et al., forthcoming). We therefore choose the length of the estimation window to obtain integer values for desired Type I error rates α . With the 200-day estimation window and the alternate hypothesis of a positive event effect, the cutoff for the 10% significance level is the 20th greatest excess return, the cutoff for the 5% significance level is the 10th greatest excess return, and the cutoff for the 1% significance level is the 2nd greatest excess return. With the 60-day estimation window, the cutoff for the 10% significance level is the 6th greatest excess return and the cutoff for the 5% significance level is the 3rd greatest excess return. The 60-day estimation window cannot differentiate between significance levels smaller than 1.67%.

A final test achieves greater power by using the theoretical prediction that the event effect should have 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). We then test the event’s effect on the summary index using the same combination of t-tests and SQ tests as before.²³ We consider two summary indices in specifications with covariates. The first averages a day’s standardized excess returns across contract lengths, and the second uses an efficient generalized least square (GLS) weighting procedure to average across contract lengths.²⁴ Those contracts that are less correlated with the others provide more “new”

²³Gelbach et al. (forthcoming) establish the validity of the SQ test for a real-valued function of the outcome vector.

²⁴The 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.

information and so receive higher weights; those contracts that are more correlated provide more redundant information and so receive lower weights. To the extent that the event affects all contract lengths more consistently than does standard noise, the summary index tests should improve our chances of detecting the event's effect in futures markets.²⁵

4 Results: The event's effect on coal and natural gas futures

We are primarily interested in whether coal futures' event day coefficients (γ^{Mon}) are significantly positive. This finding would indicate the presence of an anticipation effect in commodity markets: we would have found a highly unusual positive excess return on precisely the day predicted by theory. We consider the significance of these coefficients under standard t-tests and under the SQ test. We also compare these results to the coefficients for the day before the event (Friday) and for the day after the event (Tuesday). We have argued that Friday could have been the first phase of the event and so might have a positive coefficient. We expect Tuesday's coefficient to be statistically indistinguishable from zero.

Table 1 reports the coefficients on the event day dummy for each coal contract, along with autocorrelation-robust standard errors and significance for a t-test of the null hypothesis of a weakly negative effect. Regardless of the inclusion of covariates or of the length of the estimation window, the nearer-dated contracts show significant jumps in the coal price at the 5% level.²⁶ We see the greater power of the GARCH specification: under either estimation window, its estimated positive event effect is statistically significant at the 1% level for nearly all contracts because it produces smaller standard errors.²⁷ For the nearest-dated contracts, the event day is associated with about a 2% increase in coal futures returns. For the longest-dated contracts, the best estimate is about a 1% increase in returns. These results are stable across specifications. The central estimates of the coefficient for the day before the event (γ^{Fri}) are largely similar to those of the event day's coefficient (Table 2), though slightly smaller at the nearest dates and slightly larger at the longest dates. Some contracts' Friday coefficients are significant at the 10% level with standard error terms, and most are significant at the 1% level with GARCH error terms. In contrast, the coefficient for

²⁵The Jarque-Bera and Shapiro-Wilk tests reject normality of the summary index at the 1% and 5% levels, respectively, for coal with the 60-day estimation window. They also reject normality around the 10% level for natural gas with the 200-day estimation window (with the Shapiro-Wilk test barely above the 10% level). They fail to reject normality at standard significance levels for natural gas with the 60-day window and for coal with the 200-day window.

²⁶When including covariates, the coal contracts have an R^2 of around 0.21–0.32 in the 60-day estimation window and around 0.22–0.24 in the 200-day estimation window. The commodity index and S&P 500 equity index often have significantly positive coefficients, while the 10-year Treasury rate often has a significantly negative coefficient. Other covariates' coefficients are not typically significant.

²⁷The estimated GARCH and ARCH terms are statistically insignificant.

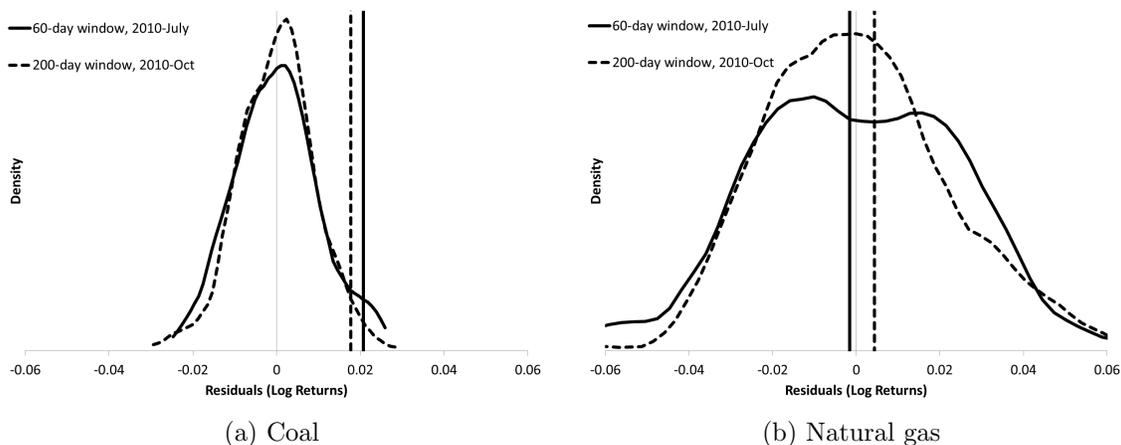


Figure 9: Kernel density plot of the residuals in the estimation window, with the coefficient (γ^{Mon}) on the Monday April 26 dummy indicated by a vertical line. Plots use the nearest-dated contract under specifications (2) and (5).

the day after the event (γ^{Tues}) does not show any significant results and its central estimates are usually negative (Table 3).

The SQ test abandons the t-test's assumption that the distribution of excess returns is normal for the weaker assumption of stationarity. Recall that the SQ test compares the event day coefficient to the distribution of excess returns realized in the estimation window. To be significant at the $(\alpha \times 100)\%$ level is to be above the $(1 - \alpha)$ -quantile of that distribution. Figure 9a plots the distribution of the nearest-dated contract's excess returns for specifications (2) and (5). The event day's excess returns for coal are far in the positive tail of this distribution. Table 4 confirms this visual evidence. For most contracts and specifications, the event day coefficient is significant at the 10% level, and it is often significant at the 5% level. In other words, the event day's excess returns are usually above the 90th percentile of the estimation window's excess returns, and they are often above the 95th percentile. Excess returns of the event day's magnitude are rarely observed in either estimation window.

Table 5 reports the event day coefficients for natural gas.²⁸ These coefficients are not statistically significant: while they are almost always positive, they are typically an order of magnitude smaller than coal's event day coefficients even as natural gas returns' greater volatility generates larger standard errors. The SQ test also fails to detect a significant event effect in gas markets. The nearest-dated contract's event day excess returns are near

²⁸When including covariates, the natural gas contracts have an R^2 of around 0.1–0.2 in either estimation window. The commodity index often has a significantly positive coefficient while the corresponding crude oil contract often has a significantly negative coefficient. Other covariates' coefficients are not typically significant.

Table 1: **Coal:** Coefficients (γ_{ci}^{Mon}) on the dummy variable for the event day.

Contract (i)	Dependent Variable: Log Futures Returns for the Indicated Contract					
	(1)	(2)	(3)	(4)	(5)	(6)
Jul-2010	0.021** (0.011)	0.021** (0.011)	0.021*** (0.0043)			
Aug-2010	0.021** (0.010)	0.021** (0.011)	0.021*** (0.0041)			
Sept-2010	0.021** (0.010)	0.021** (0.011)	0.021*** (0.0040)			
Oct-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)
Nov-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)
Dec-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)
Jan-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)
Feb-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)
Mar-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)
Apr-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)
Jun-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)
Jul-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)
Aug-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)
Sept-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)
Oct-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)
Nov-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)
Dec-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-day	60-day	60-day	200-day	200-day	200-day
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. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a t-test of the null hypothesis of a weakly negative effect. Covariates are the S&P 500, the 10-year U.S. Treasury rate, 3-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract.

Table 2: **Coal:** Coefficients (γ_{ci}^{Fri}) on the dummy variable for the day before the event.

Contract (<i>i</i>)	Dependent Variable: Log Futures Returns for the Indicated Contract					
	(1)	(2)	(3)	(4)	(5)	(6)
Jul-2010	0.017 (0.011)	0.018 (0.012)	0.018*** (0.0057)			
Aug-2010	0.017 (0.010)	0.018 (0.012)	0.018*** (0.0052)			
Sept-2010	0.017 (0.010)	0.017 (0.012)	0.017*** (0.0050)			
Oct-2010	0.014* (0.0084)	0.017* (0.0090)	0.017* (0.0097)	0.016 (0.010)	0.016* (0.0095)	0.016** (0.0074)
Nov-2010	0.014* (0.0084)	0.017* (0.0090)	0.017* (0.010)	0.016 (0.010)	0.016* (0.0095)	0.016*** (0.0032)
Dec-2010	0.014* (0.0084)	0.017* (0.0090)	0.017 (0.011)	0.016 (0.0099)	0.016* (0.0094)	0.016*** (0.0029)
Jan-2011	0.013 (0.0078)	0.015* (0.0085)	0.015*** (0.0039)	0.014 (0.0089)	0.015* (0.0085)	0.015*** (0.0024)
Feb-2011	0.013 (0.0078)	0.015* (0.0085)	0.015*** (0.0039)	0.014 (0.0089)	0.015* (0.0084)	0.015*** (0.0024)
Mar-2011	0.013 (0.0078)	0.015* (0.0085)	0.015*** (0.0039)	0.014 (0.0089)	0.015* (0.0084)	0.015*** (0.0024)
Apr-2011	0.012 (0.0077)	0.013 (0.0085)	0.013*** (0.0041)	0.012 (0.0085)	0.013 (0.0081)	0.013*** (0.0022)
May-2011	0.012 (0.0077)	0.013 (0.0085)	0.013*** (0.0041)	0.012 (0.0085)	0.013 (0.0081)	0.013*** (0.0022)
Jun-2011	0.012 (0.0077)	0.013 (0.0085)	0.013*** (0.0041)	0.012 (0.0085)	0.013 (0.0081)	0.013*** (0.0022)
Jul-2011	0.011 (0.0075)	0.013 (0.0081)	0.013*** (0.0040)	0.011 (0.0079)	0.012 (0.0076)	0.012*** (0.0019)
Aug-2011	0.011 (0.0075)	0.013 (0.0081)	0.013*** (0.0040)	0.011 (0.0079)	0.012 (0.0076)	0.012*** (0.0019)
Sept-2011	0.011 (0.0075)	0.013 (0.0081)	0.013*** (0.0040)	0.011 (0.0079)	0.012 (0.0076)	0.012*** (0.0019)
Oct-2011	0.010 (0.0076)	0.013 (0.0082)	0.013 (0.0090)	0.011 (0.0077)	0.011 (0.0073)	0.011*** (0.0019)
Nov-2011	0.010 (0.0076)	0.013 (0.0082)	0.013 (0.011)	0.011 (0.0077)	0.011 (0.0073)	0.011*** (0.0019)
Dec-2011	0.010 (0.0076)	0.013 (0.0082)	0.013 (0.017)	0.011 (0.0077)	0.011 (0.0073)	0.011*** (0.0019)
Estimation window	60-day	60-day	60-day	200-day	200-day	200-day
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. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a t-test of the null hypothesis of no effect. Covariates are the S&P 500, the 10-year U.S. Treasury rate, 3-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract.

Table 3: **Coal:** Coefficients (γ_{ci}^{Tues}) on the dummy variable for the day after the event.

Contract (<i>i</i>)	Dependent Variable: Log Futures Returns for the Indicated Contract					
	(1)	(2)	(3)	(4)	(5)	(6)
Jul-2010	-0.010 (0.011)	-0.0059 (0.012)	-0.0059 (0.0051)			
Aug-2010	-0.010 (0.010)	-0.0063 (0.012)	-0.0063 (0.0049)			
Sept-2010	-0.010 (0.010)	-0.0076 (0.011)	-0.0076 (0.0048)			
Oct-2010	-0.0094 (0.0084)	-0.0014 (0.0091)	-0.0014 (0.0063)	-0.0081 (0.010)	-0.00048 (0.0096)	-0.00048 (0.0042)
Nov-2010	-0.0094 (0.0084)	-0.0014 (0.0091)	-0.0014 (0.0069)	-0.0082 (0.010)	-0.00021 (0.0096)	-0.00021 (0.0029)
Dec-2010	-0.0094 (0.0084)	-0.0015 (0.0091)	-0.0015 (0.0075)	-0.0082 (0.0099)	0.00011 (0.0095)	0.00011 (0.0028)
Jan-2011	-0.0087 (0.0078)	-0.00063 (0.0086)	-0.00063 (0.0042)	-0.0079 (0.0089)	0.0015 (0.0086)	0.0015 (0.0027)
Feb-2011	-0.0087 (0.0078)	-0.00067 (0.0087)	-0.00067 (0.0043)	-0.0079 (0.0089)	0.0015 (0.0086)	0.0015 (0.0027)
Mar-2011	-0.0087 (0.0078)	-0.00066 (0.0087)	-0.00066 (0.0044)	-0.0079 (0.0089)	0.0016 (0.0086)	0.0016 (0.0027)
Apr-2011	-0.0084 (0.0077)	-0.003 (0.0088)	-0.003 (0.0044)	-0.0077 (0.0085)	0.00094 (0.0082)	0.00094 (0.0026)
May-2011	-0.0084 (0.0077)	-0.0032 (0.0088)	-0.0032 (0.0044)	-0.0077 (0.0085)	0.00095 (0.0082)	0.00095 (0.0026)
Jun-2011	-0.0084 (0.0077)	-0.0033 (0.0088)	-0.0033 (0.0044)	-0.0077 (0.0085)	0.00096 (0.0082)	0.00096 (0.0026)
Jul-2011	-0.0064 (0.0075)	-0.0022 (0.0084)	-0.0022 (0.0041)	-0.0058 (0.0079)	0.0016 (0.0077)	0.0016 (0.0024)
Aug-2011	-0.0064 (0.0075)	-0.0022 (0.0084)	-0.0022 (0.0041)	-0.0058 (0.0079)	0.0016 (0.0077)	0.0016 (0.0024)
Sept-2011	-0.0064 (0.0075)	-0.0022 (0.0084)	-0.0022 (0.0041)	-0.0058 (0.0079)	0.0016 (0.0077)	0.0016 (0.0024)
Oct-2011	-0.0058 (0.0076)	-0.0016 (0.0086)	-0.0016 (0.020)	-0.0053 (0.0077)	0.0020 (0.0074)	0.0020 (0.0024)
Nov-2011	-0.0058 (0.0076)	-0.0016 (0.0086)	-0.0016 (0.029)	-0.0053 (0.0077)	0.0020 (0.0074)	0.0020 (0.0024)
Dec-2011	-0.0058 (0.0076)	-0.0017 (0.0086)	-0.0017 (0.049)	-0.0053 (0.0077)	0.0020 (0.0074)	0.0020 (0.0024)
Estimation window	60-day	60-day	60-day	200-day	200-day	200-day
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. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a t-test of the null hypothesis of no effect. Covariates are the S&P 500, the 10-year U.S. Treasury rate, 3-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract.

Table 4: **Coal SQ test:** Percentile rank of the event-day residual among the estimation window residuals.

Contract (i)	(1)	(2)	(3)	(4)	(5)	(6)
Jul-2010	97**	96**	96**			
Aug-2010	97**	97**	97**			
Sept-2010	97**	100***	100***			
Oct-2010	98**	97**	97**	97**	97**	97**
Nov-2010	98**	97**	97**	97**	97**	97**
Dec-2010	98**	97**	97**	97**	97**	97**
Jan-2011	96**	97**	97**	97**	97**	97**
Feb-2011	96**	97**	97**	97**	97**	97**
Mar-2011	96**	97**	97**	97**	97**	97**
Apr-2011	94*	94*	94*	95*	96**	96**
May-2011	94*	94*	94*	95*	96**	96**
Jun-2011	94*	94*	94*	95*	96**	96**
Jul-2011	90	92*	92*	92*	93*	93*
Aug-2011	90	92*	92*	92*	93*	93*
Sept-2011	90	92*	92*	92*	93*	93*
Oct-2011	88	88	88	91*	91*	91*
Nov-2011	88	88	88	91*	91*	91*
Dec-2011	88	88	88	91*	91*	91*
Estimation window	60-day	60-day	60-day	200-day	200-day	200-day
Covariates	No	Yes	Yes	No	Yes	Yes
GARCH error	No	No	Yes	No	No	Yes

Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a null hypothesis of a weakly negative effect. The 60-day estimation window cannot distinguish between significance levels smaller than 1.67%. Covariates are the S&P 500, the 10-year U.S. Treasury rate, 3-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract.

the middle of the estimation window's distribution (Figure 9b). Indeed, Table 6 shows that the event day's returns are almost always between the 50th and 75th percentile of the estimation window's distribution. While the event day is generally associated with a slight increase in natural gas prices, this signal is statistically within the noise. However, the type of positive shock that would obscure a true negative event effect occurs in fewer than half of the estimation window's observations.²⁹

Finally, the summary index tests consider hypotheses about the sum of the contracts' event day coefficients (Table 7). The results are not sensitive to the choice of weighting scheme. Coal's index is positive and significantly different from zero at the 5% level, while the summary index for natural gas is also positive but not statistically significant. We again use the SQ test to achieve robustness against non-normal distributions. Figure 10 plots the kernel density estimator of the estimation windows' summary index distributions. As before, the event day's summary index lies far out in the positive tail for coal but sits near the middle for natural gas. Table 7 reports the percentiles and significance for the SQ test. Coal is again significant at the 5% level, meaning the event day's summary index is above the 95th percentile of the estimation window's distribution. In contrast, the event day's index values for natural gas are between the 65th and 75th percentiles. The summary index tests therefore support the stronger hypothesis that coal's contracts together demonstrate a consistently positive effect. They also make us more confident in a weak positive effect for natural gas (compared to the smaller percentiles often seen in Table 6's per-contract tests). The type of positive shock that would have obscured a true negative event effect for gas occurs in fewer than 35% of the estimation window's observations.

5 Discussion

We have seen that coal futures demonstrate a statistically significant jump in the theoretically predicted direction on the proposed event day. While the theoretical prediction for natural gas is ambiguous, the evidence suggests that the event moved natural gas futures in the same direction as coal. We next discuss the potential for confounding events. We then analyze implications for coal prices, for the cost of coal-fired electricity, and for annual coal consumption and emissions. Finally, we use the natural gas results to learn about the effect of the Senate bill on aggregate emissions.

5.1 Potentially confounding news

The primary challenge to identification is that other events might have affected returns on the event day. This can never be ruled out. The Platts Coal Trader's daily "Emissions

²⁹The coefficients on the dummy variable for Friday—the day before the event day but one for which we expect a similar shift—also reinforce the conclusion of a slight positive effect.

Table 5: **Natural Gas:** Coefficients (γ_{ci}^{Mon}) on the dummy variable for the event day.

Contract (<i>i</i>)	Dependent Variable: Log Futures Returns for the Indicated Contract					
	(1)	(2)	(3)	(4)	(5)	(6)
Jul-2010	0.0039 (0.028)	-0.0016 (0.028)	-0.0016 (0.019)			
Aug-2010	0.0043 (0.027)	0.0025 (0.027)	0.0025 (0.020)			
Sept-2010	0.0040 (0.026)	0.0032 (0.026)	0.0032 (0.020)			
Oct-2010	0.0033 (0.025)	0.0037 (0.025)	0.0037 (0.019)	0.0045 (0.024)	0.0044 (0.022)	0.0044 (0.0052)
Nov-2010	0.0026 (0.021)	0.0051 (0.021)	0.0051 (0.016)	0.0038 (0.020)	0.0040 (0.018)	0.0040 (0.0042)
Dec-2010	0.0020 (0.018)	0.0056 (0.018)	0.0056 (0.011)	0.0032 (0.017)	0.0033 (0.016)	0.0033 (0.0035)
Jan-2011	0.0021 (0.017)	0.0059 (0.017)	0.0059 (0.010)	0.0034 (0.016)	0.0034 (0.015)	0.0034 (0.0033)
Feb-2011	0.0018 (0.017)	0.0057 (0.017)	0.0057 (0.010)	0.0030 (0.015)	0.0031 (0.014)	0.0031 (0.0034)
Mar-2011	0.0012 (0.017)	0.0054 (0.017)	0.0054 (0.0095)	0.0024 (0.015)	0.0026 (0.014)	0.0026 (0.0036)
Apr-2011	0.0011 (0.016)	0.0066 (0.016)	0.0066 (0.0087)	0.0024 (0.013)	0.0033 (0.013)	0.0033 (0.0028)
May-2011	0.001 (0.016)	0.0068 (0.016)	0.0068 (0.0084)	0.0023 (0.013)	0.0035 (0.013)	0.0035 (0.0027)
Jun-2011	0.00043 (0.015)	0.0063 (0.016)	0.0063 (0.0082)	0.0018 (0.013)	0.0029 (0.013)	0.0029 (0.0027)
Jul-2011	0.00042 (0.015)	0.0066 (0.015)	0.0066 (0.0080)	0.0018 (0.013)	0.0030 (0.012)	0.0030 (0.0027)
Aug-2011	0.00045 (0.015)	0.0068 (0.015)	0.0068 (0.0079)	0.0017 (0.013)	0.0031 (0.012)	0.0031 (0.0026)
Sept-2011	0.00046 (0.015)	0.0067 (0.015)	0.0067 (0.0079)	0.0018 (0.012)	0.0032 (0.012)	0.0032 (0.0026)
Oct-2011	0.00031 (0.015)	0.0066 (0.015)	0.0066 (0.0075)	0.0016 (0.012)	0.0030 (0.012)	0.0030 (0.0025)
Nov-2011	0.00019 (0.014)	0.007 (0.014)	0.007 (0.0067)	0.0014 (0.011)	0.0029 (0.011)	0.0029 (0.0024)
Dec-2011	0.0001 (0.013)	0.0070 (0.013)	0.0070 (0.0061)	0.0011 (0.011)	0.0027 (0.010)	0.0027 (0.0022)
Jan-2012	0.000010 (0.012)	0.0071 (0.012)	0.0071 (0.0057)	0.00094 (0.010)	0.0025 (0.010)	0.0025 (0.0022)
Feb-2012	-0.00025 (0.012)	0.0068 (0.012)	0.0068 (0.0057)	0.00064 (0.010)	0.0023 (0.010)	0.0023 (0.0022)
Mar-2012	0.00017 (0.013)	0.0076 (0.013)	0.0076 (0.0056)	0.0011 (0.010)	0.0028 (0.010)	0.0028 (0.0022)
Apr-2012	-0.00092 (0.013)	0.007 (0.013)	0.007 (0.0052)	-0.000032 (0.010)	0.002 (0.010)	0.002 (0.0021)
Estimation window	60-day	60-day	60-day	200-day	200-day	200-day
Covariates	No	Yes	35 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. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a t-test of the null hypothesis of no effect. Covariates are the S&P 500, the 10-year U.S. Treasury rate, 3-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract.

Table 6: **Natural gas SQ test:** Percentile rank of the event-day residual among the estimation window residuals.

Contract (<i>i</i>)	(1)	(2)	(3)	(4)	(5)	(6)
Jul-2010	56	48	48			
Aug-2010	56	53	53			
Sept-2010	56	55	55			
Oct-2010	55	56	56	58	61	61
Nov-2010	56	58	58	59	61	61
Dec-2010	54	60	60	59	60	60
Jan-2011	55	62	62	60	60	60
Feb-2011	52	60	60	59	60	60
Mar-2011	52	59	59	58	58	58
Apr-2011	53	65	65	60	63	63
May-2011	53	66	66	61	64	64
Jun-2011	51	66	66	58	62	62
Jul-2011	51	66	66	58	63	63
Aug-2011	51	66	66	58	64	64
Sept-2011	51	66	66	59	65	65
Oct-2011	51	66	66	59	65	65
Nov-2011	50	69	69	58	65	65
Dec-2011	50	72	72	57	62	62
Jan-2012	50	73	73	57	61	61
Feb-2012	49	73	73	57	60	60
Mar-2012	49	73	73	56	61	61
Apr-2012	47	72	72	50	57	57
Estimation window	60-day	60-day	60-day	200-day	200-day	200-day
Covariates	No	Yes	Yes	No	Yes	Yes
GARCH error	No	No	Yes	No	No	Yes

Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a null hypothesis of no effect. The 60-day estimation window cannot distinguish between significance levels smaller than 1.67%. Covariates are the S&P 500, the 10-year U.S. Treasury rate, 3-month LIBOR, Baltic Dry Index, CRB Index, and the corresponding West Texas Intermediate (crude oil) futures contract.

Table 7: 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>				
Coal	1.8** (0.92)	1.7** (0.79)	1.7** (0.97)	1.7** (0.92)
Natural gas	0.54 (1.0)	0.63 (0.85)	0.30 (0.97)	0.32 (0.85)
<i>SQ tests (percentile rank)</i>				
Coal	96**	97**	96**	97**
Natural gas	67	76	66	67

The “unweighted” columns take the average of the standardized residuals across each day’s contracts, while the “weighted” columns use the efficient weighting scheme. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) levels for a null hypothesis of a weakly negative effect for coal and of no effect for natural gas.

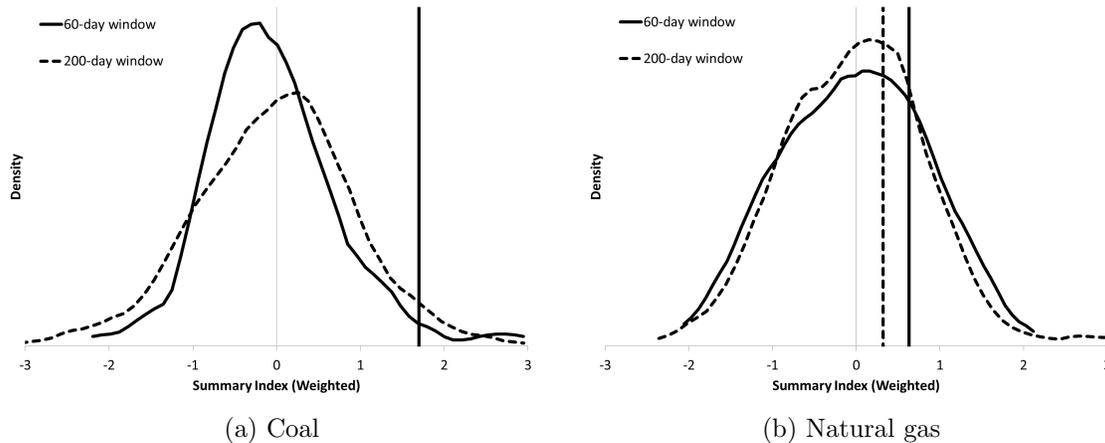


Figure 10: Kernel density plot of the weighted summary index for the estimation window under specifications (2) and (5). The summary index for Monday April 26 is indicated by a vertical line. The vertical lines for the 60-day and 200-day estimation windows overlap in the case of coal.

Roundup” regularly featured the Senate bill’s progress, demonstrating the bill’s perceived importance to coal markets. However, 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.” The logic in this series of reports implies that natural gas prices and coal prices should increase together. In contrast, our theory leads us to expect much smaller—and potentially negative—effects in the natural gas market because that commodity is less directly affected by a carbon price and would even benefit via substitution for coal. In fact, we do see an abnormally large price increase in the coal market accompanied by rather small movements in the natural gas market. Not only did our theory predict that the coal price would rise on the event day, but it outperforms contemporary market analysis in expecting only a small comovement in the natural gas market.

Our preferred specifications include covariates because they make confounding events less likely: identification is maintained even if other news appeared on the event day so long as this other news was not specific to the coal market. For instance, identification is maintained if strong German power demand (or brighter sentiment about the German economy) also affected international borrowing costs, shipping indices, commodity indices, or U.S. equity markets. In particular, the cost of shipping coal is a major component of the Baltic Dry Index, so we might expect expectations of greater transatlantic coal trade to move it. The downside to including covariates is that they can bias our estimates towards zero. If the collapse of the Senate’s climate effort also affected one of the covariates, then these specifications absorb part of the event’s effect into the explained portion of returns. However, the Senate bill’s primary near-term policy was a cap-and-trade program for the electric power sector, and its free permit allocations dampened the expected effect on retail electricity prices (EIA, 2010b). Given these facts, we do not expect that the event affected global oil markets or equity market indices. Including covariates reduces the influence of other news that might have appeared on the event day.

5.2 Magnitude of the event’s effect on coal markets

Figure 11a translates the event’s estimated effect on coal futures’ returns into its effect on futures prices. The bill’s collapse increased the nearest-dated contracts’ prices by around \$1.30 per ton of coal, with the 90% confidence interval extending from around \$0.10 to \$2.50 per ton of coal.³⁰ The effects are smaller for longer-dated contracts, with several having a

³⁰All confidence interval calculations assume that the noise introduced through β and ϵ is trivial in comparison to the noise in γ^{Mon} . Recall from Section 3 that while the Central Limit Theorem fails to hold, we cannot reject the null hypothesis of normality for the six nearest-dated coal contracts in the 60-day estimation window.

best estimate around \$1.00 per ton and the longest-dated contract having a best estimate just above \$0.60 per ton. Because the pre-event probability of passage was less than 1 and the post-event probability of passage was greater than 0, these numbers almost surely underestimate the total effect of the Senate bill. For instance, there is suggestive 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.

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 per ton for plants with less attractive outside options and by as little as \$0.09 per ton for plants with better outside options. Our central estimates for the event's effect cover a similar range. This suggests that anticipation of the Senate bill distorted coal markets at least as strongly as did railroads' market power in the wake of the Clean Air Act. Further, Cicala (2012) estimates that electricity sector deregulation lowered the price paid by coal-fired generators by \$0.25 per million Btu; our central estimates translate into between \$0.05 and \$0.10 per million Btu, depending on the state.³¹ The estimated effect of divestiture is therefore about twice as great as our estimate. It is just above our event's 90% confidence interval and around the effect obtained from combining the estimates for Friday with the event day. If, as is likely, 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 larger than that of divestiture.

Figure 11b translates these changes in futures prices into the cost of coal-fired electricity.³² Based on the nearest-dated contracts, the event increased the cost of coal-fired generation by just over 0.05 cents per kWh, with the 90% confidence interval extending from around 0.004 cents per kWh to around 0.10 cents per kWh. These changes are small compared to the average U.S. residential electricity price of 11.75 cents per kWh in April of 2010 (EIA, 2010a).

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 borne by coal consumers. Figure 11c shows that the weekend collapse of the Senate bill acted like taxing carbon dioxide 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 CO₂.³³ The projected 2013 allowance prices under the Senate's bill

³¹Calculations use the average heat content of coal consumed by each state's electricity generators (EIA, 2010a, Table C1).

³²Calculations use the CAPP contract's minimum requirement of 12,000 Btu per pound, and they assume the average tested heat rate of 10,128 Btu/kWh for coal-fired generators in 2010 (EIA, 2013, Table 8.2). CAPP futures prices can affect electricity generators' input costs when they are used to inform mark-to-market valuations used in dispatch decisions, when they influence negotiations for long-term contracts, and when they influence prices paid under long-term contracts.

³³Calculations use use 2.29 t CO₂ per short ton of coal, which follows from the CAPP 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

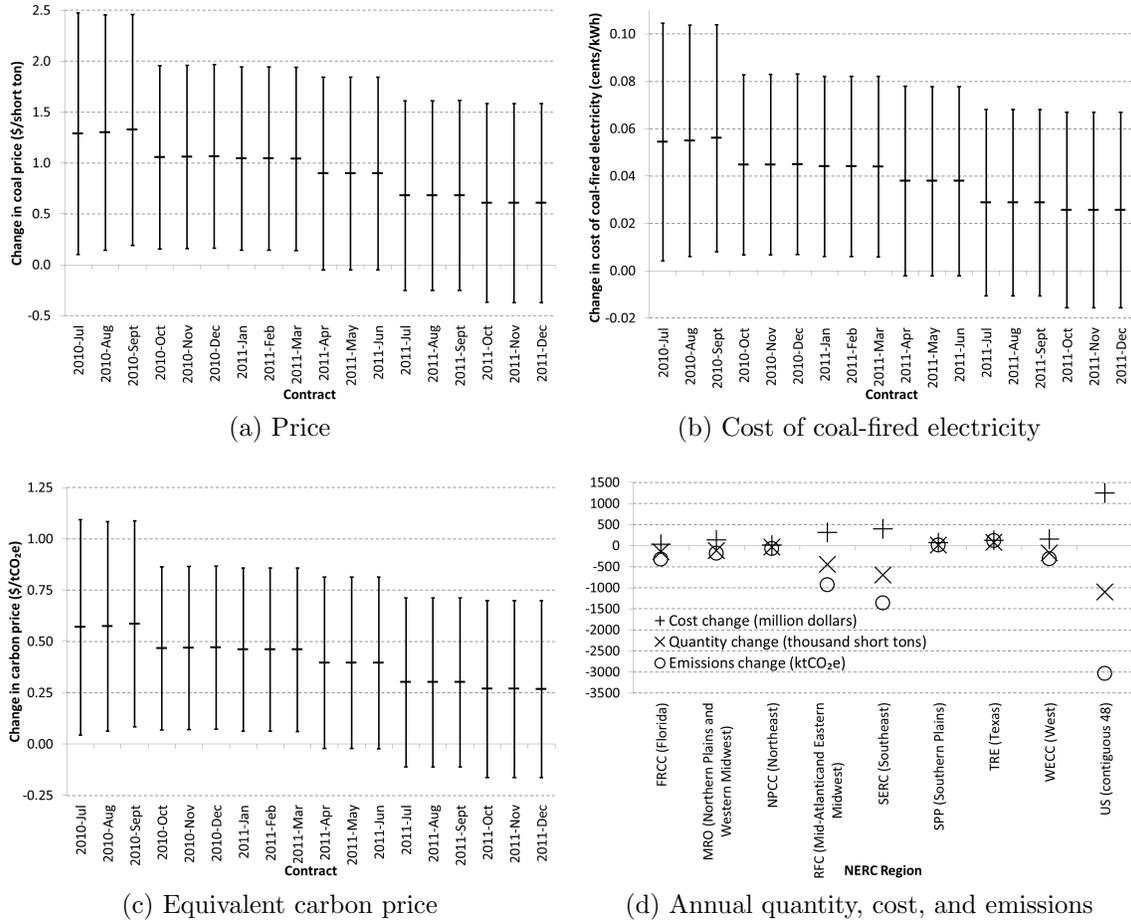


Figure 11: The estimated effect of the Senate bill’s collapse on coal futures prices and on the cost of coal-fired electricity generation. Also the carbon price that would equal the change in the futures price, and the change in coal consumption, purchase cost, and emissions if the nearest coal contract’s price change were maintained for all of 2010. Error bars give the 90% confidence interval. All plots use specification (2). “NERC Region” refers to the North American Electric Reliability Corporation’s regional reliability councils.

were \$14–25 per metric ton of CO₂ (CBO, 2010; EIA, 2010b; EPA, 2010). Based on our central estimates, anticipating the bill’s passage acted like subsidizing coal’s carbon dioxide emissions by at least 2.4–4.2% of the expected 2013 allowance prices. Further, anticipating the Senate bill’s passage probably acted like subsidizing coal’s carbon emissions by at least 12% of the U.S. government’s lower estimate of the social cost of carbon (Greenstone et al., 2011). Including the coefficient for Friday April 23 in the event effect would nearly double these results, and the anticipation effect from the bill’s actual passage would likely have been greater still.

Finally, Figure 11d plots the change in the U.S. electric sector’s coal consumption, its cost of purchasing coal, and its emissions from coal-fired plants if the event’s price change were to hold for an entire year.³⁴ Aggregate coal consumption would have fallen by around 1 million tons nationally, with the largest decreases in the Midwest and the Southeast. Meanwhile, the aggregate cost of utilities’ coal purchases would have increased by around \$1.3 billion nationally, and carbon dioxide emissions from coal-fired plants would have fallen by around 3 million metric tons nationally. Based on the event’s effect, the proposal to cap greenhouse gas emissions from 2013 may have increased greenhouse gas emissions from coal-fired plants by at least 0.05% in 2009–2012, and this number would have grown as the bill became more likely to become law.

5.3 Implications of natural gas futures for the green paradox

Even though the Senate bill’s potential passage increased greenhouse gas emissions from coal-fired power plants, our theoretical results tell us that the effect on aggregate emissions depends on how natural gas responded. It remains possible that the Senate bill was decreasing aggregate near-term emissions and would even have increased aggregate cumulative emissions. By Proposition 2, these cases both require that the Senate bill decreased the quantity of natural gas in the period prior to its implementation, meaning that our event (the bill’s collapse) would have decreased its futures price.

Analyses of the Senate’s American Power Act differed on how it would affect natural gas consumption in 2013 and beyond. The U.S. Energy Information Administration and Environmental Protection Agency used different modeling frameworks to conclude that the bill would probably have increased natural gas consumption in its first years (EIA, 2010b; EPA, 2010). This outcome would imply a decrease in natural gas consumption prior to 2013. However, analyses by two nongovernmental organizations found that the bill’s implementation would have decreased natural gas consumption (ClimateWorks Foundation, 2010;

Btu of coal (DOE, 2007, Table 1.C.3).

³⁴Calculations use the nearest-dated contract’s central estimate. Quantity and emission calculations use price elasticities for each NERC region from EIA (2012b, Table 3). Quantity and cost calculations use the electric power sector’s coal consumption by state in 2010 (EIA, 2012a, Table 26). Emission calculations use the average heat content of coal burned by each state’s electricity generators (EIA, 2010a, Table C1).

Houser et al., 2010). This outcome would be generally consistent with an increase in natural gas consumption prior to 2013.

The results in Section 4 suggest that the event increased natural gas prices. In that case, the Senate bill was increasing near-term natural gas consumption and would have decreased natural gas consumption upon implementation in 2013. These results are consistent with the unshaded region in Figure 3. A weak green paradox therefore did appear to hold: the Senate bill's anticipated carbon price was increasing near-term emissions. Moreover, a strong green paradox was unlikely to have occurred if the bill had passed (Corollary 1): the Senate bill would indeed have decreased cumulative emissions and, for a sufficiently low discount rate, would have decreased cumulative damages. However, since the bill did not end up passing, the legislative process itself potentially generated a strong green paradox. It increased emissions in 2009–2010 without establishing the carbon price that would have eventually offset that initial increase.

6 Conclusions

We analytically demonstrated how anticipated climate policies affect earlier futures prices via the direct, substitution, anticipation, and risk effects. Our model generates intertemporal linkages because extraction costs depend on cumulative extraction. The standard “green paradox” story arises naturally in a one-commodity setting: the anticipated policy increases emissions in advance of the policy. However, we found that in a two-commodity setting, substitution effects can generate patterns that contradict this standard story. Under particular combinations of elasticities and relative emission intensities, the anticipated policy can decrease emissions in advance of the policy and can actually increase cumulative emissions. In all cases that reject the weak green paradox, consumption of the low-emission product decreases in advance of the policy's implementation.

We then empirically identified the anticipation effect in coal futures markets: as predicted by theory, coal futures spiked upon the weekend collapse of the Senate's 2010 climate effort. This result is robust to the choice of estimation window, to the inclusion of covariates, to the form of the error term, and to tests of more restrictive hypotheses. While the possibility of confounding effects can never be completely ruled out, this result accords with predictions of a green paradox in coal markets. The effect on coal prices is of similar magnitude to previous findings about market power and electricity deregulation.

While the Senate's bill would have decreased future emissions from coal-fired plants, it was actually increasing them in advance of the cap-and-trade program's implementation in 2013. Because the bill's collapse probably did not decrease natural gas futures prices, the theoretically necessary condition for the bill to have decreased contemporary emissions across both products (or to have increased cumulative emissions across both products) was not met. The Senate's bill was therefore increasing aggregate emissions from 2009–2012. If

it had passed, it would have decreased cumulative emissions, but its collapse means that the overall effect of the unsuccessful policy process potentially increased cumulative emissions.

While the relevance of green paradox effects has been disputed, we have found evidence of their presence in the market for the largest contributor to climate change. The existence of the anticipation effect has implications for policy. The empirical results suggest that the proposed Senate bill, which aimed to decrease U.S. emissions to 4.75% below 2005 levels in 2013, was increasing U.S. emissions by at least 0.05% in the years leading up to 2013—and potentially by much more had the bill passed in 2010. The anticipation effect is a form of intertemporal leakage whereby the future policy regime displaces some emissions towards earlier periods. Once the capped period is reached, all further intertemporal leakage occurs between capped markets. Prior to reaching the capped period, leakage between the present and the future is an analogue of leakage between capped economies and uncapped 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 can be valuable for giving expectations, politics, 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 over \$0.50 per ton even when the bill was still uncertain to pass.

More generally, we have seen that markets are distorted by the suggestion of regulation. Major legislative proposals take time to enact, and even once enacted, policies like health care reform, financial reform, and spending cuts intentionally delay implementation by months or even years. The suggestion of regulation is likely to be particularly acute in energy markets for three reasons. First, the primary actors in these markets tend to have intertemporally-linked production costs and 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, as scientific knowledge of their effects advances, and as damages increase. This is especially true of climate change: all major actors surely expect to be regulated at some point in the not-too-distant future.

This expectation of future climate regulation must have affected energy markets for many years. 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 has probably generated more emissions than simply not pricing carbon at all. Which policy is worse in the end depends on what that carbon price

looks like, on when it occurs, and on how it shapes expectations beforehand.

7 Appendix: Proofs

7.1 Proof of Proposition 1

We seek expressions for $\partial q_1^{i*}/\partial\tau$ and $\partial q_2^{i*}/\partial\tau$, for $i \in \{H, L\}$. The Jacobian of the system of four first-order conditions in equations (2) and (3) is the Hessian H of the maximization problem:

$$H = \begin{bmatrix} U_{jj} - C_{jj} - \beta C_{jj} & -\beta C_{jj} & U_{ij} & 0 \\ -C_{jj} & U_{jj} - C_{jj} & 0 & U_{ij} \\ U_{ij} & 0 & U_{ii} - C_{ii} - \beta C_{ii} & -\beta C_{ii} \\ 0 & U_{ij} & -C_{ii} & U_{ii} - C_{ii} \end{bmatrix} \equiv \begin{bmatrix} A_j & B \\ F & A_i \end{bmatrix},$$

where, without loss of generality, the first two rows correspond to the first-order conditions for product j in periods 1 and 2 and the last two rows correspond to the first-order conditions for product i . Similarly, the first two columns correspond to partially differentiating the system with respect to the period 1 and 2 quantities of product j and the last two columns correspond to partially differentiating the system with respect to period 1 and 2 quantities of product i . The concavity of the objective ensures that the determinant of H is positive, which is a sufficient condition for a global maximum. The right-most expression rewrites H as a block matrix, where A_j , B , F , and A_i are 2×2 matrices. If commodity j were the only commodity in the market, then matrix A_j would describe the whole system. By the concavity of the corresponding objective, we know $\det(A_j) > 0$ (as is easy to confirm algebraically), which is also a sufficient condition for a global maximum in that one-commodity market.

By the Implicit Function Theorem and Cramer's Rule, we have:

$$\frac{\partial q_t^{i*}}{\partial\tau} = -\frac{\det(H_t)}{\det(H)},$$

where H_t is the matrix H with column $2+t$ replaced by the partial derivative of the first-order conditions with respect to τ . The numerator is:

$$\det(H_t) = \det\left(\begin{bmatrix} A_j & B_t \\ F & D_t \end{bmatrix}\right) = \det(A_j) \det(D_t - FA_j^{-1}B_t).$$

Column t of the 2×2 matrix B_t has a zero in the first row and $-e^j$ in the second row; its other column is the same as in B . Column t of the 2×2 matrix D_t has a zero in the first row and $-e^i$ in the second row; its other column is the same as in A_i . We have reduced the sign of the comparative static to a single determinant:

$$\frac{\partial q_t^{i*}}{\partial\tau} \propto -\det(D_t - FA_j^{-1}B_t). \quad (6)$$

Finally, inverting A_j yields:

$$A_j^{-1} = \frac{1}{\det(A_j)} \begin{bmatrix} U_{jj} - C_{jj} & \beta C_{jj} \\ C_{jj} & U_{jj} - C_{jj} - \beta C_{jj} \end{bmatrix}.$$

Begin with the comparative static for period 1. By algebraic manipulation, we find that:

$$\frac{\partial q_1^{i*}}{\partial \tau} \propto e^i \beta \{ \det(A_j) C_{ii} + (U_{ij})^2 C_{jj} \} - e^j \beta U_{ij} \{ C_{ii} [U_{jj} - C_{jj} - \beta C_{jj}] + C_{jj} [U_{ii} - C_{ii}] \}. \quad (7)$$

This yields:

$$\frac{\partial q_1^{i*}}{\partial \tau} < 0 \Leftrightarrow \frac{e^j}{e^i} > \frac{\det(A_j) C_{ii} + (U_{ij})^2 C_{jj}}{U_{ij} \{ C_{ii} [U_{jj} - C_{jj} - \beta C_{jj}] + C_{jj} [U_{ii} - C_{ii}] \}} \equiv k_1^i. \quad (8)$$

Both the numerator and denominator are positive. As $U_{ij} \rightarrow 0$, the right-hand fraction is greater than 1. For $U_{ij} < 0$, the right-hand fraction is less than 1 either as $U_{ii} \rightarrow -\infty$ or as U_{jj} and C_{jj} both $\rightarrow 0$ with $U_{ij} < U_{jj}$.

Now consider the comparative statics for the second-period quantity of product i and for the total quantity of product i . The comparative static for the period 2 quantity is

$$\begin{aligned} \frac{\partial q_2^{i*}}{\partial \tau} \propto e^i \{ \det(A_j) [U_{ii} - C_{ii} - \beta C_{ii}] - (U_{ij})^2 [U_{jj} - C_{jj}] \} \\ - e^j U_{ij} \{ [U_{jj} - C_{jj} - \beta C_{jj}] [U_{ii} - C_{ii} - \beta C_{ii}] + \beta C_{jj} C_{ii} - (U_{ij})^2 \}. \end{aligned} \quad (9)$$

This yields:

$$\frac{\partial q_2^{i*}}{\partial \tau} > 0 \Leftrightarrow \frac{e^j}{e^i} > \frac{\det(A_j) [U_{ii} - C_{ii} - \beta C_{ii}] - (U_{ij})^2 [U_{jj} - C_{jj}]}{U_{ij} \{ [U_{jj} - C_{jj} - \beta C_{jj}] [U_{ii} - C_{ii} - \beta C_{ii}] + \beta C_{jj} C_{ii} - (U_{ij})^2 \}} \equiv k_2^i. \quad (10)$$

Both the numerator and denominator are negative. The cumulative change in the quantity of product i is

$$\begin{aligned} \frac{\partial [q_1^{i*} + q_2^{i*}]}{\partial \tau} \propto e^i \{ \det(A_j) [U_{ii} - C_{ii}] - (U_{ij})^2 [U_{jj} - C_{jj} - \beta C_{jj}] \} \\ - e^j U_{ij} \{ [U_{jj} - C_{jj} - \beta C_{jj}] [U_{ii} - C_{ii}] + \beta C_{jj} U_{ii} - (U_{ij})^2 \}. \end{aligned} \quad (11)$$

For feasible U_{ij} , the condition for an increasing quantity is:

$$\begin{aligned} \frac{\partial [q_1^{i*} + q_2^{i*}]}{\partial \tau} > 0 \Leftrightarrow \frac{e^j}{e^i} > \frac{\det(A_j) [U_{ii} - C_{ii}] - (U_{ij})^2 [U_{jj} - C_{jj} - \beta C_{jj}]}{U_{ij} \{ [U_{jj} - C_{jj} - \beta C_{jj}] [U_{ii} - C_{ii}] + \beta C_{jj} U_{ii} - (U_{ij})^2 \}} \\ \equiv k_3^i. \end{aligned} \quad (12)$$

Both the numerator and denominator are positive.

Equations (8), (10), and (12) define rays from the origin in (e^i, e^j) space with slopes k_1^i , k_2^i , and k_3^i . We now consider possible relationships between the three rays. First, assume $k_2^i > k_1^i$. In that case, the region between the two rays has $\partial q_2^{i*}/\partial\tau < 0$ and $\partial q_1^{i*}/\partial\tau < 0$. It must therefore be the case that $\partial[q_1^{i*} + q_2^{i*}]/\partial\tau < 0$ in this region, which implies that $k_3^i > k_2^i$. Second, assume $k_2^i < k_1^i$. In that case, the region between the two rays has $\partial q_2^{i*}/\partial\tau > 0$ and $\partial q_1^{i*}/\partial\tau > 0$. It must therefore be the case that $\partial[q_1^{i*} + q_2^{i*}]/\partial\tau > 0$ in this region, which implies that $k_3^i < k_2^i$. Under either assumption, the ray with slope k_2^i is between the other two rays.

We have found that $k_2^i > k_3^i \Leftrightarrow k_1^i > k_2^i$. Algebraic manipulations show that $k_1^i > k_2^i$ if and only if

$$0 < -(U_{ij})^4 - \det(A_i)\det(A_j) + (U_{ij})^2 \left\{ [U_{jj} - C_{jj} - \beta C_{jj}][U_{ii} - C_{ii} - \beta C_{ii}] + 2\beta C_{jj}C_{ii} + [U_{jj} - C_{jj}][U_{ii} - C_{ii}] \right\}.$$

This is a quadratic in $(U_{ij})^2$. It is concave, and it is increasing at its negative y-intercept. To be satisfied at some feasible U_{ij} , the inequality requires that the parabola have a root in $(0, U_{ii}U_{jj})$. The smallest root is below $U_{ii}U_{jj}$ if and only if

$$2U_{ii}U_{jj} > [U_{jj} - C_{jj} - \beta C_{jj}][U_{ii} - C_{ii} - \beta C_{ii}] + 2\beta C_{jj}C_{ii} + [U_{jj} - C_{jj}][U_{ii} - C_{ii}] - \left(\{\beta^2 C_{ii}C_{jj} - \beta C_{jj}[U_{ii} - C_{ii}] - \beta C_{ii}[U_{jj} - C_{jj}]\}^2 + 4\beta\{\beta C_{ii}C_{jj}\}^2 - 4\beta\{C_{ii}(U_{jj} - C_{jj}) - C_{jj}(U_{ii} - C_{ii})\}^2 \right)^{1/2}.$$

Therefore the smallest root is below $U_{ii}U_{jj}$ only if

$$2U_{ii}U_{jj} > [U_{jj} - C_{jj} - \beta C_{jj}][U_{ii} - C_{ii} - \beta C_{ii}] + 2\beta C_{jj}C_{ii} + [U_{jj} - C_{jj}][U_{ii} - C_{ii}] - \left(\{\beta^2 C_{ii}C_{jj} - \beta C_{jj}[U_{ii} - C_{ii}] - \beta C_{ii}[U_{jj} - C_{jj}]\}^2 + 4\beta\{\beta C_{ii}C_{jj}\}^2 \right)^{1/2}.$$

Applying the triangle inequality and simplifying, the smallest root is below $U_{ii}U_{jj}$ only if

$$0 > [C_{jj} + \beta C_{jj}][C_{ii} + \beta C_{ii}] - U_{jj}[C_{ii} + \beta C_{ii}] - U_{ii}[C_{jj} + \beta C_{jj}] + C_{jj}C_{ii} - [1 + \beta]U_{jj}C_{ii} - [1 + \beta]U_{ii}C_{jj} + [4 - 2\beta^{1/2} - \beta]\beta C_{jj}C_{ii}.$$

But every term on the right-hand side is positive. Contradiction. The smallest root must instead be above $U_{ii}U_{jj}$. Therefore $k_2^i > k_1^i$, which in turn implies $k_3^i > k_2^i$.

7.2 Proof of Proposition 2

The missing factor of proportionality in each of the comparative static expressions in equations (7), (9), and (11) is $\det(A_j)^2/\det(H)$. Through summation and manipulation of previous results, we find:

$$\begin{aligned} \frac{\partial E_1}{\partial \tau} &\propto \det(A_j)^2 (e^i)^2 [\det(A_j) C_{ii} + (U_{ij})^2 C_{jj}] \\ &\quad + \det(A_i)^2 (e^j)^2 [\det(A_i) C_{jj} + (U_{ij})^2 C_{ii}] \\ &\quad - [\det(A_j)^2 + \det(A_i)^2] e^i e^j U_{ij} [C_{ii} [U_{jj} - C_{jj} - \beta C_{jj}] + C_{jj} [U_{ii} - C_{ii}]], \\ \frac{\partial E_2}{\partial \tau} &\propto \det(A_j)^2 (e^i)^2 [\det(A_j) [U_{ii} - C_{ii} - \beta C_{ii}] - (U_{ij})^2 [U_{jj} - C_{jj}]] \\ &\quad + \det(A_i)^2 (e^j)^2 [\det(A_i) [U_{jj} - C_{jj} - \beta C_{jj}] - (U_{ij})^2 [U_{ii} - C_{ii}]] \\ &\quad - [\det(A_i)^2 + \det(A_j)^2] e^i e^j U_{ij} [[U_{jj} - C_{jj} - \beta C_{jj}] [U_{ii} - C_{ii} - \beta C_{ii}] + \beta C_{jj} C_{ii} - (U_{ij})^2], \\ \frac{\partial [E_1 + E_2]}{\partial \tau} &\propto \det(A_j)^2 (e^i)^2 [\det(A_j) [U_{ii} - C_{ii}] - (U_{ij})^2 [U_{jj} - C_{jj} - \beta C_{jj}]] \\ &\quad + \det(A_i)^2 (e^j)^2 [\det(A_i) [U_{jj} - C_{jj}] - (U_{ij})^2 [U_{ii} - C_{ii} - \beta C_{ii}]] \\ &\quad - [\det(A_i)^2 + \det(A_j)^2] e^i e^j U_{ij} [[U_{jj} - C_{jj} - \beta C_{jj}] [U_{ii} - C_{ii}] + \beta C_{jj} U_{ii} - (U_{ij})^2]. \end{aligned}$$

Each of these expressions defines a quadratic in e^i . Each quadratic has real roots only if $e^j/e^i > k_t^i$ for the appropriate period t (treating $t = 3$ as the cumulative quantity) and $\det(A_i)$ is sufficiently small relative to $\det(A_j)$. $\partial E_1/\partial \tau$ is convex and is decreasing at its positive y-intercept. Define its smaller root as $e^j m_1$ and its larger root as $e^j m_6$. $\partial E_2/\partial \tau$ and $\partial [E_1 + E_2]/\partial \tau$ are concave, and they are increasing at their negative y-intercepts. Define their smaller roots as $e^j m_2$ and $e^j m_3$, respectively, and their larger roots as $e^j m_4$ and $e^j m_5$. In each of these cases, m_k is a term that does not include e^i or e^j .

There are real roots with $e^L/e^H > k_t^H$ only if $\det(A_H)$ is small relative to $\det(A_L)$. By definition, $e^L/e^H < 1$. By Proposition 1, we have $k_t^H < 1$ only if U_{HH} is sufficiently large in magnitude or if U_{LL} and C_{LL} are sufficiently small in magnitude. But large U_{HH} implies large $\det(A_H)$, and small U_{LL} and C_{LL} imply small $\det(A_L)$. Any cases with real roots must therefore occur with $e^H/e^L > k_t^L$.

We therefore have real roots in e^L only if $\det(A_L)$ is small relative to $\det(A_H)$. Consider the roots as $\det(A_L) \rightarrow 0$. The dominant terms in each quadratic become those associated with q^L . The larger root occurs at e^L slightly greater than where the q^L terms go to 0. The q^L terms in, for instance, $\partial E_1/\partial \tau$ go to 0 at $e^L = e^H/k_1^L$. Because $k_1^L < k_2^L < k_3^L$, then for $\det(A_L)$ sufficiently small, it is the case that $1/m_6 < 1/m_5 < 1/m_4$.

Assume $1/m_1 < 1/m_4$. Then there exists e^H/e^L such that $\partial E_2/\partial \tau > 0$ and $\partial E_1/\partial \tau > 0$ while $\partial [E_1 + E_2]/\partial \tau < 0$. But this is impossible. Therefore $1/m_1 > 1/m_4$.

Each quadratic's smaller root ($e^H m_1$, $e^H m_2$, and $e^H m_3$) occurs when e^L becomes small relative to $\det(A_L)$. As $\det(A_L) \rightarrow 0$, e^L small relative to $\det(A_L)$ implies that the smaller

root converges to the y-intercept (i.e., to the $(e^H)^2$ term). The intercept for $\partial E_1/\partial\tau$ is closer to zero than is the intercept for $\partial[E_1 + E_2]/\partial\tau$. Therefore, for $\det(A_L)$ sufficiently small, $1/m_1 > 1/m_3$.

Now assume that $1/m_2 > 1/m_1$. For $e^H/e^L \in (1/m_1, 1/m_2)$, we have $\partial E_1/\partial\tau > 0$, $\partial E_2/\partial\tau > 0$, and $\partial[E_1 + E_2]/\partial\tau < 0$. Contradiction. Instead assume that $1/m_2 < 1/m_3$. For $e^H/e^L \in (1/m_2, 1/m_3)$, we have $\partial E_1/\partial\tau < 0$, $\partial E_2/\partial\tau < 0$, and $\partial[E_1 + E_2]/\partial\tau > 0$. Again we have a contradiction. The only remaining option is $1/m_1 > 1/m_2 > 1/m_3$. Further, by the definition of m_3 and m_4 , we have $1/m_3 > 1/m_4$.

Define $h_i \equiv 1/m_i$. Then for $\det(A_L)$ sufficiently small, we have shown that $h_6 < h_5 < h_4 < h_3 < h_2 < h_1$. To complete the statement of the proposition, note that $\det(A_i)$ is greater when U_{ii} and C_{ii} are larger in magnitude.

7.3 Proof of Corollary 1

The cases with a weak green paradox follow directly from Proposition 2. A strong green paradox never occurs when emissions decrease in every period (i.e., when $e^H/e^L \in [h_1, h_2] \cup [h_5, h_6]$). Let \hat{E}_t and \hat{D}_t indicate emissions and damages without a tax and \bar{E}_t and \bar{D}_t indicate emissions and damages with an anticipated second-period tax. A strong green paradox occurs if and only if $\hat{D}_1 + \frac{1}{1+r}\hat{D}_2 < \bar{D}_1 + \frac{1}{1+r}\bar{D}_2$. For linear damages, this holds if and only if $\frac{1}{1+r}[\hat{E}_2 - \bar{E}_2] < \bar{E}_1 - \hat{E}_1$. When emissions decrease (increase) in the first period but increase (decrease) in the second period, both sides are negative (positive). Therefore a strong green paradox holds if and only if $r < (>) \frac{\hat{E}_2 - \bar{E}_2}{\bar{E}_1 - \hat{E}_1} - 1 \equiv \gamma$. Proposition 2 implies that the second-period change in emissions dominates the first-period change in emissions for $e^H/e^L \notin [h_1, h_6] \setminus (h_3, h_4)$, in which case $\gamma > 0$. When $e^H/e^L \in [h_2, h_3] \cup [h_4, h_5]$, a strong green paradox holds if and only if $r < \gamma$, where $\gamma < 0$.

7.4 Proof of Proposition 3

Consider a change from distribution i for τ_s to distribution j , where distribution j first-order stochastically dominates distribution i and both have support on the positive real numbers. Let \mathbb{Q}_0 indicate the degenerate distribution with the tax certainly equal to 0. Approximating $g(S_T, \tau_s)$ by a first-order Taylor expansion around $(\mathbb{E}_t^{\mathbb{Q}_0}[S_T], 0)$, the change

in the time t futures price for delivery at time T (where $t < T < s$) is:

$$\begin{aligned}
F_{t,T}^j - F_{t,T}^i &= \mathbb{E}_t^{\mathbb{Q}^j} [g(S_T, \tau_s)] - \mathbb{E}_t^{\mathbb{Q}^i} [g(S_T, \tau_s)] \\
&\approx \mathbb{E}_t^{\mathbb{Q}^j} \left[g(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0) + \frac{\partial g(S, \tau)}{\partial S} \Big|_{(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0)} (S_T - \mathbb{E}_t^{\mathbb{Q}^0} [S_T]) + \frac{\partial g(S, \tau)}{\partial \tau} \Big|_{(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0)} \tau_s \right] \\
&\quad - \mathbb{E}_t^{\mathbb{Q}^i} \left[g(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0) + \frac{\partial g(S, \tau)}{\partial S} \Big|_{(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0)} (S_T - \mathbb{E}_t^{\mathbb{Q}^0} [S_T]) + \frac{\partial g(S, \tau)}{\partial \tau} \Big|_{(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0)} \tau_s \right] \\
&= \left(\mathbb{E}_t^{\mathbb{Q}^j} [S_T] - \mathbb{E}_t^{\mathbb{Q}^i} [S_T] \right) + \frac{\partial g(S, \tau)}{\partial \tau} \Big|_{(\mathbb{E}_t^{\mathbb{Q}^0} [S_T], 0)} \left(\mathbb{E}_t^{\mathbb{Q}^j} [\tau_s] - \mathbb{E}_t^{\mathbb{Q}^i} [\tau_s] \right), \tag{13}
\end{aligned}$$

where expectations are taken jointly over τ_s and S_T and we recognize that $g(S, 0) = S$.

We now analyze $\mathbb{E}_t^{\mathbb{Q}^j} [S_T] - \mathbb{E}_t^{\mathbb{Q}^i} [S_T]$. The time t stochastic discount factor for tax distribution k and payoffs at $r \geq t$ is M_r^k . This is also the Radon-Nikodym derivative of \mathbb{Q}_k with respect to the physical measure \mathbb{P} on the time t filtration:

$$M_r^k \equiv \frac{d\mathbb{Q}_k}{d\mathbb{P}}.$$

We know that M_r^k is nonnegative and that $M_t^k = 1$. Because $\mathbb{E}_t^{\mathbb{P}} [M_r^k] = 1 \forall r \geq t$, M_r^k is a \mathbb{P} -martingale. By Girsanov's Theorem, we have:

$$dW_r^{\mathbb{P}} = \psi_r^{\mathbb{Q}_k} dr + dW_r^{\mathbb{Q}_k},$$

where dW_r^k is an increment of a Wiener process under measure k . The process $\psi^{\mathbb{Q}_k}$ is the Girsanov kernel of the measure transformation. Assume that the commodity price follows an Ito process in the absence of taxes:

$$dS_r = \mu(r) dt + \sigma(r) dW_r^{\mathbb{P}}.$$

Substituting for $dW_r^{\mathbb{P}}$ gives:

$$\begin{aligned}
dS_r &= \mu(r) dt + \sigma(r) (\psi_r^{\mathbb{Q}_k} dr + dW_r^{\mathbb{Q}_k}) \\
&= (\mu(r) - \sigma(r) \lambda_r^{\mathbb{Q}_k}) dr + \sigma(r) dW_r^{\mathbb{Q}_k}, \tag{14}
\end{aligned}$$

where we define $\lambda_r^{\mathbb{Q}_k} \equiv -\psi_r^{\mathbb{Q}_k}$ as the market price of risk for the source of randomness given by $dW_t^{\mathbb{P}}$. No-arbitrage assumptions imply that there is some positive market price of risk that makes the set of asset prices internally consistent, while complete markets would imply that the market price of risk is unique. When the commodity's price evolution is viewed

under the risk-adjusted measure \mathbb{Q}_k , the market price of risk reduces its drift while leaving its volatility term unchanged. Our expression for dS_r implies:

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}^j} [S_T] - \mathbb{E}_t^{\mathbb{Q}^i} [S_T] &= \int_0^T (\mu(r) + \sigma(r)\psi_r^{\mathbb{Q}^j}) dr - \int_0^T (\mu(r) + \sigma(r)\psi_r^{\mathbb{Q}^i}) dr \\ &= \int_0^T \sigma(r) (\lambda_r^{\mathbb{Q}^i} - \lambda_r^{\mathbb{Q}^j}) dr.\end{aligned}\quad (15)$$

If a shift in the distribution of emission prices increases the market price of risk, then the risk-adjusted measure changes to decrease the futures price via altered weighting of the no-tax commodity price.

We now analyze the similar term $\mathbb{E}_t^{\mathbb{Q}^j} [\tau_s] - \mathbb{E}_t^{\mathbb{Q}^i} [\tau_s]$, but we do not assume that taxes evolve as an Ito process. In order to differentiate between random variables under the physical measure \mathbb{P} , define τ_s^i as a tax drawn from distribution i and τ_s^j as a tax drawn from distribution j . Rewriting the risk-adjusted measures in terms of the stochastic discount factors, we have:

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}^j} [\tau_s] - \mathbb{E}_t^{\mathbb{Q}^i} [\tau_s] &= \mathbb{E}_t^{\mathbb{P}} [M_T^{\mathbb{Q}^j} \tau_s^j] - \mathbb{E}_t^{\mathbb{P}} [M_T^{\mathbb{Q}^i} \tau_s^i] \\ &= \mathbb{E}_t^{\mathbb{P}} [\tau_s^j] \mathbb{E}_t^{\mathbb{P}} [M_T^{\mathbb{Q}^j}] - \mathbb{E}_t^{\mathbb{P}} [\tau_s^i] \mathbb{E}_t^{\mathbb{P}} [M_T^{\mathbb{Q}^i}] + Cov(\tau_s^j, M_T^{\mathbb{Q}^j}) - Cov(\tau_s^i, M_T^{\mathbb{Q}^i}).\end{aligned}\quad (16)$$

By the assumption of first-order stochastic dominance, the expectation of τ_s^j is unambiguously greater than the expectation of τ_s^i . Define a parameter A by $\mathbb{E}_t^{\mathbb{P}} [\tau_s^j] = A \mathbb{E}_t^{\mathbb{P}} [\tau_s^i]$, where $A > 1$ by first-order stochastic dominance. We have the following condition for the expression in (16) to be strictly positive:

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}^j} [\tau_s] - \mathbb{E}_t^{\mathbb{Q}^i} [\tau_s] &> 0 \\ \Leftrightarrow \mathbb{E}_t^{\mathbb{P}} [\tau_s^i] \left\{ \mathbb{E}_t^{\mathbb{P}} [M_T^{\mathbb{Q}^i}] - A \mathbb{E}_t^{\mathbb{P}} [M_T^{\mathbb{Q}^j}] \right\} &< Cov(\tau_s^j, M_T^{\mathbb{Q}^j}) - Cov(\tau_s^i, M_T^{\mathbb{Q}^i}).\end{aligned}\quad (17)$$

We now establish a sufficient condition for the sign of (13) to match the sign of $\partial g(S, \tau)/\partial \tau$. First, assume that the market price of risk either does not change or changes in the direction opposite to the sign of $\partial g(S, \tau)/\partial \tau$. In that case, the expression (16) is either zero or has the same sign as $\partial g(S, \tau)/\partial \tau$, reinforcing that term in (13). Now assume, second, that the time T stochastic discount factor weakly increases and, third, that $Cov(\tau_s^k, M_T^{\mathbb{Q}^k})$ weakly increases with the shift from distribution i to j . In that case, the inequality in (17) clearly holds. Therefore the term multiplying $\partial g(S, \tau)/\partial \tau$ in (13) is positive. When we combine all three assumptions, the sign of (13) matches the sign of $\partial g(S, \tau)/\partial \tau$.

References

- Anderson, Michael L. (2008) “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*, Vol. 103, No. 484, pp. 1481–1495.
- Andrews, D. W. K. (2003) “End-of-sample instability tests,” *Econometrica*, Vol. 71, No. 6, pp. 1661–1694.
- Auerbach, Alan J. (1989) “Tax reform and adjustment costs: The impact on investment and market value,” *International Economic Review*, Vol. 30, No. 4, pp. 939–962.
- Becker, Gary S., Michael Grossman, and Kevin M. Murphy (1994) “An empirical analysis of cigarette addiction,” *The American Economic Review*, Vol. 84, No. 3, pp. 396–418.
- Binder, John J. (1998) “The event study methodology since 1969,” *Review of Quantitative Finance and Accounting*, Vol. 11, pp. 111–137.
- Bollerslev, Tim (1986) “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, Vol. 31, No. 3, pp. 307–327.
- Branson, William, Arminio Fraga, and Robert Johnson (1986) “Expected fiscal policy and the recession of 1982,” in M. H. Peston and R. E. Quandt eds. *Price, Competition and Equilibrium*, Totowa, NJ: Barnes & Noble Books, pp. 109–128.
- Broder, John M. (2010) “Graham Pulls Support for Major Senate Climate Bill,” *The New York Times*.
- Brown, Stephen J. and Jerold B. Warner (1985) “Using daily stock returns: The case of event studies,” *Journal of Financial Economics*, Vol. 14, No. 1, pp. 3–31.
- Bushnell, James, Howard Chong, and Erin T. Mansur (2013) “Profiting from regulation: Evidence from the European carbon market,” *American Economic Journal: Economic Policy*, Forthcoming.
- Busse, Meghan R. and Nathaniel O. Keohane (2007) “Market effects of environmental regulation: coal, railroads, and the 1990 Clean Air Act,” *The RAND Journal of Economics*, Vol. 38, No. 4, p. 11591179.
- Cash, Cathy (2010) “GHG bipartisan bill unveiling expected today,” *Platts Coal Trader*, p. 1.
- CBO (2010) “Discussion Draft of the American Power Act,” Technical report, Congressional Budget Office.

- Cicala, Steve (2012) “When does regulation distort costs? Lessons from fuel procurement in US electricity generation,” November.
- ClimateWorks Foundation (2010) “Analysis of the American Power Act,” Technical report.
- Conley, Timothy G. and Christopher R. Taber (2011) “Inference with “difference in differences” with a small number of policy changes,” *Review of Economics and Statistics*, Vol. 93, No. 1, pp. 113–125.
- Cowan, Richard (2010) “Global warming bill shows signs of life,” *Reuters*.
- Cowan, Richard and Thomas Ferraro (2010) “UPDATE 2—U.S. senators postpone climate bill unveiling,” *Reuters*.
- Cox, John C., Jonathan E. Ingersoll Jr., and Stephen A. Ross (1981) “The relation between forward prices and futures prices,” *Journal of Financial Economics*, Vol. 9, No. 4, pp. 321–346.
- Di Maria, Corrado, Sjak Smulders, and Edwin van der Werf (2012) “Absolute abundance and relative scarcity: Environmental policy with implementation lags,” *Ecological Economics*, Vol. 74, pp. 104–119.
- DOE (2007) “Technical Guidelines: Voluntary Reporting of Greenhouse Gases (1605(b)) Program,” Technical report, Office of Policy and International Affairs, U.S. Department of Energy.
- EIA (2010a) “Electric Power Monthly July 2010,” Technical Report DOE/EIA-0226 (2010/07), U.S. Energy Information Administration.
- (2010b) “Energy Market and Economic Impacts of the American Power Act of 2010,” Report SR/OIAF/2010-01, U.S. Energy Information Administration, Washington, DC.
- (2012a) “Annual Energy Outlook 2012,” Technical Report DOE/EIA-0383(2012), U.S. Energy Information Administration.
- (2012b) “Fuel Competition in Power Generation and Elasticities of Substitution,” Technical report, U.S. Energy Information Administration.
- (2013) “Electric Power Annual,” Technical report, U.S. Energy Information Administration.
- Eilperin, Juliet (2010a) “Kerry, Graham, Lieberman to unveil climate bill,” *The Washington Post*.

- (2010b) “Sen. Graham walks away from climate and energy bill,” April.
- EPA (2010) “EPA Analysis of the American Power Act of 2010,” Technical report, Office of Atmospheric Programs, U.S. Environmental Protection Agency.
- Fountain, John and Glenn W. Harrison (2011) “What do prediction markets predict?” *Applied Economics Letters*, Vol. 18, No. 3, pp. 267–272.
- Garrett, Major (2010) “WH Opposes Higher Gas Taxes Floated by S.C. GOP Sen. Graham in Emerging Senate Energy Bill,” April.
- Gelbach, Jonah B., Eric A. Helland, and Jonathan Klick (forthcoming) “Valid inference in single-firm, single-event studies,” *American Law and Economics Review*.
- Geman, Hélyette (2005) *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*, Chichester, West Sussex: John Wiley & Sons, Ltd.
- Gerlagh, Reyer (2011) “Too much oil,” *CESifo Economic Studies*, Vol. 57, No. 1, pp. 79–102.
- Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton (2011) “Estimating the social cost of carbon for use in U.S. federal rulemakings: a summary and interpretation,” *National Bureau of Economic Research Working Paper Series*, Vol. No. 16913.
- Gruber, Jonathan and Botond Köszegi (2001) “Is addiction “rational”? Theory and evidence,” *The Quarterly Journal of Economics*, Vol. 116, No. 4, pp. 1261–1303.
- Hakala, Scott D. (2010) “Lessons from single-company event studies: The importance of controlling for company-specific events.”
- Hall, R. E. (1971) “The dynamic effects of fiscal policy in an economy with foresight,” *The Review of Economic Studies*, Vol. 38, No. 2, pp. 229–244.
- Heal, Geoffrey (1976) “The relationship between price and extraction cost for a resource with a backstop technology,” *The Bell Journal of Economics*, Vol. 7, No. 2, pp. 371–378.
- Hoel, Michael (2012) “Carbon taxes and the green paradox,” in Robert W. Hahn and Alistair Ulph eds. *Climate Change and Common Sense: Essays in Honour of Tom Schelling*: Oxford University Press, pp. 203–224.
- Hotelling, Harold (1931) “The economics of exhaustible resources,” *The Journal of Political Economy*, Vol. 39, No. 2, pp. 137–175.
- Houser, Trevor, Shashank Mohan, and Ian Hoffman (2010) “Assessing the American Power Act: The Economic, Employment, Energy Security, and Environmental Impact of Senator Kerry and Senator Liebermans Discussion Draft,” Policy Brief 10-12, Peterson Institute for International Economics.

- Jones, Benjamin, Michael Keen, and Jon Strand (2013) “Fiscal implications of climate change,” *International Tax and Public Finance*, Vol. 20, No. 1, pp. 29–70.
- Judd, Kenneth L. (1985) “Short-run analysis of fiscal policy in a simple perfect foresight model,” *Journal of Political Economy*, Vol. 93, No. 2, pp. 298–319.
- (1987) “The welfare cost of factor taxation in a perfect-foresight model,” *Journal of Political Economy*, Vol. 95, No. 4, pp. 675–709.
- Lamdin, Douglas J. (2001) “Implementing and interpreting event studies of regulatory changes,” *Journal of Economics and Business*, Vol. 53, No. 23, pp. 171–183.
- Lange, Ian and Joshua Linn (2008) “Bush v. Gore and the effect of New Source Review on power plant emissions,” *Environmental and Resource Economics*, Vol. 40, No. 4, pp. 571–591.
- Linn, Joshua (2010) “The effect of cap-and-trade programs on firms profits: Evidence from the Nitrogen Oxides Budget Trading Program,” *Journal of Environmental Economics and Management*, Vol. 59, No. 1, pp. 1–14.
- Lizza, Ryan (2010) “As the World Burns,” *The New Yorker*.
- MacKinlay, A. Craig (1997) “Event studies in economics and finance,” *Journal of Economic Literature*, Vol. 35, No. 1, pp. 13–39.
- Manski, Charles F. (2006) “Interpreting the predictions of prediction markets,” *Economics Letters*, Vol. 91, No. 3, pp. 425–429.
- Mckenzie, Andrew M., Michael R. Thomsen, and Bruce L. Dixon (2004) “The performance of event study approaches using daily commodity futures returns,” *Journal of Futures Markets*, Vol. 24, No. 6, pp. 533–555.
- Meng, Kyle C. (2012) “The cost of potential cap-and-trade policy: An event study using prediction markets and lobbying records,” Technical report.
- Mertens, Karel and Morten O. Ravn (2011) “Understanding the aggregate effects of anticipated and unanticipated tax policy shocks,” *Review of Economic Dynamics*, Vol. 14, No. 1, pp. 27–54.
- (2012) “Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks,” *American Economic Journal: Economic Policy*, Vol. 4, No. 2, pp. 145–181.
- Perotti, Roberto (2012) “The effects of tax shocks on output: Not so large, but not small either,” *American Economic Journal: Economic Policy*, Vol. 4, No. 2, pp. 214–237.

- Romm, Joe (2010) “BREAKING: Sen. Graham walks away from climate and energy bill over immigration plans,” April.
- Sinn, Hans-Werner (2008) “Public policies against global warming: a supply side approach,” *International Tax and Public Finance*, Vol. 15, No. 4, pp. 360–394.
- (2012) *The Green Paradox: A Supply-Side Approach to Global Warming*, Cambridge, Mass: MIT Press.
- Smulders, Sjak, Yacov Tsur, and Amos Zemel (2012) “Announcing climate policy: Can a green paradox arise without scarcity?” *Journal of Environmental Economics and Management*, Vol. 64, No. 3, pp. 364–376.
- Snow, Nick (2010) “Graham leaves climate compromise effort to Kerry, Lieberman,” *Oil & Gas Journal*, Vol. 108, No. 16, pp. 37–38.
- Tankersley, Jim (2010) “Why the Senate climate bill is on life support,” April.
- Thrush, Glenn and Marin Cogan (2010) “Climate backers try to regroup,” <http://www.politico.com/news/stories/0410/36324.html>, April.
- van der Werf, Edwin and Corrado Di Maria (2012) “Imperfect environmental policy and polluting emissions: The Green Paradox and beyond,” *International Review of Environmental and Resource Economics*, Vol. 6, No. 2, pp. 153–194.
- Wolfers, Justin and Eric Zitzewitz (2006) “Five open questions about prediction markets,” Working Paper 12060, National Bureau of Economic Research.
- Yang, Shu-Chun Susan (2005) “Quantifying tax effects under policy foresight,” *Journal of Monetary Economics*, Vol. 52, No. 8, pp. 1557–1568.