American Economic Review 2017, 107(3): 748–784
https://doi.org/10.1257/aer.20150781

Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy†

By Kyle C. Meng*

This paper develops a method for forecasting the marginal abatement cost (MAC) of climate policy using three features of the failed Waxman-Markey bill. First, the MAC is revealed by the price of traded permits. Second, the permit price is estimated using a regression discontinuity design (RDD) comparing stock returns of firms on either side of the policy's free permit cutoff rule. Third, because Waxman-Markey was never implemented, I extend the RDD approach to incorporate prediction market prices which normalize estimates by policy realization probabilities. A final bounding analysis recovers a MAC range of \$5 to \$19 per ton CO_2e. (*JEL* G12, G14, Q52, Q54, Q58)

> *Overall, Waxman-Markey reduces gross domestic product by an average of \$393 billion annually between 2012 and 2035…*
>
> —The Heritage Foundation[1]

> *…under [Waxman-Markey]… the net annual economy wide cost of the cap-and-trade program in 2020 would be \$22 billion.*
>
> —Congressional Budget Office[2]

Decisions over proposed policies rely on forecasting costs and benefits. In some cases, program evaluation of similar past policies may be informative. However, when the policy of interest is relatively novel, policymakers often turn to economic

* Bren School and Department of Economics, University of California, Santa Barbara, and NBER (e-mail: kmeng@bren.ucsb.edu). I thank Ethan Kaplan, Wolfram Schlenker, and in particular Bernard Salanié for their guidance. This paper has benefited from suggestions from Joe Aldy, Doug Almond, Scott Barrett, Linda Bui, Olivier Deschênes, Ray Fisman, Olle Folke, Larry Goulder, Michael Greenstone, Rema Hanna, Geoff Heal, Solomon Hsiang, Nat Keohane, Wojciech Kopczuk, Gary Libecap, Erich Muehlegger, Suresh Naidu, Matt Neidell, Paulina Oliva, Ryan Oprea, Steve Salant, Doug Steigerwald, Tom Vogl, Gernot Wagner, Reed Walker, Marty Weitzman, Justin Wolfers, and seminar participants at Columbia University, Environmental Defense Fund, Georgia Tech, Harvard Kennedy School, IZA, NBER EEE Summer Institute, New York Fed, Rutgers, Stanford, UC Irvine, UC Santa Barbara, University of Connecticut, and University of Michigan. I am particularly grateful to Richard Goettle, Jennifer Morris, and Sergey Paltsev for sharing output from the IGEM and EPPA models. Andy Bilich provided excellent research assistance. All remaining errors are my own. I declare that I have no relevant or material financial interests that relate to the research described in this paper.

† Go to https://doi.org/10.1257/aer.20150781 to visit the article page for additional materials and author disclosure statement.

[1] Ben Lieberman, "The Economic Impact of the Waxman-Markey Cap-and-Trade Bill," *The Heritage Foundation*, June 22, 2009, heritage.org/research/testimony/the-economic-impact-of-the-waxman-markey-cap-and-trade-bill.

[2] Daniel J. Weiss, "CBO: Waxman-Markey Pollution Cuts Cost Little," *Grist*, June, 22, 2009, grist.org/politics/cbo-pollution-cuts-cost-little/.

models with poorly constrained primitive parameters. While recent studies have strengthened our understanding of climate policy benefits,[3] current estimates of climate policy costs are characterized by this dilemma. This challenge is particularly relevant for the United States which, despite being the largest cumulative greenhouse gas emitter, has yet to implement a national climate policy. Instead, recent US climate policy debates have relied on computable general equilibrium (CGE) models, which have thus far produced a wide range of cost estimates. This range is captured by the two quotes above, issued on the same day, for the same proposed policy, but informed by different CGE models.

This paper develops an alternative, empirically driven approach for forecasting the marginal abatement cost of climate policy based on market expectations. In the combined spirit of Hayek (1945) and Chetty (2009), my approach acknowledges that while primitive parameters for climate policy may be unknown to the researcher, *local* information may be held by market participants whose aggregate behavior could reveal a *sufficient statistic* for the policy's marginal abatement cost.[4] The policy of interest is the Waxman-Markey bill, a cap-and-trade climate policy that passed the US House of Representatives in 2009, failed passage in the US Senate in 2010, and is to date the US climate legislation that came closest to becoming implemented.

I recover the market expected marginal abatement cost of the Waxman-Markey bill by using three distinct features of the policy. First, under standard theoretical assumptions, the equilibrium permit price emerging from a cap-and-trade system is equal to the marginal abatement cost of the policy. Second, written into the Waxman-Markey bill was a cutoff rule granting certain firms free permits. This allows one to use a regression discontinuity design (RDD) to examine the difference in stock returns between firms just above and below the cutoff. Finally, because Waxman-Markey existed only as a probable policy during the 2009–2010 period, I extend the standard RDD approach by normalizing estimates using market beliefs that Waxman-Markey would be implemented, as captured by price changes from a prediction market tied to the policy.

Understanding the marginal abatement cost, or permit price, of the Waxman-Markey bill is important for three reasons. First, under certain conditions, the marginal abatement cost of complying with a cap-and-trade system with broad sectoral coverage such as Waxman-Markey is equal to the economy-wide US marginal abatement cost.[5] This parameter, together with recent estimates of climate policy benefits, allows for cost-benefit analyses. Second, cap-and-trade policies with permit auctioning generate government revenue that could be used to reduce distortionary

[3] See, for example, Schlenker, Hanemann, and Fisher (2005); Deschênes and Greenstone (2007, 2011); Schlenker and Roberts (2009); Feng, Oppenheimer, and Schlenker (2012); Graff Zivin and Neidell (2014); Hsiang, Meng, and Cane (2011); Dell, Jones, and Olken (2012); and Hsiang, Burke and Miguel (2013).

[4] Specifically, the permit price of a cap-and-trade policy is a sufficient statistic for marginal cost in that any vector of primitive parameters consistent with a particular permit price implies the same marginal cost. Note that unlike the welfare applications discussed in Chetty (2009), the parameter estimated in this paper captures only the policy's marginal cost, not marginal welfare. To obtain welfare effects for the policy, one needs marginal costs and benefits at every level of abatement up to the policy's constraint.

[5] The marginal abatement cost of complying with a cap-and-trade system equals the economy-wide marginal abatement cost under least-cost allocation when the permit price is binding for all sources of US emissions. Violations of this equivalence occur whenever distortions lower the residual demand for permits. Examples include the presence of binding auxiliary policies such as renewable portfolio standards and when firms can easily relocate carbon-intensive operations overseas to unregulated countries.

taxation (Bovenberg and Goulder 1996). The size of the potential welfare improvement is determined in part by the permit price. Third, the permit price of a US cap-and-trade system determines the potential benefits from linkages with other national cap-and-trade systems (Ranson and Stavins 2016).

When a policy is uncertain, neither firm value with or without the policy is directly observed. However, if market-held beliefs over the likelihood of cap-and-trade implementation can be observed via prediction market prices (Snowberg, Wolfers, and Zitzewitz 2013), then changes in such prices can be combined with stock returns in an event study to recover the market-expected firm-level effects of the policy. This suggests that the policy's permit price may be recovered by taking the difference in estimated effects between firms with and without free permits.

Unfortunately, such a broad comparison may be invalid as cap-and-trade systems generally do not allocate free permits randomly. It is, however, possible that some free permits are allocated in a manner unrelated to other firm characteristics.[6] A specific rule under Waxman-Markey indicated that firms in manufacturing subsectors with historical energy intensity greater than 5 percent would receive free permits. This suggests the basis for a regression discontinuity design (RDD) provided that firm stock returns are continuous at the 5 percent threshold under Waxman-Markey if not for the free permit rule. However, Waxman-Markey was never implemented. In an extension of the standard RDD framework to account for policy uncertainty, I formally demonstrate that identification further requires unobserved determinants of stock returns in the absence of the policy to be continuous at the 5 percent threshold.

I show that differential cumulative stock returns for firms above and below the historical energy intensity threshold closely track Waxman-Markey developments, diverging from zero at the start of 2009, peaking around the House passage of the bill, and converging back to zero thereafter during 2010. For graphical analysis, I implement two-step procedures illustrating the RD coefficient in both time-series and cross-sectional dimensions. My preferred approach estimates a one-step panel RD model which interacts time-series variation in prediction market prices with a local cross-sectional polynomial function of the RD treatment. I find robust evidence that firms in sectors with energy intensity greater than 5 percent experience a gain of 7 to 9 percentage points relative to firms in sectors with energy intensity less than 5 percent. I do not find discontinuities at placebo thresholds less or greater than 5 percent energy intensity.

Several lines of evidence suggest that, given this 5 percent threshold, sorting was unlikely. If indeed the free permit rule was unexpected, sorting in 2009 would be impossible as the energy intensity running variable is defined by data from public government sources in 2004–2006. It is unlikely that this rule was anticipated prior to 2009 as US climate legislation preceding the Waxman-Markey bill did not feature a similar free permit rule. Indeed, average stock returns for manufacturing firms in 2007 and 2008 do not exhibit a jump at the 5 percent threshold. Likewise, sector-level lobbying expenditures on climate policy as well as other covariates do not exhibit a discontinuity at the 5 percent threshold suggesting that firm activity did not determine the location of the threshold.

[6] Fowlie and Perloff (2013) use a free permit allocation rule in the RECLAIM NOx cap-and-trade program to identify the effects of free permits on equilibrium emissions.

My RDD framework estimates a reduced-form parameter capturing the jump in stock returns between firms above and below the 5 percent energy intensity threshold. To map this estimate onto my structural parameter of interest, the permit price, or marginal abatement cost of Waxman-Markey's cap-and-trade system, I consider a standard model of cap-and-trade and two additional structural assumptions, with supporting statistical and anecdotal evidence. Only one element of this mapping is neither directly observed or estimated, the cumulative number of free permits granted to treated firms under Waxman-Markey. In the spirit of Horowitz and Manski (2000), I consider *worst* and *best* case bounds by assuming that, under the policy, treated firms uniformly experience future emission decline rates equal to that of the worst and best performing manufacturing subsectors observed during years prior to the expected start of Waxman-Markey. This produces an implied lower and upper bound for the marginal abatement cost of \$5 and \$19 per ton of CO_2e in 2015, respectively. This range is narrower in spread and below the mean of CGE estimates of the same policy. Interestingly, the range of estimates recovered by my method is comparable with estimates from CGE models funded by academic and governmental institutions. Only under drastic assumptions of future emissions decline, in the range of -19 to -35 percent per year for all treated firms, do I recover marginal abatement costs in the realm produced by CGE models funded by private institutions.

A forecasting approach based on market expectations may be well-suited for estimating the marginal abatement cost of climate policy if information on key parameters, such as those dictating technological change, is largely private. The induced innovation hypothesis (Hicks 1932) suggests that climate policy could trigger significant technological advances (Jaffe, Newell, and Stavins 2003). While this has been explored theoretically (Goulder and Schneider 1999; Nordhaus 2002; Buonanno, Carraro, and Galeotti 2003; Acemoglu et al. 2012), induced technological change presents modeling difficulties for many CGE models of climate policy (Jacoby et al. 2006). My approach is agnostic about the structure, direction, and rate of technological change, and instead relies on the expectations of market participants to reveal dispersed information over the technological frontier.

Using market expectations have certain drawbacks. While my estimates capture the expected cost of the implemented policy, I am unable to confirm what markets expected to be implemented correspond exactly to the Waxman-Markey bill, especially if future policy revisions were anticipated. Furthermore, even if one knows the exact policy expected to be implemented, I could not conduct counterfactual analyses using my approach as I recover the marginal abatement cost associated with a single emission constraint and not the entire cost curve along various constraint levels.[7] As such, my method informs upon the level but not the slope of the marginal abatement cost curve and thus should complement counterfactual analyses performed by CGE models.

[7] This implies that while I am unable to estimate the total abatement cost associated with the Waxman-Markey policy, I could potentially recover an upper bound on total cost. The product of the permit price and emissions equals the total abatement cost for a firm that does not engage in any abatement response, which is greater than the total cost for any firm that does abate. I thank a referee for bringing this to my attention.

To the best of my knowledge, this paper provides the first forecast of the marginal abatement cost under cap-and-trade policy outside CGE models.[8] In doing so, it connects with several strands of literature, both in substance and methodology. Existing papers use historical fluctuations in energy prices to inform upon the potential effects of future US climate policy on employment (Deschênes 2011), competitiveness (Aldy and Pizer 2015), and welfare (Cullen 2013). In each paper, the extension to potential climate policy effects involves multiplying an estimated energy price elasticity with the expected change in energy prices under climate policy, with the latter value either assumed or itself obtained from CGE models. This paper contributes to this literature by providing an empirically estimated permit price for arguably the most important US climate policy considered to date. Cullen and Mansur (2014) directly recover an implied short-run carbon abatement cost in the absence of technological change using an estimated relationship between recent CO_2 emissions and the ratio of coal to natural gas prices. This recovered abatement cost, however, is specific only to short-run fuel-switching from coal to natural gas in the electricity sector and may not equal the marginal abatement cost of a cap-and-trade policy with national coverage over multiple sectors. Finally, this paper is intellectually similar to Anderson and Sallee (2011) and Greenstone, List, and Syverson (2012), which empirically recover the marginal cost of realized environmental regulations.

This is also among the first papers to use the prediction market event study presented by Snowberg, Wolfers, and Zitzewitz (2007) and Wolfers and Zitzewitz (2009) in a forecasting context. There is, however, a long tradition of employing traditional event study methodology to evaluate, ex post, the costs of realized regulation including event studies examining the cost of realized US (Lange and Linn 2008; Linn 2010) and EU (Bushnell, Chong, and Mansur 2013) environmental regulations. Regarding the Waxman-Markey bill, Lemoine (forthcoming) conducts a traditional event study for events related to the bill but the emphasis is not on recovering the policy's marginal abatement cost.

The remainder of the paper is structured as follows. Section I provides background on the Waxman-Markey bill and the use of prediction markets in event studies. Section II introduces a reduced-form framework for RDD under policy uncertainty. Section III presents reduced-form results along with robustness checks. Section IV provides a structural framework mapping the RD parameter onto Waxman-Markey's marginal abatement cost and compares my estimates with that of CGE models. Section V offers a brief discussion. The online Appendix includes a more detailed model of cap-and-trade, formal derivations for potential bias due to thin trading, an adjustment procedure for prediction market expiration, a data summary, further background on Waxman-Markey and related bills, and a summary of CGE models.

[8] It is possible that a market expected forecast of Waxman-Markey's marginal abatement cost may still be far off the actual value had the policy been implemented. If future economic conditions are highly uncertain, any forecast of regulatory costs would also be inaccurate. Unfortunately, because Waxman-Markey was never implemented, I cannot determine the degree of forecast error in my estimates.

I. Background

A. *The Waxman-Markey Cap-and-Trade Bill*

Over the past two decades, emissions trading, known popularly as cap-and-trade, has become an increasingly important regulatory instrument for controlling regional and global pollutants such as greenhouse gases (Stavins 1998; Aldy et al. 2010). In a typical cap-and-trade system, a limit, or cap, on cumulative emissions is set for the lifetime of the policy with the regulator issuing annual emissions permits under this cap. Regulated firms are either given or must purchase permits to cover their annual emissions. The primary appeal of cap-and-trade is that, under standard theoretical assumptions, the allocation of emissions across firms after permit trading should achieve the required cap at the lowest aggregate compliance cost (see online Appendix A.1) (Montgomery 1972; Rubin 1996). Following the success of the US SO_2 trading system introduced in the Clean Air Act Amendments of 1990 (Carlson et al. 2000; Ellerman et al. 2000), variants of cap-and-trade have been implemented domestically and internationally. Well-known systems currently in operation include the European Unions Emissions Trading System (EU-ETS), the US Regional Greenhouse Gas Initiative (RGGI), and the California cap-and-trade system.

This backdrop has made cap-and-trade the centerpiece of US domestic climate policy efforts over the last two decades. After a series of failed Senate cap-and-trade bills in the early 2000s, the Democratic-led 111th House of Representatives introduced the American Clean Energy and Security Act in the spring of 2009. Known informally as the Waxman-Markey bill after its primary sponsors, the legislation specified a declining annual limit on emissions beginning in 2012 which would eventually cover 85 percent of greenhouse-gas-emitting sectors.[9] Waxman-Markey required that covered emissions be at 83 percent, 58 percent, and 17 percent of 2005 levels by 2020, 2030, and 2050, respectively. To smooth costs over time, regulated firms were allowed unlimited permit banking and permit borrowing with some restrictions over the lifetime of the policy.[10] Importantly, in contrast to earlier periods in which Congress considered several cap-and-trade bills simultaneously, the 111th Congress only seriously deliberated over the Waxman-Markey bill and its Senate variant.

Written within the Waxman-Markey bill was a rule which granted free permits to a particular subset of firms. Specifically, Waxman-Markey considered a firm in a six-digit NAICS manufacturing sector (i.e., NAICS 31–33) eligible for free permits if for that sector recent[11] energy intensity[12] was over 5 percent and trade intensity[13]

[9] While central to Waxman-Markey, cap-and-trade was not the only component of the legislation. There were also supply-side interventions such as a renewable energy portfolio standard as well as demand-side interventions such as incentives for electric vehicles. Insofar as other policy components distort incentives toward cost-minimizing abatement options, the resulting permit price and thus marginal abatement cost would deviate from that of a hypothetical least-cost policy with the same overall abatement level.

[10] In particular, borrowing of permits one year ahead incurs an 8 percent interest rate.

[11] Section 763(b)(2)(D) required an "average of data from as many of the years of 2004, 2005, and 2006 for which such data are available" from the US Census Annual Survey of Manufacturers.

[12] Defined as "dividing the cost of purchased electricity and fuel costs of the sector by the value of the shipments of the sector" in Section 763(b)(2)(A)(ii)(II) of H.R.2454.

[13] Defined as "dividing the value of the total imports and exports of such sector by the value of the shipments plus the value of imports of such sector" in Section 401(b)(2)(A)(iii) of H.R.2454.

exceeded 15 percent.[14] These permits were granted annually according to firm-level output and industry-wide emissions intensity from 2012 to 2025 and phased out over the 2026–2035 period.[15] Importantly, this was the first time a cutoff criteria for free permits appeared in US climate legislation. All prior climate legislation[16] directed government regulators to eventually design a free permit allocation rule without providing specific guidance. This suggests that it was unlikely for firms to sort around this specific threshold for years prior to the introduction of the Waxman-Markey bill in 2009.

Waxman-Markey passed the House of Representatives on June 26, 2009, marking the first time cap-and-trade legislation had passed either Houses of Congress. Soon thereafter the Senate introduced complementary cap-and-trade bills with a free permit rule and provisions affecting permit prices that were similar to those in Waxman-Markey (see online Appendix F and Table A.4 for a detailed comparison of these bills).[17] The similarities between Waxman-Markey and Senate bills were likely due to President Obama's explicit support for the Waxman-Markey bill and because a Senate bill that deviated too much from Waxman-Markey would require the legislative burden of additional House voting.

Despite these efforts in the Senate, the overall prospects for cap-and-trade declined after House passage. With the exception of Republican Senator Lindsay Graham joining Senate cap-and-trade efforts on Nov 4, 2009, the rest of 2009 and 2010 witnessed the gradual demise of cap-and-trade. Prospects for cap-and-trade were affected by the failure to reach a new binding international agreement at the UNFCCC Copenhagen negotiations and further declined following Scott Brown's Senate victory which weakened the filibuster-proof supermajority needed by the Democrats. On April 23, 2010, Senator Lindsay Graham withdrew support for cap-and-trade. Three months later, on July 22, 2010, a little over a year after House passage of Waxman-Markey, the Senate formally dropped deliberation over a comparable cap-and-trade bill (see online Appendix E for a summary of these events). While the Waxman-Markey was ultimately never implemented, the period from May 2009 to August 2010 was marked by daily fluctuations in cap-and-trade prospects which fortunately were captured by an accompanying prediction market.

B. *Prediction Market Event Studies*

The typical prediction market contract is a bet on the realization of a particular event at a certain date. When that date is reached, holders of a contract receive \$1 if the event is realized and \$0 otherwise with contract prices fluctuating within the unit

[14] The policy also granted free permits to local distribution companies of electricity. There was, however, no cutoff for which firms got free permits and thus could not be used as an identification strategy.

[15] This is known also as *output-based allocation*. Section 764(b)(2)(A and B) required that for the period from 2012 to 2025, annual permits are granted based on the product of output over the past two years and the most recent sector-level greenhouse gas intensity (measured as ton of emissions per ton output). Section 764(a)(B) requires a complete phaseout from 2026 to 2035.

[16] Prominent examples are the 2003 McCain-Lieberman (S.R. 139), 2005 McCain-Lieberman (S.R. 1151), 2007 McCain-Liberman (S.R. 280), 2007 Lieberman-Warner (S.R. 2191), and the 2008 Boxer-Lieberman-Warner (S.R.3036) bills.

[17] In the bicameral US legislative system, a piece of legislation must pass both Houses of Congress before being sent to the President. Thus, passage of Waxman-Markey by the House of Representatives needed to be followed by a similar cap-and-trade bill approved by a Senate filibuster-proof supermajority.

interval prior to the termination date.[18] Under certain assumptions about prediction market participants,[19] the price of the contract can be interpreted as the average market belief over event realization.

Prediction market prices can be paired with stock returns in a prediction market event study. This approach provides two important advantages over traditional event study methodology. First, in the case of an eventually realized policy, prediction market prices mitigate concern over *fuzzy information* release. Traditional event studies examine abnormal stock returns in response to an unexpected release of information. Isolating the moment when markets first become aware of this information is a central challenge. This is typically manifested in the selection of an event window in which one assumes that the probability of policy realization is 0 prior to the window and 1 when the event occurs. Any early release of information may violate this assumption and result in estimates that are sensitive to event window selection (Snowberg, Wolfers, and Zitzewitz 2007).

Second, and perhaps more important, prediction market prices allow researchers to estimate abnormal returns for a probable event even if this event is never realized. In other words, the availability of prediction markets transforms event studies into a tool for policy forecasting. This potential was noted by Snowberg, Wolfers, and Zitzewitz (2013) but, to the best of my knowledge, has not been implemented in the literature. In particular, the availability of prediction markets solves a central empirical challenge: for probable but unrealized policies, neither firm values under the probable policy nor in its absence are directly observed. There are, during any legislative period, a number of important policies that fail to become law but whose costs remain of interest, perhaps to inform future legislative efforts.[20] Such a prediction market was available for cap-and-trade policy during the period 2009–2010.

C. *Cap-and-Trade Prediction Market*

From May 1, 2009 to Dec 31, 2010,[21] the online trading exchange Intrade hosted a prediction market contract on the prospects of a US cap-and-trade system. This contract was titled:[22]

> *A cap-and-trade system for emissions trading to be established before midnight ET on 31 Dec 2010.*

[18] Actual Intrade contract prices range from $0 to $10. I normalize prices to match probabilities.

[19] Wolfers and Zitzewitz (2006) show that two assumptions are required in order for prediction market prices to equal mean beliefs: (i) utility has a log form and (ii) trader wealth and beliefs are independent. For other standard utility functions, the divergence between prediction market prices and mean beliefs is shown generally to be quite small when (i) traders are risk averse; (ii) prices are within the $0.20 to $0.80 range; and (iii) the distribution of beliefs exhibit relatively low dispersion. In the case where trader wealth and belief are correlated, the prediction market price reflects the wealth weighted average belief in the trading population.

[20] Prediction markets have been offered for recent efforts to reform immigration, social security, and health care regulation in the United States. A list of all Intrade prediction markets is available here: www.intrade.com/v4/reports/special/all-intrade-markets/all-intrade-markets.xlsx.

[21] Intrade began offering this contract on March 25, 2009. However, trading began only on May 1, 2009, which marks the start of my sample period.

[22] Intrade further defined this contract by noting:

> *A cap-and-trade system will be considered established once federal legislation authorizing the creation of such a system becomes law, as reported by three independent and reliable media sources. Emissions trading under the system does not need to begin for the contract to be expired.*

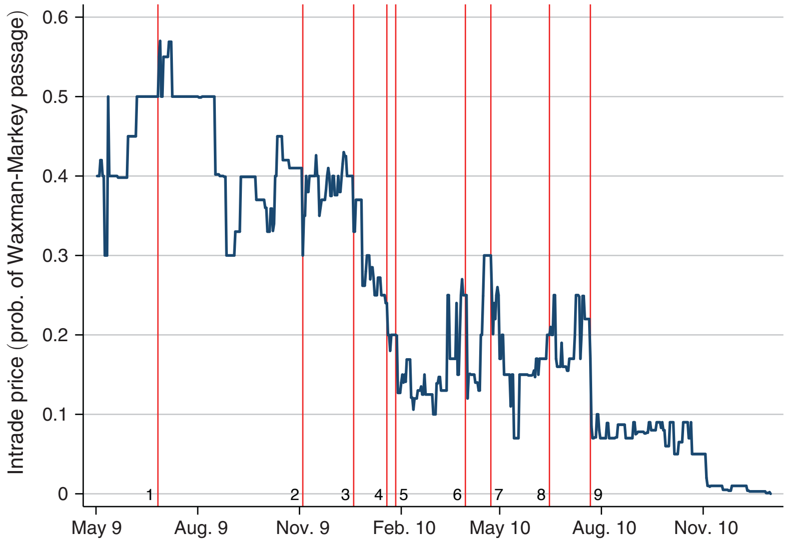

Figure 1. Cap-and-Trade Prediction Market Prices

Notes: Vertical lines mark days with events affecting cap-and-trade prospects. (1) June 26, 2009: House passes Waxman-Markey. (2) November 4, 2009: Graham joins Senate effort. (3) December 20, 2009: Copenhagen negotiations concluded. (4) January 19, 2010: Scott Brown wins Massachusetts Senate seat. (5) January 27, 2010: Graham-Kerry-Lieberman seeks non-cap-and-trade alternatives. (6) March 31, 2010: Obama supports offshore drilling. (7) April 23, 2010: Graham drops support. (8) June 15, 2010: Obama Oval Office speech. (9) July 22, 2010: Senate drops cap-and-trade legislation. See online Appendix E for further details.

Figure 1 plots the price time series for this contract. A price of \$0.50 indicates that market participants believed, on average, that cap-and-trade had a 50 percent chance of being realized before the end of 2010. Each solid vertical line identifies a political event that affected cap-and-trade prospects (summarized in online Appendix E).

Two aspects of this prediction market make it fall short of the ideal. First, the contract describes a generic cap-and-trade system without explicit mention of the Waxman-Markey bill. In Section IVA, I discuss several pieces of statistical and anecdotal evidence suggesting that market participants were likely reacting to the Waxman-Markey bill. Second, this prediction market was relatively thinly traded. During the event period, 11,260 contracts were traded for a total value of \$190,000. An average of 30 contracts were transacted every 2 days.[23] Transaction-level data acquired privately from Intrade indicates that there were 143 unique traders participating in the market.[24] To examine the potential implications of thin trading, Section IIIC conducts a series of empirical tests to examine the potential biases due to thin trading, as formalized in online Appendix B.

[23] By comparison, the prediction market used in Snowberg, Wolfers, and Zitzewitz (2007) had an average of 129 trades for every 10-minute interval during election night.

[24] While Intrade does not provide information on where traders are located, Intrade has said in a public letter to the US CFTC that "our 82,000 plus membership are predominantly resident in the United States" and that "78% of traffic to Intrade.com in the period 1 January to 30 June [2008] was from the US" (http://www.intrade.com/news/misc/CFTC_Intrade_Comment_Reg_Treatment_Event_Mkts.pdf).

II. Reduced-Form Framework: RDD under Policy Uncertainty

This section develops the econometric framework for conducting a regression discontinuity design under policy uncertainty. Identifying assumptions for the reduced-form RD parameter are first introduced graphically within a simplified setting and then formalized more generally. Section IV will introduce the structural assumptions needed to interpret the RD parameter as the marginal abatement cost of the Waxman-Markey bill.

A. *Observational Challenge: Graphical Illustration*

Figure 2 illustrates this paper's main observational challenge in its simplest form. There is one time period. The policy index p indicates whether the Waxman-Markey bill has $(p = w)$ or has not $(p = o)$ been realized. Firms are indexed by i and are in sector $j \in \{F, A\}$. Under Waxman-Markey, firms in sector $j \in F$ with energy intensity $EI_j > 0.05$ receive free permits. Conversely, firms in sector $j \in A$ with energy intensity $EI_j < 0.05$ have to purchase permits. Here, $r_i(j, p, EI_j)$ is firm-level stock returns (or the percentage change in firm value).

First, consider first the standard RDD setting for the hypothetical case where Waxman-Markey has been implemented, $p = w$, as shown in panel A of Figure 2. Solid lines indicate observed values while dashed lines indicate unobserved values. In this setting, the researcher observes the stock returns of firms with and without free permits and can estimate their respective conditional mean functions, $E[r_i | j \in F, p = w, EI_j > 0.05]$ and $E[r_i | j \in A, p = w, EI_j < 0.05]$. Identification of the discontinuity β requires that $E[r_i | j \in A, p = w, EI_j]$ be continuous in EI_j around the 5 percent threshold (Hahn, Todd, and Van der Klaauw 2001). That is, expected stock returns under Waxman-Markey should be continuous at the 5 percent energy intensity threshold if not for the assignment of free permits.

Unfortunately, Waxman-Markey was never realized. This means that the researcher observes neither $E[r_i | j \in F, p = w, EI_j]$ nor $E[r_i | j \in A, p = w, EI_j]$ over any energy intensity value. Furthermore, because Waxman-Markey was a probable policy during this time period, stock returns themselves do not directly reveal returns in the absence of Waxman-Markey, $E[r_i | p = o, EI_j]$. These three unobserved conditional mean functions are shown as dashed lines in panel B of Figure 2. Suppose, however, the researcher somehow observes the change in the probability that Waxman-Markey would be implemented at some future date. One could then estimate firm-level stock returns under Waxman-Markey *relative* to the no-policy scenario, normalize estimates according to the change in Waxman-Markey probabilities, and examine whether relative returns exhibit a discontinuity at the 5 percent energy intensity threshold. Identification of β in this setting requires an additional continuity assumption. Expected stock returns in the absence of Waxman-Markey, $E[r_i | p = o, EI_j]$, must also be continuous in EI_j around the 5 percent threshold.

B. *General Framework*

There are two reasons why, in practice, a more general framework is needed to estimate an RD parameter under policy uncertainty. First, there are more than two

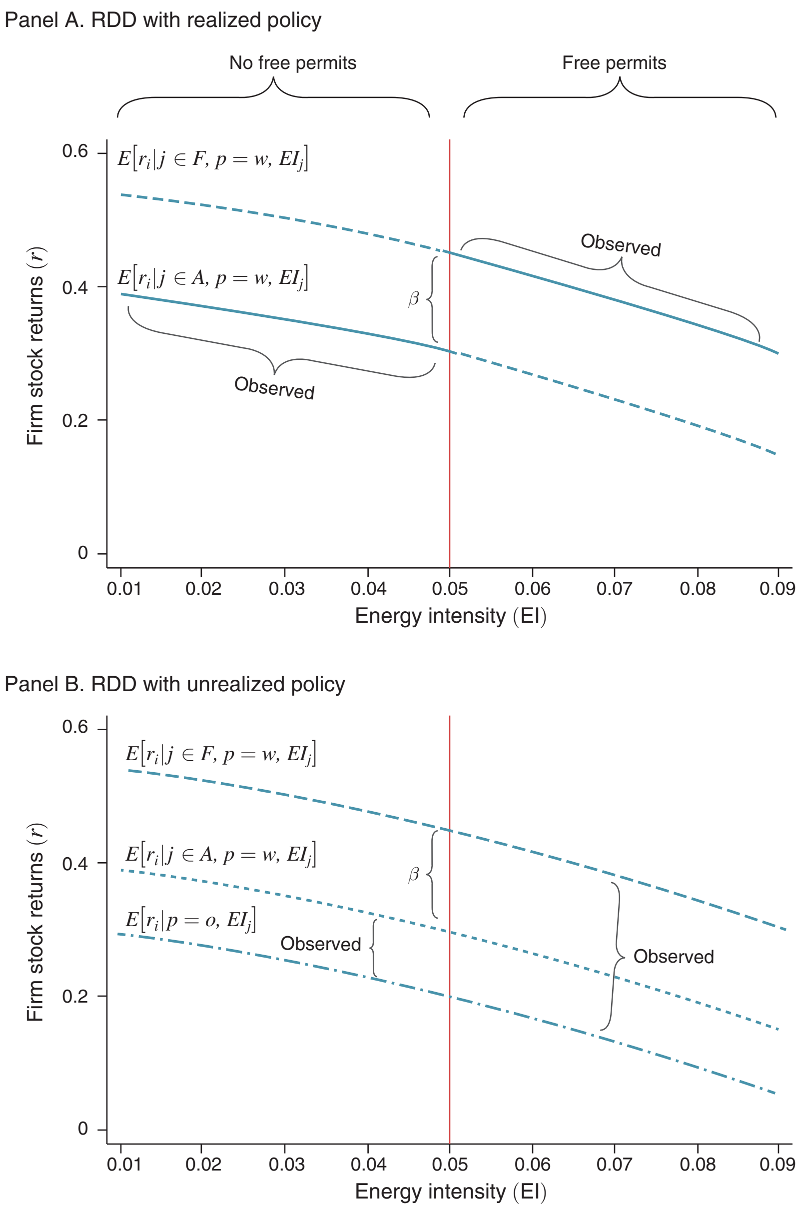

Figure 2. RDD under Realized and Unrealized Policy

Notes: Panel A shows the standard RDD setup under a realized policy. Panel B indicates the setup when policy is never realized. Solid (dashed) lines indicate observed (unobserved) values. See discussion in Section IIA.

policy states. In particular, as noted in Section IC, the cap-and-trade prediction market did not explicitly mention the Waxman-Markey bill and thus may cover multiple similar policies. Furthermore, it is possible that alternative climate policies, such as a renewable portfolio standard, were considered at the time in lieu of cap-and-trade. Second, the cap-and-trade prediction market existed continuously over many trading days such that one can gain statistical power by utilizing daily time-series variation in prediction market prices and stock returns. More broadly, this section formally demonstrates how the observational challenge presented in Section IIA could be overcome using two variables that are observed on each trading date t: the stock return of firm i in sector j, r_{ijt}, and the cap-and-trade prediction market price, θ_t.

For a representative trader, define P as the set of all policies that could be implemented on any date.[25] Next, define $W \subset P$ as the subset of cap-and-trade policies to be passed by the end of 2010 covered by the cap-and-trade prediction market. Because the cap-and-trade prediction market did not specify the Waxman-Markey bill, W may include multiple policies in addition to Waxman-Markey. The remaining subset of policies not covered by the cap-and-trade prediction market is denoted by $O \subset P$ such that $W \cup O = P$. The set O includes the no-climate policy state and any non-cap-and-trade climate policy expected to be established by the end of 2010, such as a renewable portfolio standard. It also includes all policies—cap-and-trade or otherwise—to be established after 2010. Denote $v_{ijt}(W)$ and $v_{ijt}(O)$ as expected firm value conditional on the set of policies covered and not covered by the prediction market. Firm value, v_{ijt}, is the expected value over the two policy sets:

$$(1) \qquad v_{ijt} = \theta_t v_{ijt}(W) + (1 - \theta_t) v_{ijt}(O)$$
$$= v_{ijt}(O)(1 + \gamma_{ijt}\theta_t),$$

where the second line follows after defining

$$(2) \qquad \gamma_{ijt} = \frac{v_{ijt}(W) - v_{ijt}(O)}{v_{ijt}(O)},$$

or the percentage difference in expected firm value under cap-and-trade policies covered by the prediction market relative to other policies. I henceforth refer to γ_{ijt} as the firm-level cap-and-trade effect. Taking logs and first differences of equation (1), and noting that for sufficiently small $\gamma_{ijt}\theta_t$, $\ln(1 + \gamma_{ijt}\theta_t) \approx \gamma_{ijt}\theta_t$,[26] one can write stock returns, r_{ijt}, as

$$(3) \qquad r_{ijt} = \gamma_{ijt}\theta_t - \gamma_{ijt-1}\theta_{t-1} + \Delta \ln v_{ijt}(O),$$

where Δ indicates time differences.

An RDD design can exploit the free permit assignment at the 5 percent energy intensity threshold. To demonstrate this, I decompose the firm-level cap-and-trade effect into an average effect of the permit rule and a firm- and date-varying residual term. Specifically,

$$(4) \qquad \gamma_{ijt} = \beta D_j + \eta_{ijt},$$

where $D_j = \mathbf{1}\{EI_j > 0.05\}$ is an indicator variable which equals 1 when firm i is in sector j with energy intensity greater than 5 percent. The variable β is the

[25] For example, Waxman-Markey implemented at the end of 2010 would be a different element in P than the same policy but implemented on a different date.

[26] During the sample period, average $\theta_t = 0.24$, the average estimated cap-and-trade effect is -0.02.

reduced-form parameter of interest and captures the average trading-date difference in cap-and-trade effect between firms receiving and not receiving free permits. The residual η_{ijt} captures two additional components of the cap-and-trade effect: firm and date heterogeneity in the free permit rule treatment and all other firm and date-varying determinants besides the free permit rule. Section IV presents a structural analog for equation (4) so that β can be interpreted structurally. Substituting equation (4) into equation (3), followed by adding and subtracting $\eta_{ijt}\theta_{t-1}$, yields

$$r_{ijt} = \beta D_j \Delta\theta_t + \eta_{ijt}\Delta\theta_t + \Delta\eta_{ijt}\theta_{t-1} + \Delta \ln v_{ijt}(O). \tag{5}$$

Notice that equation (5) differs from a standard RDD setup in that the cross-sectional treatment variable, D_j, is interacted with a time-series variable $\Delta\theta_t$. This implies that identification requires two continuity assumptions, one which follows standard RDD convention and another because our event of interest was never realized. Formally, Assumption 1 states:

ASSUMPTION 1 (Continuity under cap-and-trade policy):

$$E_{it}[\eta_{ijt} \mid EI_j] \textit{ is continuous at } EI_j = 0.05,$$

where the expectation is taken over firms i and dates t. This is a variant of the standard RDD continuity assumption under a certain policy, stating that the cap-and-trade effect should be continuous at the 5 percent energy intensity threshold if not for the allocation of free permits during all trading days.[27] Assumption 1 may be violated if firms could sort around the threshold or if the threshold itself was strategically placed in response to firm activity.

Next, Assumption 2 states:

ASSUMPTION 2 (Continuity under non-cap-and-trade policy):

$$E_{it}[\ln v_{ijt}(O) \mid EI_j] \textit{ is continuous at } EI_j = 0.05.$$

This assumption states that expected stock returns under all other non-cap-and-trade policies must also be continuous at the 5 percent threshold. This assumption may be invalid if the effect of alternative policies on firm returns, such as a renewable portfolio standard, also exhibits a discontinuity at 5 percent. Under these two assumptions, estimation of equation (5) for across trading dates and for firms within a narrow bandwidth of the 5 percent energy intensity threshold would produce an unbiased estimate of the reduced-form parameter of interest:

$$\hat{\beta} = \frac{\lim_{EI_j \downarrow 0.05} E_{it}[r_{ijt} \mid EI_j] - \lim_{EI_j \uparrow 0.05} E_{it}[r_{ijt} \mid EI_j]}{E_t[\Delta\theta_t]}. \tag{6}$$

[27] Because η_{ijt} includes heterogeneity in the free permit treatment effect, Assumption 1 is robust to the presence of variable treatment effects (see Hahn, Todd, and Van der Klaauw 2001).

While Assumptions 1 and 2 are sufficient to identify β, for statistical power, estimation of equation (5) requires two modifications. First, because I include firms within a certain bandwidth around the 5 percent threshold with cap-and-trade effects that may vary with energy intensity, a component of η_{ijt} is modeled as a flexible function of energy intensity. Second, to capture $\Delta \ln v_{ijt}(O)$, I employ a variety of standard controls from the finance literature as well as semiparametric controls.

III. Reduced-Form Results

This section presents my reduced-form RDD results in several parts. First, I conduct two forms of graphical analysis via two-step procedures that illustrate the RD parameter in the time-series and cross-sectional dimensions. Second, for my preferred model, I estimate a one-step panel RD estimator which interacts time-series variation in prediction market prices with a local semiparametric function of the cross-sectional RD treatment allowing for correct inference. Point estimates from this panel estimator are shown to be robust across different controls for firm value absent cap-and-trade policy, sample bandwidth around the threshold, and functional form fitted to energy intensity. Third, I consider a series of indirect tests of identifying Assumptions 1 and 2 by examining whether there are discontinuities at the 5 percent threshold across various sector-level covariates and average stock returns before and after the period 2009–2010. Last, I present additional results testing for the presence of placebo discontinuities at thresholds less than and greater than 5 percent, biases due to thin-trading in the prediction market, and biases arising from concerns over prediction market expiration.

For all results, I construct a panel of two-day stock returns derived from data obtained from the Center of Research in Security Prices (CRSP)[28] during the period from May 1, 2009 to Dec 31, 2010 with all publicly traded firms in the 31–33 NAICS manufacturing sectors.[29] In the cross-sectional dimension, I restrict my sample to firms in sectors that satisfy Waxman-Markey's rule of trade intensity exceeding 15 percent using data from the US International Trade Commission.[30] Wolfers and Zitzewitz (2006) show that with certain utility functions, a favorite-longshot bias and reverse favorite-longshot bias can occur for prediction market prices below \$0.20 and exceeding \$0.80, respectively. Addressing this concern, I restrict my sample in the time-series dimension by only including trading days when prediction prices lie between \$0.20 and \$0.80. Data sources are further detailed in online Appendix D.

[28] Two-day trading intervals account for Intrade prediction markets having later closing hours than the primary US stock exchanges as well as the occurrence of after-hours stock trading. In particular, Intrade closing prices are observed at 2 AM on weekdays and 3 AM on weekends. If after-hours stock trading were to occur, the effect of information released after 4 PM ET on trading days or over weekends would not be picked up using observed daily returns. Applying the Hadri (2000) Lagrange Multiplier test on my sample panel of stock returns failed to reject a null that all panels are stationary in favor of an alternative where some panels contain a unit root.

[29] Only three firms in my sample indicate multiple six-digit NAICS sectors over my sample period. None of the sector changes affected free permit treatment assignment.

[30] In theory, one could conduct another RDD analysis along the 15 percent trade intensity cutoff, conditional on energy intensity above 5 percent. However, in practice, there are very few firms within narrow bandwidths of the trade-intensity cutoff. For example, within a 3 percent bandwidth of the 15 percent trade intensity threshold, there are 8 publicly listed firms in my dataset. By comparison, there are 202 firms within a 3 percent bandwidth of the 5 percent energy intensity threshold.

Online Appendix Table A.1 displays six-digit NAICS sectors by energy intensity bins.[31] For conciseness, only sectors within a 3 percent bandwidth around the 5 percent energy intensity threshold are shown with bins widths corresponding to the 2 percent, 2.5 percent, and 3 percent bandwidth samples. There does not appear to be clustering of particular three-digit manufacturing sectors on either side of the threshold. Firms in paper (NAICS 322), petroleum and coal products (NAICS 324), chemical (NAICS 325), plastics and rubber products (NAICS 326), nonmetallic mineral product (NAICS 327), primary metal (NAICS 331), fabricated metal product (NAICS 332), and computer and electronic product (NAICS 334) subsectors appear on both sides of the threshold. Section IIIB tests for discontinuities in sector-level covariates at the threshold.

A. *Main RD Estimate*

Equation (5) is estimated using firms within a certain bandwidth around the 5 percent threshold. Because stock returns may vary with energy intensity for firms within the bandwidth, I follow Hahn, Todd, and Van der Klaauw (2001) and consider a local polynomial regression analog of equation (5):

$$r_{ijt} = \beta D_j \Delta\theta_t + f(EI_j)\,\Delta\theta_t + \mathbf{Z}_{it}\boldsymbol{\Psi}_i + \epsilon_{ijt}, \tag{7}$$

where as before $D_j = \mathbf{1}\{EI_j > 0.05\}$ is an indicator variable which equals 1 when firm i is in six-digit NAICS sector j with energy intensity greater than 5 percent. Here, β is the RD parameter of interest; $f(EI_j)$ is a flexible polynomial function fully interacted with D_j, allowing for different parameters on either side of the threshold. For example, under a linear specification $f(EI_j) = \alpha_1 + \alpha_2(EI_j - 0.05) + \alpha_3 D_j(EI_j - 0.05)$. Continuity of the conditional mean of the error term, $E[\epsilon_{ijt}|EI_j]$, at the 5 percent threshold is required in order to satisfy identifying Assumptions 1 and 2.

Further, $\mathbf{Z}_{it}$ is a set of controls capturing stock returns in the absence of Waxman-Markey. This is typically known in the finance literature as *normal market performance*. Six sets of controls are considered. The first three models follow standard practice from finance. The first model includes only firm fixed effects. The second model is the CAPM with firm fixed effects and an aggregate market return index multiplied by a firm-specific coefficient. The third model employs the 3-factor Fama-French model (Fama and French 1993) which augments CAPM with value-based and size-based portfolios each multiplied by firm-specific coefficients to account for common risk factors associated with book-to-market ratio and firm size. Three additional models employ more flexible time controls. My fourth model includes a full set of firm and trading date fixed effects which removes daily shocks to stock returns common to all sample firms. To explicitly allow daily oil prices to have differential effects across firms, my fifth model augments the fourth model by adding interactions between the daily percentage change in crude oil prices with

[31] My energy intensity variable uses data from the 2004–2006 US Census Annual Survey of Manufacturers (ASM), as required by the Waxman-Markey bill. I prefer this variable to the energy intensity variable constructed for the Senate requested "Interagency Report on International Competitiveness and Emission Leakage" (EPA 2009) because the latter only uses 2006 ASM data. In practice the two variables are highly correlated with a bivariate linear model producing a coefficient of 0.98 and $R^2 = 0.92$.

two-digit NAICS sector dummies. Finally, my last model includes firm fixed effects and trading date by two-digit NAICS sector fixed effects to flexibly remove any trading date shocks that may have differential effects across sectors. Before showing panel model estimates from all six versions of equation (7), I first graphically present β in the time-series and cross-sectional dimensions.

Graphical Analysis (*Time Series*).—Equation (7) can be estimated via the following two-step procedure. In the first step, the following cross-sectional local polynomial regression is estimated for every trading date t:

$$r_{ijt} = \gamma_t D_j + f_t(EI_j) + \nu_{ijt}, \tag{8}$$

where γ_t is the differential abnormal returns between firms on either side of the 5 percent energy intensity threshold on each date t. The function $f_t(EI_j)$ includes a date fixed effect; γ_t is the raw probability unadjusted daily cap-and-trade effect. In the second stage, γ_t is regressed in the time-series dimension on the daily prediction market price change:

$$\hat{\gamma}_t = \beta \Delta\theta_t + \mu_t. \tag{9}$$

Results from this two-step procedure are shown graphically in Figures 3 and 4. To generate these graphs, I first remove firm fixed effects from the panel of stock returns. Equation (8) is then estimated on the ensuing residuals for firms within a 4 percent bandwidth of the cutoff with a flexible quadratic function, $f_t(EI_j)$, in energy intensity. Figure 3 plots the cumulative daily raw cap-and-trade effect over the 2007–2011 period. Notice that Figure 3 closely tracks the timing of Waxman-Markey developments discussed in Section IA and captured by cap-and-trade prediction market prices in Figure 1. The cumulative difference between firms on either side of the threshold hover around 0 during the period 2007–2009 and dramatically increase when the 111th Congress begins at the start of 2009. The cumulative difference peak around the House passage of the bill at the start of the summer of 2009 and converge back to zero thereafter toward the end of the 111th Congress as 2010 closes. Figure 4 shows equation (9) by plotting the relationship between the raw probability unadjusted daily cap-and-trade effect and cap-and-trade prediction market price changes during the May 1, 2009 to Dec 31, 2010 sample period. A linear fit indicates a β of 7 percent.

Graphical Analysis (*Cross Section*).—I reverse the steps of the previous procedure to provide a more conventional RD graph in the cross section, similar to that illustrated in Figure 2. First, for every sample firm, I estimate the following time series regression within the period from May 1, 2009 to Dec 31, 2010:

$$r_{ijt} = \gamma_{ij} \Delta\theta_t + c_i + \nu_{ijt}, \tag{10}$$

where γ_{ij} is the firm-level cap-and-trade effect and c_i are firm fixed effects. In the second stage, I test whether there is a cross-sectional discontinuity in γ_{ij} by estimating

$$\hat{\gamma}_{ij} = \beta D_j + f(EI_j) + \mu_{ij}. \tag{11}$$

Figure 3. Cumulative Daily Probability Unadjusted Cap-and-Trade Effect, 2007–2011

Notes: Daily probability unadjusted cap-and-trade effect, γ_t , estimated using equation (8) over the period 2007–2011. Vertical lines denote the 2009–2010 period covering the 111th US Congress.

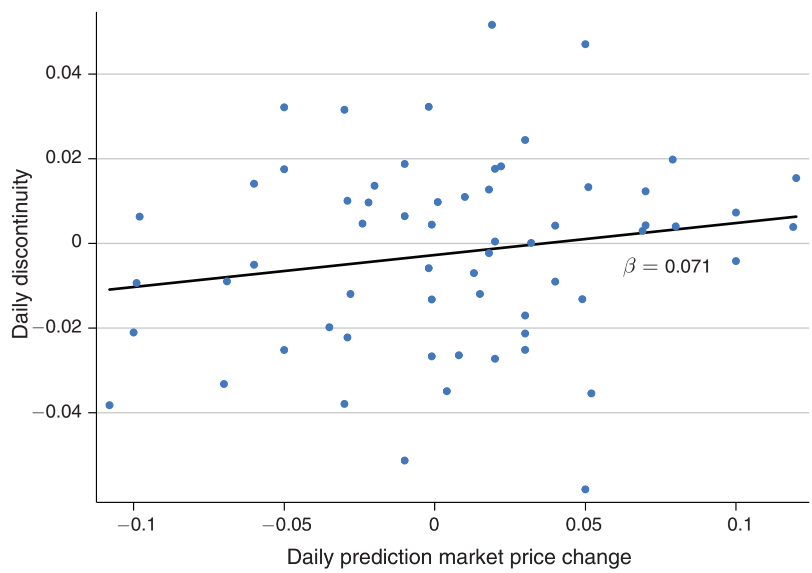

Figure 4. Daily Probability Unadjusted Cap-And-Trade Effect versus Prediction Market Price Change, 2009–2010

Notes: Scatter shows daily probability unadjusted cap-and-trade effect, γ_t, estimated using equation (8) against daily prediction market price changes within sample period. Linear fit from equation (9) shown in solid black line.

Figure 5 follows the RD literature by showing the nonparametric version of equation (11) for firms within a 4 percent bandwidth of the cutoff. To generate Figure 5, I first remove trading date fixed effects from the panel of stock returns

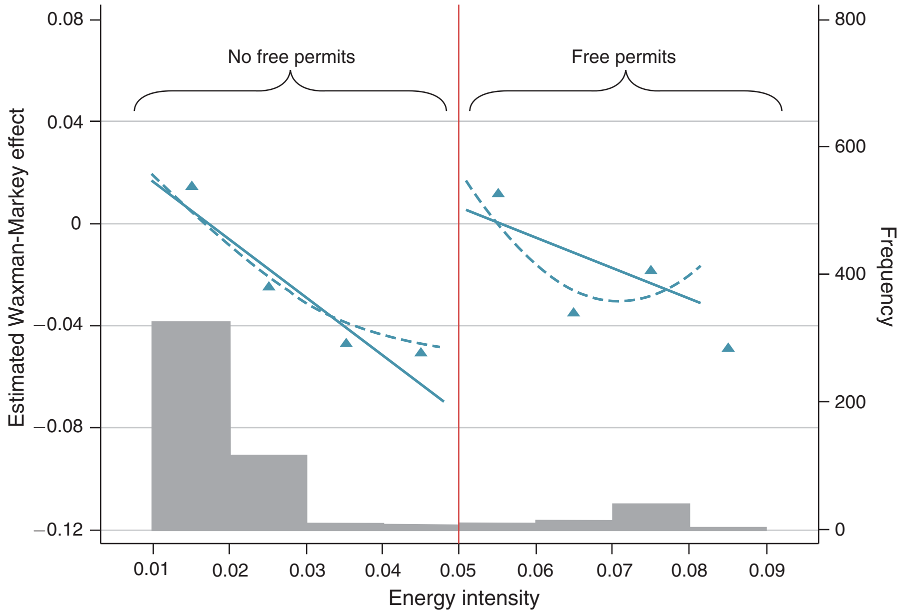

Figure 5. Discontinuity in Cap-and-Trade Effects at 5 Percent Energy Intensity

Notes: Triangles indicate local averages of firm-level cap-and-trade effects, γ_{ij}, estimated using equation (10) within a 0.01 wide bin of six-digit NAICS energy intensity. Solid (dashed) line shows a linear (quadratic) fit over the unbinned data. Distribution of sample firms by six-digit NAICS energy intensity shown in gray histogram.

before estimating equation (10). Triangles indicate the estimated cap-and-trade effect averaged over firms within separate 1-percent-wide energy intensity bins. The underlying support of sector-level energy intensity, the running variable, is shown in the gray histogram. The distribution is left-skewed with relatively fewer number of firms near the 5 percent energy intensity threshold. While this may be a concern for statistical power, it need not threaten identification provided that the firm distribution is continuous around the threshold. A density continuity test using the McCrary (2008) procedure was conducted and did not detect a discontinuity in the support of the running variable at 5 percent.[32] Solid lines show linear models fitted over the unbinned data on both sides of the 5 percent threshold while dashed lines show a similar fit using a quadratic model. Under both models, there is a discontinuity β of about 7 percent. Furthermore, it is noteworthy that the overall shape of the response function is negatively sloped on both sides of the threshold which is expected as more energy-intensive firms would experience larger losses under climate policy.

Panel Estimates.—The previous two-step RD procedures provide graphical evidence of the discontinuity at the 5 percent energy intensity threshold. However, they do not take into account sampling variability in estimates of $\hat{\gamma}_t$ and $\hat{\gamma}_{ij}$ and thus may

[32] The point estimate on a jump at 5 percent in the distribution of energy intensity is 0.78 with a standard error of 0.54.

Table 1—Discontinuity in Waxman-Markey Effects at 5 Percent Energy Intensity Threshold

		Bandwidth				
		2%	2.5%	3%	3.5%	4%
Model	Polynomial order	(1)	(2)	(3)	(4)	(5)
Panel A. Firm fixed effects						
	1	0.087	0.075	0.043	0.067	0.079
		(0.038)	(0.037)	(0.039)	(0.036)	(0.032)
	2	−0.027	0.082	0.103	0.073	0.071
		(0.033)	(0.049)	(0.043)	(0.035)	(0.035)
Panel B. CAPM						
	1	0.089	0.085	0.066	0.078	0.078
		(0.046)	(0.045)	(0.042)	(0.038)	(0.030)
	2	−0.046	0.062	0.097	0.084	0.088
		(0.035)	(0.050)	(0.051)	(0.046)	(0.045)
Panel C. 3-factor Fama-French						
	1	0.090	0.084	0.057	0.077	0.077
		(0.040)	(0.044)	(0.043)	(0.039)	(0.032)
	2	−0.040	0.067	0.104	0.082	0.090
		(0.037)	(0.045)	(0.046)	(0.041)	(0.042)
Panel D. Firm, date fixed effects						
	1	0.087	0.075	0.043	0.066	0.077
		(0.039)	(0.038)	(0.040)	(0.036)	(0.032)
	2	−0.027	0.081	0.104	0.073	0.071
		(0.033)	(0.048)	(0.043)	(0.035)	(0.035)
Panel E. Firm, date fixed effects; oil × sector						
	1	0.091	0.080	0.044	0.066	0.076
		(0.041)	(0.040)	(0.040)	(0.036)	(0.032)
	2	−0.032	0.079	0.103	0.073	0.071
		(0.033)	(0.049)	(0.043)	(0.035)	(0.035)
Panel F. Firm, date × sector fixed effects						
	1	0.064	0.050	0.042	0.071	0.069
		(0.031)	(0.036)	(0.038)	(0.036)	(0.031)
	2	0.033	0.087	0.090	0.077	0.079
		(0.058)	(0.051)	(0.039)	(0.036)	(0.038)
Number of firms		45	106	202	264	531

Notes: Each coefficient shows a separate regression discontinuity estimate of β from equation (7). Controls for normal market performance vary by panel. Polynomial order for energy intensity varies by rows within a panel. Sample bandwidth around the 5 percent threshold varies across columns. All models include 109 two-day intervals from May 1, 2009 to July 31, 2010. Robust standard errors clustered at six-digit NAICS level in parentheses.

not report correct standard errors.[33] Instead, my preferred procedure is to directly estimate equation (7) with standard errors clustered at the six-digit NAICS level to allow arbitrary forms of serial and cross-sectional correlation for firms within a given sector.

Table 1 shows estimates of β for each combination of (i) six different normal market performance controls (panels); (ii) local linear and quadratic functional forms for energy intensity (down rows within a panel); and (iii) estimation bandwidths

[33] This is akin to incorrect standard errors in a generated regressor setting.

ranging from 2 percent to 4 percent in 0.5 percent intervals around the threshold (across columns). Table 1 shows a statistically significant estimate of around 7 percent to 9 percent across these modeling choices.[34] A few patterns are worth discussing. First, point estimates are relatively unaffected by the choice of controls for normal market performance across panels of Table 1 lending support for Assumptions 1 and 2. Second, models using either a linear or quadratic function of energy intensity produce fairly similar estimates for bandwidths set at 2.5 percent, 3.5 percent, and 4 percent. The linear model fails to detect a discontinuity only within a 3 percent bandwidth though the point estimate is within the uncertainty of other models. The quadratic model fails to detect a discontinuity only within a bandwidth of 2 percent possibly due to overfitting given the reduced sample size at that bandwidth. Online Appendix Figure A.1 displays point estimates and 90 percent confidence intervals of β from the firm and date fixed effects model for finer 0.1 percent bandwidth increments using both linear (top panel) and quadratic (bottom panel) functions of energy intensity.

B. *Indirect Tests of Assumptions 1 and 2*

This section conducts several indirect tests of identifying Assumptions 1 and 2.

Sector-Level Covariates.—Assumption 1 requires that the cap-and-trade effect be continuous at 5 percent energy intensity if not for the free permit rule. Given the free permit rule, firm sorting around the 5 percent threshold is unlikely. First, as noted in Section IA, climate legislation prior to the Waxman-Markey bill did not include the 5 percent threshold such that it is unlikely that firms were anticipating this threshold based on prior experience. Second, free permit assignment is based on 2004 through 2006, publicly available, government data on sectoral energy intensity values. It is unlikely that firms in 2009 would rewrite such data. However, Assumption 1 could still be violated if there was strategic placement of the threshold by Congressional legislators, perhaps in response to firm activity.

Table 2 examines covariate balance for several six-digit NAICS sector-level variables that may influence the placement of the threshold. For each covariate, I consider both linear and quadratic models of energy intensity and different estimation bandwidths and apply the cross-sectional specification in equation (11) at the sectoral level. Panels A and B examine whether the location of the free permit rule reflected concerns over labor market effects. Panel A of Table 2 examines whether there is a discontinuity at 5 percent energy intensity in log average number of production workers for the sector. I do not detect a statistically significant discontinuity with any of my linear models and only pick up statistically significant coefficients within a 2 percent and 2.5 percent bandwidth using a quadratic function of energy intensity. Panel B models the log average number of establishments and detects only a marginally significant coefficient using a linear model of energy intensity within a 4 percent bandwidth. For neither covariate are the results sufficiently stable across modeling choices to conclude there is a discontinuity at the threshold. Firms

[34] For comparison, the average cap-and-trade effect across all listed firms estimated using the 3-factor Fama French model is −2 percent.

Table 2—Sector-Level Covariate Balance

		Bandwidth				
		2%	2.5%	3%	3.5%	4%
Covariate	Polynomial order	(1)	(2)	(3)	(4)	(5)
Panel A. log average number of workers						
	1	1.051	0.669	0.938	0.377	0.448
		(1.025)	(1.143)	(0.994)	(0.921)	(0.843)
	2	2.576	1.688	1.395	1.292	1.240
		(0.911)	(0.931)	(0.876)	(0.812)	(0.830)
Number of sectors		10	15	25	41	66
Panel B. log average number of establishments						
	1	0.643	0.487	0.755	1.050	1.106
		(0.881)	(0.899)	(0.768)	(0.659)	(0.624)
	2	1.476	1.160	0.512	0.485	0.887
		(1.058)	(0.955)	(0.907)	(0.885)	(0.826)
Number of sectors		21	29	56	87	133
Panel C. Herfindahl-Hirschman index						
	1	−0.275	−0.258	−0.053	−0.016	−0.084
		(0.274)	(0.305)	(0.255)	(0.219)	(0.213)
	2	−0.482	−0.319	−0.419	−0.432	−0.303
		(0.319)	(0.252)	(0.251)	(0.266)	(0.249)
Number of sectors		20	28	53	83	126
Panel D. log Waxman-Markey lobbying (nominal $)						
	1	3.317	2.481	1.998	1.382	1.508
		(2.406)	(1.918)	(1.469)	(1.328)	(1.117)
	2	−0.473	2.264	3.506	3.215	1.830
		(4.506)	(3.934)	(2.806)	(2.269)	(1.926)
Number of sectors		11	16	31	46	64
Panel E. log pre-2009 climate policy lobbying (nominal $)						
	1	4.020	3.604	2.267	2.442	2.122
		(3.116)	(2.524)	(1.928)	(1.721)	(1.490)
	2	4.643	2.072	5.110	3.368	2.871
		(6.713)	(4.815)	(3.504)	(3.086)	(2.549)
Number of sectors		11	15	29	37	48

Notes: Each coefficient shows a separate regression discontinuity estimate of β from estimating equation (11) at the sector level. Sector-level outcomes vary by panel. Functional forms for energy intensity vary by rows within a panel. Sample bandwidths around the 5 percent threshold vary across columns. All models include firm and date fixed effects. Robust standard errors clustered at six-digit NAICS level in parentheses.

that can pass through more of its regulatory costs onto consumers would experience lower profit losses. Panel C tests whether there is a discontinuity at 5 percent energy intensity in market concentration as measured by a six-digit NAICS-level Herfindahl-Hirschman index (HHI) constructed from average daily firm stock values in the period 2007–2008. I do not detect a jump in market concentration.

Legislators may have determined the 5 percent threshold in response to firm lobbying activity. Panel D examines the presence of a jump at 5 percent energy intensity in log lobbying expenditures on the Waxman-Markey bill (see online Appendix D for details on lobbying expenditure construction). I do not detect a statistically significant discontinuity at 5 percent energy intensity for any bandwidth and functional form in energy intensity suggesting that there was not additional lobbying pressure by firms in sectors that would have received free permits. It is, however, possible that this free permit allocation rule was determined by climate policy lobbying prior to the Waxman-Markey bill in 2009. In panel E, I extract and total lobbying expenditures across the full set of official climate legislation considered by Congress between 2003 and 2008.[35] Again, I do not detect a jump in lobbying expenditures on climate bills prior to the Waxman-Markey bill at the 5 percent threshold. Finally, it should be noted that while I do not systematically detect discontinuities in any of these covariates, some of the noisy point estimates in Table 2 are quite large. Given the small number of sectors in my estimating samples, it may be that discontinuities exist in these covariates and I am statistically underpowered in detecting them.

Firm-Level Stock Returns in 2007, 2008, and 2011.—Assumption 2 requires that conditional on controls for normal market performance, firm stock returns are continuous at 5 percent energy intensity in the absence of cap-and-trade policies covered by the prediction market. An indirect test of Assumption 2 is to examine whether stock returns in 2007 and 2008, prior to the 111th Congress, and in 2011, after the 111th Congress, exhibit discontinuities at 5 percent energy intensity. Assumption 2 would be violated if discontinuities at the 5 percent threshold existed in any alternative policy not covered by the cap-and-trade prediction market and markets believed such policy to be probable prior to or after 2009–2010. Detecting a discontinuity in stock returns prior to 2009–2010 could also imply that markets were anticipating this threshold prior to the 111th Congress.

Figure 3 already shows the discontinuity in cumulative daily stock returns at the 5 percent threshold during 2007, 2008, and 2011. Table 3 estimates the discontinuity using the cross-sectional specification in equation (11) for average daily returns in 2007, 2008, and 2011 in panels A, B, and C respectively. Rows in each panel show coefficients from models fitting a linear and quadratic function of energy intensity. Columns indicate the bandwidth size around the 5 percent threshold. I find little evidence of a discontinuity at 5 percent energy intensity across the combination of stock returns, polynomial order, and bandwidth window. Nearly every coefficient and standard error is small in magnitude and not statistically significant. I only detect a statistically significant effect in 2011 average stock returns using a quadratic function in energy intensity within a 2 percent bandwidth. This is represented graphically in online Appendix Figure A.2 which shows average daily stock returns in 2007, 2008, and 2011 for each 1 percent-wide energy intensity bin along with linear and quadratic models fitted over unbinned values.

C. *Additional Robustness Tests*

I turn now to several additional robustness tests of my main RD result.

[35] These include the 2003 McCain-Lieberman (S.R. 139); 2005 McCain-Lieberman (S.R. 1151); 2007 McCain-Lieberman (S.R. 280); 2007 Lieberman-Warner (S.R. 2191); and the 2008 Boxer-Lieberman-Warner (S.R.3036) bills.

Table 3—Firm-Level Covariate Balance

		Bandwidth				
		2%	2.5%	3%	3.5%	4%
Outcome	Polynomial order	(1)	(2)	(3)	(4)	(5)
Panel A. 2007 average stock returns						
	1	0.000	0.000	−0.001	−0.001	0.000
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	2	−0.001	0.000	0.001	0.000	−0.001
		(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
Number of firms		44	100	193	253	514
Panel B. 2008 average stock returns						
	1	−0.001	−0.000	−0.001	0.000	−0.001
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	2	−0.002	−0.001	−0.000	−0.001	−0.000
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Number of firms		45	106	202	264	530
Panel C. 2011 average stock returns						
	1	−0.000	−0.000	−0.001	−0.001	−0.001
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	2	−0.002	−0.001	−0.000	−0.000	0.001
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Number of firms		45	106	199	260	518

Notes: Each coefficient shows a separate regression discontinuity estimate of β from estimating equation (11) at the firm level. Firm-level outcomes vary by panel. Functional forms for energy intensity vary by rows within a panel. Sample bandwidths around the 5 percent threshold vary across columns. All models include firm and date fixed effects. Robust standard errors clustered at six-digit NAICS level in parentheses.

Placebo Test at Different Energy Intensity Thresholds.—Table 4 shows estimates of β from equation (7) under alternate placebo free permit thresholds at 3 percent, 4 percent, 6 percent, and 7 percent energy intensity, shown across columns. All models include both firm and date fixed effects estimated within a 4 percent estimation sample bandwidth on both sides of the treatment threshold with cluster-robust standard errors at the six-digit NAICS level. The first row of coefficients are from models fitting a local linear function of energy intensity while the second row uses a local quadratic function. For both functional forms, a statistically significant discontinuity is detected only at the correct 5 percent threshold.

Possible Consequences of Thin Trading in Prediction Market.—Section IC noted that the Intrade cap-and-trade market was thinly traded. Previous field and experimental studies have found that prediction market prices are relatively unaffected by thin trading or manipulation by individual traders (see online Appendix B for a brief literature review). Intuitively, efforts to manipulate a prediction market by some increase arbitrage opportunities for others such that distortions, if they do exist, are unlikely to last over long periods.

Because bias due to thin trading cannot be fully ruled out, online Appendix B formally considers three possible sources of deviation between the observed prediction market price and the true unobserved cap-and-trade probability: additive bias, classical measurement error, and multiplicative bias. First, I show that the inclusion of

Table 4—Placebo Tests at Different Energy Intensity Thresholds

	Energy intensity threshold				
	3%	4%	5%	6%	7%
Polynomial order	(1)	(2)	(3)	(4)	(5)
1	−0.035	−0.020	0.077	0.003	0.055
	(0.046)	(0.038)	(0.032)	(0.053)	(0.046)
2	−0.040	0.029	0.071	−0.011	0.008
	(0.072)	(0.065)	(0.035)	(0.095)	(0.061)
Number of firms	1,445	1,485	531	207	91

Notes: Each coefficient shows a separate regression discontinuity estimate of β from equation (7). Placement of placebo energy intensity threshold vary across columns. All models include firm and trading date fixed effects within a 4 percent energy intensity bandwidth. All models include 109 two-day intervals from May 1, 2009 to July 31, 2010. Robust standard errors clustered at six-digit NAICS level in parentheses.

firm fixed effects in each variant of my panel specification in equation (7) indirectly addresses additive bias. Second, I demonstrate that in the presence of classical measurement error, the interaction between my cross-sectional RD treatment variable D_j and the prediction market price change, $\Delta\theta_t$, results in lower attenuation bias than if the prediction market price change were uninteracted. Third, I derive the resulting bias in β in the presence of multiplicative bias.

Online Appendix Table A.2 provides indirect tests of whether multiplicative bias in the prediction market price affects my RD estimate. Each column shows an interacted version of equation (7):

$$(12) \quad r_{ijt} = \beta_1 D_j \Delta\theta_t + \beta_2 D_j \Delta\theta_t W_t + f(EI_j)\,\Delta\theta_t + f(EI_j)\,\Delta\theta_t W_t + \mathbf{Z}_{it}\boldsymbol{\Psi}_i + \epsilon_{ijt},$$

where W_t is a time-series variable that captures the competitiveness of the prediction market during date t. Now, β_2 is the coefficient of interest and indicates whether the estimated discontinuity is affected by the presence of potential distortions in the prediction market. Online Appendix Table A.2 shows results from models with firm and date fixed effects with both a linear and quadratic term for energy intensity. Column 1 replicates the main uninteracted coefficient from equation (7). In column 2 I include an interaction with the total trading volume for the two-day period. The interacted coefficient is close to zero and statistically insignificant while the uninteracted coefficient is similar to that in column 1 for both linear and quadratic models. Transaction-level data revealed the presence of two large traders. One large trader, a major buyer, was responsible for 38 percent of all contracts sold before the contract expired. Another large trader was responsible for 22 percent of all contracts purchased. In column 3 I examine the potential role of the two large volume traders by including an interaction with trading volume attributed to these two individuals. Again, the interacted coefficient is nearly zero and statistically insignificant while the uninteracted coefficient is similar to that in column 1. Simply considering the involvement of these two large traders, however, does not preclude other trading days in which the market was dominated by relatively few traders. Using transaction-level data with unique trader identifiers, I construct a daily buyer-based normalized Herfindahl-Hirschman Index.[36] This index captures the relative competitiveness of the prediction market for any given day and provides my most stringent test, shown in column 4. For the linear model, I detect an uninteracted term that is similar in magnitude to that in column 1 though it itself is statistically insignificant. Importantly, I do not detect a statistically significant interaction term. Coefficients for the quadratic model are also not statistically significant though point estimates are now of larger magnitude. Finally, in columns 5 and 6 I split my sample into 2009 and 2010 trading days. I do not find that estimates of β from each subsample differ much in magnitude from the full sample estimate in column 1.

Concerns over Prediction Market Expiration.—Another source of multiplicative bias may arise from Intrade contract expiration. The cap-and-trade prediction market used for this analysis expired on December 31, 2010, regardless of whether cap-and-trade regulation were to eventually pass Congress. Thus, while the prospects of cap-and-trade realization might indeed be declining in 2010, a component of the price movements shown in Figure 1 might also reflect expectations that policy realization is unlikely to occur before the end of 2010. In practice this was unlikely, as any legislation, having failed in the current Congress, is rarely reintroduced with identical features in a subsequent Congress. However, it is difficult to ascertain whether markets expected Waxman-Markey prospects to exist following the end of the 111th Congress. If so, a bias is introduced between the prediction market price and average market beliefs which increases as the expiration date nears. In online Appendix C, I detail an adjustment procedure to separate average market beliefs, the true variable of interest, from concerns over contract expiration. This procedure uses information from a similar Intrade prediction market with an earlier expiration date at the end of 2009 (see online Appendix Figure A.3). Under certain additional assumptions, I can use the period of overlap between the 2009 and 2010 expiring contracts to separate the effects of concerns over contract expiration with the true market belief in cap-and-trade prospects.

In online Appendix Table A.3, I find that adjusting for contract expiration yields estimates of β that are similar to my main result. While these estimates are slightly smaller in magnitude, they fall well within the uncertainty of my main results shown in Table 1. This is possible because whereas the adjustment procedure illustrated in online Appendix Figure A.5 inflates prediction market price levels to account for concerns of impending contract expiration, much of this adjustment is already removed from the unadjusted prediction market price after first differencing.

IV. Structural Interpretation: Marginal Abatement Cost under Waxman-Markey

Sections II and III provided the estimating framework and results for the reduced-form RD parameter β, the difference in stock returns for firms on either side of the free permit threshold under cap-and-trade policies covered by the prediction market relative to other policies. This section discusses the additional steps required for

[36] Formally, for trading day t, there are $j = 1, \ldots, J_t$ traders each purchasing s_{jt} share of all contracts transacted that day. The normalized Herfindahl-Hirschman Index is $H_t^* = \frac{H_t - 1/J_t}{1 - 1/J_t}$, where H is the Herfindahl-Hirschman Index, $H_t = \sum_j s_{jt}^2$.

mapping β onto the structural parameter of interest: the marginal abatement cost under the Waxman-Markey policy. Because the Waxman-Markey bill contained auxiliary abatement policies that may induce additional reductions in capped sectors, my recovered parameter corresponds to the cost of meeting the cap, conditional on these other policies, and may not equal the US economy-wide marginal abatement cost. For comparison purposes, most analyses obtained from CGE models will also model the presence of these auxiliary policies.

This section has four parts. First, I introduce a benchmark model of cap-and-trade regulation and consider two assumptions, one restricting variation in key features across cap-and-trade policies covered by the prediction market and another restricting variation in permit prices across trading dates. Second, I discuss a bounding analysis for the cumulative free permits expected by a treated firm to recover bounds on the marginal abatements cost under Waxman-Markey. Third, I consider two deviations from the benchmark model that may be relevant for the Waxman-Markey bill. Finally, I compare my range of implied Waxman-Markey marginal abatement costs with that estimated by prominent CGE models.

A. *Mapping Reduced Form onto Structural Parameters*

Under standard theoretical assumptions, the equilibrium permit price of a cap-and-trade policy during any period equals that period's overall marginal abatement cost of meeting the policy's aggregate cap (Montgomery 1972). Furthermore, a cap-and-trade policy with unlimited permit borrowing and banking over time results in permit prices that follow Hotelling's rule (Hotelling 1931), rising at the (exogenous) interest rate. This implies that only knowing the first period permit price, τ_0, is adequate for the time path of the overall marginal abatement cost. Thus, my structural parameter of interest is $\tau_0(w)$, the first period permit price under Waxman-Markey, $p = w$.

I consider a standard model of cap-and-trade regulation which depletes a fixed stock of aggregate emissions over the lifetime of the policy, $y \in [0, Y]$. Following Rubin (1996), output and permit markets are perfectly competitive, free permits are allocated in a lump-sum manner, and there is unlimited permit banking and borrowing over time. Section IVC explores implications when the second and third conditions are violated. For any single cap-and-trade policy covered by the prediction market, $p \in W$, optimal value for firm i in sector j during trading date t is (see online Appendix A.1 for full derivation)

$$v_{ijt}(p) = \tau_{t0}(p) A_{ij}(p) \mathbf{1}\{j \in F(p)\} + \int_0^Y e^{-\delta y} \pi(x^*_{ijty}(p))\, dy - \tau_{t0}(p) \int_0^Y x^*_{ijty}(p)\, dy, \tag{13}$$

where A_{ij} is the cumulative free permits granted to firm i over the lifetime of the policy and, as in Section IIA, F denotes the sectors receiving free permits. Here, δ is the exogenous interest rate, $\pi(\cdot)$ is a general concave instantaneous profit function, and x^*_{ijty} is optimal emissions. Notice that there are two differences between the first period permit price in equation (13) and my structural parameter of interest, the trading date-invariant first period permit price under Waxman-Markey. First, because the cap-and-trade prediction market did not specify the Waxman-Markey bill, $\tau_{t0}(p)$ is indexed by policy $p \in W$, which may contain other policies besides Waxman-Markey. Second, $\tau_{t0}(p)$ is indexed by trading date t.

To recover $\tau_0(w)$, I turn to two additional assumptions.

ASSUMPTION 3 (Homogeneity in cap-and-trade policies):

(*i*) $F(p) = F(w)\ \forall p \in W$;

(*ii*) $A_{ij}(p) = A_{ij}(w)\ \forall i, p \in W$;

(*iii*) $\tau_{t0}(p) = \tau_{t0}(w)\ \forall p \in W$.

Assumption 3 states that the various cap-and-trade policies covered by the prediction market contract exhibit the same free permit rule, firm-level cumulative free permits, and first period permit price. While it is impossible to identify each element of the set of cap-and trade policies, W, believed by market participants to be covered by the prediction market, two arguments are provided in favor of this assumption. First, as discussed in Section IA, the two most notable cap-and-trade policies considered by the 111th Senate contained key features that were identical to that found in Waxman-Markey. Specifically, the Kerry-Boxer and Kerry-Lieberman bills had identical free permit rules, consistent with Assumptions 3(i) and 3(ii). Furthermore, these Senate bills had similar provisions that would have affected permit prices such as sectoral coverage, cap schedules, international offset, and domestic offset provisions, consistent with Assumption 3(iii) (see online Appendix F and online Appendix Table A.4). Second, the placebo tests shown in Table 4 did not detect jumps in energy intensity at thresholds greater or less than 5 percent energy intensity, consistent with Assumption 3(i).[37]

For firms receiving free permits, Assumption 3 allows the cap-and-trade effect, first introduced in equation (2), to be interpreted as the Waxman-Markey effect. Applying equation (13),

$$\gamma_{ijt} = \frac{\tau_{t0}(w)A_{ij}(w)\mathbf{1}\{j \in F(w)\}}{v_{ijt}(O)} + \frac{\int_0^Y e^{-\delta y}\pi(x_{ijty}^*(w))\,dy - \tau_{0t}(w)\int_0^Y x_{ijty}^*(w)\,dy}{v_{ijt}(O)} - 1, \tag{14}$$

which is the structural analog to the reduced-form expression in equation (4). The first term captures the discounted value of cumulative free permits under Waxman-Markey. The second and third terms capture all other changes in stock returns under the policy, corresponding to η_{ijt} in equation (4).

[37] Notice that Assumption 3 is less restrictive than assuming that the cap-and-trade effect be the same for all policies covered by the prediction market.

The first period permit price in equation (14) is indexed by trading date t and reflect the role of daily aggregate shocks during 2009–2010 that may affect expectations over future abatement costs. To recover a trading date invariant first period permit price, I further assume that daily aggregate shocks do not systematically affect both the Waxman-Markey first period permit price and firm value in the absence of Waxman-Markey. Consider the following decomposition, $\tau_{t0}(w) = \tau_0(w) + \xi_t$, with $E_t(\xi_t) = 0$.

ASSUMPTION 4 (Uncorrelated permit price): $\operatorname{cov}\left(\xi_t, \frac{1}{v_{ijt}(O)}\right) = 0$.

In the context of the schematic in Figure 2, Assumption 4 requires that day-to-day vertical shifts in the conditional mean function in the absence of cap-and-trade policy, $E[r_i | p = o, EI_j]$, are uncorrelated with vertical shifts in β. In other words, aggregate determinants of baseline firm value absent cap-and-trade should not be systematically related to the demand for permits under cap-and-trade.[38] The strongest evidence in favor of this assumption comes from the stability of the RD parameter across different trading date controls, as shown down each panel of Table 1. In particular, the RD parameter is similar whether one omits trading date fixed effects (panels A–C), controls for trading date fixed effects (panel D), or controls for trading date-by-sector fixed effects (panel F). One can now write the following mapping between the reduced-form RD parameter β, which takes expectations across firms i and trading date t, and the first period permit price under Waxman-Markey:

$$\widehat{\beta} = E_{it}\left[\frac{\tau_{0t}(w) A_{ij}(w)}{v_{ijt}(O)}\right] \tag{15}$$

$$= \tau_0(w) E_i\left[A_{ij}(w) E_t\left[\frac{1 + \gamma_{ijt}\theta_t}{v_{ijt}}\right]\right] \tag{16}$$

$$= \tau_0(w)\left(E_i\left[A_{ij}(w) E_t\left[\frac{1}{v_{ijt}}\right]\right] + \widehat{\beta} E_i\left[A_{ij}(w) E_t\left[\frac{\theta_t}{\nu_{ijt}}\right]\right] + E_i\left[A_{ij}(w) E_t\left[\frac{\widehat{f(EI_j)}\,\theta_t}{v_{ijt}}\right]\right]\right), \tag{17}$$

where equation (16) employs the Law of Iterated Expectations, Assumption 4, and substitutes $v_{ijt}(O)$ using equation (1). Equation (17) applies the estimating equation

[38] As an example, consider a regulated oil-consuming manufacturing firm facing two types of aggregate shocks: a drop in oil prices and an increase in aggregate demand. The value of the firm should rise in both cases but permit prices may change in different directions. If lower oil prices induce fuel-switching away from coal-fired electricity, permit prices will drop from lower permit demand from the electricity sector. On the other hand, an increase in aggregate demand would increase emissions and permit demand, raising permit prices.

for γ_{ijt} from equation (11) (see online Appendix A.2 for detailed steps). Rearranging equation (17) in terms of $\tau_0(w)$ and replacing expectations with sample means

$$(18)\quad \tau_0(w) = \frac{\widehat{\beta}}{\frac{1}{N}\sum_i^N\left[A_{ij}(w)\frac{1}{T}\sum_i^T\left[\frac{1}{v_{ijt}}\right]\right] + \widehat{\beta}\frac{1}{N}\sum_i^N\left[A_{ij}(w)\frac{1}{T}\sum_i^T\left[\frac{\theta_t}{v_{ijt}}\right]\right] + \frac{1}{N}\sum_i^N\left[A_{ij}(w)\widehat{f(EI_j)}\frac{1}{T}\sum_i^T\left[\frac{\theta_t}{v_{ijt}}\right]\right]},$$

where N and T are the number of treated firms and trading days in the estimation sample, respectively. Observe that every term in equation (18) is either directly observed (e.g., v_{ijt}, θ_t) or estimated (e.g., $\widehat{\beta}, \widehat{f(EI_j)}$) with the exception of $A_{ij}(w)$, the firm-level cumulative number of free permits. I now turn to a bounding procedure for $A_{ij}(w)$.

B. *Bounding Cumulative Free Permits*

There are two broad approaches for obtaining $A_{ij}(w)$. One approach is to explicitly model firm-level emissions from 2012 through 2025, the period during which free permits are granted. This requires knowing the primitive parameters that determine future emissions. An alternative approach is to provide reasonable bounds for $A_{ij}(w)$ using recent emission trends. I pursue this latter approach by employing worst- and best-case scenarios as discussed by Horowitz and Manski (2000). Specifically, I look to historical data and obtain the highest and lowest rates of emission decline observed for manufacturing subsectors during the period 2006–2010.[39] I then apply either of these two rates *uniformly* to *all* treated firms in my sample from 2012 through 2025 to generate emission paths, and thus cumulative free permits. For this exercise to be valid, I am implicitly assuming that actual average emissions for sample firms during 2012 through 2025 must not decline faster (slower) than the single fastest (slowest) emission declining subsector during 2006 through 2010. The first and second columns of Table 5 display rates of emission decline for available manufacturing subsectors during 2006 through 2010. During this recent period, the petroleum refinery sector (NAICS 324110) experienced the lowest drop at -0.7 percent while the textiles sector (NAICS 313–316) experience the highest drop at -19.3 percent annually. Notice that because A_{ij} appears in the denominator of equation (18), a slower rate of emission decline implies a higher A_{ij} and thus a lower $\tau_0(w)$.[40]

Specifically, my bounding analysis follows three steps. First, I obtain 2011 average firm-level emissions for all six-digit NAICS manufacturing subsectors receiving free permits. 2006 was the last year with available subsector emissions which I divide by the number of firms in 2006 to get average firm-level emissions. To obtain 2011 subsector average firm-level emissions, I reduce 2006 values by a uniform rate of 1.5 percent annually, the rate of change for the overall NAICS 31–33

[39] It is highly unlikely that emission levels should increase on average for manufacturing firms under a cap-and-trade policy. In theory, under any strictly positive permit price, all firms, regardless of permit allocation and marginal abatement cost, should weakly decrease emissions relative to business as usual.

[40] Similarly, it is possible that markets expected future policy revisions to extend the duration of the free permit rule beyond 2035 such that A_{ij} increases, implying a lower $\tau_0(w)$. Ideally, such a scenario would not lead to values below my lower bound for $\tau_0(w)$.

Table 5—2015 Marginal Abatement Cost (2009$) Implied by RD Estimates

Assumed annual emissions rate (percent)	Corresponding sector (NAICS)	5th percentile	Mean	95th percentile
0	–	1.26	4.59	7.72
−0.70	Petroleum refining (324110)	1.33	4.88	8.19
−1.45	All manufacturing (31–33)	1.42	5.19	8.72
−5.20	Forest products (321, 322)	1.93	7.05	11.84
−7.60	Alumina and Aluminum (3313)	2.32	8.49	14.26
−11.60	Cement (327310)	3.1	11.36	19.08
−12.80	Glass (3272)	3.37	12.34	20.73
−13.90	Transport. equip. (336)	3.63	13.29	22.32
−19.30	Textiles (313–316)	5.09	18.65	31.31
−25	–	6.99	25.57	42.94
−30	–	8.95	32.78	55.04
−35	–	11.25	41.18	69.15

Notes: First column shows assumed annual emissions rate. Second column shows corresponding manufacturing subsector with 2006–2010 emissions changing at each rate. RD estimates (see equation (7) based on Table 1, panel C, row 1, and column 5); 90 percent confidence interval generated using 250 Monte Carlo draws from estimated parameter and variance-covariance matrix. Marginal abatement cost recovered using equation (18); 5 percent interest rate assumed.

manufacturing sector.[41] Second, for the 2012–2025 period under Waxman-Markey, I assume that annual emission rates for all treated firms follow that of the fastest or slowest declining manufacturing subsector during 2006–2010. Third, for the 2026–2035 period, I impose an annual reduction of 10 percent in 2026, 20 percent in 2027, up to a complete 100 percent phase-out by 2035, as specified by Waxman-Markey.

To illustrate this procedure, online Appendix Figure A.6 plots the permits received for an average firm in the plastics material and resin manufacturing subsector under a worst-case annual emission rate of −0.7 percent (petroleum refinery) and a best-case annual emission rate of −19.3 percent (textiles). Also included are the annual permit paths under annual emission rates of −1.5 percent (all manufacturing), −7.6 percent (aluminum), −13.9 percent (transport), and −35 percent for the 2012–2025 period. The vertical line marks the start of the policy such that the area under each curve represents cumulative free permits, $A_{ij}(w)$, associated with each assumed emission rate.

Table 5 displays the corresponding implied first period permit price, or marginal abatement cost, using equation (18) with $\widehat{\beta}$ and $\widehat{f(EI_j)})$ obtained from the panel data model shown in panel C, row 1, column 5 of Table 1.[42] The reported 90 percent confidence interval is obtained via Monte Carlo draws from the estimated joint variance-covariance matrix of $\widehat{\beta}$ and $\widehat{f(EI_j)})$. The marginal abatement cost, shown in 2009 US$, are for 2015 under an assumed annual interest rate δ of 5 percent.[43] I recover a lower bound marginal abatement cost of \$4.88 when assuming emissions

[41] Greenhouse gas emission levels for all six-digit NAICS subsectors are available for 2006 from EPA (2009). Later data are available from the US DOE for 2010 and cover only emissions for the overall manufacturing sector (NAICS 31–33) and for a dozen select manufacturing subsectors. Thus, I am unable to observe rates of emission change for all six-digit NAICS subsectors after 2006. See online Appendix D for more data details.

[42] The panel models using Fama-French factors and trading date fixed effects produce nearly identical point estimates and uncertainty for β. However, the panel model with date fixed effects does not separately identify the intercept in $\widehat{f(EI_j)})$ and so cannot be used to recover $\tau_0(w)$.

[43] For comparison purposes, an interest rate of 5 percent is chosen because it is the mean interest rate among the CGE models discussed in Section IVD.

from all treated firms decline under Waxman-Markey at the recent rate of the worst performing manufacturing subsector. Similarly, I recover an upper bound of \$18.65 when applying the recent rate of the best performing subsector. I recover marginal abatement costs exceeding \$25 per ton of CO_2e only when assuming emission decline rates lower than -25 percent per year.

C. *Deviations from the Benchmark Model*

I now turn to potential biases in the recovered Waxman-Markey first period permit price, $\tau_0(w)$ from two violations of the benchmark cap-and-trade model in Section IVA.

Restrictions on Permit Borrowing.—The benchmark cap-and-trade model assumes unrestricted banking and borrowing of permits. While permit banking is unlimited under Waxman-Markey, a firm would incur a $\rho = 0.08$ annual borrowing cost for each borrowed permit. For example, if a firm chooses to use one 2016-vintaged permit for 2015 compliance, it must retire (i.e., not use for compliance) 1.08 2016-vintaged permits in 2016. Similarly, if that firm uses one 2017-vintaged permit for 2015 compliance, it must retire 1.16 2017-vintaged permits in 2017.[44]

Unfortunately, knowing when the borrowing restriction would bind requires accurate forecasts of permit demand during the policy. As with Section IVB, I consider an alternative bounding approach. If the borrowing restriction never binds, then the permit price rises at the rate of interest as in the benchmark model. If, however, this restriction were to bind during any period, the permit price would rise during that period at less than the interest rate (Rubin 1996; Schennach 2000). Thus, assuming that the borrowing restriction binds in every period would generate an upper bound on the implied first period permit price.

Online Appendix A.3 formally details the implication of borrowing restrictions in an extended cap-and-trade model with vintage-specific permits. Define τ_{vyt} as the permit price of vintage v in period y on trading date t. When the borrowing restriction binds in every period, equation (15) becomes

$$\widehat{\beta} = E_{it}\left[\frac{\tau_{00}(w)}{v_{ij}(O)}\int_{y=0}^{Y} e^{-\rho y} A_{ijy}(w)\,dy\right], \tag{19}$$

which, for a given value of $\widehat{\beta}$, implies a higher first period permit price than if borrowing was unrestricted. Online Appendix Table A.6 replicates Table 5 in the main text showing the recovered first period permit price assuming that the borrowing restriction binds in all periods. The resulting lower and upper bound marginal abatement costs are \$8.41 and \$24.29 per ton CO_2e in 2015, or \$3.53 and \$5.64 higher than under unlimited permit borrowing.

[44] Technically, there are two further restrictions. First, vintaged permits 5 years beyond the current compliance period cannot be borrowed and at most 15 percent of total compliance in any year can be met using borrowed permits. For the sake of simplicity, I do not model these additional restrictions which would result in higher implied initial period permit prices. However, it is unclear if these additional constraints will ever bind during the lifetime of the policy.

Output-Based Allocation.—The benchmark cap-and-trade model assumes that permits are allocated in a lump-sum manner, under which firms face the same marginal incentives regardless of permit allocation.[45] In reality, the Waxman-Markey bill allocates free permits based on an output-based allocation rule which can affect firm-level output in the short run as well as exit decisions in the long run. The rule states that free permits received by firm i in sector j in year y is determined by

$$A_{ijy}(w) = q_{ijy}\frac{X_{jy}}{Q_{jy}}, \tag{20}$$

where q_{ijy} is output of firm i in year y and X_{jy} and Q_{jy} are total emissions and output in sector j such that X_{jy}/Q_{jy} is sector-level emissions intensity.[46]

Under output-based allocation, the allocation of free permits is no longer exogenous to the firm as it now faces an additional incentive to increase output. This implies that, all things equal, firms that receive free permits would also have higher output such that one can no longer attribute the RD parameter β entirely to the allocation of free permits. Online Appendix A.3 formally shows that in the short run without firm exits and the long run with possible exits, output-based allocation generates a downward bias in my estimates of β and my recovered permit price. Obtaining the magnitude of this bias requires knowing the concavity of the production function and expectations over future output and input prices.

D. *Comparison with CGE Estimates*

To date, multisector computable general equilibrium (CGE) models are the prevailing technique for estimating the marginal abatement cost of proposed cap-and-trade policies (see online Appendix G for a summary) and thus serve as a potential benchmark for my estimates. These models examine the cost of a specific policy through a structural representation of the economy with various primitive parameters that capture, inter alia, expected future prices, demand elasticities, and technological change. While one may attribute differences between my estimates and that of CGE models to varying information sets, it is important first to note other potential differences between these two methods.

First, the policy of interest may differ. CGE models estimate the marginal abatement cost of a known policy. The estimates from this paper correspond to the marginal cost of a policy which markets expected to be implemented. While the presence of a discontinuity at 5 percent energy intensity gives confidence that markets expected this unique feature of the Waxman-Markey bill, I cannot confirm that markets expected the realized policy to adhere strictly to the Waxman-Markey bill over its 35-year lifetime. If market expected key features to be altered at any point during implementation, my estimate would correspond to a policy that is different from Waxman-Markey.

[45] This is also referred to as the Coase Independence Theorem.

[46] Technically, Waxman-Markey requires that the average output over the previous two periods be used to determine current period free permits. I do not consider this for ease of exposition as implications are unaltered.

Second, some CGE models analyze the total costs of a stand-alone cap-and-trade policy. My estimate recovers the permit price from Waxman-Markey's cap-and-trade system in conjunction with the suite of auxiliary climate-related policies expected under Waxman-Markey. If such policies induce carbon abatement options that the cap-and-trade system alone would not induce, then my estimate corresponds to the marginal abatement cost of meeting the cap rather than the economy-wide marginal abatement cost generated by CGE models examining just the cap-and-trade component of Waxman-Markey.

Third, even if the policy of interest were the same, the unit of analysis may differ. My analysis is conducted at the firm level and thus cannot exclude non-US effects on firms with international operations nor does it capture dynamics of firm entry within a sector.[47]

Online Appendix Table A.5 summarizes the 2015 Waxman-Markey marginal abatement cost estimated by prominent CGE models (see online Appendix G for a summary of these models). For each estimate, I indicate the name and sector of the funding institution, the name and scenario of the CGE model employed, whether the model allows perfect foresight, and whether the analysis includes non-cap-and-trade components of Waxman-Markey. The unconditional mean CGE estimate of the 2015 marginal abatement cost is \$21.42 per ton CO_2e with a minimum and maximum of \$7.99 and \$37.73, respectively. Differences across a few characteristics are worth noting. First, for models that examine multiple scenarios such as the NEMS and EPPA models, assuming full usage of offsets yields lower estimates. Second, there does not seem to be a large difference in estimates between myopic CGE models and those that have perfect foresight. Third, CGE models that include other components of Waxman-Markey tend to produce higher estimates. Finally, estimates funded by private institutions are higher on average than those funded by academic or governmental institutions.

Figure 6 compares the CGE estimated marginal abatement cost shown in online Appendix Table A.5 with the lower and upper bound estimates generated from my approach. The range of my estimates falls below the unconditional mean of CGE estimates and substantially overlaps with the lower estimates generated by CGE models funded by academic and governmental institutions.[48] In order for my method to imply the higher marginal abatement costs estimated by privately funded CGE models, I need to assume annual rates of emission decline for all treated firms at or above the best performing manufacturing subsector over recent years. The overall comparison is similar when using the higher range of estimates in online Appendix Table A.6 under the assumption of fully binding borrowing restrictions.

V. Discussion

This paper develops an empirical method for forecasting the market-expected marginal abatement cost of a proposed climate policy, a central parameter for cost-benefit

[47] Ryan (2012) shows that the latter is particularly relevant for estimating the cost of the 1990 Clean Air Act Amendments on the US cement industry.

[48] My method relies heavily on the efficient market hypothesis. One implication is that market participants may be looking to CGE analyses to inform their trading behavior. Such informational dependence may explain for the close alignment between my estimates and that of certain CGE models.

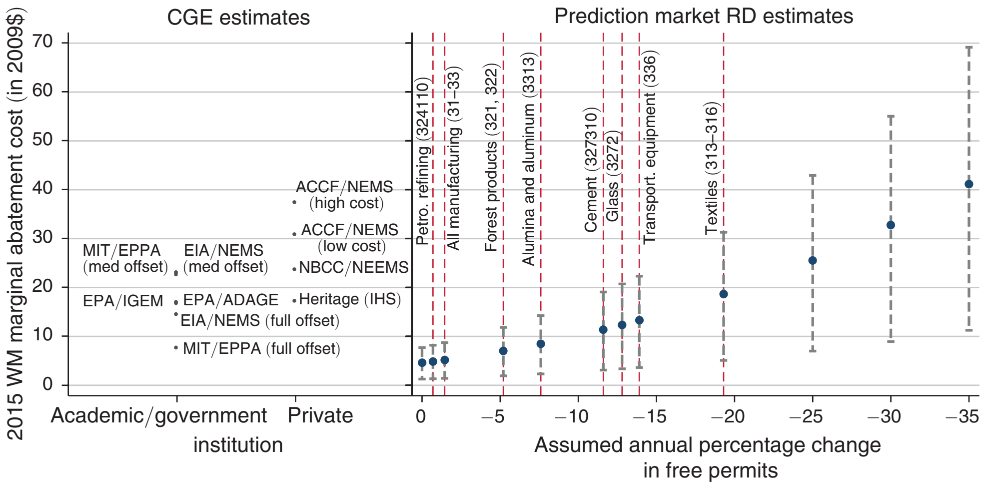

Figure 6. Comparing RD and CGE Estimates of Waxman-Markey Marginal Abatement Cost

Notes: Left panel shows estimates for the Waxman-Markey marginal abatement cost in 2015 (in 2009$) for various CGE models by sector of funding institution. Each data point is labeled by funding institution/CGE model (model scenario). Right panel shows 2015 marginal abatement cost using the prediction market RD method corresponding to the local linear model in Table 1, panel C, row 1, column 5. Uncertainty generated from Monte Carlo draws of estimated parameters. Each estimate assumes a different annual rate of change in firm-level emissions. Dashed vertical red lines show emission rates for various US manufacturing subsectors during the period 2006–2010.

analysis. This method is enabled by three features of the Waxman-Markey climate bill, the most promising, though ultimately unsuccessful, US climate policy considered to date. First, the centerpiece of the Waxman-Markey bill was a cap-and-trade system under which the equilibrium permit price, if observed, would be equal to the marginal abatement cost of the policy. Second, Waxman-Markey specified an unprecedented rule for allocating certain firms free permits which can be exploited using a regression discontinuity design. Finally, the availability of a prediction market tied to the Waxman-Markey policy allows one to recover the market-expected effect of the policy on stock returns, even if the policy was never realized. A final bounding analysis produces a range for the market-expected marginal abatement cost under Waxman-Markey in 2015 between \$5 and \$19 per ton of CO_2e.

To date, nearly all forecasts of the marginal abatement cost of climate policy are generated by computable general equilibrium models. My approach yields a range of estimates under Waxman-Markey that is both more narrow in spread and below the mean of CGE analyses of the Waxman-Markey bill. Interestingly, my estimates are within the range of the lower estimates generated by CGE analyses funded by academic and governmental institutions. This suggests that it is unlikely markets expected the higher range of cost estimates generated by privately funded CGE models.

While the method developed in this paper has the advantage of being empirically driven, it should not be viewed as a substitute for CGE models. My method recovers the marginal abatement cost of the policy that markets expected would be implemented, which may not correspond exactly to the Waxman-Markey bill.

Furthermore, while this paper may help narrow uncertainty regarding the level of the marginal abatement cost curve, it does not inform upon its slope which is necessary for optimal climate policy design. CGE models, on the other hand, structurally evaluate cap-and-trade policies for a known policy and can evaluate counterfactual policies along different abatement levels. It is possible, however, that these two methods serve complementary roles during future climate policy debates. For example, one can now conduct model selection by ruling out CGE models or a subset of the parameter space for a given model that produce estimates which greatly diverge from a market-based estimate. This opportunity for model validation may lend greater credibility to CGE-generated counterfactual analyses.

REFERENCES

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change." *American Economic Review* 102 (1): 131–66.

Aldy, Joseph E., Alan J. Krupnick, Richard G. Newell, Ian W. H. Parry, and William A. Pizer. 2010. "Designing Climate Mitigation Policy." *Journal of Economic Literature* 48 (4): 903–34.

Aldy, Joseph E., and William A. Pizer. 2015. "The Competitiveness Impacts of Climate Change Mitigation Policies." *Journal of the Association of Environmental and Resource Economists* 2 (4): 565–95.

Anderson, Soren T., and James M. Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." *American Economic Review* 101 (4): 1375–1409.

Bovenberg, A. Lans, and Lawrence H. Goulder. 1996. "Optimal Environmental Taxation in the Presence of Other Taxes: General-Equilibrium Analyses." *American Economic Review* 86 (4): 985–1000.

Buonanno, Paolo, Carlo Carraro, and Marzio Galeotti. 2003. "Endogenous Induced Technical Change and the Costs of Kyoto." *Resource and Energy Economics* 25 (1): 11–34.

Bushnell, James B., Howard Chong, and Erin T. Mansur. 2013. "Profiting from Regulation: Evidence from the European Carbon Market." *American Economic Journal: Economic Policy* 5 (4): 78–106.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer. 2000. "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy* 108 (6): 1292–1326.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods." *Annual Review of Economics* 1 (1): 451–88.

Cullen, Joseph A. 2013. "Dynamic Response to Environmental Regulation in Electricity Markets." http://www.josephcullen.com/resources/dynamicresponse.pdf.

Cullen, Joseph A., and Erin T. Mansur. 2014. "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution." National Bureau of Economic Research Working Paper 20795.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics* 4 (3): 66–95.

Deschênes, Olivier. 2011. "Climate Policy and Labor Markets." In *Design and Implementation of US Climate Policy*, edited by Don Fullerton and Catherine Wolfram, 37–49. Chicago: University of Chicago Press.

Deschênes, Olivier, and Michael Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review* 97 (1): 354–85.

Deschênes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3 (4): 152–85.

Ellerman, A. Denny, Paul L. Joskow, Richard Schmalensee, Juan-Pablo Montero, and Elizabeth M. Bailey. 2000. *Markets for Clean Air: The US Acid Rain Program.* Cambridge, UK: Cambridge University Press.

Fama, Eugene F., and Kenneth R. French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33 (1): 3–56.

Feng, Shuaizhang, Michael Oppenheimer, and Wolfram Schlenker. 2012. "Climate Change, Crop Yields, and Internal Migration in the United States." National Bureau of Economic Research Working Paper 17734.

Fowlie, Meredith, and Jeffrey M. Perloff. 2013. "Distributing Pollution Rights in Cap-and-Trade Programs: Are Outcomes Independent of Allocation?" *Review of Economics and Statistics* 95 (5): 1640–52.

Goulder, Lawrence H., and Stephen H. Schneider. 1999. "Induced Technological Change and the Attractiveness of CO2 Abatement Policies." *Resource and Energy Economics* 21 (3–4): 211–53.

Graff Zivin, Joshua, and Matthew Neidell. 2014. "Temperature and the Allocation of Time: Implications for Climate Change." *Journal of Labor Economics* 32 (1): 1–26.

Greenstone, Michael, John A. List, and Chad Syverson. 2012. "The Effects of Environmental Regulation on the Competitiveness of U.S. Manufacturing." National Bureau of Economic Research Working Paper 18392.

Hadri, Kaddour. 2000. "Testing for Stationarity in Heterogeneous Panel Data." *Econometrics Journal* 3 (2): 148–61.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica* 69 (1): 201–09.

Hayek, F. A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35 (4): 519–30.

Hicks, John R. 1932. *The Theory of Wages.* New York: Macmillan.

Horowitz, Joel L., and Charles F. Manski. 2000. "Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data." *Journal of the American Statistical Association* 95 (449): 77–84.

Hotelling, Harold. 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy* 39 (2): 137–35.

Hsiang, Solomon M., Marshall Burke, and Edward Miguel. 2013. "Quantifying the Influence of Climate on Human Conflict." *Science* 341 (6151).

Hsiang, Solomon M., Kyle C. Meng, and Mark A. Cane. 2011. "Civil Conflicts Are Associated with the Global Climate." *Nature* 476 (7361): 438–41.

Jacoby, Henry D., John M. Reilly, James R. McFarland, and Sergey Paltsev. 2006. "Technology and Technical Change in the MIT EPPA Model." *Energy Economics* 28 (5–6): 610–31.

Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins. 2003. "Technological Change and the Environment." In *Handbook of Environmental Economics*, Vol. 1, edited by Karl-Göran Mäler and Jeffrey R. Vincent, 461–516. New York: Elsevier.

Lange, Ian, and Joshua Linn. 2008. "Bush v. Gore and the Effect of New Source Review on Power Plant Emissions." *Environmental and Resource Economics* 40 (4): 571–91.

Lemoine, Derek. Forthcoming. "Green Expectations: Current Effects of Anticipated Carbon Pricing." *Review of Economics and Statistics.*

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* 59 (1): 1–14.

McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142 (2): 698–714.

Meng, Kyle C. 2017. "Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy: Dataset." *American Economic Review.* https://doi.org/10.1257/aer.20150781.

Montgomery, W. David. 1972. "Markets in Licenses and Efficient Pollution Control Programs." *Journal of Economic Theory* 5 (3): 395–418.

Nordhaus, William D. 2002. "Modeling Induced Innovation in Climate-Change Policy." In *Technological Change and the Environment*, edited by Arnulf Grübler, Nebojsa Nakicenovic, and William D. Nordhaus, 182–209. Washington, DC: Resources for the Future Press.

Ranson, Matthew, and Robert N. Stavins. 2016. "Linkage of Greenhouse Gas Emissions Trading Systems: Learning from Experience." *Climate Policy* 16 (3): 284–300.

Rubin, Jonathan D. 1996. "A Model of Intertemporal Emission Trading, Banking, and Borrowing." *Journal of Environmental Economics and Management* 31 (3): 269–86.

Ryan, Stephen P. 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80 (3): 1019–61.

Schennach, Susanne M. 2000. "The Economics of Pollution Permit Banking in the Context of Title IV of the 1990 Clean Air Act Amendments." *Journal of Environmental Economics and Management* 40 (3): 189–210.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. 2005. "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach." *American Economic Review* 95 (1): 395–406.

Schlenker, Wolfram, and Michael J. Roberts. 2009. “Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change.” *Proceedings from the National Academy of Science* 106 (37): 15594–98.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. 2007. “Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections.” *Quarterly Journal of Economics* 122 (2): 807–29.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz. 2013. “Prediction Markets for Economic Forecasting.” In *Handbook of Economic Forecasting*. Vol. 2A, edited by Graham Elliott and Allan Timmermann, 657–87. New York: Elsevier.

Stavins, Robert N. 1998. “What Can We Learn from the Grand Policy Experiment? Lessons from SO2 Allowance Trading.” *Journal of Economic Perspectives* 12 (3): 69–88.

US Environmental Protection Agency (**EPA**)**.** 2009. “Interagency Report on International Competitiveness and Emission Leakage.” https://www.epa.gov/sites/production/files/2016-07/documents/interagencyreport_competitiveness-emissionleakage.pdf.

Wolfers, Justin, and Eric Zitzewitz. 2006. “Interpreting Prediction Market Prices as Probabilities.” National Bureau of Economic Research Working Paper 12200.

Wolfers, Justin, and Eric Zitzewitz. 2009. “Using Markets to Inform Policy: The Case of the Iraq War.” *Economica* 76 (302): 225–50.