Location, location, location? What drives variation in the marginal benefits of renewable energy and demand-side efficiency

Preliminary and incomplete

Duncan Callaway, Meredith Fowlie, and Gavin McCormick *

January 1, 2015

Abstract

Greenhouse gas mitigation efforts in the electricity sector emphasize accelerated deployment of energy efficiency measures and renewable energy resources. Short-run benefits associated with incremental investments in energy efficiency or renewables manifest indirectly as reductions in the economic operating costs and emissions of marginal electricity generating units. We evaluate different renewable energy (RE) and energy efficiency (EE) technologies across regional power systems. Using standard social cost of carbon assumptions, our estimates of emissions-related benefits comprise a significant share of estimated returns on investment in some regions. On a per-MWh basis, regional variation in emissions displaced and costs avoided is more significant than variation across technologies within individual regions. This implies that the choice of location, more than the technology choice, determines the value generated by these investments. We also find that regional variation in avoided carbon benefits generates significant regional variation in the implied abatement costs associated with each technology. These results underscore the importance of designing policy incentives that accurately capture regional differences in emissions-related returns on RE and EE investments.

*The authors would like to acknowledge the generous support of the UC Berkeley Hellman Faculty Fund.

1 Introduction

Investments in renewable energy and demand-side energy efficiency improvements are playing a crucial role in efforts to reduce greenhouse gas emissions from the power sector which accounts for an estimated 40 percent of domestic CO_2 emissions. Renewable energy investment in U.S. has increased nearly 250 percent since 2004, reaching 36.7 billion in 2013.[1] Annual investments in energy efficiency, estimated as the extra cost for efficient goods and services relative to the average goods and services, are also on the rise in the U.S., estimated to be approximately $90 billion per year in 2010 (Laitner, 2013).[2].

Over the short to medium-run, return on incremental investments in renewable energy (RE) and energy efficiency (EE) manifest indirectly in two ways. First, RE generation and EE savings reduce operating costs at incumbent electricity generators. These cost-related benefits can generally be captured privately in the form of revenues from electricity sold (in the case of RE) or reduced energy expenditures (in the case of EE). Benefits also manifest in the form of avoided emissions at marginal generating units on the system. Emissions-related benefits are, to a large extent, external to electricity market transactions. In particular, greenhouse gas emissions remain untaxed in much of the power sector.

Much of the recent investment in renewable energy cannot be rationalized on the basis of private returns alone. Numerous policies and programs are currently in place to incentivize investments that are socially – but not privately – cost effective. External, uncompensated benefits associated with reduced greenhouse gas emissions serve as an important justification for these policy interventions.

The primary objective of this paper is to estimate the CO_2 emissions-related benefits generated by incremental investments in renewable energy and energy efficiency, and to assess the economic significance of these benefits. We are particularly interested in analyzing variation in these emissions impacts along spatial and technological dimensions. If variation in external, uncompensated benefits is economically significant, policy incentives should be designed to reflect this variation.

Many of the policies and programs used to accelerate investments in RE and EE do not explicitly account for variation in external benefits across RE and EE resources.[3] Conceptu-

[1]Michael Liebreich, Bloomberg New Energy Finance Summit (London: Bloomberg New Energy Finance, 2013), available at http://about.bnef.com/summit/content/uploads/ sites/3/2013/12/2013-04-23-BNEF-Summit-2013-keynote- presentation-Michael-Liebreich-BNEF-Chief-Executive.pdf; Pew Charitable Trusts, Whos Winning the Clean Energy Race? (2014), available at http://www.pewenvironment.org/uploadedFiles/PEG/Publications/Report/clen-whos- winning-the-clean-energy-race-2013.pdf.

[2]Laitner, Skip (2013). "Calculating the Nation's Annual Energy Efficiency Investments", ACEEE

[3]There are several important policies that currently serve to accelerate investment in RE and EE. Twenty-

ally, this variation can be partitioned into regional variation (arising from differences in the emissions profiles of regional power systems) and technological variation (arising from differences in temporal profiles of production or savings different RE and EE resources). Several recent studies have explored regional differences in the quantity of emissions displaced per unit of renewable energy generated (e.g. Callaway and Fowlie (2009), Siler-Evans, Azevedo and Morgan (2012), Kaffine, McBee and Lieskovsky (2013), Graff Zivin, Kotchen and Mansur (2014)). Novan (2014) finds that even within a single region (Texas), output from different renewable energy technologies can provide different external benefits due to differences in the correlation between the emissions profile of the power system and the production profile of the renewable energy resource. Novan argues that, if governments continue to subsidize RE and EE via production and investment-based policies, more emphasis should be placed on designing policies that more accurately reflect variation in external benefits across regions and technologies.

In principle, increasing the accuracy with which policy incentives reflect variation in external benefits should improve the allocative efficiency of policy outcomes. But increased accuracy comes at a cost of increased complexity in terms of both policy design and implementation. This begs the question: How economically significant is the variation in emissions displaced by EE and RE resources? And what dimensions (e.g. spatial or technological) are most important from a policy design perspective?

We use detailed hourly data from six major independent system operators (ISOs) in the United States over the period 2010-2012, together with detailed, site-specific profiles of renewable energy production potential and energy efficiency savings potential, to estimate the impacts of incremental RE and EE investments on power system emissions and operating costs. Emissions displacement and avoided operating cost are evaluated on the same empirical footing in order to facilitate a comparative assessment of the economic returns on alternative renewable energy and energy efficiency investments. We explore the extent to which variation in emissions displacement across regions and technologies drives variation in marginal returns on investment and marginal abatement costs.

There are several important findings. The first pertains to regional variation in emissions displacement benefits. We document statistically significant regional variation in marginal

nine states have adopted renewable portfolio standards which mandate minimum levels of renewable generation. Twenty states have efficiency standards which establish specific targets for demand-side energy savings. Incentives offered under these programs are based on electricity production. In addition, the federal government has established minimum efficiency standards for certain appliances and buildings, and provides sizable tax credits for renewable energy and energy efficiency improvements.

operating emissions rates (i.e. the emissions intensity of marginal producers). The quantity of emissions displaced per MWh of renewable energy generation (or per MWh of energy saved in the case of EE investments) also varies significantly across regions. For example, the quantity of emissions displaced, on average, by a MWh of renewable electricity in the midwest is more than double the average rate in California where marginal producers tend to be relatively clean gas-fired plants.

In contrast, emissions displacement (on a per-MWh basis) does not vary significantly across technologies in most of the regions we analyze. The reason is that marginal operating emissions rates are relatively homogeneous within regions across hours, days, or seasons. So resources with very different profiles displace very similar quantities of emissions per MWh. This finding has implications for the design of policy incentives. Policy incentives that are regionally differentiated – but neglect to capture intra-regional variation in resource profiles – can capture the vast majority of the variation in emissions displacement in our data.

A third finding is that emissions-related benefits can comprise a large share of the short-run returns on RE and EE investments. Our measure of the short-run value generated per MWh of RE or EE is comprised of both the avoided operating costs (e.g. fuel) and the value of avoided CO_2 emissions. Using a social cost of \$38 /ton of CO_2, emissions related benefits account for anywhere between one quarter and one half of the total estimated value per MWh in regions where emissions are not subject to a binding cap.[4] In contrast, variation in emissions displacement benefits across technologies within a region has little economic significance.

Finally, we assess the extent to which variation in emissions displacement rates drives variation in abatement costs. We combine our estimates of avoided emissions and avoided operating costs with estimates of investment costs in order to compute the implied cost per ton of avoided emissions across resources and regions. Variation in investment costs across technology types drives much of the variation in abatement costs, although regional variation in emissions displacement benefits has a significant role to play.

Taken together, these findings underscore the importance of designing policy incentives to accurately capture regional variation in external, emissions-related benefits. Within a region, variation in emissions displacement across resources and technologies is less likely to be economically significant.

The paper proceeds as follows. Section 2 provides a conceptual framework for the analysis.

[4]The Regional Greenhouse Gas Initiative (RGGI) imposed a binding cap on GHG emissions from the power sector during our study period. In RGGI states, RE and EE investments should have no impact on aggregate emissions, so emissions displacement benefits are assumed to be zero.

Section 3 summarizes the data. Section 4 estimates marginal operating emissions rates across time and space. Section 5 estimates marginal emissions displacement rates across regions and technologies. Section 6 relates estimates of emissions displacement to a more comprehensive measure of economic value. Section 7 estimates region and resource-specific marginal abatement costs. Section 8 concludes.

2 Conceptual framework

An overarching goal of this paper is to estimate the returns on investment in renewable energy and energy efficiency over the short-run, and to summarize the variability in these values along spatial, temporal, and technological dimensions. This is a short-run analysis in that we condition on the existing infrastructure of regional electric power systems. This is a marginal analysis in that we focus on incremental investments in grid-connected renewables and energy efficiency.

Returns on these investments manifest indirectly as avoided emissions and reductions in variable operating costs at marginal generating units on the system. Our analysis will emphasize the former; policies designed to accelerate investment in renewable energy and energy efficiency are largely rationalized on the basis of these uncompensated external benefits. To put these estimates in context, however, we also take a more comprehensive look at relative costs and benefits. More precisely, we assess the extent to which variation in emissions displacement benefits affects returns on investment (inclusive of avoided operating costs) and marginal abatement costs across regions and technologies.

Our analysis proceeds in four steps. In this section we organize these steps within a simple conceptual framework; in Sections 4–7 we describe each step in more detail and discuss estimation results.

2.1 Marginal operating emissions rate

On the operating margin, environmental benefits associated with additional renewable generation or efficiency improvements are determined by the emission intensities of the marginal units displaced. Modeling the relationship between marginal changes in system operating conditions and emissions is a critical first step in our analysis.

We specify an emissions equation, $EM_r(y_{rt}, x_{rt})$ which defines system-wide emissions in region r as a function of factors we can observe. The y_{rt} denotes the total production from generators that respond to marginal changes in production from RE, or savings from EE, in region r at time t. Other observable factors that can affect system operating conditions, such as weather, are captured by x_{rt}. Differentiating with respect to net load, we obtain an expression for the marginal operating emissions rate (MOER):

$$\phi_{rt} \equiv \frac{\partial EM_r(y_{rt}, x_{rt})}{\partial y_{rt}} \tag{1}$$

This partial derivative captures the system-wide emissions associated with the last megawatt produced by dispatchable units. Figure 1 serves to illustrate how this MOER can vary across hours and seasons, with this example specific to New York State. In this case there is significant seasonal and diurnal variation in MOERs.[5] Electricity consumption levels are higher in the day than at night, and higher on average in New York in the summer. The diurnal variation in the winter MOER curve is likely the result of coal more often on the margin at night, while combined cycle gas turbines are more often on the margin during the day.

2.2 Marginal emissions displacement rate

The next step in our analysis estimates the quantity of emissions displaced by a given renewable or efficiency technology j. Our approach turns on two features that distinguish grid connected wind, solar, and demand-side efficiency technologies from combustion-based generation resources. First, wind, solar and energy efficiency savings are variable and "non-dispatchable". Second, because the variable costs of wind, solar and efficiency savings are negligible when compared to combustion generators, they will almost always cause a reduction in output from emitting units. Taking these two factors together, if one neglects changes in transmission and distribution line losses, the quantity of emissions displaced by a resource is given by the product of the hourly MOER and the technology's electricity production or savings in that hour. Therefore it is essential to capture the hourly correlation between MOER and resource profile; we will explain this mathematically below.

Figure 2 plots average hourly resource availability for a representative wind site, a representative solar PV installation, a generic residential lighting upgrade, and a generic commercial lighting upgrade (all located in New York state). Realized resource availability varies around these average values. Referring to Figures 1 and 2 together, one can see that different resources have potentially significantly different capability to displace emissions from

[5] To put these rates in perspective, using the U.S. Energy Information Administration estimates for prime mover heat rates in 2012 and CO_2 emissions by fuel type, coal plants emit roughly 2075 pounds of CO_2 per MWh (assuming bituminous coal), combined cycle gas turbines (CCGT) emit 892 pounds per MWh and simple cycle gas turbines (SCGT) emit 1346 pounds per MWh.

Figure 1: Seasonal marginal operating emissions rate profile (NYISO)

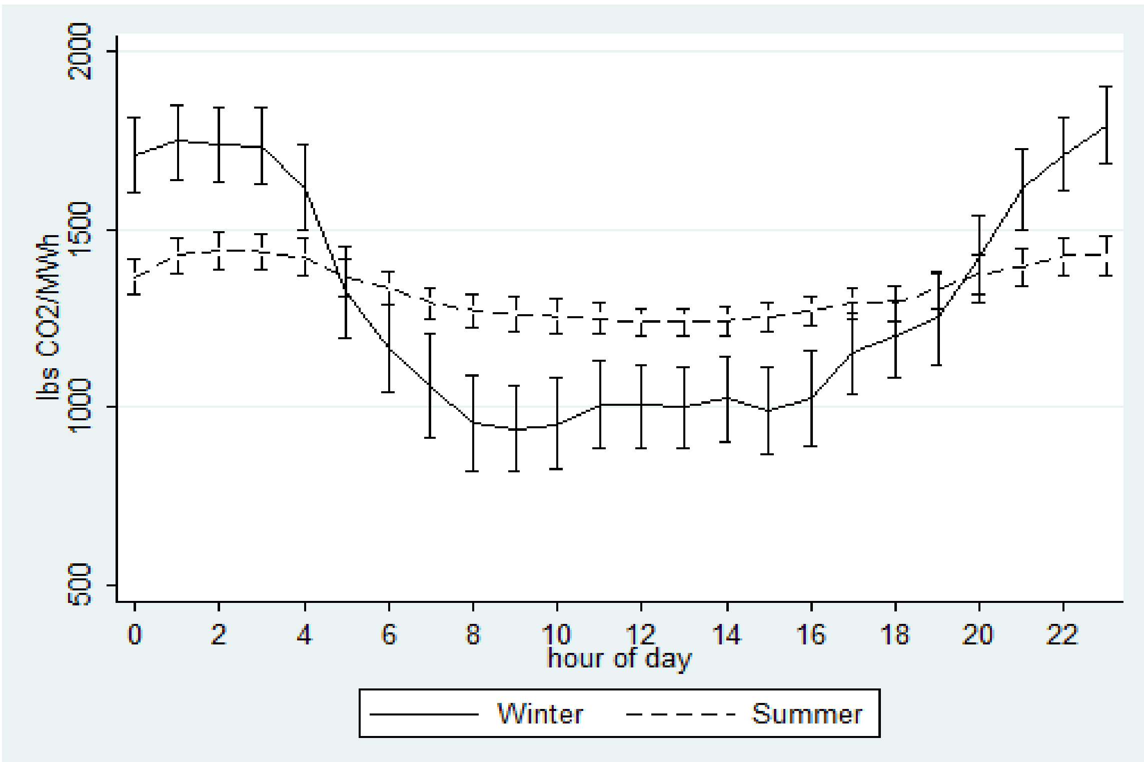

Notes: This figure illustrates hour-specific estimates of the marginal operating emissions rate in New York by season. Bars denote 95 percent confidence intervals. Our approach to constructing point estimates and associated confidence intervals is explained in detail in Section 4.

combustion generators.

Figure 2: Resource-specific production profiles

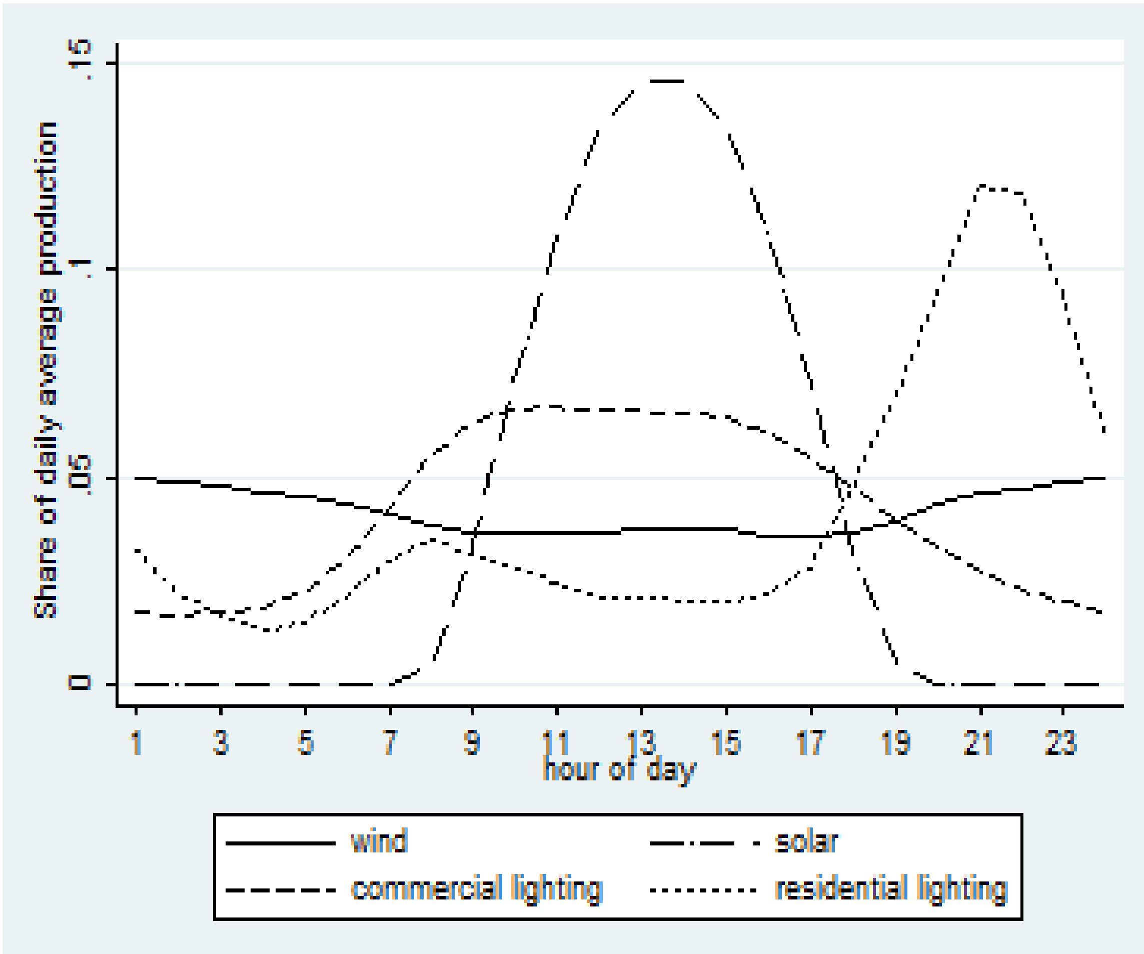

Notes: This figure plots the share of energy generated (or saved) on an average winter day in New York by hour of day. See the data appendix for a discussion of data sources.

To facilitate direct comparisons of emissions displacement across resource profiles, we define δ_{rj} as the marginal emissions displacement rate (MEDR) specific to each technology

j and region r over time horizon T:

$$\delta_{rj} \equiv E\left[\sum_{t=1}^{T}\left(\omega_{rjt}\phi_{rt}\right)\right] \tag{2}$$

$$= T \cdot E\left[\omega_{rjt}\right] \cdot E\left[\phi_{rt}\right] + T \cdot cov(\omega_{rj}, \phi_r) \tag{3}$$

$$= \overline{\phi_r} + T \cdot cov(\phi_r, \omega_{rj}) \tag{4}$$

where the weights ω_{rjt} represent the energy produced (or saved) by a resource in time interval t normalized by the total production from the resource over the time horizon T.[6] Variation in the MOER and resource profile parameters ω_{rjt} consists of both systematic and random components; $E[\cdot]$ denotes an expected, or average, value.

A key implication of Equation (4) is that the average quantity of emissions displaced per MWh generated or saved by a resource is determined not only by the average MOER in the region, $\overline{\phi_r}$, but also the correlation between the resource production profile and the marginal operating emissions rate. This highlights the importance of capturing both regional and technological variation in our analysis. As MOER and resource profiles vary across regions, the ability of a particular technology to deliver emissions displacement benefits can also change.

2.3 Marginal economic value

Having summarized variation in emissions displacement rates across regions and technologies, we turn to an assessment of the economic significance of this variation. To put our estimates of emissions displacement into perspective, we introduce a measure of marginal value that accounts for both avoided environmental damages as well as its impact on power system direct operating costs.

$$E\left[MB_{rj}\right] = E\left[\sum_{t=1}^{T}\left(\tau\omega_{rjt}\phi_{rt}\right)\right] + E\left[\sum_{t=1}^{T}\left(\omega_{rjt}\lambda_{rt}\right)\right] \tag{5}$$

$$= \underbrace{\tau(\overline{\phi_r} + Tcov(\phi_r, \omega_{rj}))}_{\text{Emissions displacement value}} + \underbrace{\overline{\lambda_r} + Tcov(\lambda_r, \omega_{rj})}_{\text{Avoided operating costs}} \tag{6}$$

[6]For ease of exposition, we ignore variation in resource profiles for each technology within a region. Below we show empirical support for the assumption that within-region variation is not a significant driver of variation in resource value

The τ parameter captures the monetary value of the health and environmental damages avoided per unit of displaced emissions.[7] The λ parameter represents the cost of the last MWh produced by dispatchable units over a particular hour. If ω_{rj} is positively (negatively) correlated with the marginal cost of supplying load, this will positively (negatively) influence the marginal economic value of the renewable or efficiency resource.

2.4 Marginal abatement cost

To assess the economic implications of external emissions benefits, we should ideally account for benefits *and* costs of investing in renewable energy generation and energy efficiency. We combine our estimates of marginal benefits introduced in the previous section with estimates of levelized investment costs to compute the marginal abatement cost – i.e. the net cost per ton of CO_2 emissions avoided – as follows:

$$E\left[MAC_{rj}\right] = \frac{LCOE_{rj} - (\overline{\lambda_r} + Tcov(\lambda_r, \omega_{rj}))}{(\overline{\phi_r} + Tcov(\phi_r, \omega_{rj}))} \tag{7}$$

The numerator is the net cost per MWh: the levelized cost of electricity (LCOE) net of the value of avoided fuel costs. The LCOE is a common benchmarking tool used to assess the relative cost-effectiveness of different energy technologies. Conceptually, it measures the constant (in real terms) price per unit of electricity generated that would equate the net present value of revenue from the plant's output with the net present value of the cost of production. Dividing the net cost per MWh by the quantity of emissions displaced per MWh (in the denominator) yields a cost per ton of emissions avoided. If this value exceeds the social cost of carbon emissions, the investment cannot be rationalized on the basis of the emissions externality alone.[8]

The marginal abatement cost measure summarized by Equation (7) provides a simple metric that facilitates a comparison the relative merits of wind, solar and efficiency investments on the basis of the investment required to avoid carbon emissions. Framing comparisons in this way, we can evaluate the extent to which an accurate internalization of emissions benefits would alter the rank order – and level – of investment in RE and EE resources across regions and technologies.

[7] We assume this damage value is constant over the relevant range of emissions levels.

[8] In a regional electricity market that has imposed an emissions cap, if this cost per ton exceeds the prevailing permit price, the investment is not cost effective unless there are other market failures (e.g. learning by doing) in play.

3 Data

The data sets used in this paper include combustion generator emissions and production; hourly production from wind and solar generators; hourly savings from efficiency measures; marginal electricity prices; and estimates of levelized costs for each renewable and efficiency resource. Except where otherwise noted, the period of analysis is each hour of the period 2010-2012. Before we describe these data sets in detail, we discuss three research design choices that shape our data set construction.

3.1 Regional unit of analysis

We define the regions r to be the six major independent system operators (ISOs) in the United States: ISO New England (ISONE), the New York ISO (NYISO), the PJM Interconnection, the Midcontinent Independent System Operator (MISO), the Electric Reliability Council of Texas (ERCOT) and the California ISO (CAISO).[9] We chose these regions for two central reasons. First, they coordinate large-scale pooled electricity markets to economically balance local load with supply on daily, hourly and sub-hourly time scales.[10] Second, as balancing authority areas (BAAs), these ISOs coordinate local generation via ancillary services – most notably frequency regulation – to balance net load forecast errors on a second-to-second basis, after all electricity markets have cleared.

The choice of ISOs as the region of analysis contrasts with other papers that use North American Electric Reliability Corporation (NERC) regions (Siler-Evans et al., 2012; Graff Zivin et al., 2014). NERC regions are used for monitoring expansion plans and assessing historical reliability performance, but do not define the footprint of any single pooled market or BAA. In fact, in some cases NERC regions are much larger than ISOs (CAISO is a very small part of the WECC NERC region), and in other cases ISOs straddle multiple NERC regions (for example PJM straddles the MRO, RFC and SERC NERC regions; MISO straddles the MRO and RFC NERC regions).

[9] Because the Sacramento Municipal Utility District and Los Angeles Department of Water and Power are surrounded completely by CAISO, we include generators in those footprints in our analysis; therefore we refer to the total region of analysis as California.

[10] On March 1, 2014 the Southwest Power Pool began coordinating daily, hourly and sub-hourly markets via its Integrated Marketplace. At the time of writing this paper there was insufficient data to include it in our analysis.

3.2 Marginal generating units

Several recent studies have estimated the emissions impacts of grid-connected renewable electricity generation in a variety of contexts. Empirical strategies vary in terms of the degree of complexity, data requirements, and identification strategies. In this study, we use observable variation in production at grid-connected thermal power plants to proxy for the effects of adding a new grid-connected renewable energy resource or efficiency improvements.

To estimate marginal operating emissions rates, we regress hourly CO_2 emissions in a regional power system on generation at thermal power plants in the region. An alternative approach would regress emissions on the sum of generation from combustion-based sources as well as production from non-emitting inframarginal sources such as renewables, large hydro and nuclear. However, though inframarginal, the output of these non-emitting resources might change for reasons unrelated to operator dispatch decisions. For example, wind and solar change production in response to available natural resources,[11] and hydro plant operators typically self-schedule their output, in part to manage environmental constraints. To the extent the output of these non-marginal generators are correlated (but not causally) with changes in net load, marginal emissions analyses that regress emissions on net load will be biased.

3.3 Power flows between regions

Our base specification does not include transfers across ISO boundaries for two reasons. First, there is no straightforward way to allocate flow into a region to specific types of generation (specifically emitting versus non-emitting generation) and for this reason, as discussed in the previous paragraph, using a measure of total imports in a regression suffers the same problems that load-based regression estimates suffer (namely changes in non-marginal, non-emitting generation can bias MOER estimates). Second, market barriers (for example the challenge of congestion management across ISOs, a lack of standard definitions for energy products, and fees for importing and exporting across ISO interchanges) prevent generators outside an ISO's footprint from efficiently participating in that ISO's market. Though ongoing efforts between ISOs to address these so-called "seams" issues may eventually resolve these inefficiencies,[12] generators within an ISO's footprint currently face fewer hurdles to being

[11] In extreme conditions wind and solar can be curtailed, however, to the extent this curtailment occurs in existing systems it is generally driven by transmission congestion rather than a system-wide excess of supply. In that sense those curtailed resources are locally marginal, but not at the aggregate level.

[12] `http://www.isorto.org/Documents/Report/2010IRCMetricsReport_2005-2009.pdf`, last accessed December 28, 2014.

dispatched to meet a change in net load.[13]

We note that this modeling choice specifies a mechanism for how generation is dispatched, but given that choice we are not omitting data. That is, we will focus on estimating the marginal emissions rates of combustion-fired generation within an ISO footprint – with an accurate accounting of local generation and emissions – and our key assumption is that the output of those local generators will change in response to renewables or efficiency. This is in contrast to load-based estimates that measure the marginal change in local emissions with respect to a marginal change in local load (e.g. Siler-Evans et al. (2012)); that approach suffers from accounting errors because local emissions may in fact be changing in response to remote changes in load, or remote emissions may be changing in response to local changes in load. A credible modeling alternative, similar to that employed in Graff Zivin et al. (2014), is to regress total emissions within an entire interconnection (i.e. an aggregation of ISOs and utilities that interchange power) on a vector whose elements comprise load in each sub-region of the interconnection. This implicitly captures exchange between areas, however, we still prefer the generation-based approach used in this paper for two reasons. First, it avoids the load-based estimate problem of accounting for non-marginal changes in non-emitting generation mentioned above. Second, load in neighboring regions tends to be highly collinear, which complicates the interpretation of the estimated coefficients.

3.4 Data Sources

Combustion generator production and emissions data. We obtained generation (in MWh per hour) and CO_2 emissions (in pounds per hour) for combustion-fired plants that report to EPAs Continuous Emissions Monitoring System (CEMS) dataset. We use plant latitude and longitude to locate the plants within ISOs using a spatial database of the footprints of each ISO[14]. We exclude combined heat and power units and so-called "self generating" units which are unlikely to ever be called upon to follow load).

Generation cost. To capture marginal fuel and operating costs (λ_r in the previous section) we collected real-time hourly average locational marginal prices for each region[15]. In ISONE, we set λ_r equal to the Internal Hub real-time LMP in ISONE. For all other regions,

[13] In October 2014, CAISO and the utility PacifiCorp began operating a shared energy imbalance market to facilitate inter-hour economic adjustments in flow between the regions. This market is designed to reduce the inefficiency of inter-regional trade and may need to be considered in future analyses of this type.

[14] `http://www.ventyx.com/en/solutions/business-operations/business-products/velocity-suite`, last accessed December 28, 2014.

[15] `http://www.gdfsuezenergyresources.com/index.php?id=712`, last accessed December 28 2014.

we set λ_r equal to the unweighted spatial average of each region's hourly LMPs. Note that ERCOT's nodal market began on December 1, 2010; we dropped all preceding dates from our analysis of ERCOT. LMPs in ISOs containing states that participate in the Regional Greenhouse Gas Initiative (RGGI) will be influenced by the marginal carbon abatement cost in RGGI; in Section 7 we will adjust those LMPs by the carbon market price, which we take to be the permit price in 2012, \$1.93/ton.[16].

Wind production data. We obtained simulated wind production data from the National Renewable Energy Laboratory's (NREL) Eastern Wind dataset[17] and Western Wind dataset[18]. NREL and its partners produced these datasets with a combination of meso-scale wind speed models and the production characteristics of hypothetical wind farms. The resulting simulated datasets cover more than 30,000 sites across the United States, have a temporal resolution of 6 minutes (which we used to construct hourly averages), and span 3 years from 2004-2006.[19] We used the latitude and longitude of each simulated wind site to locate the production within each ISO, and normalized these data to hourly energy production per megawatt of installed capacity for each site.

Levelized cost of wind energy (LCOE) are constructed from data provided by Lawrence Berkeley National Laboratory (LBNL) Wiser et al. (2014). This report analyzes power purchase agreements (PPAs) from a large sample of wind installations to produce annual average levelized prices per megawatt-hour of wind. We used the 2012 data for each of four regions: Great Lakes, Interior, West and Northeast. Prices and ISOs that we assigned to each region are in the Appendix. Because all projects should have received the federal production tax credit (PTC) we set the total LCOE equal to the sum of PPA prices and the 2012 PTC (\$22/MWh). Assuming the wind industry is competitive, these prices are representative of total social costs per MWh. Table 1 summarizes the wind LCOE data along with LCOEs for other technologies, described below.

Solar production data are from NRELs PV WATTS simulation tool [20]. This software applies PV performance modeling to typical meteorological year weather data to estimate the hourly average production of a solar array installed at thousands of different sites. We

[16] `https://www.rggi.org`, last accessed December 28, 2014.

[17] `http://www.nrel.gov/electricity/transmission/eastern_wind_dataset.html`, last accessed December 28, 2014.

[18] `http://www.nrel.gov/electricity/transmission/western_wind.html`, last accessed December 28, 2014.

[19] Data of this extent are not available in the years that we collected combustion generator data (2010-2012). Though correlation between wind speed and electricity load is very weak, using wind data from different years than those used to construct MOERs could introduce small errors in our analysis.

[20] `http://pvwatts.nrel.gov`, last accessed December 28, 2014.

Table 1: LCOEs for all technologies and regions.

technology	CAISO	ERCOT	ISONE	MISO	NYISO	PJM
Utility scale solar	$123.49	$133.47	$140.80	$150.41	$150.27	$151.73
Utility scale wind	$87.36	$52.58	$75.11	$64.39	$75.11	$75.11
Residential lighting	$25.25	$25.25	$25.25	$25.25	$25.25	$25.25
Commercial lighting	$4.12	$4.12	$4.12	$4.12	$4.12	$4.12

replicate this average year for each year in our analysis. We used the default assumptions for a fixed PV array facing south, with a tilt angle set equal to the sites latitude. Because solar production is highly spatially correlated on hourly time scales we use only two sites per region and chose sites to be relatively far apart and such that one was in a location with very good resource potential for that region; these sites are listed in the Appendix.

Solar LCOE data are constructed from data from LBNL Barbose et al. (2014). We used the 2012 installed cost for >5MW utility scale systems ($2.97 per watt).[21] The LBNL data are reported prior to receipt of any direct financial incentives or tax credits, therefore assuming the PV industry is competitive, these prices are representative of total social costs per MWh. We use the same cost model as in (Baker et al., 2013), namely: we assume that the inverter is replaced every 10 years at a cost of $0.20/W but declining at 2% annually in real terms; assume a project life of 30 years; assume a panel degradation rate of 0.5% per year; and assume a real discount rate of 3%. We computed LCOE for each of the two sites per region and averaged the result within each region; the resulting LCOE are in Table 1.

Energy efficiency "production" data. For efficiency, we will focus only on commercial and residential lighting efficiency We obtained a year of simulated hourly consumption data for typical residential and commercial buildings (details on the data are can be found in (Wei et al., 2012)). These hourly profiles vary by hour of day, weekdays/weekends and season. Commercial lighting consumption is concentrated in business hours and residential lighting energy is concentrated in evening hours. We assumed that a unit of energy saved from lighting efficiency would be distributed in proportion to the hourly consumption data, and treat those saved units of energy in a given hour as equivalent to energy produced from a wind or solar generator. We assumed the hourly consumption profiles, conditioned on season, and weekend / weekday, would be the same in each year of our analysis.

Energy efficiency LCOE data are from the US Department of Energy Appliance

[21]We note that residential scale solar installed prices vary systematically across the country, and large-scale systems likely do as well, however the data available comprise only a single nation-wide number. However, local resource potential drives levelized cost, and this was factored into our analysis.

and Equipment Standard Programs 2011 Fluorescent Lamp Ballast rulemaking. For each appliance efficiency rulemaking under consideration, the DOE releases a technical support document including either a Life Cycle Cost Assessment or a National Impacts Analysis which provide estimates of the energy savings and costs associated with different efficiency levels (EL) under consideration. For both residential and commercial categories, we focused on general service fluorescent lamps (GSFL). DOE estimates that there are more than 2 billion of these lamps in service in the US residential and commercial sectors (DOE, 2009), with most (92%) in the commercial sector [22]. DOE documents GSFL lamp characteristics extensively for rulemaking purposes. The current DOE standard is 88 lumens per watt for the lamp-ballast system (DOE, 2009). For each sector, we chose the baseline as the technology with the lowest installed cost in that sector that also meets the standard. We defined the efficiency option as the technology with the second lowest installed cost that also meets the standard and is more efficient than the baseline. The resulting technology choices were different for the residential and commercial sectors. We calculated a levelized cost of energy saved by the efficiency option over a fifteen-year period (to reflect ballast lifetimes DOE (2009)) at a 3% discount rate. Further detail on efficiency calculations, including the chosen technologies and their costs, in the appendix.

4 Marginal operating emissions rates

Our primary empirical challenge in this section is to isolate the variation in generation that most closely mimics the system-wide response to an incremental investment in EE or RE. To capture this variation, we estimate the following equation:

$$E_{rkt} = \alpha_{rkhs} + \phi_{rkhs} G_{rkt} + e_{rkt}, \tag{8}$$

where E_{rkt} measures electricity production at dispatchable, fossil-fueled sources in region r and hour t. As noted in Section 3, we exclude production at resources that do not typically vary in output to follow load (such as biomass, landfill gas, wind, hydro, solar, and nuclear).

Much of the unit scheduling that determines how electricity generating units are dispatched occurs day-ahead. Therefore, in any given hour, system-wide emissions are a function of not only contemporaneous operating conditions, but also the forecast conditions

[22] Though this suggests the number in the residential sector is relatively small, at roughly 35 W per lamp, the residential sector alone has over 6.5 GW of lamps installed. DOE estimates these lamps are used 791 hours per year, suggesting roughly 5 TWh of end-use electricity consumption per year

throughout the day that determined how units were dispatched day-ahead. We want to isolate the variation in system-wide emissions that could conceivably be caused by marginal increases in RE or EE. In other words, we want to be careful to use variation across days with similar load profiles to estimate the coefficients in Equation (8). Comparing system-wide emissions across days with very different load profiles can confound the effects of small differences in net load (such as those associated with incremental EE and RE investments) with the effects of larger differences in underlying system dispatch. To address this issue, we cluster days within a region and season that share very similar generation profiles. More precisely, we use a k-means clustering algorithm to cluster daily observations (within a region and season) over the period 2010-2012 into groups of days with very similar load profiles and peak loads. This algorithm, which is explained in more detail in the appendix, gives rise to clusters of days within each season and region denoted by k. The α and ϕ coefficient values are allowed to vary across clusters to reflect differences in underlying operating conditions that are distinct from the kind of variation generated by incremental changes in RE and EE.

The α parameter captures the average emissions level observed in region r, season s, hour h, and load profile type k. Differencing out these average values helps to control for the effect of systematic differences in system operating conditions across regions, hours, or seasons that will persist independent of incremental RE or EE investments.

We are primarily interested in the ϕ coefficients which are estimated separately for each hour of the day to capture systematic, within-day variation in marginal operating conditions. To capture systematic, seasonal variation in MOERs, these hour-specific coefficients are estimated separately for summer and winter seasons (denoted s). We define our seasons to match the seasonal NOx emissions regulations which switch on in May and switch off in October and which affect the marginal operating costs of fossil-fuelled generating units.

Standard errors throughout the analysis are estimated using a block bootstrap. We generate 1,000 permutations of the main data set as follows. For each each region and season, we select (with replacement) a set of days which preserves the observed composition of week days and week-end days within that region-season over the study period (2010-2012). We keep 24 hour blocks within each day together because electricity grid operations are optimized day-ahead for the following day. Each simulated data set is used to repeatedly estimate Equation 8.

Figure 3: Seasonal marginal operating emissions rates

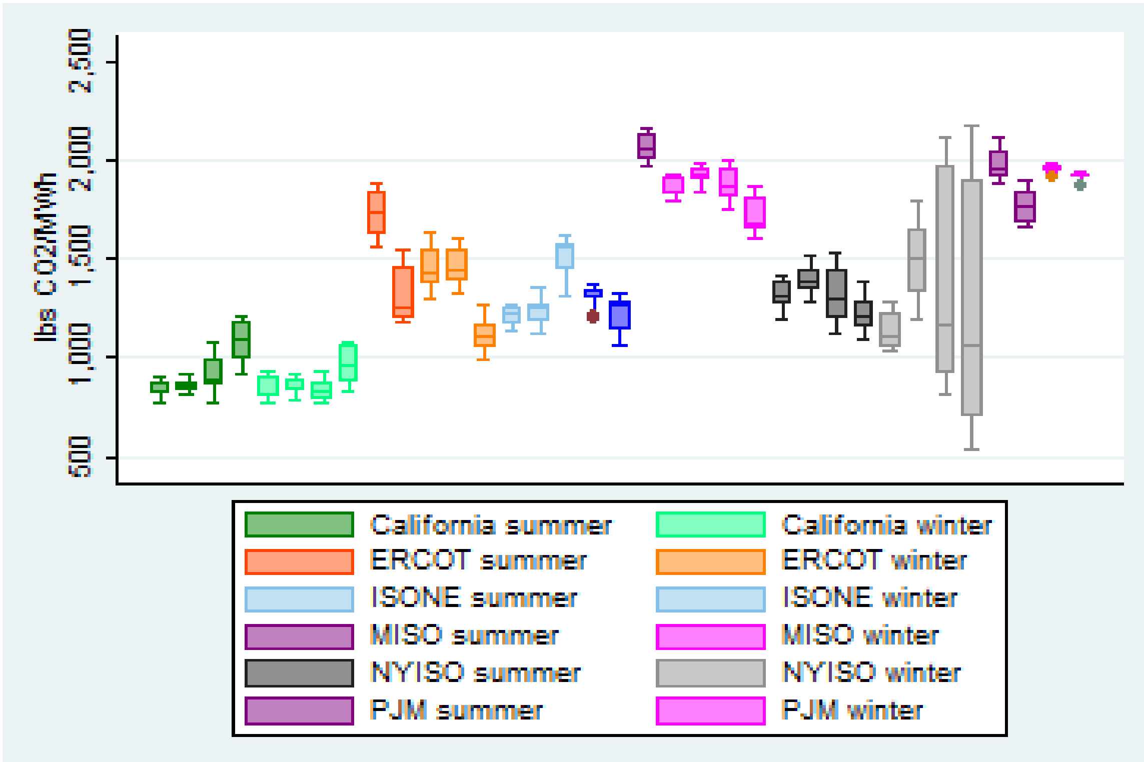

This figure illustrates the range of hour-specific estimates of the marginal operating emissions rate by season and cluster.

4.1 Estimation results

With six regions, twenty-four hours, two seasons, and an average of three clusters per region and season, the empirical strategy summarized above yields no fewer than point estimates of 840 ϕ_{hrsk}.

Figure 3 summarizes the range of variation in these MOER point estimates. Each box corresponds to a region, season, and cluster. Within a region and season, clusters are displayed in increasing order of generation. Fifty percent of the corresponding, hour-specific point estimates fall within the range of the box. The line within each box denotes the median value. The whiskers add (or subtract) 1.5 times the interquartile range to the third (or first) quartile.

The figure illustrates striking variation in marginal operating emissions rates across regions. California's fossil fuel mix is dominated by natural gas. Variation in emissions rates is therefore driven primarily by variation in fuel efficiency, with more fuel efficient plants preceding less efficient plants in the merit order. Consequently, we see MOER estimates increasing with generation levels in both the summer and winter. In contrast, coal accounts for a significant fraction of the fuel mix in PJM and MISO regions. The MOER estimates in these regions indicates coal is marginal in some fraction of hours. Intuitively, MOER estimates are decreasing with generation levels in these regions (because coal units will often precede cleaner gas units in the merit order).

Region-hour-season-specific MOER estimates are reported in the appendix along with bootstrapped confidence intervals. [23]

5 Marginal Emissions Displacement Rates

In this section, we estimate the region and technology specific marginal emissions displacement rates, δ_{rj} introduced in Section 2. These values are estimated empirically as:

$$\delta_{rjt} = \frac{\sum_{t=1}^{T} (\phi_{rkhs} \cdot q_{rjt})}{\sum_{t=1}^{T} (q_{rjt})}. \tag{9}$$

For each hour of the data period (indexed by t), we multiply the quantity of simulated renewable energy production (or energy demand reductions in the case of efficiency) q_{rjt} with the corresponding regional MOER estimate. The numerator in Equation 9 is the estimate of the pounds of CO_2 displaced by renewable energy generation – or avoided due to demand reductions – over time period T. Dividing by the sum of energy produced (or saved) yields an estimate of the average quantity of emissions displaced per MWh. For each region-technology pairing, confidence intervals are estimated using the block bootstrap described above.

5.1 Estimation results

Equation 4 from Section 2 shows how variation in these MEDR values will depend on both variation in regional average MOERs and any correlation between MOER profiles and resource profiles. Figure 4 starts to unpack this variation along regional and technological dimensions.[24]

[23]The appendix also summarizes an exercise in which summer marginal emissions rates are re-estimated using a data set that excludes those generating units that report in summer-only. We find that dropping these units does not significantly affect the summer estimates.

[24]Because bootstrapping the full suite of results is a computationally intensive process, we choose two representative PV sites and 20 wind sites per region to incorporate in the subsequent analysis. So, in the

Figure 4: Marginal emissions displacement rates

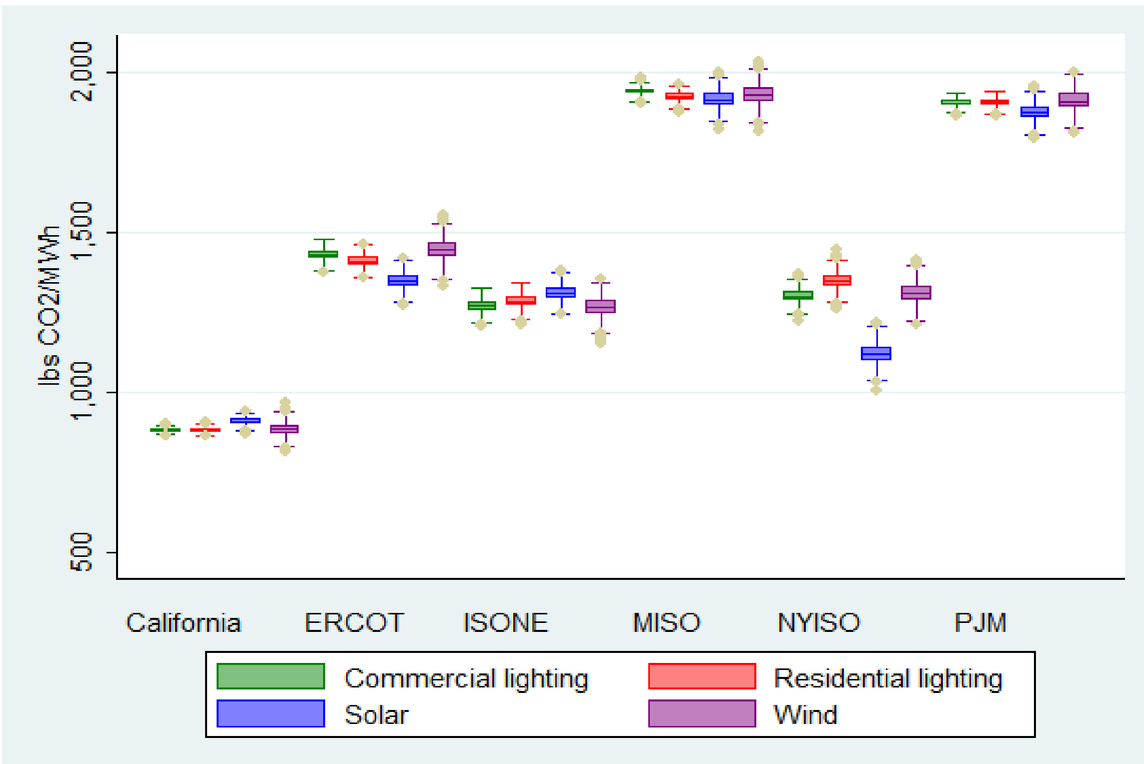

Notes: This figure illustrates the range of resource-specific marginal emissions displacement rates estimated by region. The top and bottom of each box represent the upper and lower quartile values, respectively. Whiskers denote 1.5 times the interquartile range beyond the 25th and 75th percentile values. Resource profiles for lighting efficiency improvements capture generic seasonal and hourly variation in energy savings. Solar and wind profiles vary within and across days according to simulated meteorological conditions and are site specific. Subsets of sites from each region are used to estimate regional MEDRs.

These box plots illustrate substantial variation in emissions displacement rates across regions. In California, marginal operating emissions rates are low and relatively constant over the course of a day. Emissions displacement rates in this region are accordingly low and do not vary significantly across technologies with different resource profiles. In New York, there is more inter-temporal variation in the MOER profile, and thus more variation in emissions displacement rates across technologies with different resource profiles. Solar PV resource availability is negatively correlated with the MOER profile (see Figures 1 and 2); solar is therefore associated with a relatively low δ. In contrast, demand reductions associated with residential lighting improvements are positively correlated with diurnal MOER profiles. This results in a relatively high average MEDR for this technology.

Figure 4 summarizes emissions displacement estimates for generic efficiency profiles and a representative set of PV and wind resource installations from each region. However, we observe thousands of PV and wind sites within each region. Whereas PV production tends to be highly correlated across sites within a region, spatial variation in elevation, topography, and vegetation can generate significant variation in wind patterns across relatively short distances. Significant intra-regional differences in wind resource profiles could imply significant variation in emissions displacement rates and within technology. To assess the extent of the intra-regional variation in emissions displacement across wind sites, we generate MEDR point estimates for all wind sites in the data. Results are summarized in Appendix A.4. We find very limited intra-region, intra-technology variation in emissions displacement values. Subsequent analysis will therefore focus on variation across regions and technologies.

5.2 Analysis of variance in emissions displacement rates

Suppose a policy maker is looking to design policy incentives that compensate external emissions displacement benefits summarized by Figure 4. The nature of the variation in MEDRs across resources should inform the design of these policy incentives. The Figure clearly shows that inter-regional variation in simulated marginal emissions rates dominates within region variation. From a policy design perspective, this suggests that much of the variation in emissions displacement benefits across resources could be captured using production-based incentives that vary significantly across- but not within-regions.

In what follows, we assess the extent to which alternate regional measures of emissions intensity capture variation in MEDRs across regions and resource types. For each bootstrap

case of solar PV for example, Figure 4 summarizes the results of 2,000 bootstrap repetitions in each region. We discuss within region, within technology variation in more detail below.

repetition (indexed by b), we compute two regional summary statistics for each year of data: the regional average emissions rate (averaged across all fossil units) and the regional average MOER. We then use a regression-based approach to analyze the correlation between annual MEDR estimates, the δ_{rjb}, and these summary measures.

As a point of reference, Column (1) of Table 2 reports the average δ value of 1442 lbs CO_2 per MWh. Column (2) regresses δ_{rb} on a regional measure of the average annual emissions intensity (the average emissions rate of fossil generators within a region). This average emissions rate captures a significant share of the variation in MOERs; the R-squared is 0.83. Column (3) restricts the coefficient on the regional average to equal one and includes a set of technology-specific indicator variables. These technology-specific coefficients measure the average difference – by technology type – between this regional average proxy and the δ_{rb}. The regional average emissions rate exceeds the average δ_{rb} for all technology types. This implies that, if resource owners were compensated based on a regional average emissions rate, avoided emissions would be over-compensated. This over-compensation is most significant for solar and commercial lighting, both of which peak during the middle of the day when the MOERs tend to be lowest.

Columns (4) and (5) of Table 2 regress the δ_{rj} on the regional annual average MOER. This marginal (versus average) rate is a superior proxy for the emissions displaced per unit of renewable energy generation or energy saved; the R-squared increases to 0.97. Emissions displaced by residential lighting improvements are slightly underestimated on average. Emissions displaced by commercial lighting and solar PV are significantly over-estimated. These technology-specific interactions report average deviations from the regional proxy. Appendix A.4 reports the results from estimating a more fully saturated model. Intuitively, deviations from the regional average are relatively small on average in regions like California and MISO where MOER profiles are quite flat. In contrast, deviations are more significant in New York. For example, the average MOER over-estimates emissions displaced per MWh of electricity generated by solar PV by approximately 190 lbs on average (a 17 percent increase above the estimated MEDR for this technology in this region).

6 Marginal Economic Value

In the preceding section, we document statistically significant variation in marginal emissions displacement rates across regions and technologies. In this section, we begin to assess the economic implications of this variation. More specifically, we assess how differences in

Table 2: Regression-based decomposition of variance in marginal emissions displacement

Dependent variable is MEDR	(1)	(2)	(3)	(4)	(5)
Constant	1441.6**	193.9**	.	14.289**	.
	(9.36)	(1.29)		(0.41)	
Average emissions rate (regional by year)		0.85**	1		
		(<0.01)			
Marginal operating emissions rate (regional by year)				0.98**	1
				(<0.01)	
Wind			-11.56**		0.46
			(0.54)		(0.49)
Solar PV			-53.89**		-41.97**
			(13.06)		(10.37)
Residential lighting			-9.34		2.67
			(8.32)		(4.98)
Commercial lighting			-32.39		-20.38
			(14.64)		(12.26)
R-squared		0.83	0.84	0.97	0.98
Observations	72	72	72	72	72

Note: Bootstrapped standard errors.
* Significant at the 5 percent level
** Significant at the 1 percent level

emissions displacement drive differences in social returns on investment across regions and technologies.

We use a monetary measure of marginal economic value (summarized by Equation 6 from Section 2) that includes both the value of avoided emissions (often external to market transactions) and the operating costs (e.g. fuel costs) associated with generation displaced marginal units. This marginal value measure is implemented empirically as:

$$MB_{rj} = \tau \cdot \delta_{rj} + \frac{\sum_{t=1}^{T} (\lambda_{rt} \cdot q_{rjt})}{\sum_{t=1}^{T} (q_{rjt})}. \tag{10}$$

To construct the first term on the right hand side, we multiply the MEDR estimates by the social cost of carbon denoted τ. We assume a value of \$38 per ton CO_2 (in 2011 dollars).[25] Notably, this value will not fully manifest in cases where power sector emissions are subject to a binding cap. We return to this point below.

To estimate the second term on the right hand side of Equation 10, we need regional and hourly measures of the variable operating costs at marginal dispatchable generating units (λ_{rjt}). We use region-specific, real-time locational marginal prices (LMPs) as a proxy. These prices reflect the marginal cost of supplying (at least cost) the next increment of electricity to a particular location given the supply and demand bids submitted by market participants and the physical constraints on the system.

Our approach to estimating the value of avoided operating costs parallels our approach to estimating avoided emissions. In each hour we multiply the megawatt-hours of simulated renewable energy production (or energy demand reductions in the case of efficiency improvements) with the corresponding regional LMP value. Aggregating these avoided costs across all hours and dividing by the sum of energy produced (or saved) yields a region and technology-specific estimate of the average marginal value per MWh. For each region-technology pair, confidence intervals are estimated using the block bootstrap described above.

6.1 Avoided operating costs

Figure 5 summarizes region and technology-specific estimates of avoided operating costs per MWh. Variation within a region and across technologies is driven by differences in

[25] This is approximately equal to the value associated with a 3 percent discount rate: U.S. Interagency Working Group on Social Cost of Carbon. 2013. Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866. `http://www.whitehouse.gov/sites/default/files/omb/inforeg/social_cost_of_carbon_for_ria_2013_update.pdf`, last accessed December 20, 2014.

Figure 5: Avoided operating costs per MWh by technology type and region.

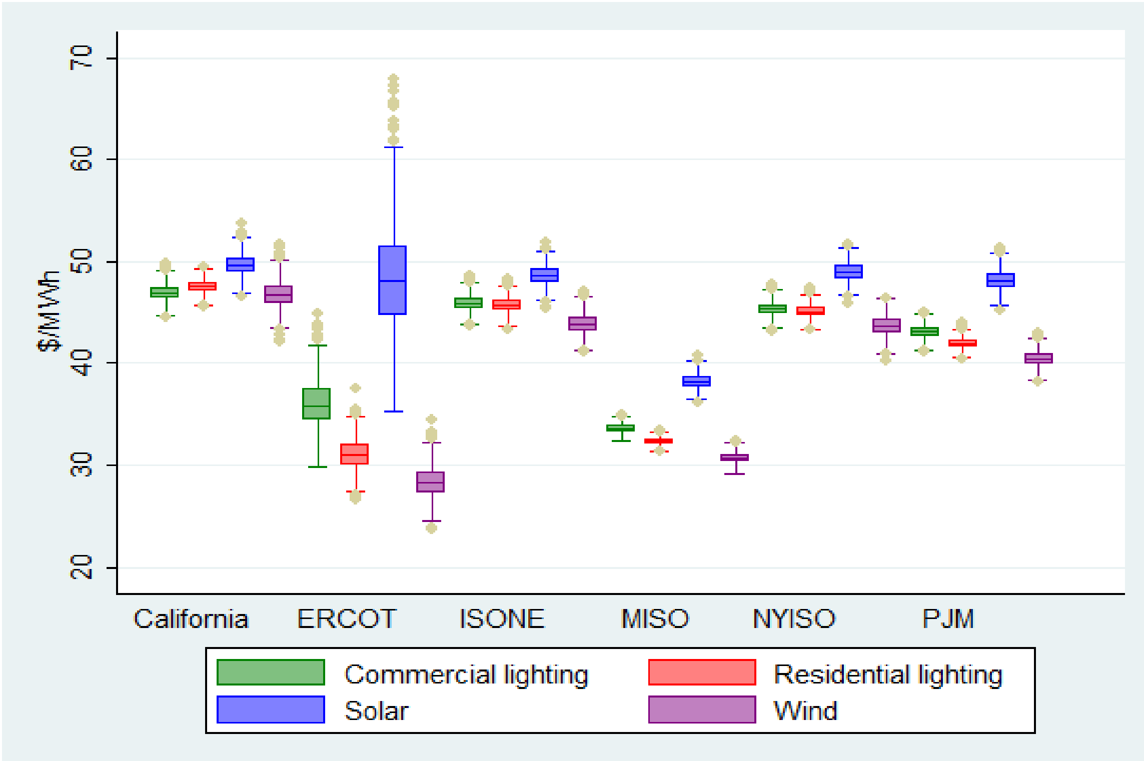

the temporal correlation between resource profiles and marginal operating costs. Intuitively, across all regions, solar PV is associated with the highest value estimates. Of all the resources we consider, solar PV production peaks are most coincident with peak load. In contrast, wind production tends to be negatively correlated with demand.

Cross-region comparisons of the λ_{rjt} should be made carefully. Differences in marginal prices across regions can reflect, among other factors, differences in market structure and associated incentives that govern the bidding behavior of the market. In our context, there are two institutional considerations that warrant careful consideration.

The first pertains to regional differences in resource adequacy and procurement. In contrast to other regions, ERCOT does not presently have a direct mechanism to procure generation capacity; all generator revenue comes from transactions for energy. As a consequence, energy prices in ERCOT are allowed to rise to very high levels to reflect scarcity of generation and incentivize construction of new capacity. In other words, ERCOT prices capture the cost to build new generation capacity in addition to fuel costs. Figure 5 shows that estimates of avoided operating costs are relatively more volatile in ERCOT. In the other regions, load serving entities are required to contract with generators solely on the basis of their existing capacity before any energy is transacted. These transactions are intended to ensure that sufficient generation capacity is built (or kept operating) to maintain system reliability. The resulting payments to owners of generation form part of their total revenue, meaning they only need to capture a portion of their revenue in energy markets to be profitable. MEV comparisons between ERCOT and other ISOs should made in this context.

A second consideration pertains to regional differences in emissions regulations. New York, New England, and some states in the PJM participate in the Regional Greenhouse Gas Initiative (RGGI). This initiative imposes a cap on CO_2 emissions from electricity generation; electricity producers must hold permits to offset emissions. This binding cap has two important implications for our analysis. First, the price of electricity generation in these three regions reflects the carbon permit price in addition to other variable operating costs. Second, because the emissions cap is binding, marginal increases in renewable generation capacity or energy efficiency will not reduce carbon emissions in RGGI. In other words, there are no external emissions displacement benefits of RE or EE in these regions; benefits manifest as avoided abatement costs on the margin (captured by the wholesale prices).

6.2 The marginal social value of RE and EE resources

Figure 6 summarizes the point estimates of marginal economic value in terms of emissions displacement benefits (red) and the value of avoided operating cost components (blue) by region and technology.

In the figure, there are no emissions displacement benefits associated with RE and EE investments in New York and New England. These markets are completely covered by the binding cap on CO_2 emissions. Wholesale market prices in these regions reflect fuel, operations, and emissions abatement costs incurred on the system operating margin. In the remaining regions where all (or in the case of PJM, a majority) of emissions are uncapped, the value of emissions displaced per MWh is added to the value of displaced operating costs. In California, where natural gas is on the margin in a majority of hours, emissions displacement benefits comprise a relatively small share (approximately a quarter) of marginal social value. Emissions displacement benefits comprise a larger share (as much as one half) of marginal social value estimates in PJM and MISO.

An overarching implication of Figure 6 is that, in regions where emissions are uncapped, external emissions displacement value comprise an important source of marginal social value

Figure 6: Marginal social value by technology type and region.

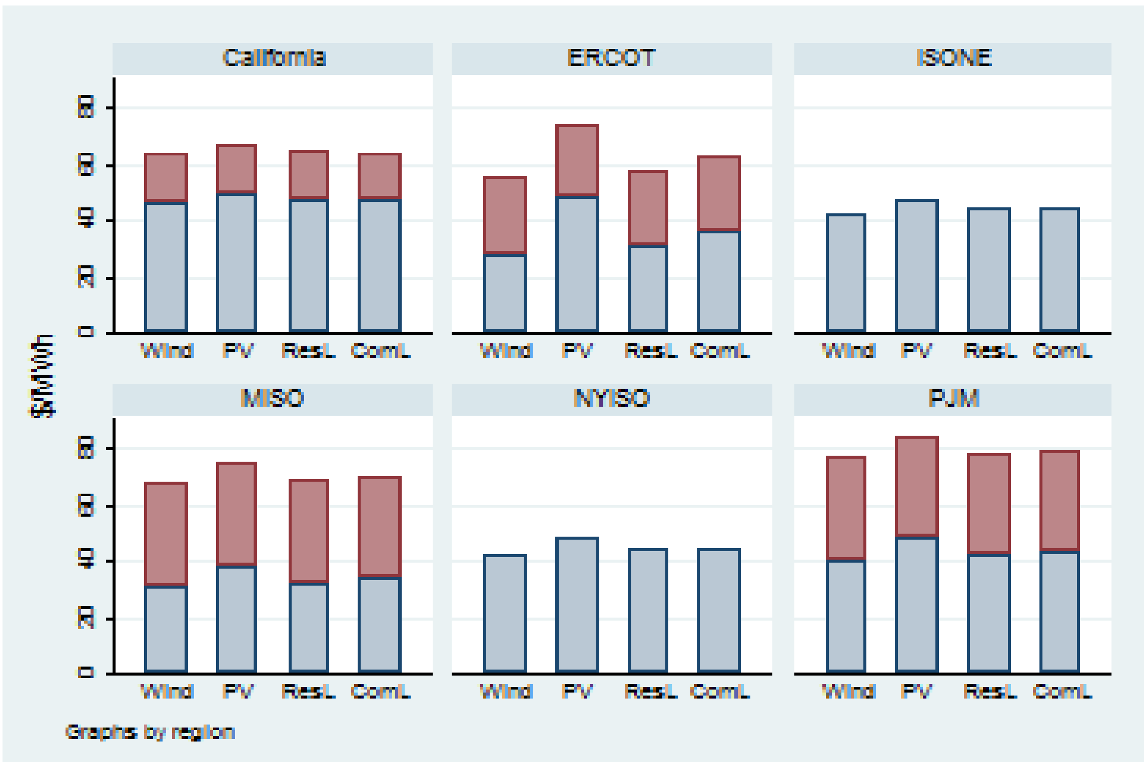

Notes: This figure summarizes point estimates of emissions displacement (in red; measured in monetary terms) and operating costs (in blue) displaced per MWh of renewable energy generated or demand-side electricity saved.

– ranging from roughly 25% of value in California to over half in MISO. By summing operating costs and the cost of carbon, we also see that, in the non-RGGI states, total social value becomes more similar across regions. That is, the regions with low operating costs tend to be those with the highest avoided carbon benefits. We also find that variation in operating costs can dominate the ranking of total social value – for example, solar PV ranks last with respect to emissions displacement value in MISO, ERCOT, and PJM, but first in terms of total social value due to the operating costs displaced during peak hours.

7 Marginal abatement cost

The final step in our analysis incorporates estimates of both benefits and costs of EE and RE investments. Using the notation from Section 2, where j corresponds to technology type, r denotes region, and t denotes time in hours, we compute the region- and technology-specific MACs as follows:

$$\mathrm{MAC}_{rj} = \frac{\mathrm{LCOE}_{rj} - \sum_{t=1}^{T} (\lambda'_{rt}\omega_{rjt})}{\delta_{rj}}$$

where λ'_{rt} is the marginal operating cost, and LCOE_{rj} is the levelized cost of energy introduced in Section 3. As in Section 2, ω_{rjt} is the ratio of a technology's production (or savings) in a region in period t to the production over the total period of analysis, and δ_{rj} is the mean MEDR computed in Section 5.

For CAISO, MISO, ERCOT and PJM, we set λ'_{rt} directly equal to the marginal fuel and operating cost λ_{rt} introduced in Section 3. For NYISO and ISONE, where all states are subject to a binding cap on power sector carbon emissions, we subtract the product of the emissions permit price and the corresponding MOER from the ISO LMP data to construct the marginal energy cost: $\lambda'_{rt} = \lambda_{rt} - P_{RGGI}\delta_{rj}$. We set P_{RGGI} equal to the average RGGI permit price in 2012(\$1.93/ton CO_2). [26]

Although we have made some adjustments to the λ parameter to account for carbon permit prices, there are several remaining caveats to consider. The first relates to levelized cost of energy calculations. As described in Section 3, wind LCOEs are built from records of region-specific power purchase agreements and solar LCOEs are built from historical in-

[26] We did not adjust PJM prices because, although Maryland, Delaware and New Jersey participated in RGGI for some or all of our study period, they are only a fraction of the larger PJM market, and we are using a spatially weighted LMP.

stalled costs (\$/watt) combined with region-specific solar resource potential. On the other hand, our lighting efficiency LCOE calculations are built from DOE's engineering-economic forecasts of the retail price of and costs to install new technologies, their power consumption, and estimates of how much they are operated on an annual basis (see Appendix for details). There are a number of reasons to argue that other factors should be included in the true levelized cost of lighting efficiency, including the economic impact on manufacturers and differences between the utility and the performance of the baseline technology and the more efficient option. In sum, considering that the levelized cost estimates for each technology are derived from disparate data sources, we cannot draw precise comparisons across technologies. However, as we shall see in the results, the marginal abatement cost differences across technologies tend to be very large, suggesting that, in most cases, the ranking of technologies is likely to be robust to methods for computing LCOE.

Second, as we discussed above, the ERCOT LMP also includes the effect of supply scarcity, whereas all other ISOs have capacity markets that tend to suppress wholesale energy prices. This inflates the average energy price in ERCOT somewhat, which has the effect of reducing ERCOT's MACs relative to what they would be in the presence of a capacity market.

Finally, note that for ISONE and NYISO – the two regions fully within RGGI – the MAC should not be interpreted as a marginal *abatement* cost (implying a reduction in carbon if the technology were deployed) but rather a marginal *compliance* cost. In other words, our estimates represent the cost to comply with RGGI's binding cap if one chose to use the technologies we are studying.

In light of these caveats, we will not show our results with confidence intervals because, although our bootstrap approach captures the variability from marginal operating emissions rates and wind, solar and lighting profiles, we cannot fully characterize the remaining sources of uncertainty in the MAC calculations, especially with respect to the LCOE. The latter uncertainty may well dominate any uncertainty we are able to capture. Therefore we present our results as point estimates only, with strong caveats on interpreting small differences in MAC values.

Figure 7 summarizes the results. The figure shows striking variation across technologies. Owing to its relatively high installed cost, solar PV is associated with the highest MAC.[27] Wind, still not competitive with wholesale electricity prices during our time period– at least

[27]It is worth noting that recent estimates of signed, but not completed, utility scale solar PPAs suggest that the cost of utility scale solar is dropping precipitously, in some cases to that of wind or even lower Bolinger and Weaver (2014).

Figure 7: Marginal abatement costs across technologies and regions

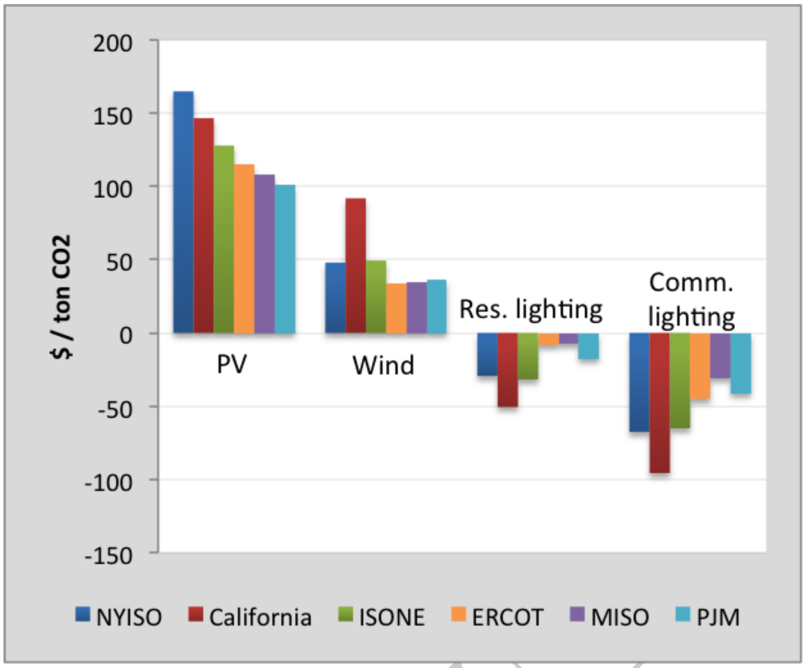

when measured with the average PPA prices in the data we used – is associated with positive costs per ton of CO_2 offset. However, in several of the regions we consider, estimated costs are similar to standard estimates of the social cost of carbon.

Turning to the investments in lighting efficiency, cost estimates are *negative* due to LCOE estimates produced by DOE. One interpretation is that the efficiency standards should be more stringently in order to equate marginal costs with marginal social returns. Note that our estimates take as given the engineering estimates of energy savings and technology costs. If, for example, assumed utilization rates are too high or assumed implementation costs are too low, this would reduce cost effectiveness. Recall also that we are using a 3% discount rate to calculate the LCOE, again reflecting our focus on total social costs rather than private consumer costs. A higher discount rate would be more realistic for private decisions and would push residential lighting to have positive cost in some regions.[28]

Keeping these caveats in mind, the basic ranking of options – PV with the highest MAC and commercial efficiency with the lowest – is likely to be robust to modifications to total social LCOEs for efficiency, owing to the large difference between them. For lighting technology, commercial MAC are superior to residential; though slightly different technologies are used for the LCOE calculations in commercial and residential, this difference is largely

[28]For example, a 7 percent discount rate would raise the LCOE by roughly 20 percent.

driven by the difference in operating hours (roughly 9 hours per day in the commercial sector versus two hours for residential).

We now turn our attention to variation in MAC estimates across regions and within technologies. Figure 7 technology rankings are identical across all regions. But within technologies, regional rankings vary. For solar PV, MAC estimates are highest in NYISO and California – but for very different reasons. In the case of New York, though the MEDR and economic value are average relative to other regions, the solar LCOE is low owing to a low solar resource. California, on the other hand, has a relatively good solar LCOE and economic value, but the MEDR is the lowest of the regions we are studying. At the other end, perhaps surprisingly, PJM and MISO are the best regions for solar in spite of their relatively low solar LCOE, due to the high avoided emissions.

The ranking of regions changes modestly for wind. In particular, due to historically high PPA prices there, California overtakes New York as the worst region. However the ranking remains intact otherwise.

For efficiency, the regional rankings within technologies takes on a different interpretation. Variation in negative cost estimates is driven predominantly by regional variation in MEDR (the difference between LCOE and avoided energy value is relatively similar across regions), with the lowest MEDR regions having the most negative MAC. This result is mathematically intuitive – smaller emissions reductions in the denominator results in a larger (in absolute value) fraction. In other words, relatively few tons of CO_2 are avoided for each unit of electricity saved.

In regions where greenhouse gas emissions are subject to a binding cap, these marginal abatement cost estimates can be compared to carbon allowance prices which reflect the shadow value of the constraint imposed by the emissions cap. In the period of our analysis, RGGI prices were below \$2 per ton of carbon. In contrast, we find that, in NYISO and ISONE (regions that fully participate in RGGI), PV's associated abatement cost is over \$120 per ton and wind's is about \$50 per ton. On the basis of carbon displacement potential alone, and with current technology costs and electricity supply infrastructure, PV and wind do not appear to be cost-effective technologies for carbon abatement vis a vis other abatement options. On the other hand, using data from the DOE to compute a levelized cost of energy for lighting efficiency, we find that the associated abatement cost is negative – a result driven by the fact that avoided energy costs exceed the LCOEs we use. This result suggests that further increasing efficiency standards would be very cost effective tool for reducing carbon emissions – or, in the case of ISONE and NYISO, reducing the costs of complying with the

RGGI cap.

8 Discussion

In this paper we have estimated the carbon displacement benefits associated with different RE and EE technologies across several regions in the United States. Within each region we find limited temporal variation in marginal operating emissions rates. As a result, technology-specific estimates of emissions displacement benefits vary little within regions, even though the timing of energy production does vary significantly across technologies.

In contrast, we find significant variation in emissions-related benefits *across* regions. For example, marginal emissions displacement rates in MISO and PJM's are more than double those in California across all technologies. Valuing carbon at $38 per ton, we find that carbon displacement potential is an important driver of marginal returns on investment. In regions where power sector emissions are not capped, emissions displacement value ranges from 50 to 100 percent of the avoided operating cost benefits of a given technology.

In addition to marginal economic value, we also examine marginal cost per ton of CO_2 – a measure that reflects regional and technological variation in levelized cost of energy, avoided operating cost and carbon displacement potential. Here we find that the variation in levelized cost of energy across technologies has a very strong influence on the marginal abatement cost. Cost estimates for solar PV exceed $100 per ton CO_2, whereas estimated costs associated with lighting efficiency improvements are negative. Within technologies, variation in emissions displacement rates across regions differs by technology. For example, abatement costs associated with PV are highest in NY and lowest in PJM, and costs associated with wind are highest in California and lowest in Texas. The range of variation within technologies is also economically large – marginal abatement costs for PV are 65 percent greater in NY than PJM and abatement costs for wind are nearly three times greater in California than Texas.

Overall, these results underscore the importance of designing policies that capture variation in emissions displacement benefits across regions. For example, the economic value of PV to a potential developer will appear to be highest in California under a production-based policy applied uniformly across regions. But from a social welfare perspective, our analysis suggests that PV will be more valuable in ERCOT, MISO and parts of PJM. In contrast, the efficiency gains associated with designing policies to capture intra-regional variation in emissions-related benefits appear small (if the primary policy objective is to efficiently internalize the carbon externality). [29]

These results should be interpreted with some important caveats in mind. First, this is a short-run analysis that conditions on the power system structure, the policy environment, and the technology characteristics we observe in recent years. Second, this is a marginal analysis; we evaluate impacts of relatively small, incremental increases in renewable energy and efficiency investment. Our approach is not well suited to evaluating long run impacts, nor should our estimates be used to value returns on large, non-incremental investments.

These caveats notwithstanding, our results do highlight the general importance of designing policies that accurately reflect regional differences in emissions displacement potential. The future of carbon regulation in the power sector is highly uncertain. However, under the proposed Clean Power Plan, RE and EE investments are virtually certain to play a very significant role in achieving emissions reduction targets. States are currently considering a wide range of possible compliance options, including production-based policies such as renewable portfolio standards and efficiency portfolio standards. If emissions targets are to be met cost effectively, it will be critical that policy incentives are designed to capture regional variation in emissions-related returns on carbon mitigation investments.

A Appendix

A.1 Data

Wind PPA data. In the 2013 Wind Technology Market Report, LBNL reports on a data set of 343 power purchase agreements (PPAs) totalling nearly 30 GW of installed wind capacity. LBNL collected these data from multiple sources, including FERCs Electronic Quarterly Reports, FERC Form 1, avoided-cost data filed by utilities, pre-offering research conducted by bond rating agencies, and a Berkeley Lab collection of PPAs. Figure 8 shows a summary of the full data set. There is a clear downward trend in wind prices following the 2009 peak. Though there are relatively few data in 2012-2013, their averages fall in line with the overall trends in the data set.

These PPAs bundle together the sale of electricity, capacity and renewable energy certificates and the receipt of federal incentives (e.g. the production tax credit, investment tax

[29] We note that our conclusions are based on interpreting performance-based (i.e. per-MWh) incentives. Though we did not directly relate our findings to capacity-based incentives, our findings regarding the relative insignificance of intra-regional variation in the emissions externality likely generalizes to capacity-based incentives.

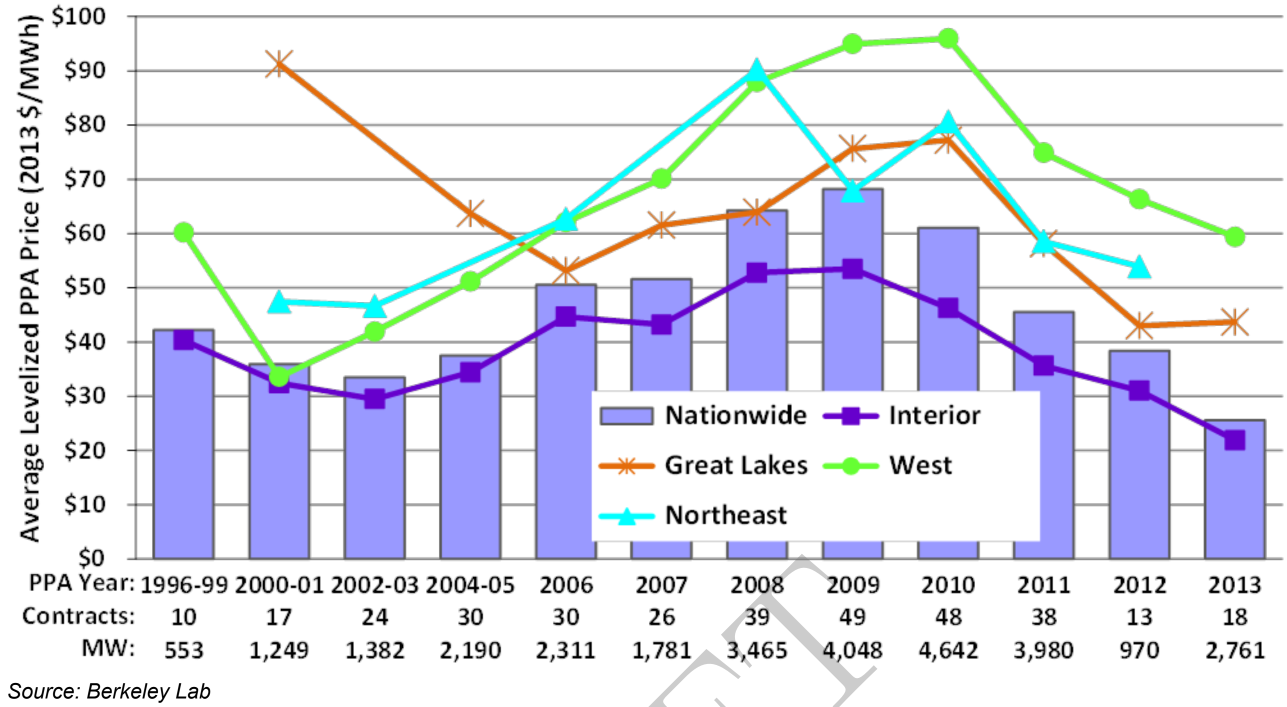

Figure 8:

credit or treasury grant). Neglecting the influence of policies at the state and local level as well as local market characteristics on PPAs, and assuming a competitive wind market, the PPA plus federal incentives will be representative of the levelized cost of wind power, and we treat them as such in this paper.

Solar production data. We chose 2 sites per region: Butte, CA and China Lake, CA (California); Austin, TX and Marfa, TX (ERCOT); Boston, MA and Concord, NH (ISONE), Rapid City, SD and Lansing, MI (MISO); Binghamton, NY and New York, NY (NYISO); Mansfield, OH and Virginia Beach, VA (PJM).

Efficiency LCOE calculations We used DOE estimates of technology costs and energy consumption to compute efficiency LCOEs. The key assumptinos are in Table 3, taken from the DOE's Technical Support Document for the General Service Fluorescent and Incandescent Reflector Lamps Energy Conservation Standard (Navigant Consulting, 2009).

The current DOE standard for general service fluorescent lamps is 88 lumens per watt for the lamp-ballast system (DOE, 2009). We gathered data on technology costs and energy consumption from the National Impacts Analysis for the current standard . In addition to technical assessments of the lumens per watt and installed cost (including retail price to

Table 3: Data used to calculate efficiency LCOE.

	residential	commercial
baseline technology	0.75 ballast factor, 32 watt	0.78 ballast factor, 32 watt
efficiency option	0.75 ballast factor, 30 watt	0.75 ballast factor, 32 watt
baseline cost	\$52.96	\$62.87
baseline energy	39.2 kWh	224.1 kWh
efficiency cost	\$53.55	\$63.31
efficiency energy	37.3 kWh	215.4 kWh
LCOE	\$15.38	\$2.51

Notes: (1) All lamps are electronic ballast. (2) We assumed a lamp and ballast replacement (due to failure of existing lamp and ballast) for both residential and commercial. (3) All lamps are T8. (4) Levelized cost computed by dividing cost difference between baseline and efficient option by the energy saved times an annuity factor for 15 years at 3% discount rate (=12.3).

consumer, taxes and installation labor) for each technology, DOE assumes residential lamps will be operated 791 hours per year, and commercial lamps for 3,435 hours per year. For each sector, we chose the baseline as the technology with the lowest installed cost in that sector that also meets the current standard. We defined the efficiency option as the technology with the lowest installed cost from among technologies that are more efficient than the baseline. We calculated a levelized cost of energy saved by the efficiency option over a fifteen-year period (per DOE's estimates that ballast lifetime is 15 years (DOE, 2009)), at a 3% discount rate.

A.2 k-means clustering

We cluster calendar days in our data using a k-means clustering algorithm. Within a given region and season, every day of the period 2010-2012 is given a 24-dimensional value based on the megawatt-hours per hour of fossil fuel generation in that hour. An additional value is added for the megawatt-hours per hour at peak that day (which may occur at different times). We then k-means cluster these 25-dimensional values. Thus, we are matching days on both the shape of electricity demand and the quantity in that day. We seed the clusters by initially matching entirely based on peak load in that day. In practice, the clustering algorithm frequently returns results which also closely map to clustering based on peak quantity, with load shape having relatively less influence on the cluster assignment.

We determine the number of clusters using the following algorithm. For each region and season combination, we cluster each of 11 different ways from 2 clusters, through 12. We then calculate MOERs using these methods. Beginning with the 12 cluster approach, we then check for a statistically significant difference between any of the MOERs that result. If there are no statistically significant differences between any of the clusters, we drop to the next lowest number of clusters. Thus, we use the smallest number of clusters which provides more informational content than the next smallest number of clusters.

Figure summarizes the results of this exercise. Each figure plots the average generation profile for each season-region-cluster triad. Bars denote 95-percentile confidence intervals. These graphs illustrate significant variation in load profiles even within a region-season.

A simpler approach to capturing this variation in load profiles would be to use the calendar month to proxy for intra-seasonal variation in load profiles. Figure X illustrates how our season-specific cluster composition varies in space and time. In each region and season, green denotes the first cluster associated with low load levels. Higher numbered clusters correspond with higher average load profile days. The figures show how our approach leads to a very different grouping of days as compared to a by-month grouping. We argue that our approach does a better job at controlling for the effects of load profile differences across days within a region-season.

A.3 MOER estimates

MOER panel figures

Also include (in future drafts). MOER table of estimates

MOER table of estimates dropping units reporting summer only.

A.4 Intra-regional variation in emissions displacmement across wind sites

We are interested in assessing the potential significance of variation in wind energy production profiles within a region. We start by estimating marginal emissions displacement rates for the over 30,000 wind sites in the data. Figure X arranges these sites in ascending order of estimated MEDR values. The figure suggests minimal variation in emissions displacement across sites.This is not altogether surprising given the limited variation in MOERs within most regions.

To put the variation summarized in the figure above into context, we systematically compare the MEDR estimates in either tail of these regional distributions. More precisely, in each region we select the ten sites straddling the 2.5 percentile value and ten sites straddling

Figure 9: Marginal operating emissions rates

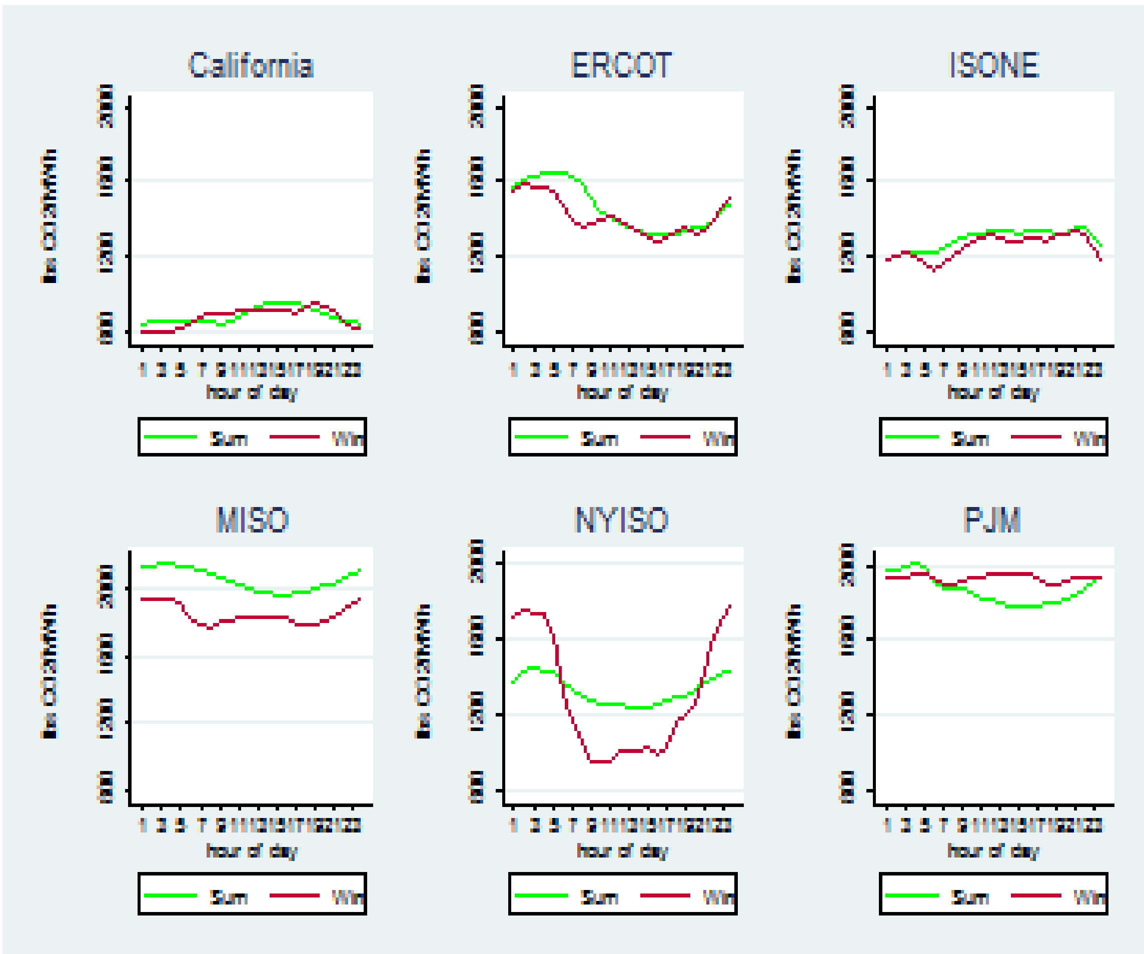

MOER estimates by season and region.

Figure 10: Site-specific marginal emissions displacement

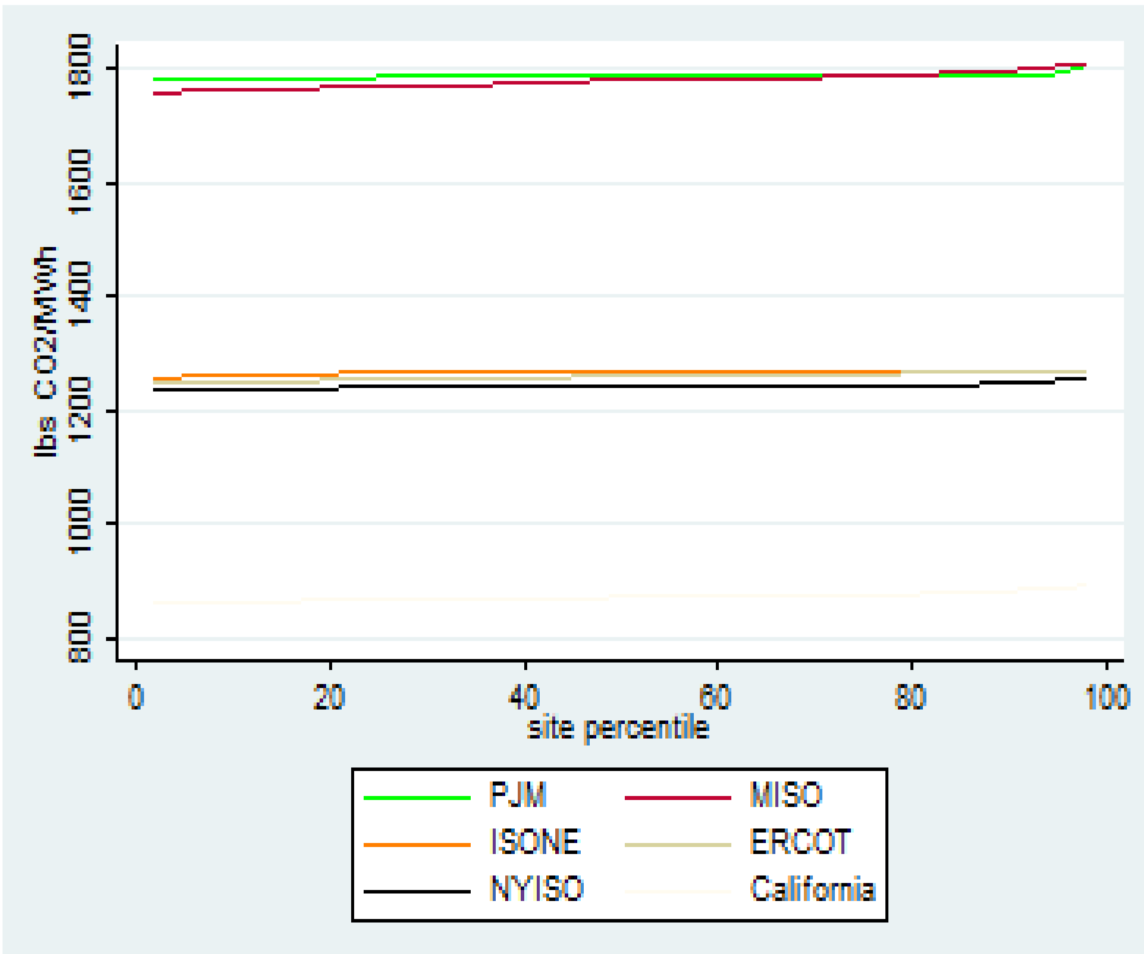

See the data appendix for a discussion of data sources.

the 97.5th percentile value. We bootstrap within region differences in MEDR across all possible pairwise comparisons between low and high ranked sites.

The table below summarizes these differences. The average pairwise difference between sites with high and low emissions displacement estimates, normalized by the average MEDR across all sites in the region, varies from 1 to 3 percent. We also report a more extreme difference. We take the two most different sites in each region, bootstrap the difference, and report the 95th percentile difference in MEDRs. The table shows that even this extreme measure of the difference in emissions displacement rates across sites within a region is small relative to the average MEDR (averaged across all sites in the region).

Region	Mean difference (lbs/MWh)	Extreme difference (lbs/MWh)	Mean difference as share of regional average	Extreme difference as share of regional average
ISONE	13	61	1%	5%
ERCOT	26	124	2%	9%
MISO	31	137	2%	7%
NYISO	28	89	2%	7%
PJM	14	105	1%	5%
California	24	141	3%	16%

Based on these results, we conclude that intra-regional variation in emissions displacement rates across wind sites is very small. We thus use only a subset of the sites (twenty sites from each region) to summarize the variation in emissions displacement values across regions and technologies.

A.5 Variation in marginal emissions displacement rates

The paper reports the average deviation of simulated marginal emissions displacement rates from regional and annual averages by technology. These interaction terms mask some regional variation. The table below reports the coefficient estimates from a fully saturated model (i.e. one including technology-region interactions).

References

Baker, E., Fowlie, M., Lemoine, D., and Reynolds, S. S. (2013). The economics of solar electricity. *Annual Review of Resource Economics*, 5(1):387–426.

Technology	California	ERCOT	ISONE	MISO	NYISO	PJM
Commercial	14.04	-30.25	14.22	-19.24	-90.77	-10.28
lighting	(7.84)	(12.18)	(11.39)	(16.21)	(13.42)	(16.18)
Solar	39.89	-86.57	42.45	-24.39	-190.19	-33.03
PV	(9.15)	(15.08)	(15.84)	(23.89)	(23.38)	(21.82)
Residential	2.20	-22.90	15.31	-16.73	37.67	0.53
lighting	(3.30)	(5.95)	(5.85)	(6.44)	(12.35)	(5.78)
Wind	3.29	11.23	-3.49	-13.02	0.26	4.50
	(20.19)	(27.60)	(21.25)	(31.45)	(22.39)	(28.35)

Barbose, G., Weaver, S., and Darghouth, N. (2014). Tracking the sun vii: An historical summary of the installed price of photovoltaics in the united states from 1998 to 2013. Technical report, Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA (US).

Bolinger, M. and Weaver, S. (2014). Utility-scale solar 2013: An empirical analysis of project cost, performance, and pricing trends in the united states.

DOE (2009). Energy conservation program: Energy conservation standards and test procedures for general service fluorescent lamps and incandescent reflector lamps; final rule. Technical report, Department of Energy.

Graff Zivin, J. S., Kotchen, M. J., and Mansur, E. T. (2014). Spatial and temporal heterogeneity of marginal emissions: implications for electric cars and other electricity-shifting policies. *Journal of Economic Behavior & Organization*.

Navigant Consulting, I. (2009). Technical support document: Development of the energy conservation standards for general service fluorescent lamps and incandescent reflector lamps. Technical report.

Siler-Evans, K., Azevedo, I. L., and Morgan, M. G. (2012). Marginal emissions factors for the us electricity system. *Environmental science & technology*, 46(9):4742–4748.

Wei, M., Nelson, J., Ting, M., Yang, C., Kammen, D., Jones, C., Mileva, A., Johnston, J., and Bharvirkar, R. (2012). Californias carbon challenge: Scenarios for achieving 80% emissions reductions in 2050. *Lawrence Berkeley National Laboratory, UC Berkeley, UC Davis, and Itron to the California Energy Commission. URL.*

Wiser, R., Bolinger, M., and Berkeley, L. (2014). 2013 wind technologies market report.