Online Appendix: Imperfect Markets versus Imperfect Regulation in U.S. Electricity Generation

Steve Cicala

November 13, 2019

A Data Appendix: For Online Publication

A.1 Power Control Area Definitions

The definition of a Power Control Area (PCA), or Balancing Authority (BA) is somewhat flexible, varying across regions, regulatory agencies, and over time. For the purposes of this paper, I am interested in identifying decision-making units responsible for allocating production to generating units to keep supply and demand for electricity in balance at all moments in time. As a practical matter, the method I use to measure the value of reservoir hydropower (discussed below) relies on the opportunity cost of water based on offsetting fossil power. I therefore classify PCAs as those recognized as such in load reporting by FERC in 1999, also reporting control of fossil-fired units based on a combination of reporting in FERC Form 714 (Part II, Schedule 1), and the 1999 configuration of the grid based on EPA's eGRID database.[25] This results in the consolidation of a number of "planning areas" that report their own load, though they do not dispatch plants, as well as a few hydropower-only PCAs in the Pacific Northwest (see Tables A.1, A.2, and A.3). I use county-level approximations of these 1999 PCA configurations when using demographic, meteorological, and economic variables to predict load (see Section A.2). To construct these maps, I begin with the 1999 EIA Form 861, "Annual Electric Power Industry Report", which connects local utilities to PCAs, and reports the counties in which respondent utilities have generation equipment. I then use individual service territory maps (via internet search) to refine these boundaries.

Much like the Neighborhood Change Database, the goal is to create a time-invariant characterization of the grid, which has indeed changed over time. New generation units, for example, are associated with their contemporary ISO rather than historical PCA. To determine the 1999 PCA in which new generation *would have* been located (ignoring differential investment issues), I first use local utility association: many of these utilities are unchanged in spite of changes to the bulk electricity system. If that utility belonged to a 1999 PCA, the plant inherits that association. If no other information is available, the 1999 PCA maps are used to associate new generation with historical areas.

A.2 Load Data

Hourly data on electricity usage (load) are compiled from a combination of the Federal Energy Regulatory Commission (FERC) Form 714, local system operators, and the North American Electricity Reliability Corporation (NERC), depending on data availability. While this data, in theory, are publicly-available from a straightforward download from the FERC site, this is emphatically not the case in practice. Until 2006, there was no required submission format for hourly load data, so that each PCA's data might be submitted in anything from an Excel file to free-form text, often without a codebook. In addition, there is no standard procedure for accounting for daylight savings time: some areas ignore it completely, others report zero at the start and double-report the final hour, etc. Annual reports are missing altogether for some PCA-years, or are reported as part of the load of an adjacent area (again, often without documentation). A number of smaller areas that do not own generation (and are therefore *planning* areas, rather than *control* areas) are combined with the neighboring PCA that conducts dispatch. To avoid PCAs composed entirely of estimated hydro generation, a handful of areas in the Pacific Northwest are combined as well. Areas that join an ISO

[25] eGRID is used as the starting point, then corrections are made by hand based on FERC reporting because these forms are only available as (occasionally handwritten) pdfs of plant names, rather than EIA facility codes.

often have their load included in the ISO total, and may not be available as a single PCA.

When missing, hourly load data are estimated separately for each PCA using LASSO to uncover the best functional form in a disciplined manner. One benefit of consumers' insulation from electricity market conditions is that electricity load can estimated extremely well as a function of time (of year, week and day), population, weather, and economic conditions. The day of week/year variables used in prediction are a set of trigonometric functions with varying periodicity over the course of the week and year to account for regular calendar fluctuations. Temperature variables measure heating and cooling degrees (degrees above or below 65^oF) on the daily minimum, maximum, and hourly temperatures, as well as dew point. This data come from NOAA stations and PRISM Climate Group (2004), and collapses county-level data with population weights for PCA-wide measures. Economic data include unemployment rates as well as electricity-intensive employment in manufacturing and mining sectors aggregated from the county level to the approximate footprint of the PCA in 1999 as with the meteorological data.

These variables are used in a LASSO estimation procedure to avoid over-fitting by including a regularization term in the standard OLS procedure that sets less important predictors to exactly zero rather than fit on noise. When estimated using data for even years, it produces estimates that have a mean absolute deviation of less than 5% when validated against odd years.[26]

Western Interconnection

Load data for the Western Interconnection come from both FERC and the Western Electricity Coordinating Council (WECC), depending upon availability. PCAs as of 1999, and constituent load-reporting areas are reported in Table A.1. The abundance of hydropower requires the consolidation of a number of Public Utility Districts in Washington and Irrigation Districts in California to arrive at a level of aggregation such that reservoir power is offsetting fossil power. In addition, there is relatively inconsistent reporting of load in the former Southern California Edison territory, though total load from the California Independent System Operator (CAISO) is well-reported, as is load from the other territories in the CAISO footprint. Southern California Edison load is therefore calculated from the remaining CAISO load after subtracting off load from Pacific Gas & Electric, San Diego Gas & Electric, and their respective constituent load areas.

Texas Interconnection

The Electric Reliability Council of Texas (ERCOT) is a separate interconnection that consolidated ten PCAs into a single market on 31 July, 2001. After a period of only reporting total ERCOT load in 2001 and 2002, the ISO began reporting load by eight "weather zones" that do not cleanly overlap with the original PCAs. The ERCOT total is consistently reported throughout the sample period. I therefore run LASSO using the ERCOT total and the population-weighted characteristics for the entire ISO, then use the ISO-derived coefficients projected upon the PCA-level characteristics to predict PCA-level load. Final estimates are scaled by the ratio of observed ERCOT load to the sum of predicted PCA loads to ensure that the totals match those observed in the data. This method delivers estimates of load in 1999 and 2000 for the original PCAs that have an absolute mean deviation from the true loads of about 6%, in line with the out of sample estimates delivered by estimating fixed footprints to years without load data.

[26]Using only PCA-specific means yields an error of about 20%, which is reduced to 15% by using PCA-hour means, and no other explanatory variables.

Eastern Interconnection

Load data for the Eastern Interconnection varies in the consistency of reporting. The Northeastern ISOs in New York and New England did not consolidate multiple PCAs upon transition to markets, but simply changed the method for allocating output over a fixed territory–load reporting is consistent throughout. PCAs in the Pennsylvania-Jersey-Maryland (PJM) market have delegated load reporting to the ISO, but PJM has helpfully preserved the original footprints as the basis for more detailed reporting available through their website. The Southwest Power Pool (SPP) has also made hourly load data available by original PCAs in spite of aggregate reporting to FERC. This has, unfortunately, not been the case for the Midwest ISO (MISO), which declines to release the disaggregated data which was previously publicly available before the ISO took over load reporting in 2009. Instead, the most disaggregated load available from MISO is broken down by three large regions spanning many former PCAs each. Fortunately MISO began market dispatch three years before taking over load reporting, so demand is largely observed through the transition to markets. Determining loads from constituent PCAs in MISO since 2009 required a bit of reverse-engineering. MISO publishes anonymized bid data in their day-ahead market, which includes the price and quantity cleared. They also publish location-specific prices that clear the day-ahead market. Because identifiers in the bid data are relatively persistent, there is a time path of cleared prices forms a unique identifier that can be matched to exact locations in published prices (congestion creates node-specific prices). This allows me to place cleared quantities in their respective load zones, which correspond to former PCAs. The overlapping years of separate FERC reporting and published day-ahead quantities confirm an exceptionally high-quality match using this method.

Table A.1: Power Control Areas in the Western Interconnection

1999 PCA	Constituent Load	Unreported Periods	ISO	Market Date
Avista	Avista Corp			
Arizona Public Service	Arizona Public Service Co			
	WAPA Lower Colorado River	2001-2007		
Bonneville Power Authority	Bonneville Power Authority			
	PUD 1 of Chelan County			
	PUD 1 of Douglas County			
	PUD 2 of Grant County			
El Paso Electric	El Paso Electric Company			
Imperial Irrigation District	Imperial Irrigation District			
Idaho Power	Idaho Power Company			
Los Angeles Dept Of Water & Power	Los Angeles, City Of			
Montana Power	Montana Power Company			
Nevada Power	Nevada Power Company			
Pacificorp	Pacificorp			
Pacific Gas & Electric	Pacific Gas & Electric	2011-2012	CAISO	1 April 1998
	Modesto Irrigation District		CAISO	1 April 1998
	WAPA Sierra Nevada Region		CAISO	1 April 1998
	City of Redding		CAISO	1 April 1998
	Sacramento Municipal Utility District		CAISO	1 April 1998
Portland General Electric	Portland General Electric Co			
Public Service Co Of New Mexico	Public Service Co Of New Mexico			
Public Service Co Of Colorado	Public Service Co Of Colorado			
Puget Sound Energy	Puget Sound Energy			
	Seattle Department of Lighting			
	Tacoma Power			
Southern California Edison	Southern California Edison		CAISO	1 April 1998
	City of Vernon		CAISO	1 April 1998
	CA Dept of Water Resources		CAISO	1 April 1998
	City of Anaheim		CAISO	1 April 1998
	City of Santa Clara		CAISO	1 April 1998
	City of Riverside		CAISO	1 April 1998
	City of Pasadena		CAISO	1 April 1998
San Diego Gas & Electric	San Diego Gas & Electric		CAISO	1 April 1998
Sierra Pacific Power	Sierra Pacific Power Co			
Salt River Project	Salt River Project			
Tucson Electric Power	Tucson Electric Power Co			
WAPA Colorado-Missouri	WAPA Colorado-Missouri			

Table A.2: Power Control Areas in the Texas Interconnection

1999 PCA	Constituent Load	ISO	Market Date
Central And South West Services	Central And South West Services (AEP)	ERCOT	31 July 2001
	South Texas Electric Cooperative	ERCOT	31 July 2001
	Brownsville Public Utilities Board	ERCOT	31 July 2001
Lower Colorado River Authority	Lower Colorado River Authority	ERCOT	31 July 2001
	Austin Energy	ERCOT	31 July 2001
Reliant Energy	Reliant Energy	ERCOT	31 July 2001
San Antonio Public Service Board	San Antonio Public Service Board	ERCOT	31 July 2001
Texas Municipal Power Pool	Texas Municipal Power Pool	ERCOT	31 July 2001
TXU Energy	Texas Utilities	ERCOT	31 July 2001
	Texas-New Mexico Power Company	ERCOT	31 July 2001

Table A.3: Power Control Areas in the Eastern Interconnection

NERC Region	1999 PCA	Constituent Load	Unreported Periods	ISO	Market Date
ECAR	American Electric Power	American Electric Power		PJM	1 October 2004
ECAR		Buckeye Power		PJM	1 October 2004
ECAR		American Municipal Power - Ohio		PJM	1 October 2004
ECAR	Allegheny Power Service	Allegheny Power Service		PJM	1 April 2002
ECAR	Big Rivers Electric	Big Rivers Electric Corp		MISO	1 December 2010
ECAR	Cinergy	Cinergy		MISO	1 April 2005
ECAR	Consumers Energy	Consumers Energy		MISO	1 April 2005
ECAR	Detroit Edison	Detroit Edison Co		MISO	1 April 2005
ECAR	Duquesne Light	Duquesne Light Company		PJM	1 January 2005
ECAR	Dayton Power & Light Co	Dayton Power & Light Co		PJM	1 October 2004
ECAR	East Kentucky Power Coop	East Kentucky Power Coop			
ECAR	Firstenergy	Firstenergy		MISO	1 April 2005
ECAR	Hoosier Energy	Hoosier Energy		MISO	1 April 2005
ECAR	Indianapolis Power & Light	Indianapolis Power & Light		MISO	1 April 2005
ECAR	Louisville Gas & Electric	Louisville Gas & Electric		MISO	1 April 2005
ECAR	Northern Indiana Public Service	Northern Indiana Public Service		MISO	1 April 2005
ECAR	Southern Indiana G & E	Southern Indiana G & E		MISO	1 April 2005
FRCC	Florida Municipal Power Agency	Florida Municipal Power Agency			
FRCC		Orlando Utilities	2003, 2004		
FRCC		City of Lakeland	2004		
FRCC	Florida Power	Florida Power Corporation			
FRCC	Florida Power & Light	Florida Power & Light			
FRCC	Gainesville Regional Utilities	City Of Gainesville			
FRCC	Jacksonville Electric Authority	Jacksonville Electric Authority			
FRCC	Seminole Electric Coop	Seminole Electric Coop	2003		
FRCC	City Of Tallahassee	City Of Tallahassee			
FRCC	Tampa Electric	Tampa Electric Company	2004		
MAAC	Pennsylvania-Jersey-Maryland	Pennsylvania-Jersey-Maryland		PJM	1 April 1997
MAIN	Alliant East	Alliant East		MISO	1 April 2005
MAIN	Ameren	Ameren	2007-2012	MISO	1 April 2005
MAIN		Columbia Water & Light	2004	MISO	1 April 2005
MAIN	Commonwealth Edison	Commonwealth Edison Co		PJM	1 May 2004
MAIN	Central Illinois Light	Central Illinois Light Co	2007-2012	MISO	1 April 2005
MAIN	Illinois Power	Illinois Power	2009-2012	MISO	1 April 2005
MAIN	Southern Illinois Power Coop	Southern Illinois Power Coop		MISO	1 April 2005
MAIN	Springfield (IL) Water Light & Power	City of Springfield, IL		MISO	1 April 2005
MAIN	Wisconsin Energy	Wisconsin Electric Power		MISO	1 April 2005
MAIN	Wisconsin Public Service	Wisconsin Public Service	2004	MISO	1 April 2005
MAIN		Madison Gas & Electric Co	2004	MISO	1 April 2005
MAPP	Alliant West	Alliant West		MISO	1 April 2005
MAPP	WAPA Upper Missouri	WAPA Upper Missouri East Basin	2000, 2001, 2004, 2005	MISO	1 April 2005
MAPP		Basin Electric Power Cooperative	2000, 2001, 2004, 2005		
MAPP		WAPA Upper Missouri West Basin	2000, 2001, 2004, 2005	MISO	1 April 2005
MAPP	Dairyland Power Coop	Dairyland Power Coop	2000, 2001, 2004	MISO	1 June 2010
MAPP	Great River Energy	Great River Energy	2000-2002, 2004	MISO	1 April 2005
MAPP	Midamerican Energy	Midamerican Energy		MISO	1 September 2009
MAPP		Muscatine Power & Water	2000, 2001, 2004, 2006-2009	MISO	1 September 2009
MAPP	Minnesota Power & Light	Minnesota Power & Light	2000-2004, 2009-2012	MISO	1 April 2005
MAPP	Nebraska Public Power District	Nebraska Public Power Dist	2004	SPP	1 April 2009
MAPP		Lincoln Electric System	2004	SPP	1 April 2009
MAPP	Northern States Power	Northern States Power Co	2001, 2004	MISO	1 April 2005
MAPP		Southern MN Municipal Power	1999, 2000, 2004, 2006	MISO	1 April 2005
MAPP	Omaha Public Power District	Omaha Public Power District	2000, 2001, 2004	MISO	1 April 2005
MAPP	Otter Tail Power	Otter Tail Power	2001, 2004	MISO	1 April 2005
MAPP		Minnkota Power Cooperative	1999-2001, 2004	MISO	1 April 2005
NPCC	New England Power Pool	New England Power Pool		NEISO	1 May 1999
NPCC	New York Power Pool	New York Power Pool		NYISO	18 November 1999
SERC	Alabama Electric Cooperative	Alabama Electric Cooperative			
SERC	Associated Electric Cooperative	Associated Electric Cooperative			
SERC	Carolina Power & Light	Carolina Power & Light			
SERC	Duke Energy	Duke Energy			
SERC	South Carolina Electric & Gas	South Carolina Electric & Gas			
SERC	South Carolina Pub Serv Auth	South Carolina Public Service Authority			
SERC	South Mississippi Electric Power	South Mississippi Electric Power	2001		
SERC	Southern	Southern Co			
SERC		Oglethorpe Power			
SERC	Tennessee Valley Authority	Tennessee Valley Authority			
SERC	Dominion Virginia Power	Dominion Virginia Power		PJM	1 May 2005
SPP	CLECO	Central Louisiana Electric Co	2002	SPP	1 February 2007
SPP		Lafayette Utility System		SPP	1 February 2007
SPP	Empire District Electric	Empire District Electric		SPP	1 February 2007
SPP	Entergy	Entergy			
SPP	Grand River Dam Authority	Grand River Dam Authority	2002	SPP	1 February 2007
SPP	Kansas City Power & Light	Kansas City Power & Light	2002, 2006	SPP	1 February 2007
SPP		City of Independence	2002, 2006	SPP	1 February 2007
SPP	Kansas City Board of Public Utilities	Kansas City Board of Public Utilities	2002	SPP	1 February 2007
SPP	Louisiana Energy & Power Authority	Louisiana Energy & Power Auth			
SPP		Louisiana Generating	2000, 2007-2012		
SPP	Aquila Networks - MPS	Missouri Public Service Co		SPP	1 January 2010
SPP	Oklahoma Gas & Electric Co	Oklahoma Gas & Electric Co	2002	SPP	1 February 2007
SPP	Sunflower Electric Cooperative	Sunflower Electric Cooperative		SPP	1 February 2007
SPP	Public Service of OK (SWEPCO)	Public Service of OK (SWEPCO)	2002, 2006	SPP	1 February 2007
SPP	Southwestern Power Admin	Southwestern Power Admin		SPP	1 February 2007
SPP		Southwestern Public Service	2002	SPP	1 February 2007
SPP		Golden Spread Electric Cooperative	2003	SPP	1 February 2007
SPP	Western Farmers Elec Coop	Western Electric Farmers Coop	2002	SPP	1 February 2007
SPP		Oklahoma Municipal Power Authority	2000, 2001, 2003-2005	SPP	1 February 2007
SPP	Aquila Networks - WPK	Aquila - WestPlains	2007-2012	SPP	1 February 2007
SPP	Western Resources	Western Resources	2002	SPP	1 February 2007

A.3 Generation

EPA Continuous Emissions Monitor System (CEMS) Data

Roughly 95% of fossil-powered generation in the United States from 1999-2012 is reported at the boiler-hour level in the EPA CEMS data. While reporting for the largest units begins in 1996, comprehensive reporting does not begin until 1999. This data system was designed to monitor emissions for compliance with NO_x and SO_2 programs of the 1990 Amendments to the Clean Air Act. I adjust for net-to-gross ratios, as CEMS reports gross generation for unit i in hour t, but power station usage (typically about 6%) must be subtracted to measure how much power is being sent to the grid. The adjustment is a unit-level version of that used by Cullen and Mansur (n.d.) to measure hourly net generation, except at the interconnection-fuel level, rather than generating unit:

$$\text{Net Generation}_{it} = \frac{\text{EIA Net Generation}_{im}}{\sum_{t \in m} \text{CEMS Gross Generation}_{it}} * \text{CEMS Gross Generation}_{it}$$

The hourly gross load data from CEMS are merged to the monthly EIA data on net generation and heat rate (at the unit level from EIA-767 when possible, otherwise at the plant-prime mover level from EIA-906/920/923), and scaled by the ratio of monthly net generation to monthly gross load. This accomplishes two tasks: First, it ensures output represents net generation at the hourly level, smoothing start-up and ramping costs over the month. Second, a number of units (especially Combined Cycle units) only include the steam portion of the unit in CEMS, leaving the generation from the second cycle unreported. This scaling treats the (hourly) unreported generation as dispatched at the same rate as the main unit. Figure A.1 plots the kernel density estimates of the net-to-gross ratios of CEMS units. While a density of the scaling is overwhelmingly concentrated at about 94% (the gross-to-net scale), combined cycle units feature a bimodal distribution with a second (much smaller) peak around 1.4 (reflecting the contribution of the unmatched cycles). The top and bottom percentile of scales is trimmed for outliers (outside of 0.2 and 2), and estimated scales based on a regression of observable unit characteristics are used instead.

Nuclear Generation

Daily output from nuclear-powered units is reported to the Nuclear Regulatory Commission (NRC) as a share of potential output for each generating unit. The exact potential output is calculated by taking the ratio of the monthly total net generation reported on EIA Forms 906/920/923 to the sum of these daily potential outputs:[27]

$$\text{Net Generation}_{it} = \frac{\text{EIA Net Generation}_{im}}{\sum_{t \in m} \text{NRC Share of Capacity}_{it}} * \text{NRC Share of Capacity}_{it}$$

Output levels are determined by multiplying the calculated potential output by daily share of output generated, and distributing the generation over the hours of the day. There is minimal potential for error in this last step because nuclear units are typically running at maximum capacity, down for maintenance, or transitioning between the two over the course of days.

[27] Note that the ratio of monthly net generation to the sum of capacity shares is equal to the unit's capacity:$\text{Capacity}_{im} = \frac{\text{EIA Net Generation}_{im}}{\sum_{t \in m} \text{NRC Share of Capacity}_{it}}$

Figure A.1: Net-Gross Scaling of CEMS Units with Monthly Data

(A) All CEMS units

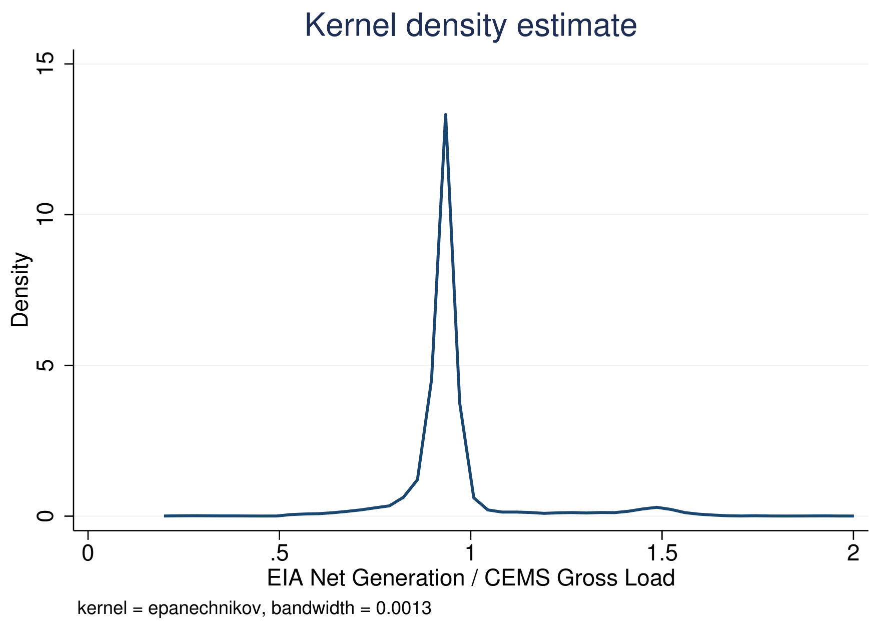

(B) Combined Cycle Units in CEMS

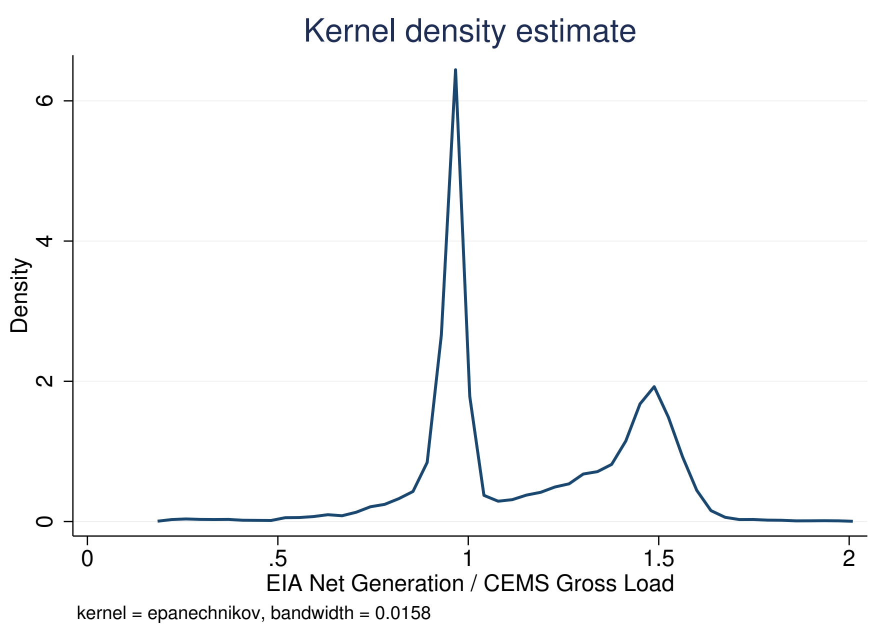

Hydro Generation

Monthly hydro generation is reported on Form EIA-906/920/923. I use discharge and/or streamflow data to distribute this aggregated generation across the hours of the month. Inquiries at individual hydropower administrators yielded exact hourly turbine discharge numbers. These sources include the Tennessee Valley Authority, U.S. Army Corps of Engineers, and U.S. Bureau of Reclamation. Over one third of hydropower generation is collected from these sources.

When exact numbers were not available, I infer hourly generation from streamflow data collected by the U.S. Geological Survey's streamgage network. To do so, I use plant coordinates from EIA to locate generators along the National Hydrography Database's stream network. I then use network analysis software to identify the nearest downstream gage that collects streamflow data (in many cases the dam's discharge monitor *is* the USGS streamflow recorder). I then allocate the month's generation to hours based on the share of monthly streamflow released in that hour in the same spirit as the generation allocation methods described above. A validation exercise that regresses hourly generation for those dams where it is directly observed from administrative agencies on predicted generation based on the USGS streamgage method has a coefficient of 0.96 and an R^2 of 0.84.

Fossil Generation Reported Monthly

Between 2% and 6% of generation from fossil-powered units does not appear in the CEMS data. For these units, I use monthly data at the unit level from EIA-767 (mostly small steam-powered boilers), or plant-prime mover level from EIA-906/920/923 (mostly small gas turbines and internal combustion generators). To allocate this production to hours of the month, monthly generation is allocated over hours by ranking the hours of the month by load, and producing at maximum capacity in the highest load hours up to the total reported monthly generation (or analogously for annual generation if that is the level of EIA-906/920/9233 reporting).

For example, if a turbine with a capacity of 10MW reports in EIA-906/920/923 that it produced 10MWh in a month, it is assumed that it only produced in the hour of maximum load that month. If it produced 50MW, it is assumed to have produced 10MWh in the five highest demand hours that month, and was idle otherwise. This approach is motivated by demand sweeping through the merit order, so that generating units are only briefly marginal, and therefore typically producing at maximum capacity, if at all. In this case, generation is proportional to the number of hours in which demand is sufficiently high for the unit to be "in" the merit order.

This algorithm is applied consistently throughout the sample period, and is invariant to the institutions used to allocate production.

Other Generation

Monthly wind generation by farm is reported in EIA-906/920/923. I merge this data with hourly wind speed data from nearby weather stations, as reported in National Oceanic and Atmospheric Administration's (NOAA) Integrated Surface Database. I estimate the potential output of the farm based on a cut-in wind speed of 3 m/s, a cut-out speed of 20 m/s, a rated wind speed of 10 m/s (at which point the farm produces at maximum capacity), and an increasing cubic between the cut-in and rated wind speeds. I then use predicted power to distribute the observed monthly total over the month:

$$\text{Net Generation}_{it} = \frac{\text{EIA Net Generation}_{im}}{\sum_{t\in m}\text{Predicted Wind Generation}_{it}} * \text{Predicted Wind Generation}_{it}$$

For areas that report hourly wind generation in the ISO footprint (ERCOT since 2007, SPP and MISO since 2008), I use the hourly analog for wind farms in the ISO n's footprint:

$$\text{Net Generation}_{it} = \frac{\text{ISO Wind Generation}_{nt}}{\sum_{i \in n} \text{Predicted Wind Generation}_{it}} * \text{Predicted Wind Generation}_{it}$$

For solar generation, I match the county of each installation reporting in EIA-906/920/923 with hourly location-specific data on global horizontal irradiance (GHI) from NREL's Solar Radiation Database. Monthly net generation is allocated to hours in proportion to the share of the month's total GHI at that site, in that hour.

Finally, geothermal generation is also reported in EIA-906/920/923 at the plant-month level. Geothermal plants are used as baseload, run at maximum possible capacity at zero marginal cost. The monthly generation is evenly distributed over the hours of the month.

A.4 External Validation of Fit

Figure 6 shows how these various sources come together to produce a picture of supply and demand that match relatively well. Panel (b) shows idiosyncratically high imports relative to exports between PCAs in 2000, prospective net exports around 2004, and a relatively steady growth in the gap between imports and exports since that time. The differences between imports and exports between PCAs in this figure is either because something is missing in my calculations of load and generation, or that the US has changed its net exports with respect to Mexico and Canada over time.

To validate the net electricity disposition of the US, I collected monthly data on imports and exports from Mexican and Canadian statistical agencies. For Mexico, this data is available back to 2005,[28] and shows that net imports are generally around zero and rarely the equivalent of 200MW-months. In Canada, this data is published by the National Energy Board's Commodity Tracking System.[29] Monthly net imports from Canada are plotted against the difference between load and generation in the data in Figure A.2. It shows that the data are largely in sync, with differences rarely greater than the equivalent output of a few GW power plants. The two series generally match trends in how net imports have changed over time.

A.5 Heat Rates and Capacities

With a Leontief production function, the ratio of output to heat inputs measures the productivity of generation unit. A substantial literature has developed in industrial organization to consistently measure (Hicks-neutral) productivity, which is typically unobserved and time-invariant. It is possible to measure supply curves in the electricity setting because unit productivity is (more or less) time invariant and capacities are known whether the unit is operating or not. This simplifies matters quite a bit.

Heat rates when operating are observed at the unit-month level in EIA-767, and EIA-906/920/923 at the plant-prime mover-month. When not operating, I use estimated heat rates based on regressions including unit-specific trends.

Heat rates for cogeneration units are a bit trickier: these are units that also provide useful steam energy, making it economical to run even if would not be operating otherwise. These units tend

[28] https://datos.gob.mx/busca/dataset/comercio-exterior-importacion-y-exportacion-de-energia-electrica
[29] https://apps.neb-one.gc.ca/CommodityStatistics/Statistics.aspx?language=english

Figure A.2: Monthly Net Imports from Canadian Statistics versus Load-Generation Mismatch

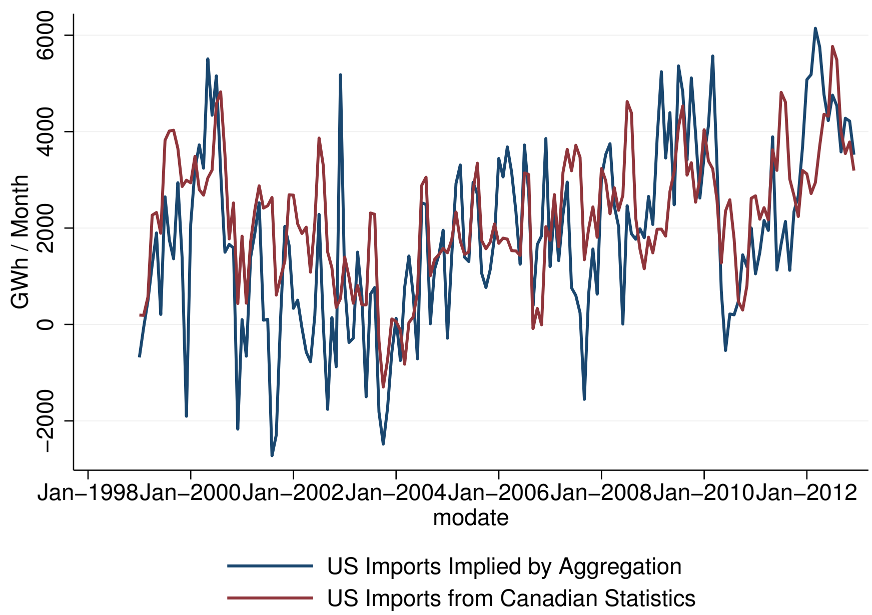

to have higher heat rates, which unaccounted for, would show up as out-of-merit generation. This then becomes a question of how much of the cost of fuel should be attributed to electricity versus steam production. The approach I use is to estimate what the unit's heat rate would be if not for cogeneration based on its vintage, capacity, etc. Fuel usage in excess of this estimated heat rate is attributed to steam generation, and not counted as a cost of running the unit for electricity production. The share of generation sent to the grid by cogeneration and industrial units is reported in EIA-906/920/923.

Unit capacities are reported in EIA-860 (as well as EIA-767), but to ensure out of merit costs are always positive, I use the maximum net generation observed from CEMS units in a year to measure the capacity of the unit.[30] For nuclear units, I calculate capacity as described above: the ratio of monthly output to the share of capacity utilized, as reported by the NRC. Because energy inputs are a fundamental constraint on production from renewables units, I assume that wind and hydro units are never withheld, and are therefore their observed production is their maximum capacity at that moment in time. Differential capacity factors of wind farms, for example, do not contribute to out of merit calculations. For hydro reservoirs this means that dry years, for example, similarly do not show up as economical units sitting idle.

A.6 Fuel Prices

Coal Prices

As in Cicala (2015), this paper uses detailed data on coal deliveries to power plants from the Energy Information Administration (Forms EIA-423, "Monthly Report of Cost and Quality of Fuels for Electric Plants," and EIA-923, "Power Plant Operations Report") and Federal Energy Regulatory Commission (Form FERC-423, "Monthly Report of Cost and Quality of Fuels for Electric Plants"). These are shipment-level data, reported monthly for nearly all of the coal burned for the production of electricity in the United States (all facilities with a combined capacity greater than 50MW are required to report). The reader is referred to the Online Appendices of that paper for more details.

For this paper, the extensive use of bilateral contracts for coal procurement is potentially problematic: the merit order is determined by spot prices, not contract prices. This is because it is the opportunity cost of coal that determines its value when allocating production to plants. On the other hand, minimum purchase volume provisions and resale frictions, can make the opportunity cost of coal essentially free–the high contract prices is a sunk cost and the plant is stuck with the coal.

This is a conceptually important distinction, though in practice the main results of the paper are largely invariant to using the observed contract prices instead of estimated spot prices.

The approach I use to estimate spot prices is to separate the delivered price of coal delivered to plant i in region d and month m from mine county origin o into mine-mouth and shipping costs using hedonic regressions that include the characteristics of the coal and a third order polynomial in distance shipped:

$$c_{oim} = X_{oim}\left[\chi\left(spot_{oim}\right)\beta_m^s + \chi\left(contract_{oim}\right)\beta_m^c\right] + f_{rm}(distance_{oim}) + \varepsilon_{iom}$$

[30] This effectively defines the merit order frontier as the lowest possible cost of production, inclusive of the ability to produce slightly higher than the rated limit when necessary, or less than the rated limit due to constraints not considered in nameplate rating reporting. Results are substantially unchanged when using EIA-860 nameplate capacity instead of observed capacities for CEMS units.

I estimate separate coefficients for the price of each characteristic based on whether it was delivered under a long-term or spot contract. I then apply the spot characteristic prices to all deliveries to predict what prices would have been under a spot contract delivery.

Oil Prices

Oil-burning units also report deliveries in forms EIA-423/923. I estimate spot prices at the state-month level separately by fuel type (diesel, residual fuel oil, etc.), or wider geographic if state-month deliveries are not reported.

Natural Gas Prices

Spot prices for natural gas deliveries are based on daily prices from 65 major trading hubs across the country with consistent prices series from 1999-2012. These data come from the Platts, Bloomberg, and NGI. I use plant coordinates and natural gas pipeline network shapefiles from EIA to locate plants along the pipeline network. I then use network analysis software to connect each plant to the nearest pricing hub. These hubs are not uniformly distributed across the US, as illustrated in Figure A.3, which displays the connections between power plants and hubs. The southeast generally lacks pricing points in particular. Missing daily prices for weekends and holidays are based on carrying the most recent trading day forward.

Figure A.3: Gas Hub - Power Plant Links

B Supplementary Results: For Online Publication

B.1 Testing Potential Confounders

The validity of the research design used in the paper requires assumptions about costs, capacities, and operations. This appendix section tests a number of these assumptions in the context of correlated timing with the introduction of market dispatch.

If generating capacities or fuel prices are themselves directly affected by market dispatch in the short run, controlling for the merit order cost of meeting load, $C^*_{pt}(L_{pt})$ is endogenous. Figure B.1A presents the event study specification used for the main results with the log of PCA generating capacity as the outcome variable. While there is a small uptick in capacity the year before markets open, all magnitudes are small and there is nothing persistent that mirrors the pattern of outcomes in the main results. Figure B.2 evaluates fuel prices for coal and gas as the outcome variables, and similarly does not find a pattern that suggests market dispatch is directly affecting these variables.

As described in the paper, the adoption of market mechanisms to determine production was a separate development from state-led restructuring that allowed generators to become the residual claimants of operational revenues. A natural concern is that these non-utility generators have a greater incentive to operate when economical, and that their expansion was correlated with the adoption of market dispatch. To gauge the extent to which this other policy change might bias estimates of market dispatch's impact, Figure B.1B presents an event study-type figure analogous to the main results, but with the logarithm of non-utility generation capacity as the dependent variable. The figure shows that there has indeed been differential growth of non-utility capacity in areas that adopted market dispatch. However, the trend through the onset of treatment is smooth and continues at the same pace throughout the four-year period relative to market adoption. The divestiture and/or installation of new non-utility generation does not line up with the onset of treatment.

An additional concern is that a PCA operating closer to the idealized merit order as I have defined it might actually lead to more costly operations, as it fails to economize on intertemporal costs–such as ramping and start-up costs. Figure B.3 evaluates the extent to which the cost reductions I find in the body of the paper might be offset with larger intertemporal costs. I define ramping as the absolute value of a generator's output from hour-to-hour (looking at ramping up versus ramping down separately does not affect the conclusion). Cold starts are quantified as the nameplate capacity of a unit that goes from zero to positive output in a given hour. Because of the large number of hours with zero cold starts, I add one to the outcome before taking logs here. While market dispatch appears to have had no discernible impact on cold starts, there is a modest uptick in ramping during the first 6 months of market operation. This is consistent with generators learning that the maintenance costs of ramping generators are not covered by running a unit every moment the wholesale price exceeds the marginal fuel cost, akin to highway versus city driving for a car. In any case, this is a transitory change that does not mirror the cost reductions in the main outcomes.

Figure B.1: Treatment Effects by Months to Market: Installed Capacity

(A) *Log*(Total Capacity)

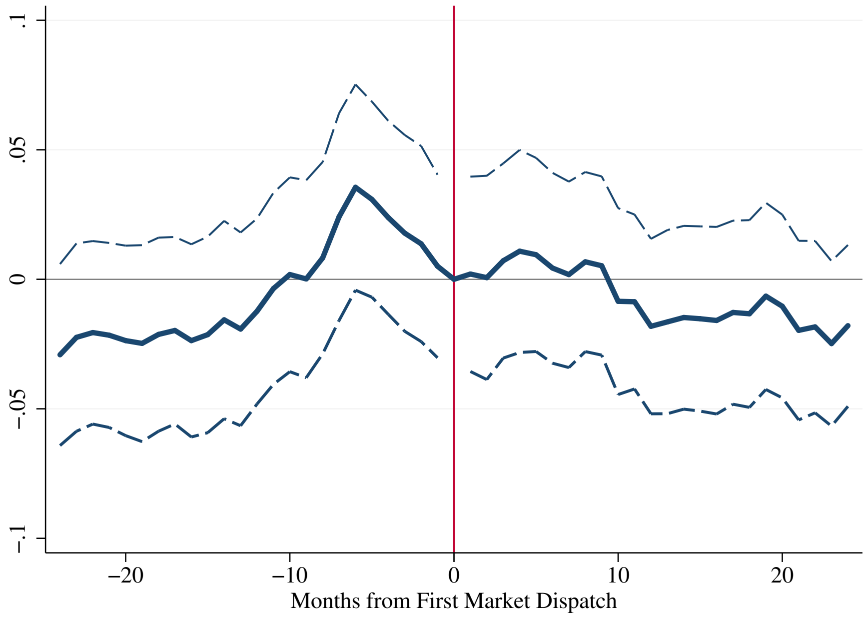

(B) *Log*(Non-Utility Capacity)

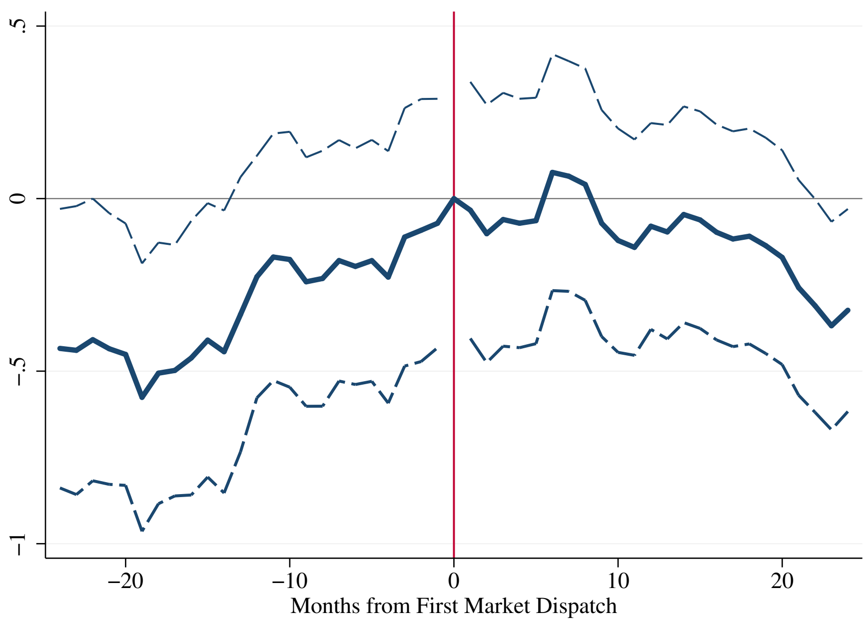

Note: These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (4) of Table 3, where observations are weighted by mean PCA load in 1999. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.2: Treatment Effects by Months to Market: Fuel Prices

(A) *Log*(Coal Price per MMBTU)

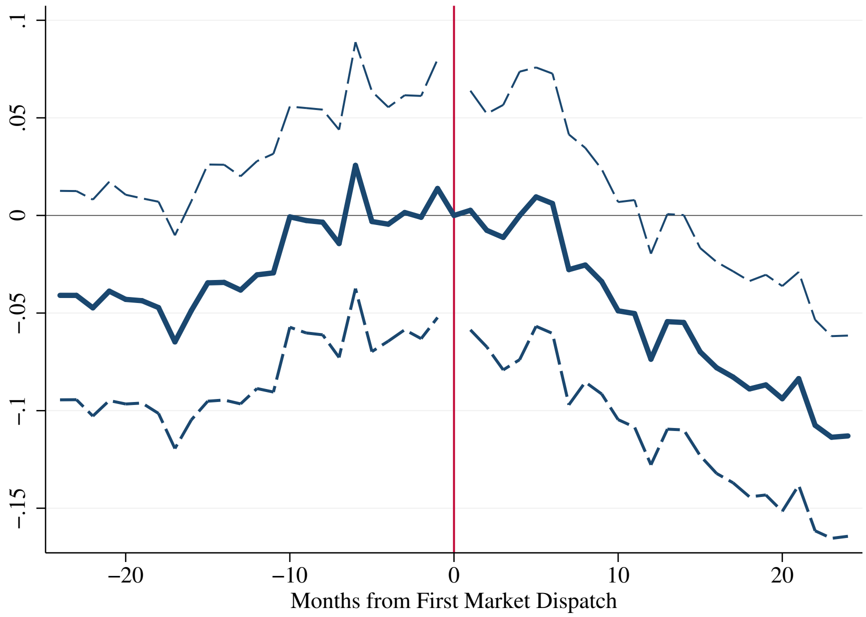

(B) *Log*(Gas Price per MMBTU)

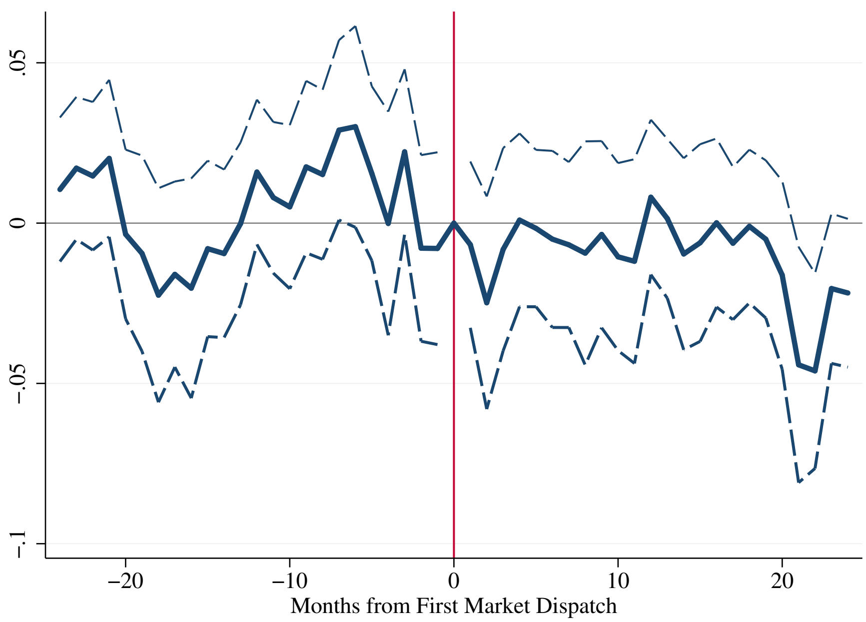

Note: These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch, PCA-specific controls for load, date-hour-region and PCA fixed effects. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.3: Treatment Effects by Months to Market: Cold Starts and Ramping

(A) *Log*(Cold Start MW)

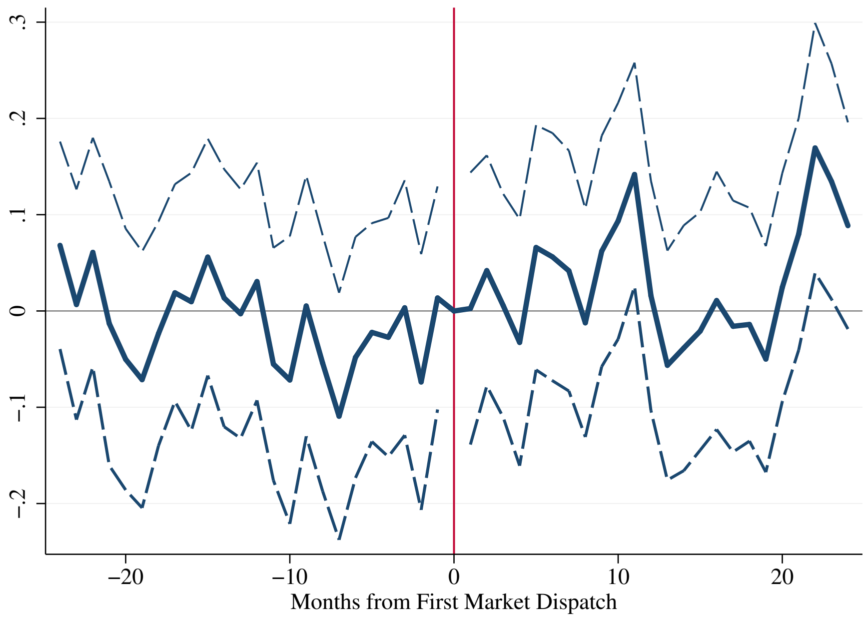

(B) *Log*(Ramping MW)

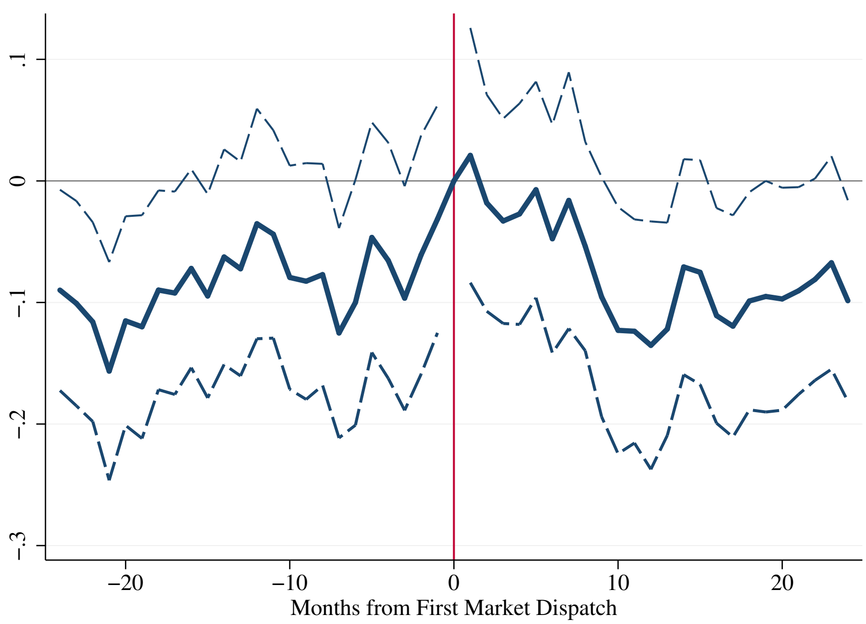

Note: These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (4) of Table 3, where observations are weighted by mean PCA load in 1999. Change in $Log(Load_{pt})$ from the prior hour is also included as a control. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

B.2 Sensitivity to Alternative Specifications

This section presents alternative specifications of the main results, and sensitivity to alternative assumptions regarding generation costs. Tables B.1 and B.2 begin with coarse fixed effects (year and PCA), and progressively add finer controls. It first adds PCA-Month of Year to account for PCA-specific, time invariant seasonal fluctuations. It then adds daily instead of annual fixed effects, and finally region-specific daily fixed effects. Column (1) of Tables 3 and 2 are one level finer still, moving from daily-region to date-hour-region.

Tables B.5 and B.6 replicate Tables 3 and 2 from the body, but use unweighted OLS. The results are qualitatively similar, but create noise in the event-study figures, as small areas with only a few generators can only have relatively large proportional changes.

Figures B.4 and B.5 present the event study analogs of the Column (1) specifications of Tables 3 and 2: there are no controls for load or the merit order cost of meeting load. The figures are qualitatively similar to their counterparts in the body of the text, but there are trends that the more comprehensive specifications help remove. It is clear that the increase in estimated trade volumes from Columns (1) and (4) in Table B.4 is due to the removal of a downward-sloping pre-trend.

The remaining figures present event-study results under alternative assumptions about fuel pricing. Figure B.7 estimates fuel prices without the restricted-access portions of the EIA-data on fuel deliveries. Figure B.8 uses the restricted-access data, but calculates the merit order using the prices of fuel delivered under contract instead of estimating spot market prices as in the body of the paper. For these figures I continue to use the daily hub prices for natural gas. Figure B.9 presents results under an alternative pricing assumption for the value of hydropower. In the paper, hydropower from reservoirs is valued at the marginal cost according to the merit order. Here, I instead suppose that all hydropower were zero marginal cost. For each of these alternative assumptions or data sources the overall results are largely unaffected.

Table B.1: Sensitivity to Fixed Effects: Quantities

A. *Log*(Trade Volume)

	(1)	(2)	(3)	(4)
Market Dispatch	0.174***	0.177***	0.180***	0.168***
	(0.029)	(0.028)	(0.028)	(0.032)
Time FE	Year	Year	Date	Date-Region
PCA FE	PCA	PCA-Month of Year	PCA-Month of Year	PCA-Month of Year
Clusters	16464	16464	16464	16464
PCAs	98	98	98	98
R^2	0.446	0.463	0.468	0.484
Obs.	12004719	12004719	12004719	12004719

B. *Log*(MWh Out of Merit)

	(1)	(2)	(3)	(4)
Market Dispatch	–0.090***	–0.087***	–0.086***	–0.072***
	(0.014)	(0.012)	(0.012)	(0.013)
Time FE	Year	Year	Date	Date-Region
PCA FE	PCA	PCA-Month of Year	PCA-Month of Year	PCA-Month of Year
Clusters	16440	16440	16440	16440
PCAs	98	98	98	98
R^2	0.866	0.878	0.881	0.886
Obs.	11625543	11625543	11625543	11625543

Note: Observations weighted by mean PCA load in 1999. All specifications include annual event-time dummies beyond 2 years from treatment. Standard errors clustered by PCA-Month in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table B.2: Sensitivity to Fixed Effects: Allocative Efficiency

A. *Log*(Gains from Trade)

	(1)	(2)	(3)	(4)
Market Dispatch	0.486***	0.483***	0.481***	0.436***
	(0.070)	(0.066)	(0.066)	(0.070)
Time FE	Year	Year	Date	Date-Region
PCA FE	PCA	PCA-Month of Year	PCA-Month of Year	PCA-Month of Year
Clusters	16412	16412	16412	16412
PCAs	98	98	98	98
R^2	0.338	0.374	0.389	0.417
Obs.	8477369	8477369	8477369	8477369

B. *Log*(Out of Merit Costs)

	(1)	(2)	(3)	(4)
Market Dispatch	–0.081***	–0.084***	–0.093***	–0.130***
	(0.030)	(0.028)	(0.024)	(0.028)
Time FE	Year	Year	Date	Date-Region
PCA FE	PCA	PCA-Month of Year	PCA-Month of Year	PCA-Month of Year
Clusters	16437	16437	16437	16437
PCAs	98	98	98	98
R^2	0.809	0.822	0.838	0.852
Obs.	11618837	11618837	11618837	11618837

Note: Observations weighted by mean PCA load in 1999. All specifications include annual event-time dummies beyond 2 years from treatment. Standard errors clustered by PCA-Month in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table B.3: Inverse Hyperbolic Sine Transformed Outcomes: Quantities

A. *asinh*(Trade Volume)	(1)	(2)	(3)	(4)
Market Dispatch	0.168***	0.149***	0.212***	0.227***
	(0.033)	(0.033)	(0.031)	(0.032)
1^{st} Neighbor				0.044
Market Dispatch				(0.036)
2^{nd} Neighbor				0.010
Market Dispatch				(0.033)
asinh(Load)		Yes	Yes	Yes
asinh(Load Merit Cost)			Yes	Yes
Clusters	16464	16464	16464	16464
PCAs	98	98	98	98
R^2	0.532	0.563	0.578	0.579
Obs.	12028128	12028128	12028128	12028128

B. *asinh*(MWh Out of Merit)	(1)	(2)	(3)	(4)
Market Dispatch	–0.066***	–0.068***	–0.052***	–0.053***
	(0.014)	(0.013)	(0.014)	(0.014)
1^{st} Neighbor				–0.030*
Market Dispatch				(0.016)
2^{nd} Neighbor				0.025*
Market Dispatch				(0.014)
asinh(Load)		Yes	Yes	Yes
asinh(Load Merit Cost)			Yes	Yes
Clusters	16464	16464	16464	16464
PCAs	98	98	98	98
R^2	0.880	0.887	0.891	0.892
Obs.	12028128	12028128	12028128	12028128

Note: All specifications include PCA-Month of Year and Region-Date-Hour fixed effects. Controls for the logarithm of load L_{pt} and its merit order cost $C^*_{pt}(L_{pt})$ are estimated with separate slopes by PCA-Month of Year. Standard errors clustered by PCA-Month in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table B.4: Inverse Hyperbolic Sine Transformed Outcomes: Allocative Efficiency

A. *asinh*(Gains from Trade)

	(1)	(2)	(3)	(4)
Market Dispatch	0.585***	0.583***	0.616***	0.617***
	(0.100)	(0.102)	(0.095)	(0.097)
1^{st} Neighbor				0.221*
Market Dispatch				(0.118)
2^{nd} Neighbor				–0.097
Market Dispatch				(0.109)
asinh(Load)		Yes	Yes	Yes
asinh(Load Merit Cost)			Yes	Yes
Clusters	16464	16464	16464	16464
PCAs	98	98	98	98
R^2	0.403	0.451	0.478	0.479
Obs.	12028128	12028128	12028128	12028128

B. *asinh*(Out of Merit Costs)

	(1)	(2)	(3)	(4)
Market Dispatch	–0.117***	–0.104***	–0.148***	–0.171***
	(0.029)	(0.028)	(0.025)	(0.026)
1^{st} Neighbor				–0.013
Market Dispatch				(0.032)
2^{nd} Neighbor				–0.010
Market Dispatch				(0.025)
asinh(Load		Yes	Yes	Yes
asinh(Load Merit Cost			Yes	Yes
Clusters	16464	16464	16464	16464
PCAs	98	98	98	98
R^2	0.848	0.857	0.865	0.866
Obs.	12028128	12028128	12028128	12028128

Note: All specifications include PCA-Month of Year and Region-Date-Hour fixed effects. Controls for the logarithm of load L_{pt} and its merit order cost $C^*_{pt}(L_{pt})$ are estimated with separate slopes by PCA-Month of Year. Standard errors clustered by PCA-Month in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table B.5: Unweighted Main Results: Quantities

A. *Log*(Trade Volume)

	(1)	(2)	(3)	(4)
Market Dispatch	0.136***	0.136***	0.188***	0.207***
	(0.020)	(0.020)	(0.020)	(0.021)
1^{st} Neighbor				0.096***
Market Dispatch				(0.028)
2^{nd} Neighbor				0.089***
Market Dispatch				(0.023)
$Log(L_{pt})$		Yes	Yes	Yes
$Log(C^*_{pt}(L_{pt}))$			Yes	Yes
Clusters	16464	16464	16464	16464
PCAs	98	98	98	98
R^2	0.577	0.610	0.626	0.626
Obs.	12004719	12004719	12004719	12004719

B. *Log*(MWh Out of Merit)

	(1)	(2)	(3)	(4)
Market Dispatch	–0.121***	–0.130***	–0.128***	–0.112***
	(0.017)	(0.017)	(0.017)	(0.018)
1^{st} Neighbor				0.013
Market Dispatch				(0.019)
2^{nd} Neighbor				0.035**
Market Dispatch				(0.016)
$Log(L_{pt})$		Yes	Yes	Yes
$Log(C^*_{pt}(L_{pt}))$			Yes	Yes
Clusters	16440	16440	16440	16440
PCAs	98	98	98	98
R^2	0.842	0.849	0.857	0.857
Obs.	11625543	11625543	11625543	11625543

Note: All specifications include PCA-Month of Year and Region-Date-Hour fixed effects, as well as annual event-time dummies beyond two years from treatment. Controls for the logarithm of load L_{pt} and its merit order cost $C^*_{pt}(L_{pt})$ are estimated with separate slopes by PCA-Month of Year. Standard errors clustered by PCA-Month in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Table B.6: Unweighted Main Results: Allocative Efficiency

A. *Log*(Gains from Trade)	(1)	(2)	(3)	(4)
Market Dispatch	0.333***	0.362***	0.397***	0.432***
	(0.043)	(0.042)	(0.042)	(0.045)
1^{st} Neighbor				0.271***
Market Dispatch				(0.059)
2^{nd} Neighbor				0.023
Market Dispatch				(0.049)
$Log(L_{pt})$		Yes	Yes	Yes
$Log(C^*_{pt}(L_{pt}))$			Yes	Yes
Clusters	16412	16412	16412	16412
PCAs	98	98	98	98
R^2	0.514	0.572	0.600	0.600
Obs.	8475828	8475828	8475828	8475828

B. *Log*(Out of Merit Costs)	(1)	(2)	(3)	(4)
Market Dispatch	–0.250***	–0.237***	–0.246***	–0.221***
	(0.031)	(0.030)	(0.029)	(0.030)
1^{st} Neighbor				0.201***
Market Dispatch				(0.036)
2^{nd} Neighbor				–0.003
Market Dispatch				(0.030)
Log(Load		Yes	Yes	Yes
Log(Load Merit Cost			Yes	Yes
Clusters	16437	16437	16437	16437
PCAs	98	98	98	98
R^2	0.798	0.810	0.823	0.823
Obs.	11618837	11618837	11618837	11618837

Note: All specifications include PCA-Month of Year and Region-Date-Hour fixed effects, as well as annual event-time dummies beyond two years from treatment. Controls for the logarithm of load L_{pt} and its merit order cost $C^*_{pt}(L_{pt})$ are estimated with separate slopes by PCA-Month of Year. Standard errors clustered by PCA-Month in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Figure B.4: Main Quantity Results: No Load or Cost Controls

(A) *Log*(Trade Volume)

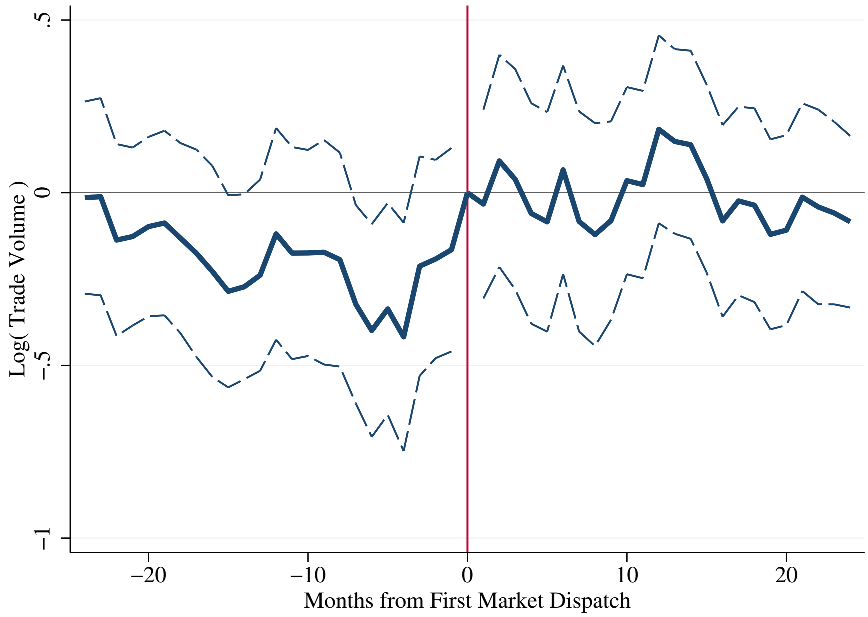

(B) *Log*(MWh Out of Merit)

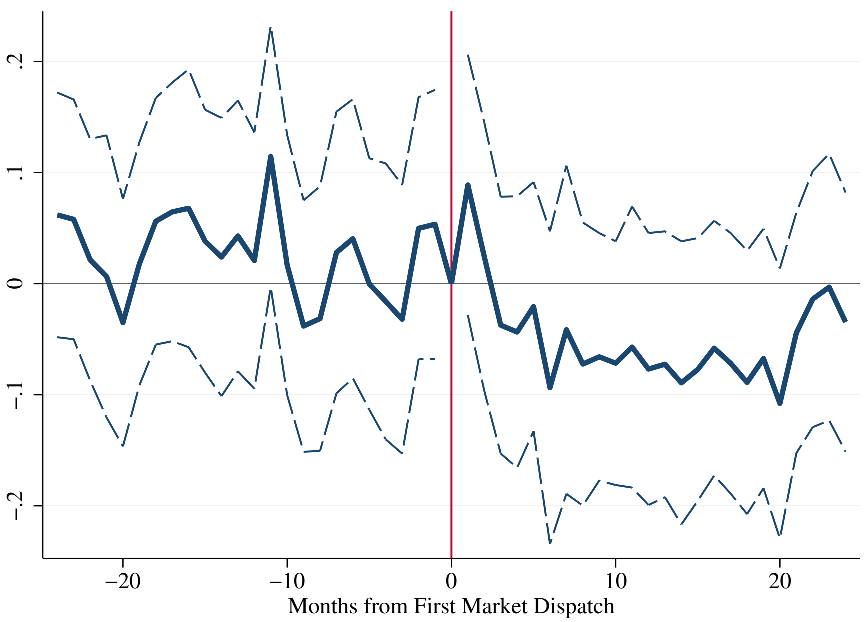

Note: These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (1) of Table 3, where observations are weighted by mean PCA load in 1999. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.5: Main Allocative Efficiency Results: No Load or Cost Controls

(A) *Log*(Gains from Trade)

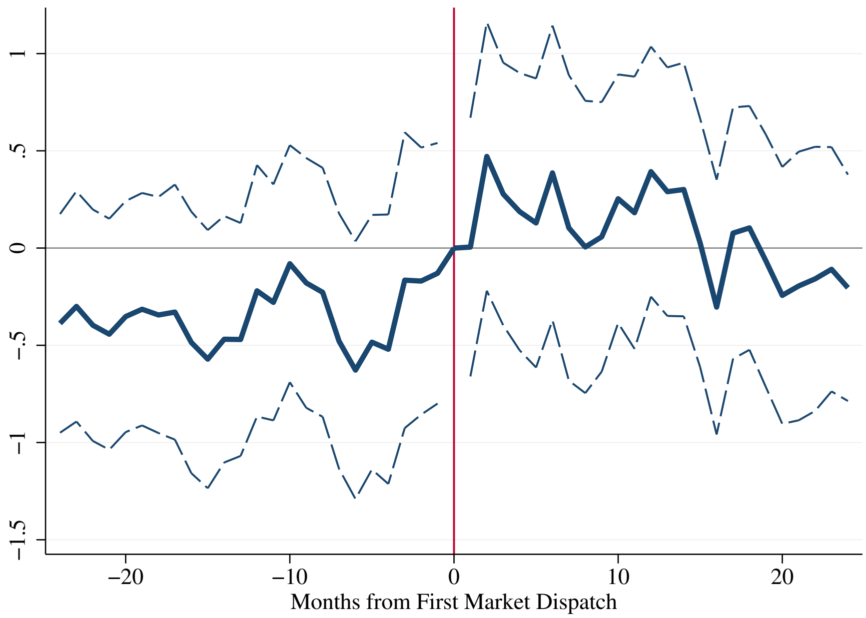

(B) *Log*(Out of Merit Costs)

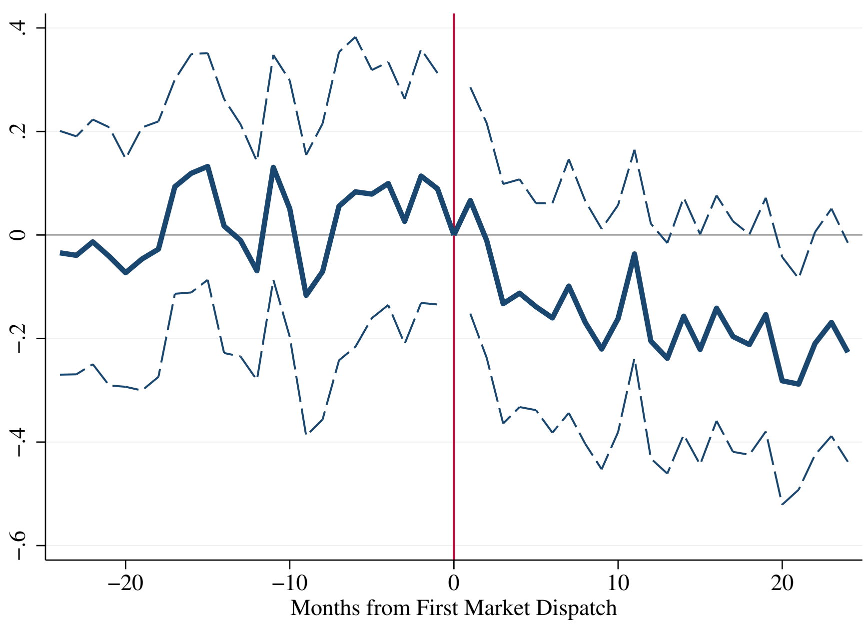

Note: These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (1) of Table 2, where observations are weighted by mean PCA load in 1999 but there are no controls for load or the merit order cost of meeting load. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.6: Main Allocative Efficiency Results: Inverse Hyperbolic Sine Transformation

(A) *asinh*(Gains from Trade)

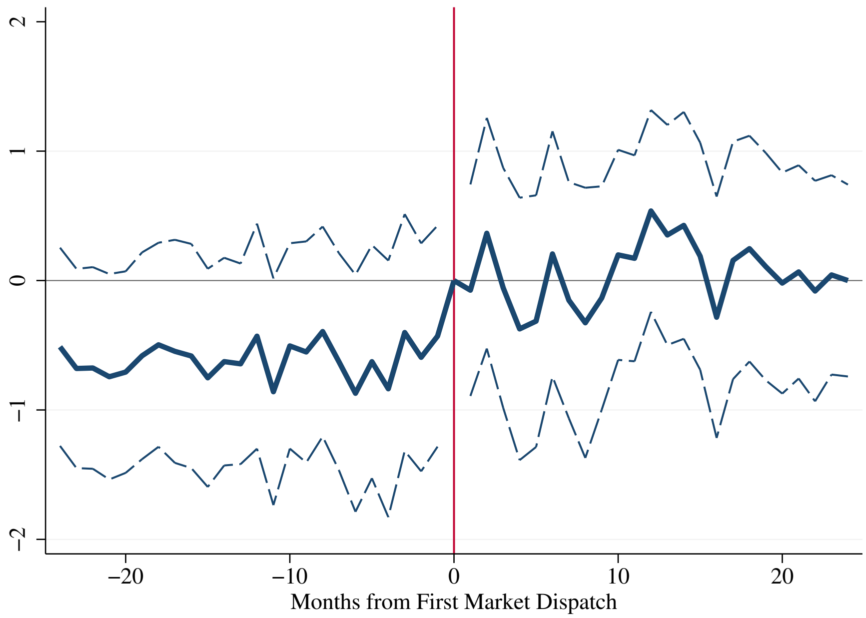

(B) *asinh*(Out of Merit Costs)

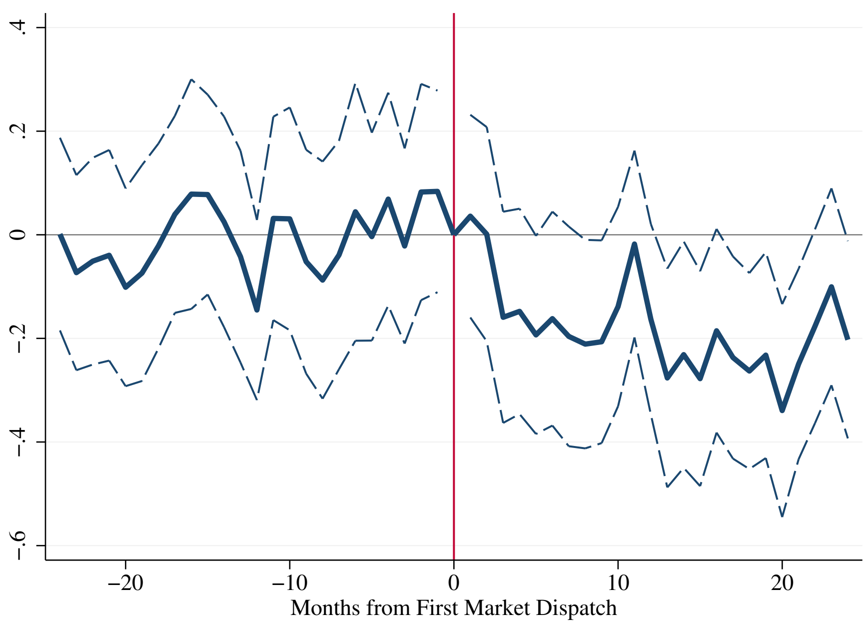

Note: These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (1) of Table 2, where observations are weighted by mean PCA load in 1999 but there are no controls for load or the merit order cost of meeting load. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.7: Main Allocative Efficiency Results: Publicly-Available Fuel Prices

(A) *Log*(Gains from Trade)

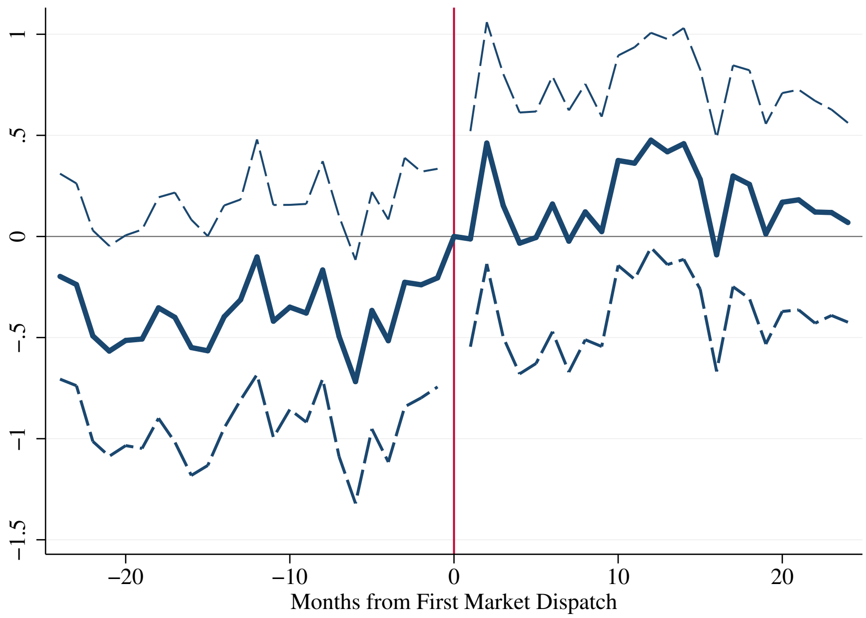

(B) *Log*(Out of Merit Costs)

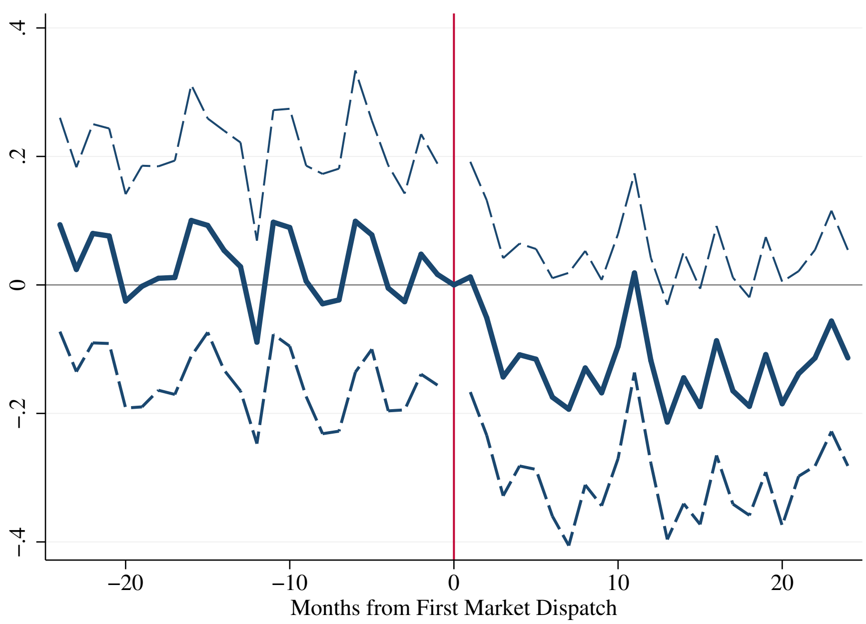

Note: These figures are based on estimated fuel prices using publicly-available data from EIA for coal, oil, and biomass. Logged outcomes are regressed on a set of indicator variables for each month until (after) the transition to market dispatch, date-hour-region and PCA-month of year fixed effects. This corresponds with column (4) of Table 2, where observations are weighted by mean PCA load in 1999. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.8: Main Allocative Efficiency Results: No Spot Market Price Estimation

(A) *Log*(Gains from Trade)

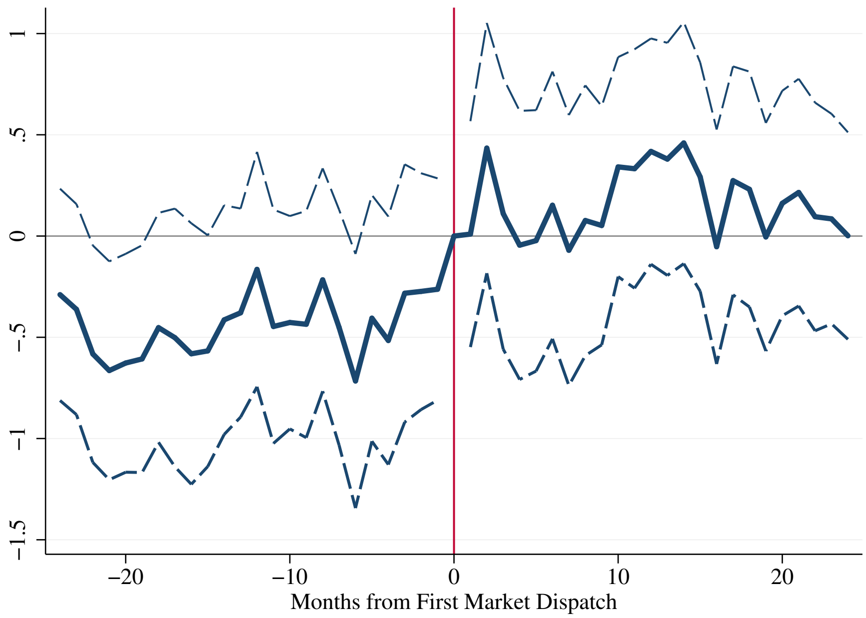

(B) *Log*(Out of Merit Costs)

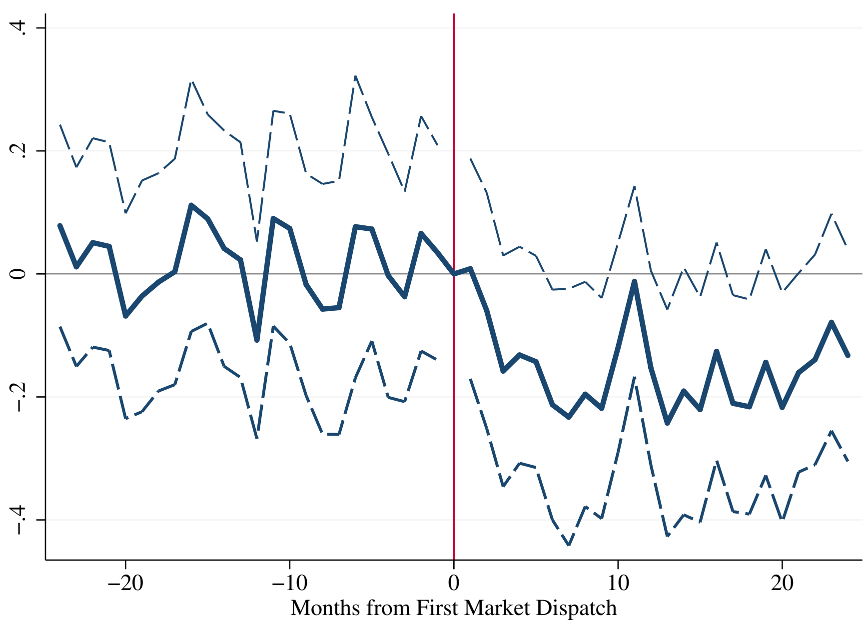

Note: Fuel prices used in the underlying data for these figures do not adjust for contract versus spot purchases of coal. These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (4) of Table 2, where observations are weighted by mean PCA load in 1999. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.

Figure B.9: Main Allocative Efficiency Results: Zero Cost Reservoir Hydropower

(A) *Log*(Gains from Trade)

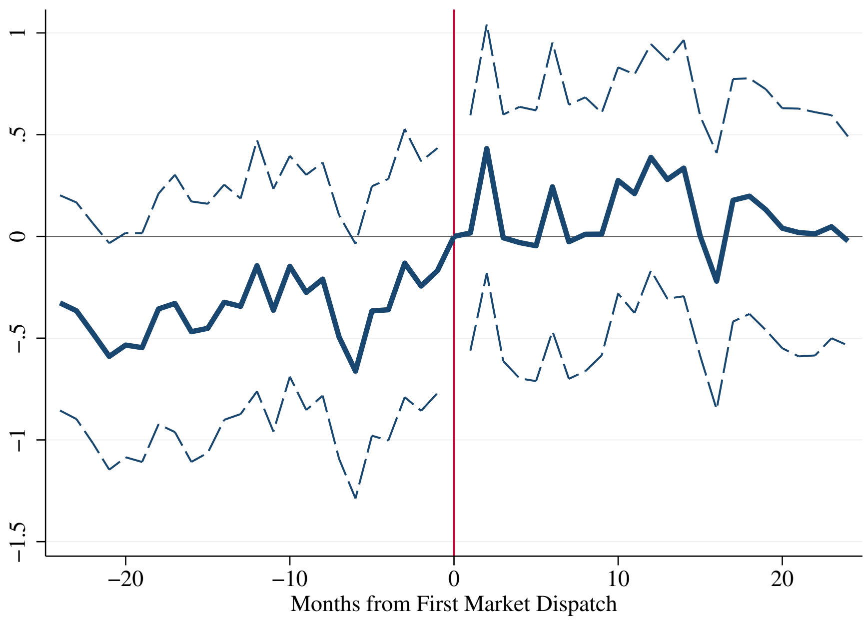

(B) *Log*(Out of Merit Costs)

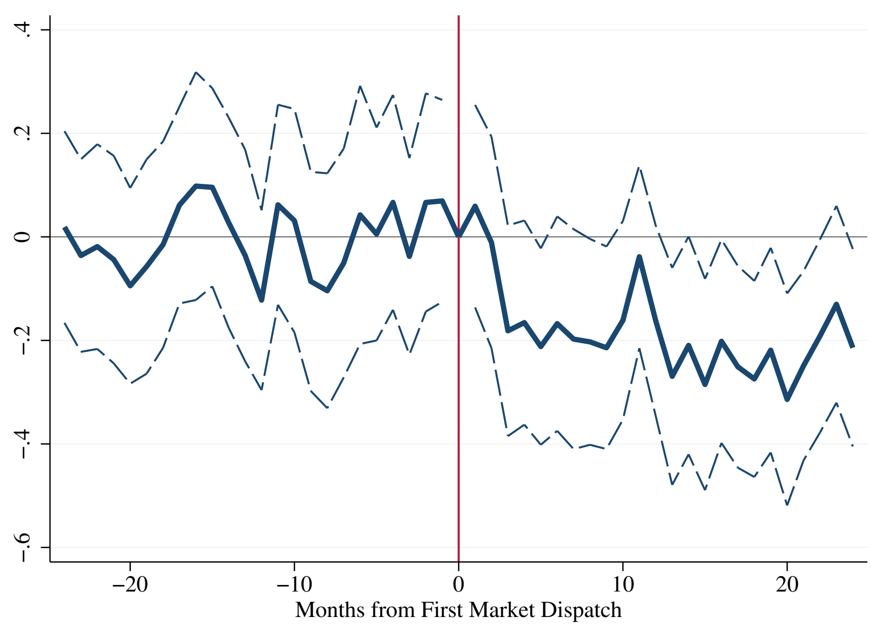

Note: In constructing outcomes for these figures all hydropower is assumed to be delivered at zero marginal cost, instead of the marginal cost of generation according to the merit order, as in the body of the paper. These figures are based on regressing logged outcomes on a set of indicator variables for each month until (after) the transition to market dispatch. The specification corresponds with column (4) of Table 2, where observations are weighted by mean PCA load in 1999. The month prior to treatment is normalized to zero. 95% Confidence intervals in dashed lines are based on clustering at the PCA-month level.