

Machine Learning from Schools about Energy Efficiency

Fiona Burlig
University of Chicago

Christopher Knittel
MIT

David Rapson
UC Davis

Mar Reguant
Northwestern University

Catherine Wolfram*
UC Berkeley

September 29, 2017

Appendix: For online publication

*Burlig: Department of Economics and Energy Policy Institute, University of Chicago, burlig@uchicago.edu. Knittel: Sloan School of Management and Center for Energy and Environmental Policy Research, MIT and NBER, knittel@mit.edu. Rapson: Department of Economics, UC Davis, dsrapson@ucdavis.edu. Reguant: Department of Economics, Northwestern University, CEPR and NBER, mar.reguant@northwestern.edu. Wolfram: Haas School of Business and Energy Institute at Haas, UC Berkeley and NBER, cwolfram@berkeley.edu.

Table A1: Panel fixed effects results by intervention type (levels)

	(1)	(2)	(3)	(4)	(5)
Any intervention	-1.829 (0.392)	-1.908 (0.391)	-2.224 (0.392)	-0.750 (0.364)	-0.894 (0.363)
Observations	19,113,926	19,113,926	19,113,066	19,113,066	19,113,926
HVAC interventions	-1.800 (0.477)	-1.873 (0.479)	-2.296 (0.476)	-1.123 (0.451)	-1.199 (0.459)
Observations	14,841,496	14,841,496	14,840,772	14,840,772	14,841,496
Lighting interventions	-2.594 (0.646)	-2.677 (0.640)	-2.973 (0.639)	-1.459 (0.590)	-1.564 (0.584)
Observations	12,850,033	12,850,032	12,849,420	12,849,420	12,850,032
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	Yes	No
Month of Sample Ctrl.	No	No	No	Yes	No
Month of Sample FE	No	No	No	No	Yes

Notes: This table reports results from estimating Equation (3.1), with hourly energy consumption in levels (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. A coefficient of -0.894, for example, means that going from 0 to 100% of a school’s expected savings delivers energy savings of approximately 0.894 kWh per hour on average. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period.

Table A2: Panel fixed effects results by intervention type (trimming)

	(1)	(2)	(3)
Any intervention	-0.021 (0.009)	-0.016 (0.008)	-0.012 (0.008)
Observations	19,503,680	19,113,434	18,723,432
HVAC interventions	-0.026 (0.011)	-0.018 (0.010)	-0.014 (0.010)
Observations	15,144,112	14,840,990	14,538,259
Lighting interventions	-0.031 (0.012)	-0.025 (0.011)	-0.023 (0.010)
Observations	13,111,944	12,849,541	12,587,456
Trimming			
Dependent variable (1, 99)		X	
Dependent variable (2, 98)			X

Notes: This table reports results from estimating Equation (3.1), with hourly energy consumption in logs (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. A coefficient of -0.045, for example, means that going from 0 to 100% of a school’s expected savings delivers energy savings of approximately 4.5 percent on average. Standard errors, clustered at the school level, are in parentheses. This table presents three types of trimming of the dependent variable: Column (1) does not trim at all; Column (2) trims the sample to exclude observations below the 1st or above the 99th percentile, as in the main text; and Column (3) trims the sample to exclude observations below the 2nd or above the 98th percentile. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period.

Table A3: Panel fixed effects results by hour-block (Any intervention)

	(1)	(2)	(3)	(4)	(5)
Aggregate	-0.043 (0.007)	-0.043 (0.007)	-0.047 (0.007)	-0.014 (0.008)	-0.016 (0.008)
Midn. to 3 AM x Treat	-0.061 (0.009)	-0.051 (0.008)	-0.054 (0.008)	-0.021 (0.009)	-0.024 (0.009)
3 AM to 6 AM x Treat	-0.061 (0.009)	-0.048 (0.008)	-0.052 (0.008)	-0.019 (0.009)	-0.021 (0.009)
6 AM to 9 AM x Treat	-0.020 (0.008)	-0.024 (0.007)	-0.029 (0.006)	0.004 (0.008)	0.003 (0.009)
9 AM to Noon x Treat	-0.034 (0.008)	-0.048 (0.006)	-0.051 (0.006)	-0.018 (0.008)	-0.021 (0.008)
Noon to 3 PM x Treat	-0.026 (0.009)	-0.044 (0.007)	-0.047 (0.007)	-0.014 (0.009)	-0.017 (0.009)
3 PM to 6 PM x Treat	-0.042 (0.009)	-0.040 (0.008)	-0.047 (0.008)	-0.014 (0.009)	-0.013 (0.010)
6 PM to 9 PM x Treat	-0.051 (0.009)	-0.043 (0.008)	-0.049 (0.008)	-0.016 (0.009)	-0.016 (0.009)
9 PM to Midn. x Treat	-0.055 (0.009)	-0.051 (0.008)	-0.054 (0.008)	-0.021 (0.009)	-0.024 (0.009)
Observations	19,113,321	19,113,321	19,112,451	19,112,451	19,113,321
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: This table reports hour-block-specific results from estimating Equation (3.1), with the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. A coefficient of -0.045, for example, means that going from 0 to 100% of a school’s expected savings delivers energy savings of approximately 4.5% on average. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. The results from this table are presented graphically in Panel A of Figure 2 in the main text.

Table A4: Panel fixed effects results by hour-block (HVAC interventions)

	(1)	(2)	(3)	(4)	(5)
Aggregate	-0.043 (0.009)	-0.042 (0.009)	-0.048 (0.009)	-0.017 (0.010)	-0.018 (0.010)
Midn. to 3 AM x Treat	-0.068 (0.012)	-0.053 (0.010)	-0.058 (0.010)	-0.027 (0.011)	-0.029 (0.011)
3 AM to 6 AM x Treat	-0.069 (0.011)	-0.051 (0.010)	-0.055 (0.010)	-0.024 (0.011)	-0.026 (0.011)
6 AM to 9 AM x Treat	-0.016 (0.010)	-0.016 (0.008)	-0.024 (0.008)	0.006 (0.010)	0.008 (0.010)
9 AM to Noon x Treat	-0.033 (0.010)	-0.043 (0.008)	-0.048 (0.007)	-0.017 (0.009)	-0.019 (0.010)
Noon to 3 PM x Treat	-0.022 (0.010)	-0.040 (0.008)	-0.043 (0.008)	-0.013 (0.010)	-0.015 (0.010)
3 PM to 6 PM x Treat	-0.029 (0.011)	-0.037 (0.010)	-0.045 (0.010)	-0.014 (0.011)	-0.012 (0.012)
6 PM to 9 PM x Treat	-0.051 (0.011)	-0.048 (0.010)	-0.054 (0.010)	-0.024 (0.011)	-0.023 (0.011)
9 PM to Midn. x Treat	-0.054 (0.011)	-0.055 (0.009)	-0.059 (0.009)	-0.028 (0.010)	-0.030 (0.011)
Observations	14,840,991	14,840,990	14,840,264	14,840,264	14,840,990
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: This table reports hour-block-specific results from estimating Equation (3.1), with the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. A coefficient of -0.045, for example, means that going from 0 to 100% of a school’s expected savings delivers energy savings of approximately 4.5% on average. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. These regressions include only schools that underwent an HVAC upgrade or never underwent an upgrade during our sample period. The results from this table are presented graphically in Panel B of Figure 2 in the main text.

Table A5: Panel fixed effects results by hour-block (Lighting interventions)

	(1)	(2)	(3)	(4)	(5)
Aggregate	-0.052 (0.009)	-0.051 (0.009)	-0.055 (0.009)	-0.024 (0.011)	-0.025 (0.011)
Midn. to 3 AM x Treat	-0.062 (0.013)	-0.050 (0.011)	-0.052 (0.011)	-0.020 (0.012)	-0.023 (0.012)
3 AM to 6 AM x Treat	-0.063 (0.012)	-0.050 (0.011)	-0.054 (0.011)	-0.022 (0.012)	-0.023 (0.012)
6 AM to 9 AM x Treat	-0.028 (0.010)	-0.036 (0.008)	-0.043 (0.008)	-0.011 (0.011)	-0.009 (0.011)
9 AM to Noon x Treat	-0.046 (0.011)	-0.065 (0.008)	-0.068 (0.008)	-0.036 (0.011)	-0.039 (0.010)
Noon to 3 PM x Treat	-0.032 (0.011)	-0.059 (0.009)	-0.061 (0.009)	-0.029 (0.011)	-0.033 (0.011)
3 PM to 6 PM x Treat	-0.053 (0.012)	-0.053 (0.011)	-0.058 (0.011)	-0.026 (0.012)	-0.026 (0.013)
6 PM to 9 PM x Treat	-0.063 (0.011)	-0.051 (0.010)	-0.056 (0.011)	-0.024 (0.012)	-0.024 (0.012)
9 PM to Midn. x Treat	-0.069 (0.012)	-0.055 (0.011)	-0.056 (0.011)	-0.024 (0.012)	-0.028 (0.012)
Observations	12,849,512	12,849,511	12,848,896	12,848,896	12,849,511
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: This table reports hour-block-specific results from estimating Equation (3.1), with the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. A coefficient of -0.045, for example, means that going from 0 to 100% of a school’s expected savings delivers energy savings of approximately 4.5% on average. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. These regressions include only schools that underwent a lighting upgrade or never underwent an upgrade during our sample period. The results from this table are presented graphically in Panel C of Figure 2 in the main text.

Table A6: Matching results – HVAC interventions

	(1)	(2)	(3)	(4)	(5)
Any district	-0.010 (0.026)	-0.010 (0.026)	-0.005 (0.027)	-0.007 (0.027)	-0.011 (0.026)
Same district	0.008 (0.022)	0.008 (0.022)	0.013 (0.024)	0.016 (0.021)	0.014 (0.021)
Opposite district	-0.067 (0.015)	-0.067 (0.015)	-0.068 (0.016)	-0.034 (0.017)	-0.039 (0.015)
Observations	2,379,037	2,379,033	2,378,466	2,378,466	2,379,033
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	Yes	No
Month of Sample Ctrl.	No	No	No	Yes	No
Month of Sample FE	No	No	No	No	Yes

Notes: This table reports results from estimating Equation (3.1), with the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. As above, the independent variable is defined as the fraction of total expected savings that have been installed by time t . The untreated group in these regressions is chosen via nearest-neighbor matching. In particular, we match one untreated school to each treated school. Each row in the table employs a different restriction on which schools are allowed to be matched to any given treatment school. “Any district” matches allow any untreated school to be matched to a treatment school; “same district” matches are restricted to untreated schools in the same school district, and “opposite district” matches are restricted to untreated schools from different districts. In each case, the matching variables are the mean, maximum, and standard deviation of electricity consumption in each three-hour block (e.g., 9 AM-Noon) from the pre-treatment period; demographic variables measured at the census block level, including the poverty rate, log of per capita income, school-level variables (enrollment; age of the school; grades taught; an academic performance index; and climate). These estimates are relatively sensitive to which schools are included. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. This table is analogous to Table 3 in the main text, but presents results for HVAC upgrades only.

Table A7: Matching results – Lighting interventions

	(1)	(2)	(3)	(4)	(5)
Any district	-0.060 (0.021)	-0.060 (0.021)	-0.056 (0.022)	-0.025 (0.023)	-0.034 (0.020)
Same district	-0.046 (0.016)	-0.046 (0.016)	-0.045 (0.017)	-0.029 (0.019)	-0.027 (0.017)
Opposite district	-0.046 (0.012)	-0.046 (0.012)	-0.049 (0.013)	-0.005 (0.020)	-0.004 (0.019)
Observations	1,914,567	1,914,563	1,914,147	1,914,147	1,914,563
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	Yes	No
Month of Sample Ctrl.	No	No	No	Yes	No
Month of Sample FE	No	No	No	No	Yes

Notes: This table reports results from estimating Equation (3.1), with the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. As above, the independent variable is defined as the fraction of total expected savings that have been installed by time t . The untreated group in these regressions is chosen via nearest-neighbor matching. In particular, we match one untreated school to each treated school. Each row in the table employs a different restriction on which schools are allowed to be matched to any given treatment school. “Any district” matches allow any untreated school to be matched to a treatment school; “same district” matches are restricted to untreated schools in the same school district, and “opposite district” matches are restricted to untreated schools from different districts. In each case, the matching variables are the mean, maximum, and standard deviation of electricity consumption in each three-hour block (e.g., 9 AM-Noon) from the pre-treatment period; demographic variables measured at the census block level, including the poverty rate, log of per capita income, school-level variables (enrollment; age of the school; grades taught; an academic performance index; and climate). These estimates are relatively sensitive to which schools are included. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. This table is analogous to Table 3 in the main text, but presents results for lighting upgrades only.

Table A8: Machine learning results by intervention type (levels)

	(1)	(2)	(3)	(4)	(5)
Any intervention	-2.868 (0.368)	-2.835 (0.368)	-2.808 (0.386)	-1.711 (0.348)	-1.878 (0.334)
Observations	19,113,602	19,113,602	19,112,748	19,112,748	19,113,602
HVAC interventions	-2.856 (0.471)	-2.821 (0.469)	-2.847 (0.495)	-1.912 (0.444)	-1.989 (0.423)
Observations	14,841,226	14,841,226	14,840,509	14,840,509	14,841,226
Lighting interventions	-3.407 (0.519)	-3.386 (0.519)	-3.381 (0.541)	-2.191 (0.510)	-2.436 (0.481)
Observations	12,849,702	12,849,702	12,849,090	12,849,090	12,849,702
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	Yes	No
Month of Sample Ctrl.	No	No	No	Yes	No
Month of Sample FE	No	No	No	No	Yes

Notes: This table reports results from estimating Equation (3.1), with the prediction errors in hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period. All regressions include a control for being in the post-training period for the machine learning.

Table A9: Machine learning results by intervention type (trimming)

	(1)	(2)	(3)
Any intervention	-0.022 (0.012)	-0.031 (0.008)	-0.032 (0.007)
Observations	19,503,680	19,113,602	18,723,528
HVAC interventions	-0.019 (0.013)	-0.032 (0.009)	-0.033 (0.008)
Observations	15,144,112	14,841,225	14,538,344
Lighting interventions	-0.039 (0.013)	-0.041 (0.010)	-0.040 (0.009)
Observations	13,111,944	12,849,701	12,587,462
Trimming			
Dependent variable (1, 99)		X	
Dependent variable (2, 98)			X

Notes: This table reports results from estimating Equation (3.1), with the prediction errors of the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. A coefficient of -0.045, for example, means that going from 0 to 100% of a school’s expected savings delivers energy savings of approximately 4.5 percent on average. Standard errors, clustered at the school level, are in parentheses. This table presents three types of trimming of the dependent variable: Column (1) does not trim at all; Column (2) trims the sample to exclude observations below the 1st or above the 99th percentile, as in the main text; and Column (3) trims the sample to exclude observations below the 2nd or above the 98th percentile. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period.

Table A10: Machine learning results by hour-block (Any intervention)

	(1)	(2)	(3)	(4)	(5)
Aggregate	-0.044 (0.007)	-0.045 (0.007)	-0.045 (0.008)	-0.029 (0.008)	-0.031 (0.008)
Midn. to 3 AM x Treat	-0.041 (0.009)	-0.040 (0.008)	-0.038 (0.009)	-0.023 (0.009)	-0.027 (0.009)
3 AM to 6 AM x Treat	-0.032 (0.009)	-0.036 (0.009)	-0.036 (0.009)	-0.020 (0.009)	-0.023 (0.009)
6 AM to 9 AM x Treat	-0.009 (0.008)	-0.019 (0.007)	-0.022 (0.007)	-0.006 (0.008)	-0.005 (0.008)
9 AM to Noon x Treat	-0.051 (0.007)	-0.057 (0.007)	-0.057 (0.007)	-0.041 (0.008)	-0.043 (0.008)
Noon to 3 PM x Treat	-0.057 (0.008)	-0.060 (0.008)	-0.061 (0.008)	-0.045 (0.008)	-0.047 (0.008)
3 PM to 6 PM x Treat	-0.057 (0.009)	-0.054 (0.009)	-0.055 (0.009)	-0.039 (0.009)	-0.040 (0.009)
6 PM to 9 PM x Treat	-0.052 (0.009)	-0.044 (0.009)	-0.043 (0.009)	-0.027 (0.009)	-0.031 (0.009)
9 PM to Midn. x Treat	-0.053 (0.009)	-0.046 (0.008)	-0.045 (0.009)	-0.029 (0.009)	-0.033 (0.009)
Observations	19,113,592	19,113,592	19,112,654	19,112,654	19,113,592
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: This table reports hour-block-specific results from estimating Equation (3.1), with the prediction errors of the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period. All regressions include a control for being in the post-training period for the machine learning. The results from this table are presented graphically in Panel A of Figure 7 in the main text.

Table A11: Machine learning results by hour-block (HVAC interventions)

	(1)	(2)	(3)	(4)	(5)
Aggregate	-0.044 (0.008)	-0.045 (0.008)	-0.046 (0.008)	-0.032 (0.009)	-0.032 (0.009)
Midn. to 3 AM x Treat	-0.045 (0.010)	-0.044 (0.010)	-0.043 (0.010)	-0.030 (0.010)	-0.032 (0.010)
3 AM to 6 AM x Treat	-0.038 (0.010)	-0.040 (0.010)	-0.040 (0.010)	-0.026 (0.010)	-0.027 (0.010)
6 AM to 9 AM x Treat	-0.004 (0.009)	-0.012 (0.008)	-0.017 (0.008)	-0.004 (0.008)	0.000 (0.008)
9 AM to Noon x Treat	-0.049 (0.008)	-0.052 (0.008)	-0.054 (0.008)	-0.040 (0.008)	-0.040 (0.008)
Noon to 3 PM x Treat	-0.053 (0.009)	-0.056 (0.009)	-0.057 (0.009)	-0.043 (0.009)	-0.044 (0.009)
3 PM to 6 PM x Treat	-0.045 (0.011)	-0.050 (0.010)	-0.052 (0.010)	-0.038 (0.010)	-0.037 (0.010)
6 PM to 9 PM x Treat	-0.055 (0.010)	-0.048 (0.010)	-0.048 (0.010)	-0.035 (0.010)	-0.035 (0.010)
9 PM to Midn. x Treat	-0.062 (0.011)	-0.053 (0.009)	-0.051 (0.010)	-0.037 (0.010)	-0.040 (0.010)
Observations	14,841,212	14,841,212	14,840,417	14,840,417	14,841,212
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: This table reports hour-block-specific results from estimating Equation (3.1), with the prediction errors of the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period. All regressions include a control for being in the post-training period for the machine learning. These regressions include only schools that underwent an HVAC upgrade or never underwent an upgrade during our sample period. The results from this table are presented graphically in Panel B of Figure 2 in the main text.

Table A12: Machine learning results by hour-block (Lighting interventions)

	(1)	(2)	(3)	(4)	(5)
Aggregate	-0.052 (0.009)	-0.052 (0.009)	-0.052 (0.009)	-0.036 (0.010)	-0.041 (0.010)
Midn. to 3 AM x Treat	-0.030 (0.011)	-0.035 (0.011)	-0.032 (0.011)	-0.017 (0.011)	-0.024 (0.011)
3 AM to 6 AM x Treat	-0.022 (0.012)	-0.035 (0.011)	-0.035 (0.011)	-0.020 (0.011)	-0.025 (0.011)
6 AM to 9 AM x Treat	-0.016 (0.010)	-0.032 (0.008)	-0.035 (0.008)	-0.020 (0.009)	-0.021 (0.009)
9 AM to Noon x Treat	-0.078 (0.009)	-0.075 (0.008)	-0.075 (0.009)	-0.060 (0.010)	-0.065 (0.009)
Noon to 3 PM x Treat	-0.086 (0.010)	-0.076 (0.009)	-0.074 (0.009)	-0.059 (0.010)	-0.066 (0.010)
3 PM to 6 PM x Treat	-0.080 (0.012)	-0.066 (0.011)	-0.065 (0.011)	-0.051 (0.012)	-0.055 (0.011)
6 PM to 9 PM x Treat	-0.059 (0.011)	-0.051 (0.011)	-0.050 (0.011)	-0.035 (0.011)	-0.041 (0.011)
9 PM to Midn. x Treat	-0.046 (0.011)	-0.046 (0.011)	-0.043 (0.011)	-0.028 (0.011)	-0.035 (0.011)
Observations	12,849,685	12,849,685	12,849,003	12,849,003	12,849,685
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: This table reports hour-block-specific results from estimating Equation (3.1), with the prediction errors of the log of hourly energy consumption (averaged across “blocks” of three hours) as the dependent variable. The independent variable is defined as the fraction of total expected savings that have been installed by time t , and takes on values from 0 to 1 in the treated schools, and 0 always in the control schools. Standard errors, clustered at the school level, are in parentheses. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period. All regressions include a control for being in the post-training period for the machine learning. These regressions include only schools that underwent a lighting upgrade or never underwent an upgrade during our sample period. The results from this table are presented graphically in Panel C of Figure 2 in the main text.

Table A13: Machine learning results: Pre- vs. post-period training (Any intervention)

	(1)	(2)	(3)	(4)	(5)
Trained on pre	-0.044	-0.045	-0.045	-0.029	-0.031
	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
Observations	19,113,602	19,113,602	19,112,664	19,112,664	19,113,602
Trained on post	-0.059	-0.058	-0.061	-0.034	-0.032
	(0.008)	(0.008)	(0.008)	(0.008)	(0.009)
Observations	19,026,970	19,026,970	19,026,079	19,026,079	19,026,970
Pooled	-0.052	-0.052	-0.053	-0.032	-0.032
	(0.007)	(0.007)	(0.007)	(0.008)	(0.007)
Observations	38,140,574	38,140,574	38,138,654	38,138,654	38,140,574
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: In this table, we present three variations on our prediction procedure. The “Trained on pre” panel presents results where the pre-treatment period is used to train the machine learning model, which we then forecast into the post-treatment periods to estimate treatment effects. In the “Trained on post” panel, we reverse this procedure, training the model on the post-treatment period, and projecting into the pre-treatment period (scaled such that treatment effects have the same sign). Finally, the “pooled” panel uses both trained-on-pre and trained-on-post predictions. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Standard errors are clustered at the school level.

Table A14: Machine learning results: Pre- vs. post-period training (HVAC interventions)

	(1)	(2)	(3)	(4)	(5)
Trained on pre	-0.044	-0.045	-0.046	-0.032	-0.032
	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)
Observations	14,841,225	14,841,225	14,840,437	14,840,437	14,841,225
Trained on post	-0.057	-0.056	-0.059	-0.032	-0.031
	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)
Observations	14,759,556	14,759,556	14,758,793	14,758,793	14,759,556
Pooled	-0.050	-0.050	-0.052	-0.032	-0.031
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Observations	29,603,837	29,603,835	29,602,206	29,602,206	29,603,835
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

Notes: In this table, we present three variations on our prediction procedure. The “Trained on pre” panel presents results where the pre-treatment period is used to train the machine learning model, which we then forecast into the post-treatment periods to estimate treatment effects. In the “Trained on post” panel, we reverse this procedure, training the model on the post-treatment period, and projecting into the pre-treatment period (scaled such that treatment effects have the same sign). Finally, the “pooled” panel uses both trained-on-pre and trained-on-post predictions. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Standard errors are clustered at the school level. These regressions include only schools that underwent an HVAC upgrade or never underwent an upgrade during our sample period.

Table A15: Machine learning results: Pre- vs. post-period training (Lighting interventions)

	(1)	(2)	(3)	(4)	(5)
Trained on pre	-0.052 (0.009)	-0.052 (0.009)	-0.052 (0.009)	-0.036 (0.010)	-0.041 (0.010)
Observations	12,849,701	12,849,701	12,849,024	12,849,024	12,849,701
Trained on post	-0.067 (0.010)	-0.066 (0.010)	-0.064 (0.010)	-0.042 (0.011)	-0.050 (0.011)
Observations	12,746,680	12,746,680	12,746,047	12,746,047	12,746,680
Pooled	-0.061 (0.009)	-0.061 (0.009)	-0.059 (0.009)	-0.040 (0.010)	-0.047 (0.010)
Observations	25,599,442	25,599,441	25,598,071	25,598,071	25,599,441
School FE, Block FE	Yes	Yes	Yes	Yes	Yes
School-Block FE	No	Yes	Yes	Yes	Yes
School-Block-Month FE	No	No	Yes	No	Yes
Month of Sample FE	No	No	No	Yes	No
Month of Sample Ctrl.	No	No	No	No	Yes

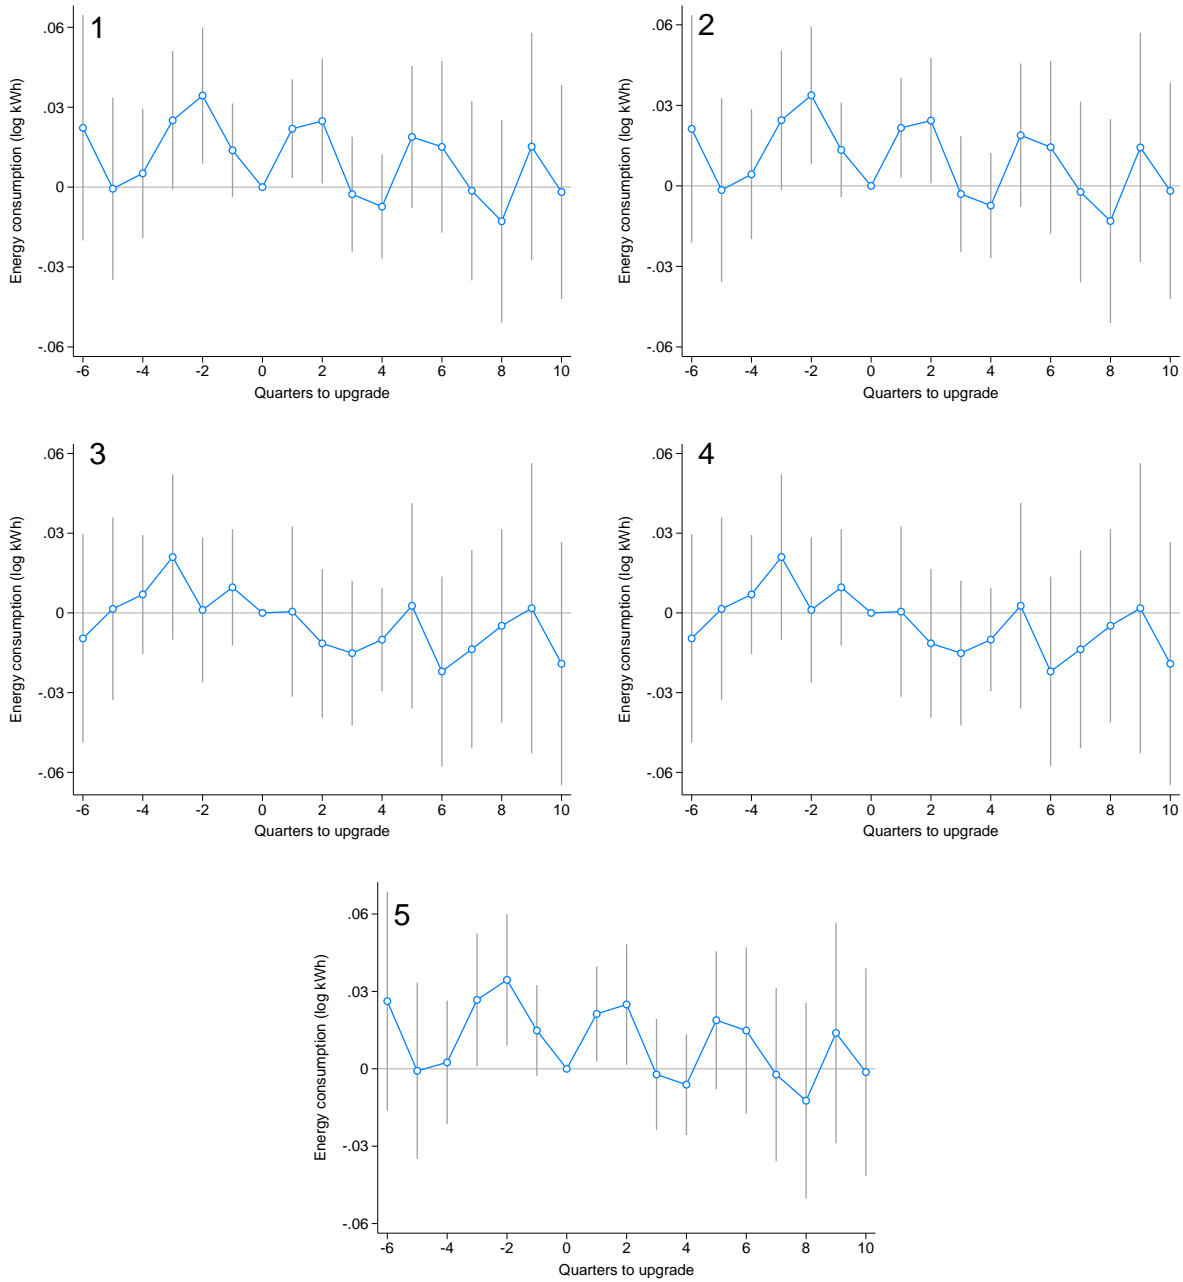
Notes: In this table, we present three variations on our prediction procedure. The “Trained on pre” panel presents results where the pre-treatment period is used to train the machine learning model, which we then forecast into the post-treatment periods to estimate treatment effects. In the “Trained on post” panel, we reverse this procedure, training the model on the post-treatment period, and projecting into the pre-treatment period (scaled such that treatment effects have the same sign). Finally, the “pooled” panel uses both trained-on-pre and trained-on-post predictions. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Standard errors are clustered at the school level. These regressions include only schools that underwent a lighting upgrade or never underwent an upgrade during our sample period.

Table A16: Machine learning results: double LASSO

	(1) Any intervention	(2) HVAC interventions	(3) Lighting interventions
Treatment (aggregate)	-0.029 (0.007)	-0.030 (0.008)	-0.036 (0.009)
Midn. to 3 AM x Treat	-0.021 (0.008)	-0.027 (0.009)	-0.017 (0.011)
3 AM to 6 AM x Treat	-0.019 (0.008)	-0.024 (0.009)	-0.020 (0.011)
6 AM to 9 AM x Treat	-0.012 (0.006)	-0.008 (0.007)	-0.025 (0.009)
9 AM to Noon x Treat	-0.040 (0.006)	-0.036 (0.007)	-0.057 (0.008)
Noon to 3 PM x Treat	-0.043 (0.006)	-0.040 (0.007)	-0.058 (0.009)
3 PM to 6 PM x Treat	-0.040 (0.007)	-0.038 (0.009)	-0.052 (0.010)
6 PM to 9 PM x Treat	-0.029 (0.008)	-0.035 (0.009)	-0.036 (0.010)
9 PM to Midn. x Treat	-0.026 (0.008)	-0.032 (0.009)	-0.027 (0.011)
Observations	18,955,076	14,739,256	12,676,852

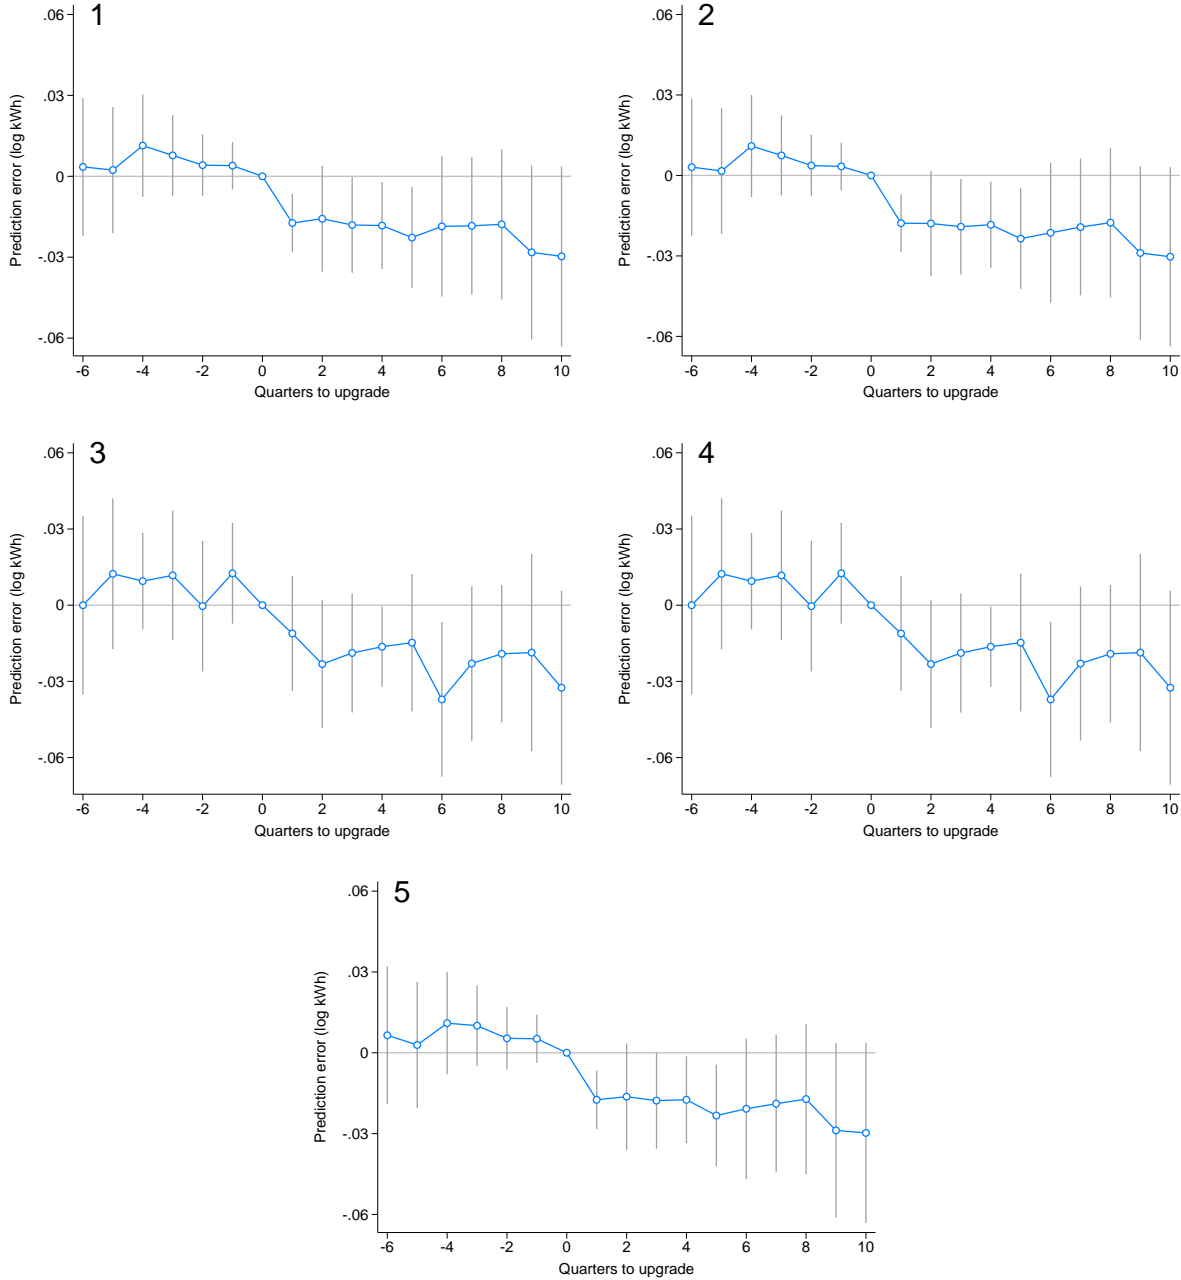
Notes: This table presents results from our extension of Belloni, Chernozhukov, and Hansen (2014)'s double selection approach. We first estimate a LASSO to predict the timing of treatment, next estimate a second LASSO to predict electricity consumption, and finally estimate a third LASSO with time as the dependent variable, in order to accommodate trends. Similar to the standard double selection procedure, we then regress energy consumption on treatment timing and the union of the non-zero-coefficient variables from all three LASSOs. To make this computationally tractable, we residualize each dependent variable by the full set of controls, and implement the final step by regressing residualized prediction errors on residualized treatment date error and residualized time error. These procedures are mathematically equivalent, via Frisch-Waugh-Lovell. All samples are trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Standard errors are clustered at the school level. Regressions for HVAC and light interventions include only schools with an intervention of this type and pure untreated schools that never underwent energy efficiency interventions of any kind during the sample period.

Figure A1: Panel fixed effects event study – all specifications



Notes: This figure displays point estimates and 95 percent confidence intervals from event study regressions of energy consumption before and after an energy efficiency upgrade. We estimate Equation (3.2) with the log of hourly electricity consumption (in kWh, averaged by three hour block) as the dependent variable. We normalize time relative to the quarter each school undertook its first upgrade. Each panel corresponds to the like-numbered column of Table 2, and includes both treated and untreated schools. Standard errors are clustered by school, and the sample has been trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Even with flexible controls, these estimates display strong patterns - perhaps reflecting seasonality in upgrade timing. We also do not see strong evidence of a shift in energy consumption as a result of energy efficiency upgrades.

Figure A2: Machine learning event study – all specifications



Notes: This figure displays point estimates and 95 percent confidence intervals from event study regressions of energy consumption before and after an energy efficiency upgrade. We estimate Equation (3.2) with prediction errors based on log electricity consumption in kWh (averaged across three-hour “blocks”) as the dependent variable. We normalize time relative to the quarter each school undertook its first upgrade. Each panel corresponds to the like-numbered column of Table 4, and includes both treated and untreated schools. Standard errors are clustered by school, and the sample has been trimmed to exclude observations below the 1st or above the 99th percentile of the dependent variable. Unlike the regression estimates displayed in Figure 3, there is a clear change in energy consumption after the installation of energy efficiency upgrades, which persists more than a year after the upgrade. Furthermore, we fail to reject differential energy consumption between treated and untreated schools prior to the upgrades.