

Longer-Run Evidence on Whether Building Energy Codes Reduce Residential Energy Consumption

Matthew J. Kotchen

Abstract: This paper provides an ex post evaluation of how changes to a building energy code affect energy consumption. Using residential billing data for electricity and natural gas over 11 years, the analysis is based on comparisons between residences constructed just before and just after a building code change in Florida. While an earlier study using 3 years of data for the same residences showed savings for both electricity and natural gas, new results show an enduring savings for natural gas only. These findings underscore the importance of accounting for all sources of energy consumption when conducting evaluations of building codes. More broadly, the results provide a counterpoint to the growing literature casting doubt on whether ex ante forecasts of energy efficiency policies and investments can provide useful information about actual energy savings. Indeed, more than a decade after Florida's energy code change, the measured energy savings still meets or exceeds the forecasted amount.

JEL Codes: Q47, Q48

Keywords: Building codes, Energy, Evaluation, Regulation

ARE BUILDING ENERGY CODES effective at saving energy? The answer to this question is important given the growing reliance on building codes as a central part of energy and climate policy in the United States and abroad. The promotion of building energy codes is a priority at the US Department of Energy, which estimates energy expenditure savings in the hundreds of billions of dollars by 2030, with emission reductions equivalent to taking millions of vehicles off the road (US DOE 2011). In the European Union, all member countries must comply with the revised Energy Performance

Matthew Kotchen is at Yale University and the National Bureau of Economic Research (matthew.kotchen@yale.edu). Without implicating them for any errors or omissions, I am grateful to Grant Jacobsen, Arik Levinson, Erin Mansur, and Joe Shapiro for helpful comments and discussion. I also thank two anonymous reviewers and the editor Stephen Holland for constructive suggestions on an earlier version of the paper.

Received February 15, 2016; Accepted June 1, 2016; Published online January 25, 2017.

JAERE, volume 4, number 1. © 2017 by The Association of Environmental and Resource Economists. All rights reserved. 2333-5955/2017/0401-0004\$10.00 <http://dx.doi.org/10.1086/689703>

of Buildings Directive that seeks to promote energy efficiency and help meet the EU's greenhouse gas emissions targets (European Union 2012). Despite the growing emphasis on building codes as a regulatory instrument, our understanding of the actual impacts on energy consumption remains thin. There are only a handful of peer reviewed studies that seek to evaluate the extent to which building energy codes affect construction practices and energy consumption, and among these studies the results are quite mixed.¹

In one of the more recent papers, Jacobsen and Kotchen (2013) study the effects of a change in Florida's building code using residential billing data for electricity and natural gas. Their study is based on a comparison between residences in Gainesville, Florida, that were constructed just before and just after an increase in the stringency of Florida's energy code in 2002. Jacobsen and Kotchen find that the code change caused a 4% decrease in electricity consumption and a 6% decrease in natural gas consumption, and these savings were close to those predicted *ex ante* for the code change. They also find that energy consumption in post-code change residences is less responsive to weather shocks in ways consistent with greater energy efficiency and that the social and private payback periods for code compliance range between 3.5 and 6.4 years, respectively.

Even more recently, Levinson (2014) studies the effect of building energy codes on electricity consumption in California. Despite the fact that California is often perceived as a model state promoting energy-efficient buildings, Levinson (2014) finds little evidence that building codes have anything close to the anticipated effect on residential electricity consumption.² An important aspect of Levinson's contribution is the careful attention given to building age as a distinct feature from building vintage with respect to codes. He shows that newer homes consume less electricity simply because they are new, and this observation can be problematic for reliably estimating the effect of building codes. Hence careful methods are needed to separate age from vintage effects, that is, older and newer homes from those built before or after a building code change.

The potential importance of a newness effect raises questions about Jacobsen and Kotchen's research design and findings in Florida. Did the post-code-change residences consume less energy simply because they were new? And, over time, will the energy consumption of post-code-change residences more closely resemble the pre-code-change residences as the differences in age reflect less of a newness effect? Indeed, Levinson (2014, 8) writes, "I suspect if we revisited those Gainesville homes today, 10 years later,

1. Studies that focus explicitly on the effects of regulatory building codes include Horowitz and Haeri (1990), Jaffe and Stavins (1995), Aroonruengsawat et al. (2012), Costa and Kahn (2010), Jacobsen and Kotchen (2013), and Levinson (2014).

2. Levinson's (2014) paper has received a substantial amount of media attention given its counterintuitive finding. Readers may find of interest a 30-minute Freakonomics podcast dedicated to the paper at <http://freakonomics.com/tag/arik-levinson/>.

we would find no difference in energy use for households built before and after the 2002 code change.”

This conjecture is the starting point for the present paper. I test whether Jacobsen and Kotchen’s findings still hold 11 years after the building code change. This is important for at least two reasons. First, Jacobsen and Kotchen’s approach provides a clear identification strategy for estimating building code effects, and their results are some of the very few that find an effect. Knowing whether the results endure is therefore highly policy relevant. Second, regarding energy efficiency investments more generally, a growing literature finds that *ex ante* engineering studies significantly overestimate realized savings in *ex post* evaluations.³ Yet Jacobsen and Kotchen’s study provides a counterpoint in the important context of building codes. Because they find that the engineering forecasts are in line with the estimated savings, the question is whether or not the results hold up.

This paper also contributes with new insight about the effect of building codes on energy consumption over time and how future studies should approach *ex post* evaluations. The results do not yield a simple yes or no to an enduring building code effect. There are differences between the results for electricity and natural gas. While the initial estimates of the building code effect on electricity consumption diminish over time, the original results for natural gas underestimate the longer-term energy savings. Nevertheless, the overall net effect on energy consumption in a combined measure of million British thermal units indicates a rather consistent level of code-induced energy savings over the 11-year period. Together, the results highlight the importance of not focusing exclusively on electricity consumption—as most studies have done—and for waiting a few years after an energy code change to begin evaluation.

1. EMPIRICAL SETTING AND DATA COLLECTION

Florida’s residential building code that took effect in March 2002 included provisions to strengthen the energy efficiency of newly constructed houses. The new requirements focused on energy used for space heating, space cooling, and water heating.⁴ To evaluate the effect of the code on residential energy consumption, Jacobsen and Kotchen use monthly billing data on electricity and natural gas consumption for residences in Gainesville, Florida. They focus on residences constructed 3 years before and 3 years after the building code change. A unique feature of their data set is that billing data were combined with information about the physical characteristics of each residence, in addition to monthly weather data.

3. Examples include Dubin, Meidema, and Chandran (1986), Metcalf and Hassett (1999), and Fowle, Greenstone, and Wolfram (2015). The general issues are also reviewed in Allcott and Greenstone (2012), Gillingham and Palmer (2014), and Gerarden et al. (2015).

4. See Jacobsen and Kotchen (2013) for a detailed description of specific changes to the 2002 code and methods of compliance.

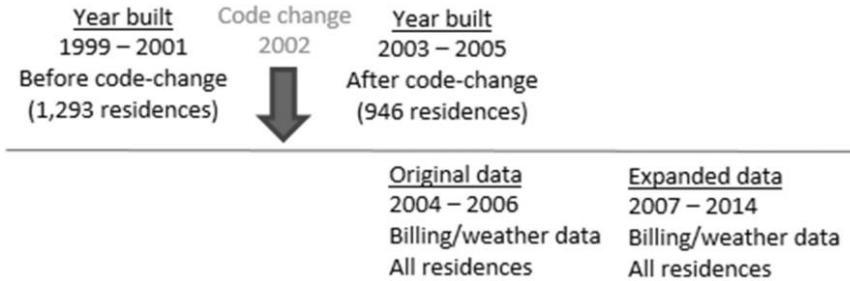


Figure 1. Before and after code change residences, and the original and expanded data set

Figure 1 illustrates the basic research design, the original data set, and the expanded data set used in the present paper. Before code change residences (BCCRs) were built in the years 1999–2001, and after code change residences (ACCRs) were built in the years 2003–5.⁵ Unless otherwise indicated, residences built in 2002 are excluded from the analysis because they cannot be clearly identified as having been subject to the before or after code change requirements. Jacobsen and Kotchen used billing data for the years 2004–6 to test for differences in energy consumption between the before and after code change residences. In this paper, I use an expanded data set that includes billing and weather data for the additional years 2007–14, yielding a total of 11 years of energy consumption data after the building code change. Importantly, the additional billing data are for the identical set of before and after code change residences.

Most of the expanded data are from the same sources. Billing data were obtained from Gainesville-green.com, and weather data on monthly average cooling degree days (ACDD) and average heating degree days (AHDD) were obtained for the weather station located at the Gainesville Regional Airport.⁶ The new data were merged with the original to form a monthly panel from January 2004 through December 2014 for 2,239

5. Residences are partitioned as BCCRs or ACCRs based on the observed effective year built (EYB) as it relates to the timing of energy code compliance at the time the residence received its permits (for further details, see Jacobsen and Kotchen 2013). To ensure that only newly constructed residences are included, any residence is dropped from the analysis if it had a utility bill on record prior to the EYB, as this would indicate remodeling rather than new construction.

6. One difference in the source of data was necessary because natural gas data were not fully up to date at Gainesville-green.com. I obtained natural gas billing data from October 2013 through December 2014 directly from the Gainesville Regional Utilities (GRU), which makes all billing data publicly available for the two most recent years and ultimately provides the source for updating Gainesville-green.com.

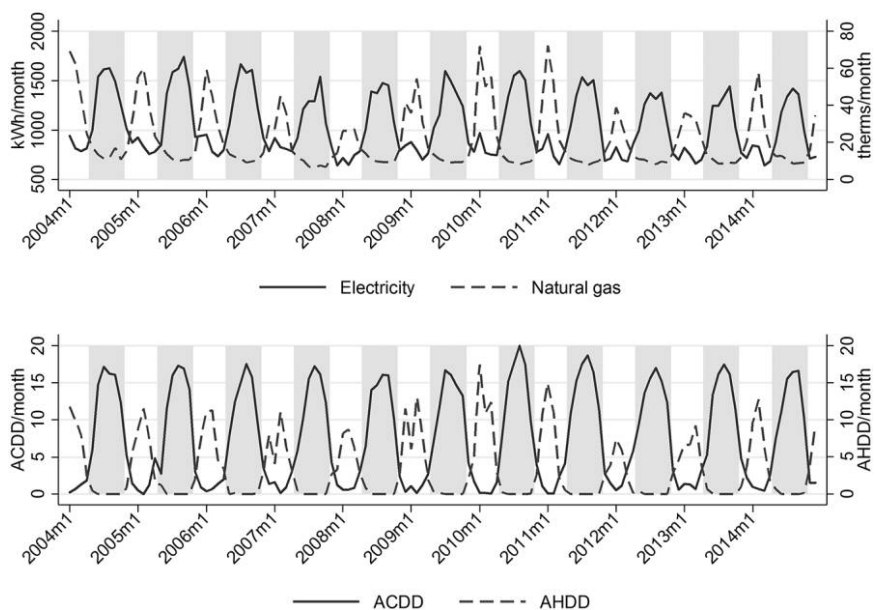


Figure 2. Average monthly electricity and natural gas consumption and cooling and heating degree days from 2004 to 2014, with months May–October shaded each year.

residences, of which there are 1,293 BCCRs and 946 ACCRs.⁷ The time-invariant variables on the physical characteristics of each residence are square footage; number of bathrooms, bedrooms, and floors; and indicators for central air conditioning and a shingled roof.⁸

The top panel of figure 2 shows the average monthly electricity and natural gas consumption for all residences from January 2004 through December 2014. The hotter months of the year—May through October—are shaded in the figure. There is a clear pattern where electricity consumption is higher during the summer months when de-

7. Two areas of missing billing data are for the majority of residences between April 2007 and February 2008, and natural gas data for all residences in July 2014. Attempts to obtain the missing data from GRU have been unsuccessful. In what follows, I include all of the available data in the analysis, but all results are robust to dropping the months where there are missing data.

8. Summary statistics for residence characteristics remain unchanged from Jacobsen and Kotchen’s original analysis (see Jacobsen and Kotchen 2013, tables 1, 2). The only difference with a potentially meaningful magnitude and statistical significance is that ACCRs are 4.5% smaller.

mand for air conditioning is greater, and natural gas consumption is higher during winter months when demand for heating is greater. The bottom panel of figure 2 shows the patterns of monthly ACDD and AHDD based on the standard 65° Fahrenheit threshold. Comparisons between the panels of figure 2 reveal how electricity demand closely follows cooling degree days, while natural gas follows heating degree days.

2. BEFORE AND AFTER CODE CHANGE COMPARISONS

2.1. Overall Differences

The key estimates from Jacobsen and Kotchen are regression-based, average differences in electricity and natural gas consumption between before and after code change residences. The preferred specification is

$$Y_{ijt} = \delta \text{CodeChange}_i + \beta \mathbf{X}_i + v_{jt} + \varepsilon_{ijt}, \quad (1)$$

where the dependent variable represents consumption of either electricity (kilowatt hours, kWhs) or natural gas (therms) in residence i , zip code j , and month t ; CodeChange_i is an indicator for whether a residence was built after the code change; \mathbf{X}_i is a vector of the observable residence characteristics; v_{jt} represents zip code by month-year fixed effects; and ε_{ijt} is an error term. The residence characteristics included in the model are the natural log of square footage, an indicator for central air conditioning, an indicator for a shingled roof, and categorical variables for the number of bathrooms, bedrooms, and stories.⁹ When reporting results for this model, and all others throughout the paper, I use standard errors that are clustered at the residence level.¹⁰

Table 1 reproduces Jacobsen and Kotchen's results with the original 2004–6 data for electricity and natural gas in the first column. I report the coefficients of interest, the δ 's, along with the mean monthly consumption and percentage difference between before- and after-code-change residences. Jacobsen and Kotchen found that ACCRs consumed 4.27% less electricity and 6.67% less natural gas. The second column of table 1 reports analogous results using the expanded data set through 2014. With the longer span of data, the difference in electricity consumption between before- and after-code-change residences is no longer statistically significant and has a point estimate very close to zero. Levinson's (2014) conjecture is therefore consistent with the data for electricity; however, the pattern is very different for natural gas. The difference in natural gas consumption between before- and after-code-change residences for the ini-

9. Jacobsen and Kotchen also estimate models that include simply month-year fixed effects. Unless otherwise indicated, I focus throughout this paper on models with zip code by month-year fixed effects because they are less restrictive and therefore preferred.

10. In an appendix, available online, I discuss alternative specifications to estimate overall building code differences that account explicitly for residence age. I report selected results and explain why the models are not the preferred specifications for this particular analysis.

Table 1. Average Differences in Energy Consumption between Before- and After-Code-Change Residences

	Original 2004–6	All Data 2004–14
Electricity:		
Code change	-48.922*** (20.295)	-.502 (15.651)
Mean kWh/month	1,146.4	1,049.1
Percent difference	-4.27%	-.05%
Observations	64,471	256,635
Natural gas:		
Code change	-1.572** (.704)	-2.875*** (.537)
Mean therms/month	23.6	21.3
Percent difference	-6.67%	-13.5%
Observations	64,471	254,467
Combined mmBtu:		
Code change	-.324*** (.109)	-.290*** (.080)
Mean mmBtu/month	6.3	5.7
Percent difference	-5.14%	-5.09%
Observations	64,471	254,467

Note. Each coefficient is from a different regression model that includes residence characteristics and zip code × month-year fixed effects. Standard errors in parentheses are clustered at the residence level.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

tial 2004–6 period provides an underestimate of the difference over the whole period through 2014. Using the full series of data, I find that ACCRs use 13.5% less natural gas on average—double the initial estimate, and with a high level of statistical significance.

The bottom panel of table 1 reports the results of new models that combine electricity and natural gas into a single measure of overall energy consumption, quantified as millions of British thermal units (mmBtu).¹¹ Combining electricity and natural gas into a single measure of energy consumption has the advantage of estimating an overall effect that accounts for potential substitution between energy sources. The focus on overall energy use is also more appropriate for a performance-based code, such as Florida’s, where builders can trade off among different energy options. Because com-

11. The exact conversion is based on $\text{mmBtu} = 0.0034121416 \times \text{kWhs} + 0.1 \times \text{therms}$.

pliance is based on an overall rating, rather than conforming to specific requirements, there is no reason to expect that consumption of both electricity and natural gas would necessarily decrease. These results indicate that for the initial period of 2004–6, ACCRs consume 5.14% less energy overall, and the difference is highly statistically significant. Then, when using all the data through 2014, the average difference remains nearly identical, at 5.09% with the same level of statistical significance.

The focus here is on energy consumption because building energy codes have saving energy as the primary objective. But other outcomes are of interest as well, such as cost savings and carbon dioxide (CO₂) emissions. Combining electricity and natural gas consumption into mmBTUs is essentially a weighting between the two sources of residential energy use. Combining consumption into monthly energy expenditures and CO₂ are alternative approaches that I consider to estimate overall average differences between BCCRs and ACCRs.¹²

Table 2 reports the key coefficients from estimating (1) with the different dependent variables and all of the data from 2004 to 2014. The percentage changes in expenditures for electricity and natural gas are very close to the estimates for the change in consumption. They differ only because the rate schedules are nonlinear. The CO₂ emissions are, however, a linear function of consumption, so these percentage changes are identical to those in table 1. The combined estimates for expenditures and emissions have a lower magnitude than those for consumption and are not statistically different from zero. This difference occurs because the expenditure and emission weighting places relatively more weight on electricity than the weighting for mmBTU. Hence, while there is statistically significant evidence of energy savings, the evidence is weaker for savings on utility bills and reducing CO₂ emissions.

2.2. Differences by Effective Year Built

The previous estimates are for overall differences in energy consumption between BCCRs and ACCRs. As shown in figure 1, residences are partitioned into the two groups by their effective year built (EYB). The overall estimate can therefore be decomposed further into average differences by EYB. This is useful for examining potential differences

12. Expenditures are estimated based on GRU's variable, pre-tax price schedules in 2014. After several attempts, I was not able to obtain price schedules for all years, but I have been assured that prices did not change significantly over the 10-year period. Hence using 2014 prices should provide reasonable estimates. For electricity, the increasing block rate is \$0.031 for 0–250 kWh, \$0.042 for 251–750 kWh, and \$0.084 for over 750 kWh, along with \$0.078 for all kWh as a fuel adjustment charge. For natural gas, the rate is \$.502 per therm, plus a gas plant recovery charge of \$0.0556 and a \$.46 fuel adjustment per therm. CO₂ emissions for electricity are based on the marginal emission rate of electricity from Graff Zivin, Kotchen, and Mansur (2014) of 1.29 lbs/kWh for the FRCC NERC region, while for natural gas it is the physical relationship of 11.7 lbs/therm.

Table 2. Average Differences in Variable Energy Expenditures and CO₂ Emissions between Before- and After-Code-Change Residences

	Expenditures	CO ₂ Emissions
Electricity:		
Code change	-.066 (2.433)	-.648 (20.190)
Monthly mean	140.1	1,353.4
Percent difference	-.05%	-.05%
Natural gas:		
Code change	-2.926*** (.547)	-33.636*** (6.288)
Monthly mean	21.7	249.0
Percent difference	-13.4%	-13.5%
Combined:		
Code change	-3.022 (2.553)	-34.545 (21.811)
Monthly mean	161.3	1,599.2
Percent difference	-1.87%	-2.2%

Note. Each coefficient is from a different regression model that includes residence characteristics and zip code × month-year fixed effects. Models have the same number of observations as those indicated in table 1. Standard errors in parentheses are clustered at the residence level.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

in energy consumption by the age of residences within and between the two groups. Specifically, the model is

$$Y_{ijt} = \delta EYB_i + \beta X_i + v_{jt} + \varepsilon_{ijt}, \tag{2}$$

which differs from (1) because EYB_i is a categorical variable for each effective year built from 1999 through 2005. Jacobsen and Kotchen estimate the same model, and as in the original analysis, I include residences built in 2002 as the omitted category.

Figure 3 illustrates the estimated δ 's graphically for electricity, natural gas, and combined energy consumption. There is no trend in electricity consumption by effective year built. This result is consistent with no observable, enduring effect of the energy code change on electricity consumption. With natural gas, however, the results are again different. While there is no trend for the BCCRs, consumption is declining for the ACCRs. Yet it appears that it may take a year or so for the decline to begin after the code change, raising the possibility that residences with an effective year built in 2003 (and maybe some in 2004) might still have been subject to the before-code-change

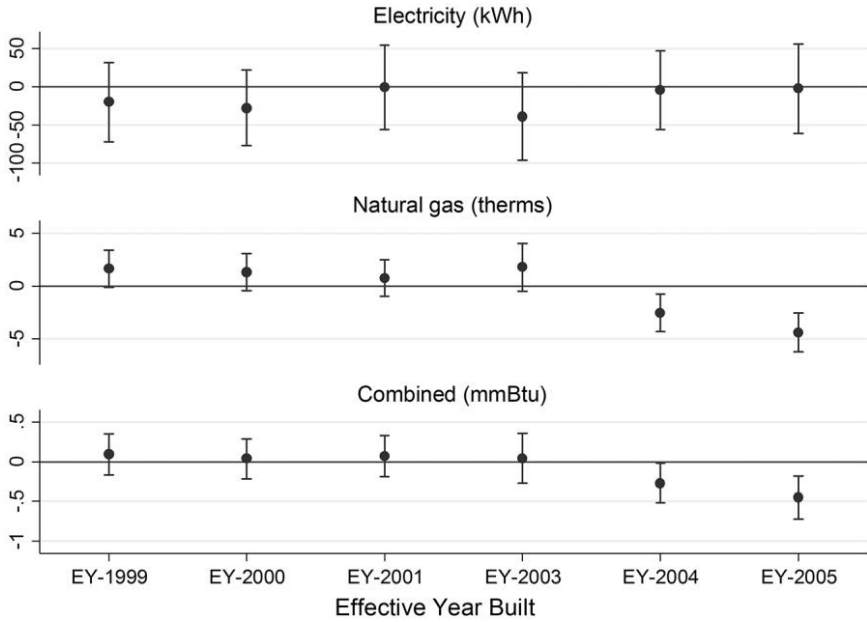


Figure 3. Overall differences in electricity, natural gas, and combined energy consumption among residences by effective year built, relative to those built in 2002, with 95% confidence intervals.

requirements. The results for combined energy consumption, which is effectively a weighted average of electricity and natural gas, more closely resemble those for natural gas. The important takeaway from the figure is that the trend in energy consumption appears to break around the time of the energy code change, suggesting that the code change had an effect on energy consumption, rather than the estimates simply capturing an age effect.¹³

2.3. Yearly Differences

While the energy code change is associated with lower energy consumption overall, as measured with mmBtu, the results in table 1 clearly indicate that something different is happening over time with electricity compared to natural gas. This motivates closer scrutiny of how the estimated building code effects differ over time. Accordingly, I estimate additional models of the form

13. Jacobsen and Kotchen's figures 3 and 4 present results for electricity and natural gas using only the first 3 years of consumption data. Those reported here differ by clearly showing no code change effect on electricity, but a more clear effect on natural gas. Jacobsen and Kotchen did not provide separate estimates for combined energy consumption.

$$Y_{ijt} = \delta \text{CodeChange}_i \times \text{Year}_t + \beta X_i + v_{jt} + \varepsilon_{ijt}, \quad (3)$$

where Year_t is a categorical variable for each year in the sample. Estimation of this model yields distinct energy code effects for each year, where the δ 's capture differences in the annual averages between BCCRs and ACCRs. I estimate separate models for electricity, natural gas, and mmBtu and report the results graphically.

Figure 4 includes all three sets of results. The δ 's are scaled as a percentage difference from average consumption for the corresponding energy measure and year. It is worth keeping in mind that when the estimates are based on annual differences, there is less statistical power, yet the trend in point estimates is of primary concern here. The results for electricity illustrate an upward trend and clearly show how focusing on the first three years of data provides an overestimate of the electricity savings.¹⁴ Indeed, ACCRs appear to consume even more electricity in the more recent years, despite initially consuming less. But the trend clearly differs for natural gas. After the first two years, the ACCRs consume significantly less natural gas, with point estimates ranging between 10% and 20% less, and the difference endures over time. Finally, the results for overall energy consumption indicate that ACCRs uniformly consume less energy, but the differences vary from year to year, with point estimates ranging between 2% and 10% less. In this case, we can see visually how the short-run estimate over three years is actually quite close to the 11-year estimate.

What might explain the pattern of differences for electricity and natural gas over the 11-year period? With respect to electricity, Levinson (2014) posits that such a pattern might arise because ACCRs are simply newer, and the newness effect might dissipate over time. Although better insulated when newly constructed, residences may quickly lose some of the tight envelope that increases initial efficiency. New heating and cooling systems are more likely to have clean and efficient air filters that are subsequently changed less regularly. It may also take the occupants of new residences time to fully move in and acquire the same set of appliances as those in older residences.

An alternative explanation for the electricity results is that, despite the small differences in the effective year built between BCCRs and ACCRs, occupants of the ACCRs may be different, perhaps younger and more likely to add members to a growing family. Younger families may also have relatively increasing demands for goods that use electricity in ways unaffected by the building code (e.g., televisions and electronics). In this case, however, the relative increase in electricity consumption in ACCRs over time would not mean the code change had no effect. Instead, it would raise questions about the identification assumption underlying the regression models for estimating average differences, a topic to which I return in section 3.

14. Recall that ACCRs are continuously entering the data set in 2004 and 2005. Estimates for these years are therefore based on a smaller set of ACCRs with monthly observations not uniformly spread over months of the year.

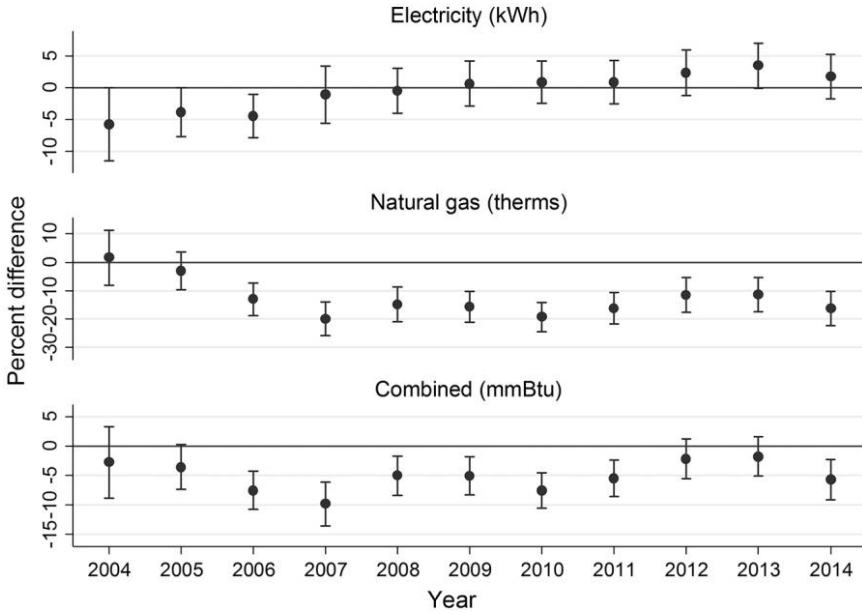


Figure 4. Annual differences in electricity, natural gas, and combined energy consumption between before- and after-code-change residences, with 95% confidence intervals.

The natural gas results show a consistent pattern of energy savings from the building code after the first two years. While there may be an initial adjustment period for natural gas consumption in new residences, perhaps because new occupants are learning to use systems, I prefer an explanation that simply discounts estimates for the first two years because all ACCRs are not yet fully included in the data set. That natural gas shows a consistent energy code savings beginning in 2006 is surely consistent and plausible given the focus of the code change on space and water heating, the main drivers of natural gas consumption other than cooking.

2.4. Monthly Differences

A fourth set of models examines differences in electricity and natural gas consumption over months of the year. Jacobsen and Kotchen conducted this analysis to show that ACCRs consumed less electricity in the summer and less natural gas in the winter, because of lower energy demand for air conditioning and heating, respectively. Here I examine the patterns of monthly consumption with models of the form

$$Y_{jt} = \delta \text{CodeChange}_i \times \text{Month}_t + \beta X_i + v_{jt} + \varepsilon_{ijt}, \tag{4}$$

where *Month_t* is a categorical variable for month of the year. In this case, the δ 's provide estimates of the differences in consumption between before- and after-code-change residences for each month of the year.

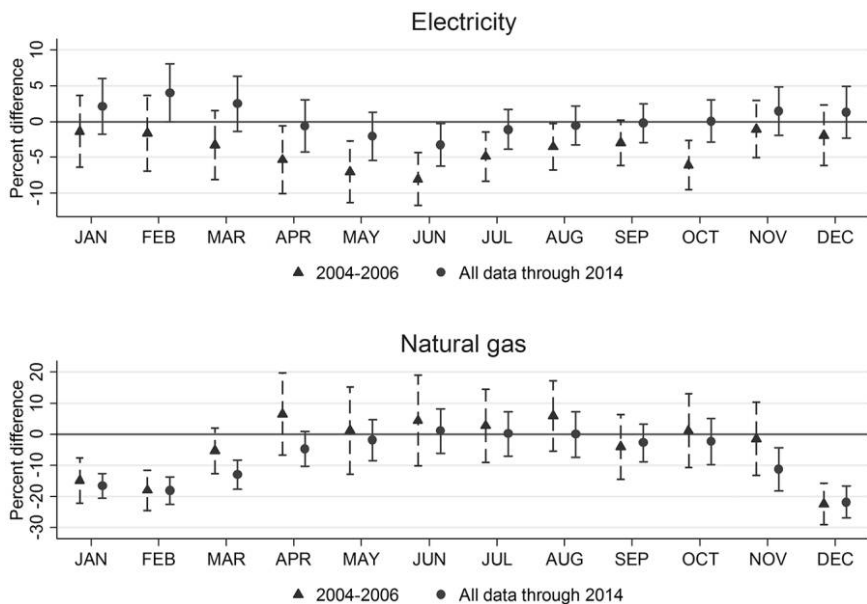


Figure 5. Monthly differences in electricity and natural gas consumption between before and after code change residences, 2004–6 and all data through 2014, with 95% confidence intervals.

Figure 5 illustrates the results graphically for electricity and natural gas.¹⁵ I report results using the original 2004–6 data and the full set of data through 2014. The top panel shows the pattern for electricity described by Jacobsen and Kotchen: using data 3 years after the code change, there are no differences in electricity consumption between before- and after-code-change residences during the colder and winter months, but the ACCRs consume less electricity during the hotter and summer months when demand for air conditioning is greater.¹⁶ When using all the data through 2014, the profile of electricity consumption remains the same, yet the difference between before- and after-code-change residences tends to be greater across all months of the year. We have already established in figure 4 that electricity consumption is increasing over time in the ACCRs compared to the BCCRs, and figure 5 shows that the increase

15. I do not report the results for mmBtu because the reason for estimating monthly differences is to gain insight into the seasonal patterns of consumption for the original sources of energy demand, electricity and natural gas. But, of course, similarly formatted results for mmBtu are available upon request.

16. The results presented here are based on models with zip code by month-year fixed effects rather than simply month-year fixed effects as shown in Jacobsen and Kotchen’s figures 1 and 2.

occurs because of greater baseline demand rather than a seasonal effect.¹⁷ The natural gas results in the bottom panel reveal lower consumption of natural gas in ACCRs in the colder and winter months, indicating greater efficiency with heating. With natural gas, however, there does not appear to be an increasing or decreasing trend over months of the year when using the full set of data, and this is consistent with the ACCRs consuming uniformly less natural gas in figure 4.

3. DIFFERENCES IN WEATHER RESPONSIVENESS

An advantage of the analysis in the previous section is that it yields an estimate of the average effect of the building code change on residential energy consumption. A disadvantage of the approach is potential vulnerability to omitted variable bias. If there is some unobserved variable that is correlated with energy consumption and the BCCR-ACCR categorization, for reasons unrelated to the building code change, then the estimates could be biased. As discussed previously, examples include differences in families that purchased homes in Gainesville a few years later, and differences in the stock of appliances in newer residences.

To address these concerns, I estimate models to test for differences in the way BCCRs and ACCRs adjust energy consumption in response to weather shocks. The approach is essentially a difference-in-differences strategy that takes account of unobservable, time-invariant heterogeneity with the inclusion of residence fixed effects. While Jacobsen and Kotchen estimate similarly specified fixed effects models for electricity and natural gas consumption separately, I focus here on the combined measure of mmBtu.¹⁸ The combined measure is preferable because, as mentioned previously, it yields an overall effect that accounts for potential substitution between electricity and natural gas.

The first model takes the form

$$Y_{it} = \delta \text{CodeChange}_i \times [\text{ACDD}_t, \text{AHDD}_t] + \beta [\text{ACDD}_t, \text{AHDD}_t] + \text{Month}_t + \text{Year}_t + \mu_i + \varepsilon_{it} \quad (5)$$

where the dependent variable is monthly mmBtu consumption, the code change indicator is interacted with the weather variables, Month_t is a set of month of year dum-

17. Although I do not find evidence for it here, a seasonal effect of comparatively increasing demand for electricity in the summer by ACCRs would be consistent with an air conditioning rebound effect.

18. Though not reported here, I also estimated separate models for electricity and natural gas. The specifications are identical to those in J&K's analysis but with the expanded data set through 2014. The main results continue to hold, whereby ACCRs are less responsive in electricity due to ACDD and less responsive in natural gas due to AHDD. These results are available upon request.

mies, $Year_t$ is a set of year dummies, and μ_i is a residence fixed effect. The model provides estimates of how temperature variation affects monthly energy consumption in the BCCRs (the β 's) and how the estimates differ in the ACCRs (the δ 's).

Two other specifications provides alternative estimates of the difference in differences—that is, the δ 's—but control for the direct weather effects more flexibly. One specification is

$$Y_{it} = \delta CodeChange_i \times [ACDD_t, AHDD_t] + v_t + \mu_i + \varepsilon_{it} \tag{6}$$

where v_t represents month-year fixed effects. Note that the weather variables do not enter on their own because with one weather station, they are not separately identified from the month-year fixed effects. The other specification is the same except for the inclusion of zip code by month-year fixed effects,, as specified in previous models.

Table 3 reports the results of all three models using all of the data through 2014. The results of specification (5), reported in the first column, show the unsurprising result that greater ACDD and AHDD in a month increases energy consumption. This is reflected in the positive and statistically significant estimates of β_{ACDD} and

Table 3. Difference-In-Differences Estimates of Energy Consumption due to Weather Variability, 2004–14

	mmBtu		
	(1)	(2)	(3)
Code change \times ACDD (δ_{ACDD})	-.014*** (.004)	-.013*** (.004)	-.009** (.004)
Code change \times AHDD (δ_{AHDD})	-.118*** (.011)	-.112*** (.011)	-.093*** (.010)
ACDD (β_{ACDD})	.115*** (.004)
AHDD (δ_{AHDD})	.403*** (.007)
Month dummies	Yes	No	No
Year dummies	Yes	No	No
Month-year dummies	No	Yes	No
Zip code \times month-year dummies	No	No	Yes
Residence fixed effects	Yes	Yes	Yes
R-squared (within)	.273	.286	.331
Observations	254,467	254,467	245,467

Note. Standard errors are clustered at the residence level.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

β_{AHDD} , as well as the linear combinations $\beta_{ACDD} + \delta_{ACDD}$ and $\beta_{AHDD} + \delta_{AHDD}$. The results also show that the before- and after-code-change residences respond differently, as indicated by the negative and statistically significant estimates of δ_{ACDD} and δ_{AHDD} . Specifically, the ACCRs increase energy consumption by less than the BCCRs, and this is consistent with the ACCRs being more energy efficient. The same results hold with similar magnitudes and statistical significance for the more flexible models in columns 2 and 3.¹⁹

To get a sense for the magnitude of how much BCCRs and ACCRs differ in their responsiveness, it is useful to focus on the results in column 1. Specifically, the ratios of coefficients $\delta_{ACDD}/\beta_{ACDD} = -0.123$ and $\delta_{AHDD}/\beta_{AHDD} = -0.292$ have useful interpretations.²⁰ The first implies that a unit increase in the monthly ACDD causes the ACCRs to increase energy consumption by 12% less on average than the BCCRs. The second implies that a unit increase in the monthly AHDD causes the ACCRs to increase energy consumption by 29% less on average than the BCCRs. These magnitudes suggest that the building code change had a substantial effect on tempering energy consumption in response to more extreme temperatures.

Because we have seen important differences in how the building code affects energy consumption over time, it is worth considering how the estimated differences in responsiveness to weather shocks might change over time. The final model that I estimate expands on specification (5) as follows:

$$Y_{it} = \delta \text{CodeChange}_i \times [ACDD_t, AHDD_t] \times \text{Year}_t + \beta [ACDD_t, AHDD_t] \times \text{Year}_t + \text{Month}_t + \text{Year}_t + \mu_i + \varepsilon_{it} \quad (7)$$

where Year_t is interacted with the weather variables to estimate the coefficients of interest separately for each year. Then, after estimating the model, I derive the ratio δ_{kt}/β_{kt} for $k = ACDD, AHDD$ and $t = 2004 \dots 2014$ along with the 95% confidence intervals (see n. 20 for details) and report the results graphically in figure 6.

19. Studying California residences over decades, Chong (2012) and Levinson (2014) find evidence that more recently constructed houses subject to more stringent building codes increase electricity consumption more in response to hotter weather. Chong (2012) explains his results for Riverside, California, as increased air conditioning ownership having outweighed other energy-saving impacts. While Levinson (2014) is able to control for air conditioning ownership over a broader geographic area, the same explanation could underlie his results because his models control for average air conditioning effects without allowing electricity responsiveness to vary differently among residences with and without air conditioning.

20. These ratios are nonlinear combinations of two coefficients. Test statistics are derived using the delta method, and both ratios are statistically different from zero at the 99% level. The 95% confidence intervals for the ACDD and AHDD ratios are $(-0.186, -0.059)$ and $(-0.340, -0.243)$, respectively.

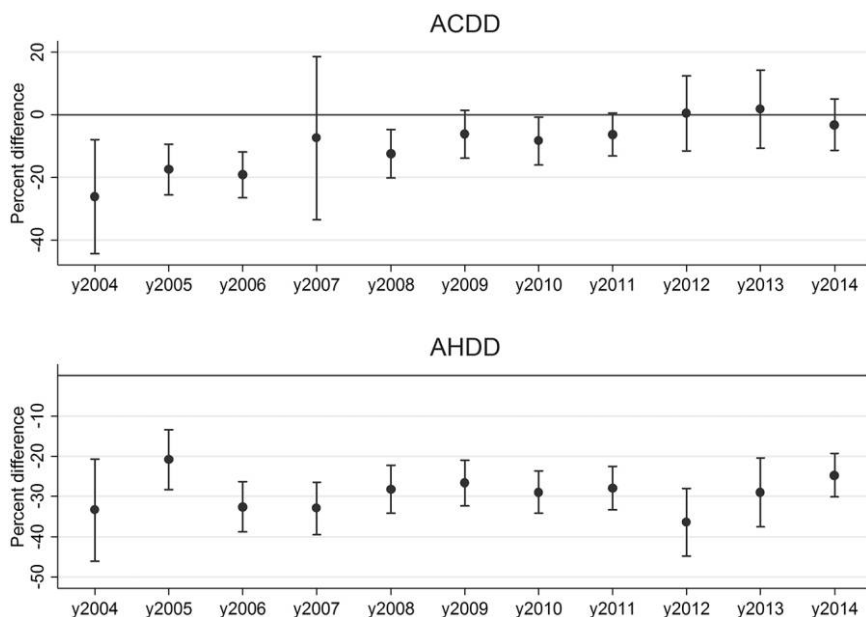


Figure 6. Annual difference in differences in energy consumption responsiveness between before- and after-code-change residences due to weather variability, with 95% confidence intervals.

Despite the very different empirical strategy from that in section 2, there is a now familiar pattern to the results. Soon after the building code change, ACCRs increase energy consumption significantly less in response to more ACDD, and this suggests greater efficiency with air conditioning.²¹ But the difference between before- and after-code-change residences appears to dissipate over time, until there is no evidence of an energy code affect about 8 years later. This pattern is clearly consistent with a relative newness effect in ACCRs that is not enduring. In contrast, the difference in responsiveness to AHDD appears roughly constant over the 11 years and is always statistically different from zero, with point estimates ranging between 20% and 30%. This result suggests that the energy code had real effects on the efficiency of residences for heating, and this tracks the previous findings for natural gas.

4. CONCLUSION

This paper considers the question of whether building energy codes actually save energy. Using more than a decade’s worth of billing data for residences in Gainesville,

21. The large confidence interval for the point estimate in 2007 is most likely the result of having fewer observations to estimate an effect for that year. Recall the missing data mentioned in n. 7.

Florida, the answer is yes, but the results differ for electricity and natural gas. Despite what appears to be an initial code change effect that reduces electricity consumption, the difference between before- and after-code-change residences disappears after a few years. In contrast, the code change has a significant and enduring effect on natural gas consumption, causing a reduction of more than 10%. Both the electricity and natural gas results are consistent with the way before- and after-code-change residences respond to weather shocks. In particular, ACCRs increase natural gas consumption by nearly 30% less than BCCRs in response to marginally colder weather.

Comparison of results from the present paper to those in Jacobsen and Kotchen's original study yields two important methodological insights. First, future studies that take advantage of the discontinuity design of comparing before- and after-code-change residences should wait several years after the code change before using billing data for analysis. Second, evaluation of building codes should not focus exclusively on electricity consumption, which has been the case in most previous studies. Energy codes seek to improve the efficiency of space heating, space cooling, and water heating. While natural gas is used primarily for these purposes, a growing share of residential electricity consumption is for other uses, including appliances, electronics, and televisions. There is also the potential for substitution among energy sources for home heating and many appliances (e.g., clothes dryers and hot water heaters). Unfortunately data on energy sources for different uses are not available for this study, but the estimates of combined energy consumption seek to account for such substitution. In addition to estimating the overall energy effect, I consider the effect on energy expenditures and CO₂ emissions, which are effectively different weighting approaches that place more emphasis on electricity than natural gas. While there is statistically significant evidence of overall energy savings, the evidence is weaker for savings on utility bills and reducing CO₂ emissions.

Finally, the results of this study have broader implications for the evaluation of energy efficiency policies. While a growing number of studies find that engineering forecasts significantly overestimate realized savings of efficiency investments, this does not appear to be the case with Florida's building code. Forecasts predicted that the 2002 code change would translate into a 2% increase in residential energy efficiency. The revised estimate here is very close to the forecast: a 2.9% savings in overall energy use. Nevertheless, this savings does not necessarily imply that building codes are an efficient or even desirable policy. Following the same approach outlined by Jacobsen and Kotchen, the revised estimates imply social and private payback rates of about 10 and 16 years (up from 4 to 6), respectively. Whether Florida homeowners would find this private payback rate desirable, and how the social net benefits from building codes might compare to other policy instruments to promote energy efficiency are open and important questions for future research.

REFERENCES

- Allcott, Hunt, and Michael Greenstone. 2012. Is there an energy efficiency gap? *Journal of Economic Perspectives* 6:3–28.
- Aroonruengsawat, Anin, Maximilian Auffhammer, and Alan Sanstad. 2012. The impact of state level building codes on residential electricity consumption. *Energy Journal* 33:31–52.
- Chong, Howard. 2012. Building vintage and electricity use: Old homes use less electricity in hot weather. *European Economic Review* 56:906–30.
- Costa, Dora L., and Matthew E. Kahn. 2010. Why has California's residential electricity consumption been so flat since the 1980s? A microeconomic approach. NBER Working paper 15978, National Bureau of Economic Research, Cambridge, MA.
- Dubin, Jeffrey, Allen K. Meidema, and Ram V. Chandran. 1986. Price effects of energy-efficient technologies: A study of residential demand for heating and cooling. *RAND Journal of Economics* 17:310–25.
- European Union. 2012. Directive 2012/27/EU of the European Parliament and of the Council. *Official Journal of the European Union* 55:1–56.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2015. Do energy efficiency investments deliver? Evidence from the weatherization assistance program. NBER Working paper 21331, National Bureau of Economic Research, Cambridge, MA.
- Gerarden, Todd D., Richard G. Newell, Robert N. Stavins, and Robert C. Stowe. 2015. An assessment of the energy-efficiency gap and its implications for climate-change policy. NBER Working paper 20905, National Bureau of Economic Research, Cambridge, MA.
- Gillingham, Kenneth, and Karen Palmer. 2014. Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Review of Environmental Economics and Policy* 8:18–38.
- Graff Zivin, Josh S., Matthew J. Kotchen, and Erin T. Mansur. 2014. Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies. *Journal of Economic Behavior and Organization* 107:248–68.
- Horowitz, Marvin J., and Hossein Haeri. 1990. Economic efficiency v. energy efficiency: Do model conservation standards make good sense? *Energy Economics* 12:122–31.
- Jacobsen, Grant D., and Matthew J. Kotchen. 2013. Are building codes effective at saving energy? Evidence from residential billing data in Florida. *Review of Economics and Statistics* 95:34–49.
- Jaffe, Adam B., and Robert N. Stavins. 1995. Dynamic incentives in environmental regulations: The effects of alternative policy instruments on technology diffusion. *Journal of Environmental Economics and Management* 29:43–63.
- Levinson, Arik. 2014. How much energy do building energy codes really save? Evidence from California. NBER Working paper 20797, National Bureau of Economic Research, Cambridge, MA.
- Metcalfe, Gilbert E., and Kevin A. Hassett. 1999. Measuring the energy savings from home improvement investments: Evidence from monthly billing data. *Review of Economics and Statistics* 81:516–28.
- US DOE (US Department of Energy). 2011. Building energy codes resource guide for policy makers. Report PNNL-SA-81023.