

The Distributional Effects of Building Energy Codes*

Chris Bruegge, Tatyana Deryugina, and Erica Myers†

Abstract. We evaluate the distributional consequences of building energy codes for home characteristics, energy use, and home value. We exploit spatial variation in California’s code strictness created by building climate zones, combined with information on over 350,000 homes located within 3 kilometers of climate zone borders. Our key findings are that stricter codes create a non-trivial reduction in homes’ square footage and the number of bedrooms at the lower end of the income distribution. On a per-dwelling basis, we observe energy use reductions only in the second lowest income quintile, driven by decreases in square footage. Energy use per square foot actually *increases* in the bottom quintile. Home values of lower-income households fall, while those of high-income households rise, suggesting that building energy codes result in more undesirable distortions for lower-income households.

* We thank the editor, two anonymous referees, Matt Kotchen and participants in the Energy Policy Tradeoffs between Economic Efficiency and Distributional Equity conference for their helpful comments.

† Bruegge: Stanford University, cbruegge@stanford.edu. Deryugina: University of Illinois and NBER, deryugin@illinois.edu. Myers: University of Illinois and E2e, ecmyers@illinois.edu.

I. Introduction

Building energy codes that set minimum efficiency requirements for new construction are widely employed throughout the United States. Typically, the imposition of a building code is motivated by either an externality (e.g., an individual will not take into account their neighbor's house catching fire when choosing the level of fire safety), informational barriers (e.g., it is difficult to observe how sturdy a building is), or other market failures such as inattention on behalf of home buyers. Energy building codes have in part been justified by the significant environmental externalities associated with energy use. However, no state currently taxes building energy use, which is a more cost-effective way to reduce energy consumption in the absence of informational or landlord-tenant market failures.

The adoption of standards such as building energy codes may also be motivated by distributional concerns (Ito and Sallee 2018). That is, while building energy codes are less cost-effective, they may achieve a more preferable distributional outcome than would be feasible with energy taxes. This could be the case if the value of building codes, net of their costs, represented a larger fraction of household income in poor areas than in wealthy areas. Codes could cause even higher absolute levels of savings in low-income areas if standards are not binding for homes occupied by wealthier individuals, although such savings may not be welfare-improving if they are costly. These conjectures have not been tested empirically, and furthermore there is even debate about the extent to which building energy codes reduce energy use at all (Jacobsen and Kotchen 2013; Levinson 2016; Kotchen 2017; Novan et al. 2017).

We adopt a novel approach to studying both the effectiveness and distributional properties of building energy standards by exploiting spatial and temporal variation in the stringency of California's codes. California has 16 distinct climate zones, and climate zones whose "typical meteorological year" is more extreme generally face more stringent energy code requirements. Thus, homes that are built in the same geographic area and at the same time but on opposite sides of a climate zone border will face different levels of code stringency. Moreover, the introduction of climate zones in 1982 provides a second source of variation that allows us include granular spatial controls, eliminating time-invariant characteristics that might differ across the boundary.

We obtain account-level electricity and natural gas billing data for the years 2009-2015 from four California utilities, which together serve the vast majority of homes in the state, and combine these with data on home characteristics, occupant characteristics, and home values from ReferenceUSA (RefUSA). We identify about 354,000 homes that are within 3 kilometers of a climate zone border and use them to study the impact of building energy codes on home characteristics, energy use, and home values. We adopt a difference-in-differences approach, comparing cross-border differences among homes built prior to the introduction of energy building codes (1947-1977) to differences among homes built after climate zones were introduced (1982-2006).³ Our identification assumption is that, absent building energy code climate zones, differences in energy use for post-1981 homes on either side of a climate zone border would have been statistically indistinguishable from differences in energy use for pre-1977 homes. We are able to provide some empirical evidence for this assumption by examining trends in energy use for pre-1977 homes on different sides of post-1981 climate zone borders. Importantly, our approach allows us to flexibly control for changes in building practices over time with vintage fixed effects, something previous studies have been unable to do. To elucidate the distributional effects of building energy codes, we estimate how the effect of building codes on energy use and housing prices varies by income quintile.

We find that stricter energy codes reduce the square footage and the number of bedrooms in homes occupied by households in the bottom two income quintiles by 4-6%, a channel through which the codes might reduce total energy consumption. The top three quintiles also show a small (0.6-1.6%) reduction in the number of bedrooms and the third and fourth quintile show a similarly small decline in square footage (0.6-0.7%). These results suggest that the codes create more undesirable distortions to attributes in low-income homes. To our knowledge, we are the first to estimate how builders respond to building energy codes along the dimension of home attributes. Our finding that building codes alter home characteristics also highlights the difficulty of only using intertemporal variation for identification in this and similar settings, especially over longer time periods. If there are secular trends in building characteristics and a possibility that building energy codes affect these characteristics, both of which are very likely, then it is impossible to accurately estimate the effects of building codes, either net of changes in

³ Between 1977 and 1982, building requirements varied slightly according to local heating and cooling degree days. However, the modern climate zones did not yet exist.

characteristics or not. If the researcher controls for home characteristics, the estimated effects will not represent the causal effect of building codes because the codes affect attributes;⁴ and if the researcher does not control for home characteristics, secular trends will obfuscate the true effects of building codes. Thus, cross-sectional variation is essential for proper identification.

When we consider energy usage, we find that, for the average border in our sample, building energy codes do not significantly (in a statistical sense) reduce the total or per-square foot natural gas consumption. However, electricity usage falls by 1% on a per-house basis and by 0.6% on a per-square foot basis. These averages are an order of magnitude smaller than a back-of-the-envelope engineering calculation of how cross-border differences should translate into energy use reductions. While our data do not allow us to test for mechanisms that may be behind these shortfalls, the results are consistent with recent work showing that engineering estimates sometimes overstate real-world gains, for example because the engineering estimates are biased, work is not carried out properly, or because individual behavior does not align with engineering assumptions (e.g., Davis et al. 2014; Hanna et al. 2016; Allcott and Greenstone 2017; Fowlie et al. 2018). The latter could happen if households respond to energy efficiency improvements by increasing their energy use – the so-called “rebound effect” – or take other actions that offset the energy reductions that would have otherwise materialized.

The average treatment effect masks substantial heterogeneity across income groups. Specifically, we find that households in the second quintile of the income distribution experience the largest decreases in total energy use (about 4%). Our estimates for the bottom quintile are very imprecise, and we cannot rule out that they experience a comparable or even slightly larger decline in total energy use. For the other three quintiles, however, we can rule out energy reductions of 1.5% or more with 95% confidence. Considering gas and electricity separately, we find that the second income quintile experiences similar decreases in both (4.5-4.7%), with estimates for the bottom quintile again being statistically imprecise. For the top three quintiles, we see significant reductions in electricity use of 1.1-1.2%, but no meaningful changes in natural gas use. On a per-square-foot basis, we only see (small) reductions in electricity use in the top income quintile, and no significant decline in natural gas or total energy use anywhere in the income distribution. Energy use per square foot in the bottom quintile actually increases,

⁴ Attributes are a “bad control” in the sense of Angrist and Pischke (2009). Because codes affect attributes, holding the attributes constant in a regression equation will not allow the researcher to capture the causal effect of variation in code stringency. Angrist and Pischke explain that this bias could be either positive or negative.

possibly because some energy demand (e.g., appliances) is independent of square footage, and square footage in these homes falls. Although our empirical strategy uses different variation in building energy codes than the existing literature, our conclusion about their effects on energy use is qualitatively similar: building energy codes lead to some energy use reductions, but *ex post* savings fall substantially short of *ex ante* predictions.

Using the same methodology, we also estimate the degree to which the energy code differences are capitalized into housing prices. Because of our difference-in-differences approach, these estimates will be purged of price differences due to construction costs, land scarcity, and other local factors that could affect the value of an existing home. Thus, any estimated differences should be driven by differences in the attributes of homes subject to stricter building energy codes, including their energy efficiency. We find that housing prices increase by about 2%, on average and on a per-square-foot basis. The increases are concentrated in the top two income quintiles, and are too large to be explained by the net present value of measured electricity savings, suggesting that building energy codes also lead to changes in unobservable (to us) characteristics that higher-income households value. By contrast, housing prices in the bottom two quintiles fall by 8-12%; about half of this effect can be explained by the decline in square footage. Because the per-square-foot price changes are not accompanied by per-square-foot changes in energy use, our interpretation is that there are also changes in home characteristics other than square footage, such as the number of bedrooms and other attributes that are unobservable to us.

Our paper makes contributions along several dimensions. First, our identification strategy is unique in this literature. Part of the difficulty in evaluating the effectiveness of building energy codes is that most of the existing empirical studies focus on changes in building codes over time (Jacobsen and Kotchen 2013; Levinson 2016; Kotchen 2017). As such, they lack a comparison group of homes that were built at same time but do not face the same building energy code. One exception is Aroonruengsawat et al. (2012), who use state-level panel data from the US and find that building energy codes reduce electricity use by as much as 5%. Because their estimation strategy is based largely on differential timing in the introduction of energy codes across states, the accuracy of their results depends crucially on their ability to control for potential confounding factors, such as heterogeneous trends and non-linear relationships between weather and electricity use. By contrast, our estimates rely on a much weaker identification assumption:

we only require parallel trends near climate zone borders. In some cases, these borders even bisect cities, lending more credence to our identification assumption. We are also able to consider natural gas use, something Aroonruengsawat et al. (2012) do not observe. Unlike any of the existing studies, we also examine capitalization, distributional consequences, and whether building codes affect home characteristics, such as square footage. Our findings that home characteristics respond to building energy codes also highlight the importance of using cross-sectional variation for identification.

We also contribute to the growing literature on the distributional effects of energy policy.⁵ To date, this research has focused on carbon taxes, gasoline taxes, and fuel economy standards, and there is virtually no research on the distributional consequences of building energy codes. However, building energy codes are a very pervasive type of energy policy; only eight relatively sparsely populated states do not have statewide energy codes for residential buildings (DOE 2017). Thus, their distributional consequences are worth studying. We show that lower-income households suffer the largest distortions to square footage and the number of bedrooms and a decline in home value, while higher-income households experience very small distortions in these characteristics and an overall increase in home value. These findings suggest that builders use different approaches to comply with building energy codes in different parts of the income distribution. Ito and Sallee (2018) show that standards which distort attributes may improve efficiency under some conditions, but that similar benefits can be obtained with a compliance trading scheme that does not result in attribute distortion. They speculate that attribute-based regulation may be ultimately motivated by distributional considerations. However, while the lower-income households in our sample obtain the largest monetary savings from building energy codes, this reduction is brought about by builders reducing square footage. Absent other market failures, such as asymmetric information, building energy codes are thus unlikely to disproportionately benefit lower-income homeowners.

The rest of the paper is organized as follows. In Section II, we provide background on energy building codes in California and discuss the ways in which energy building codes might affect building attributes and energy use. In Section III, we outline our estimation strategy and the data we use. Section IV presents our results, and Section V discusses and concludes.

⁵ e.g., West 2004; West and Williams 2004; Hassett et al. 2009; Jacobsen 2013; Williams et al. 2014; Borenstein and Davis 2016. See Bento 2013 for a review.

II. Background

A. California building energy codes

California was one of many states to implement a statewide energy efficiency code in the 1970s (Aroonruengsawat et al. 2012). It officially adopted such a code in 1978, although some building energy standards that were adopted in the code began to be enforced in late December of 1976 (CEC 1978). Because we focus on energy use in single-family homes, we restrict the discussion below only to requirements that apply to such dwellings.

In 1982, California introduced 16 “climate zones” (see Figure 1) and specified slightly different energy efficiency requirements for each one.⁶ The stringency of the code in each climate zone is determined by the zone’s “typical meteorological year”. Originally, these were developed from weather data in a “representative city” in each climate zone.⁷ In later years, other weather stations were added to more accurately capture the typical temperature distribution faced by population centers in each climate zone. In general, zones with more moderate climates face fewer restrictions than zones with more extreme climates.

Builders in each climate zone can comply with California’s building energy codes in one of two ways. First, they can utilize the so-called “performance method” by demonstrating that the proposed building is expected to use less energy than the maximum allowed for a “standard” home of the same size in that climate zone. Because the performance method offers a lot of flexibility, it is used by the vast majority of new home builders to demonstrate compliance.⁸ Alternatively, builders can adopt a prescriptive set of requirements for how the building must be built (called “alternative packages”). Although this second method is not frequently used, the maximum allowed energy use under the performance method is determined by the projected

⁶ In 1995, the California Energy Commission conducted a review of the climate zones and changed the classification of several cities (California Energy Commission 1995). This re-classification affects relatively few homes in our sample and we omit them from the analysis. In principle, these changes create the perfect natural experiment because the California Energy Commission points out that the reclassification was due to mistakes in the original boundaries. Unfortunately there are not enough homes affected by the change to detect an effect on energy usage, so we omit these homes from our sample rather than try to exploit the change for identification.

⁷ See http://www.energy.ca.gov/maps/renewable/building_climate_zones.html. We further discuss the typical meteorological year data in Section III.

⁸ We contacted plans examiners in three jurisdictions across the state of California (Ryan Pursley from the City of Concord Building Office, Joe Espinosa from the Palo Alto Building Office, and Leslie Edwards from the Kern County Building Office) and Michael Kunz from Title 24 Express, a company that helps builders certify energy building code compliance. All of them confirmed that new homes almost exclusively utilize the performance method for compliance. Moreover, several building energy code documents themselves state that the majority of builders use the performance method.

energy usage of a home that meets the prescriptive standards. Therefore, the alternative packages determine the building code strictness, so we briefly discuss them first.

Alternative packages offer builders several possible combinations of minimum requirements for key energy-relevant home characteristics, including (a) the amount of ceiling, wall, and other types of insulation, (b) glazing (i.e., windows and glass doors) energy efficiency, (c) the share of a home's wall area that can consist of glazing, and (d) in some climate zones, shading minimums. The 1982 building codes contained three such packages for each climate zone (A, B, and C); starting in 1983, there have been five (A-E). Alternative packages thus offer much less choice and flexibility than the performance method, which allows builders to employ a lot more possible combinations of energy efficiency characteristics.

The performance method has evolved slightly since it was first introduced. In 1982-1983, the maximum allowed energy use was a simple per-square-foot "energy budget", expressed in thousands of BTUs per square foot of conditioned space per year for space conditioning and thousands of BTUs per dwelling unit per year for water heating. There were separate energy budgets for heating and cooling and for each climate zone (see Online Appendix Table A1). Whether or not a particular building met the required energy budget was determined by software that projects energy use as a function of building characteristics and location.

In later building code vintages, simulations continued to be the basis for determining a proposed home's energy use for compliance using the performance method. However, the per-square-foot energy budget ceased to be constant, but instead became determined by a "standard design". The standard design is a home that has the same square footage as the proposed home and has the prescriptive requirements of a particular "alternative package". For example, in building energy code vintages 1995 and later, the characteristics of alternative package D are used to establish the performance standards for each climate zone. Thus, if package D in one climate zone contains characteristics that translate into lower energy use than package D in another climate zone, then a builder will have to utilize more energy-efficiency measures in the former compared to the latter, all else equal.

Building energy codes similar to California's can be thought of as changing the marginal costs of different attributes (see Section A in the Online Appendix for a simple conceptual framework). For example, a larger house may become costlier to build in the presence of stricter building energy codes, as it requires more additional insulation material than a smaller house.

Alternatively, some attributes may become implicitly cheaper if they relax an energy use constraint imposed by building energy codes. Identifying and describing all the assumptions of the energy modeling software that may penalize or reward certain building attributes is beyond the scope of this paper. The most obvious incentive to distort attributes is the one we allude to above: building codes increase the marginal cost of each square foot by requiring better insulation. Of course, under the performance standard, builders are not necessarily required to install the insulation levels described in the prescriptive packages. But because the reference home is assumed to have that level of insulation, producing equivalent energy reductions through other means will on average be more difficult for a larger house.

In California, there is an additional implicit size penalty because reference homes used in compliance calculations have square footprints and windows that are evenly distributed across all four walls. This modeling assumption makes it increasingly difficult for large homes to meet the standards, because energy use due to heat loss is proportional to surface area, not floor area, and holding floor area fixed, more rectangular homes will have a larger surface area.⁹ Thus, for homes without square footprints, the difficulty of meeting the performance standards increases with home size.

Because the stringency of California’s code thus depends on home size, California’s building energy code is an example of “attributed-based” regulation, where requirements for one characteristic (e.g., energy efficiency) depend on the level of a “secondary attribute” (e.g., home size). Ito and Saltee (2018) develop and analyze a general theoretical framework of such regulation. They show that attribute basing generally leads to a distortion in the attribute on which the policy is based. Such a distortion can be welfare-improving (relative to a non-attribute-based standard) if it helps equalize the marginal costs of compliance across goods or achieves distributional outcomes that would otherwise be unattainable. However, attribute-basing could also reduce the effectiveness of a standard if it encourages larger homes that use more energy.

⁹ To see this, let $SA(\text{Square}, A)$ be the wall surface area of a home with a square footprint and square footage A . Define $SA(\text{Rectangle}, A, c)$ to be the wall surface area of a home with a rectangular footprint, with square footage A , and whose length is c times its width. Note that $SA(\text{Square}, A)/SA(\text{Rectangle}, A) = 4\sqrt{c}/(2 + 2c) \in (0, 1]$. The limit of this fraction is 0 as c approaches 0 or infinity, and the ratio reaches its maximum value of 1 when $c = 1$, i.e. the rectangular home is actually a square. Thus, non-square homes will always be penalized in this dimension of the performance standard. To the extent that $(A, c) > 0$, larger homes will be penalized more than smaller homes. The assumption that $corr(A, c) > 0$ would be true under many circumstances, including the desire to have natural window light available in all rooms of the house.

To apply the Ito and Sallee intuition to building codes, notice that the features of California’s code discussed above implicitly penalize size, which leads to a further reduction in energy use. The opposite might be true depending on the details of any particular code. For example, if larger homes have lower compliance costs on a per square-foot basis because they can take advantage of economies of scale in making fixed investments like efficient HVAC systems, then building energy codes may create incentives to build larger homes. Overall, California’s building energy codes are similar in structure to that of other states, 49 of which use some version of the International Energy Conservation Code (IECC), statewide or locally (ICC 2017).¹⁰ Like California, the IECC also has a set of prescriptive standards that builders can choose to follow and a performance standard whereby builders can use software to demonstrate that the proposed home is expected to use less energy than a standard design under the prevailing environmental conditions. Thus, most of our conclusions about the incentives created by building energy codes are also applicable outside of California.

B. Building energy code differences across climate zones

Our identifying variation comes from the fact that the stringency of California’s code depends on a home’s climate zone. Therefore, prescriptive standards (and consequently the maximum energy consumption set by the performance standard) vary across space. All else equal, the amount of heat a building gains or loses depends on the difference between indoor and outdoor temperatures. Thus, insulation and other energy efficiency measures matter more in more extreme climates, and alternative packages in more extreme climates – as measured by each zone’s typical temperatures – have stricter requirements.

Unfortunately, comparison of climates of neighboring zones is not always straightforward, as one zone may have more heating degree days (defined as the annual sum of degrees under 65F) while the other has more cooling degree days (defined as the annual sum of degrees over 65F). This leads to cases in which one side of a border has more stringent insulation requirements but the other side requires better window shading, for example. We thus adopt a statistical approach, asking how neighboring zones’ building code strictness is related to their archetypal climate differences *on average*. In Online Appendix Section B and Online Appendix Table A2, we

¹⁰ See U.S. Department of Energy’s website (<https://www.energycodes.gov/adoption/state-code-adoption-tracking-analysis>) for a list of states’ current energy building codes.

demonstrate that differences in bordering zones' typical meteorological year temperatures are strongly related to differences in their building code strictness, as measured by the prescriptive packages. That is, larger climate differences between neighboring zones, as measured by heating and cooling degree days, are associated with stricter building requirements imposed on the zone with the more extreme climate. We also perform a back-of-the-envelope calculation and demonstrate that a home with the more energy-efficient insulation characteristics (as described by the prescriptive packages) will use about 12-15% less energy than the less energy-efficient home, not counting appliances. Because the performance packages are based on the prescriptive packages, similar conclusions should apply to the former. Thus, the cross-border differences in building code requirements that we identify are non-trivial.

The extent to which building energy codes affect energy consumption across climate zones also depends on how binding they are, i.e., how different actual homes are from homes that would have been built in absence of such codes, in terms of insulation, HVAC efficiency, and so on. This counterfactual is not something we can observe directly, but the less binding building energy codes are, the lower the observed difference in energy use across climate zone borders will be. Thus, our estimates of building energy codes' effects on energy use is partly a test of how binding such codes are.

III. Empirical strategy

A. Regression approach

It may seem natural to use neighboring zones' differences in prescriptive standards to see how building code differences translate into differences in energy use. However, prescriptive standards are multidimensional, and it is not straightforward to combine their features into a single "strictness" variable. Instead of adopting an *ad hoc* formula, we use neighboring zones' archetypal climate differences as independent variables that proxy for and are strongly correlated with building code strictness. It is important to remember that these climate differences are those of zones' representative cities, not differences in climates just across the border. Climate zone boundaries were drawn to capture areas whose climates are similar to a representative city in each zone.¹¹ Although in some cases boundaries follow mountain ridges or valley streams, in

¹¹ See http://www.energy.ca.gov/maps/renewable/building_climate_zones.html for a list of these cities. Since 2013, climate zones have been enforced at the zip code level and in many cases do not cross city boundaries.

general the locations of the borders were determined for administrative convenience rather than in response to abrupt changes in topology. Thus, climates around these borders should be similar. We are unable to empirically verify this similarity using weather data because the number of weather stations that are sufficiently close to climate zone borders and to each other is too small.¹² However, in Online Appendix Figure A1, we show that the per-square-foot energy consumption of older (1947-1977) homes on different sides of climate zone borders is statistically indistinguishable across different calendar months, relative to January, strongly suggesting that the climate is very similar across the borders in our sample.

It is possible, however, that there are important differences in neighborhood, demographic, and other characteristics that change discretely at zip code or city boundaries, which often – but not always – coincide with climate zone borders. To control for any remaining time-invariant dissimilarities, including any small differences in climate, we adopt a difference-in-differences approach rather than simply comparing post-1982 homes on different sides of a climate zone border.¹³ Specifically, we compare cross-border differences between homes built in 1947-1977 to cross-border differences between homes built in 1982-2006 by estimating the following equation:

$$\begin{aligned} \ln(Y_{imy}) = & \beta_1 Post_i * DiffCDD_i + \beta_2 Post_i * DiffHDD_i \\ & + \beta_3 DiffCDD_i + \beta_4 DiffHDD_i + \alpha_{bv} + \alpha_{bmy} + \alpha_z + \mathbf{L}'_i \boldsymbol{\alpha}_b + \varepsilon_i, \end{aligned} \quad (1)$$

where Y_{imy} is the electricity use, natural gas use, or total energy use in dwelling i , month m , and year y . The variable $DiffCDD_i$ is the difference in cooling degree days between the typical meteorological year in the climate zone of dwelling i and that of the nearest climate zone. $DiffHDD_i$ is the corresponding difference in heating degree days. The variable $Post_i$ is an

¹² Specifically, of 1,105 weather stations operating in California at some point between 1950 and 2016, only 47 are located (a) within 5 kilometers of a climate zone border and (b) within 5 kilometers of another weather station. However, for 28 of these stations the nearby station is not on the other side of the border, leaving us with only 19 qualifying stations.

¹³ Alternatively, one could use the post-1982 homes and the spatial variation in code stringency induced by the climate zones. This would require climate zone boundaries to be unrelated to any unobserved factors that also influence energy consumption, home attributes, or occupant preferences. While it appears that some climate zone borders are determined arbitrarily by freeways or arbitrary jurisdictional boundaries, we have tested for smoothness in demographics and home attributes in neighboring climate zones prior to the introduction of the codes and determined that only a small number of borders appear to display pre-code continuity in the distribution of these variables. Thus, we prefer our weaker identification assumption that permits us to use a larger, more representative sample of California households and provides results that are more likely to be externally valid than a regression-discontinuity type analysis of a small number of borders.

indicator for the dwelling being built in 1982 or later. We control for the fact that there may be a secular trend in the energy efficiency of new homes with vintage-by-border (year-built-by-border) fixed effects, α_{bv} . To control for seasonality common to newer and older homes in a particular area, we include border-by-month-of-sample fixed effects, α_{bmy} . Finally, we also include zip code and border-specific linear latitude and longitude controls ($\mathbf{L}'_i\alpha_b$). Standard errors are two-way clustered by dwelling and by month-of-sample. We also estimate the effects of building energy codes on time-invariant home characteristics (square footage, number of bedrooms, price, and price per square foot). In these specifications, we omit the border-by-month-of-sample fixed effects and use robust standard errors, as we only have one observation per home.

To simplify interpretation, we report linear combinations of β_1 and β_2 at the average cooling and heating degree day difference for treated homes. To calculate the difference in degree days at this “average” border, we compute the weighted mean of heating and cooling degree day differences, using zones that have more cooling degree days than their neighbor(s) as the reference.¹⁴ We weight by the number of “treated” homes built in or after 1982, when climate zones are in place. The average difference in heating degree days is -53.5, and the average difference in cooling degree days is 990.9. This approach is not the only reasonable way of combining the coefficients, but it has the advantage of capturing any correlation between heating and cooling degree-day differences as well as the distribution of homes in our sample.

To estimate heterogeneity in the impact of building energy codes by income, we allow β_1 and β_2 to vary by the household’s estimated income quintile:

$$\begin{aligned} \ln(Y_{imy}) = & \sum_{q=1}^5 \beta_{q1}(Post_i * DiffCDD_i * 1[QInc_i = q]) & (2) \\ & + \sum_{q=1}^5 \beta_{q2}(Post_i * DiffHDD_i * 1[QInc_i = q]) \\ & + \beta_3 DiffCDD_i + \beta_4 DiffHDD_i \\ & + \alpha_{qb} + \alpha_{bv} + \alpha_{bmy} + \alpha_z + \mathbf{L}'_i\alpha_b + \varepsilon_i, \end{aligned}$$

¹⁴ If we averaged without a reference point, the average would primarily be driven by differences in the number of homes on either side of the border. This is because climate differences are symmetric across the zone border: if homes on one side have $diffCDD = X$, the homes on the other side mechanically have $diffCDD = -X$.

where $1[QInc_i = q]$ is an indicator equal to 1 if household i 's estimated income quintile is q . Quintiles are defined using the 2010 American Community Survey microdata for California. In addition to the controls in equation (1), we add income-quintile-by-border fixed effects (α_{qb}), although our results are very similar if we only retain the original controls. We have also replicated our estimation with income deciles instead of income quintiles, concluding that there is no meaningful advantage to further subdividing income groups in this way.

Our identifying assumption is that, conditional on the fixed effects above, differences between cross-border homes that are *not* driven by energy building codes are on average the same for pre-1977 and post-1981 homes. Because we are able to difference out any neighborhood amenities and other time-invariant factors with our difference-in-differences specification, the main threat to our identification strategy is systematically different development that occurs on either side of the border and that is unrelated to building energy codes. For example, if new development on one side of the border consists of larger or more expensive homes and new development on the other side of the border is smaller or lower-income, we would attribute differences in energy consumption and home attributes to the building codes when in fact these differences are due to differential location-specific trends. We assess whether such a situation is likely in Section IV.

B. Data and estimation sample

Home characteristics, home prices, and occupant demographics. We obtain housing characteristics data from ReferenceUSA (RefUSA). The dataset contains detailed characteristics for over 6 million single-family homes, including the exact premise address, latitude/longitude coordinates, when the home was built, its square footage and number of bedrooms, and the estimated home value (based on assessor data).

RefUSA also provides an estimate of the socioeconomic characteristics of each dwelling's occupants. The data include basic demographic characteristics, such as the age of the head of the household, whether the head of the household is married, and the estimated number of children residing in the dwelling. Importantly, RefUSA also contains a measure of the household's income, as inferred from various sources, including census block characteristics,

credit card purchases, magazine subscriptions, and other data that can be linked to a mailing address.¹⁵

Using the latitude and longitude provided by RefUSA, we calculate each home's own climate zone and its distance to the nearest climate zone. We then restrict the sample to single-family homes that are within 3 kilometers of a climate zone border. Because we are interested in considering total energy use as well as energy use per square foot, we eliminate homes for which we do not have square footage information. We also omit homes located in cities whose climate zone changed in 1995.

To maximize the comparability of the treated and control group, we exclude homes built prior to 1947 or after 2006, as there is evidence that new homes use less energy in the first few years of their existence than in the longer run (Levinson 2016; Kotchen 2017). For our main regressions, we also drop homes built in 1977-1981. While there was some variation in building requirements by heating and cooling degree days between 1977 and 1981, most of the variation was at the city level and did not correspond to climate zone boundaries.

Energy use. We obtain monthly premise-level electricity and natural gas usage data from four major California utilities: San Diego Gas & Electric (SDG&E), Pacific Gas and Electric (PG&E), Southern California Edison (SCE), and Southern California Gas (SCG). Together, they serve almost all of California, with the exception of the very northern part of the state, the Sacramento area, and a few cities that have their own electric and gas utilities (see Figure 2). Thus, every climate zone border has the possibility of being represented in our sample although in practice some borders do not have any homes located nearby.

Our energy usage data span the time period of January 2009 through July 2015, allowing us to obtain a fairly precise measure of each premise's expected electricity and natural gas usage. Electricity usage is measured in kilowatt-hours (kWh) and natural gas usage is measured in therms. A kilowatt-hour is equivalent to about 3,400 British Thermal Units (BTUs), and a therm is equal to 100,000 BTUs. To gauge the change in total energy usage, we convert both therms and kilowatt-hours to thousands of BTUs and add them together.

We use the address provided by the utility to match each home in the 3-kilometer RefUSA sample described above to its energy use. To minimize false matches, which could

¹⁵ The exact algorithm is proprietary.

introduce measurement error in our measure of building code strictness, we only retain cases where addresses match perfectly (including the street number, street name, city, and zip code) or where the only difference in the addresses is an abbreviated street suffix (e.g., “Ave” instead of “Avenue”, “Rd” instead of “Road”, etc.).¹⁶ We then drop a few homes that match to more than one utility providing the same type of energy. To maintain consistency across specifications, we restrict our sample to homes where we observe both electricity and natural gas usage. Using the most recent RefUSA record from the years 2006-2012, we are able to obtain income data for about 80 percent of the border homes for which we have energy use. Our final sample contains 353,597 single-family homes built between 1947 and 2006 that are located within 3 kilometers of a climate zone border.

Figure 3 shows the geographic locations of these homes, including larger maps of three metropolitan areas where we have the largest density of border homes: San Francisco (28,592 homes), Los Angeles (184,140 homes), and San Diego (46,652 homes). These homes represent almost three quarters of our sample. Homes built before modern climate zones were established are in light gray, while homes built after are in dark gray (in case of overlap, light gray is superimposed on dark gray). Black lines correspond to climate zone boundaries. In many places there is significant geographic overlap between older and newer homes (although this is sometimes difficult to see visually due to the high home density and overlaps of new and old homes) and good balance in homes on different sides of the boundary. The areas where there are few or no older/newer homes or few/no homes just on the opposite side of the boundary demonstrate the necessity of the geographic controls in equation (1) for proper inference. In Online Appendix Figure A2, we show the spatial distribution of incomes in our sample. Although our sample is quite rich on average, and there are some visible clusters of high-income homes, each of the three major metropolitan areas and most climate zone borders have both low- and high-income households.

To increase the precision of our estimates, we would like to eliminate variation in energy use that is due to seasonal patterns in a particular area. Including geographically-varying calendar month fixed effects is the most straightforward way of accomplishing this. However, households’ billing cycles vary substantially in our sample. We thus apply a simple

¹⁶ We found that attempting to match more homes by accounting for misspellings and other errors did not significantly increase the number of matches, at least for homes in the San Diego Gas and Electric utility area.

transformation to allocate energy use during a billing cycle to a particular calendar month.¹⁷ Specifically, we calculate the proportion of days in each billing cycle that fell into each calendar month and then allocate that proportion of overall energy use to that month. For example, if a household used 100 kilowatt-hours (kWh) during a 31-day billing cycle that started on January 10th 2015 and ended on February 9th 2015, we would allocate $\frac{22}{31} * 100 = 71$ kWh of usage to January 2015 and $\frac{9}{31} * 100 = 21$ kWh of usage to February 2015. Finally, we combine all energy use allocated to a particular month-year to arrive at each household's mean daily energy use during a calendar month.

Heating and cooling degree-days. We obtain the “typical meteorological year” for each zone from the California Energy Commission.¹⁸ Among other variables, the data report a “typical” temperature for each hour of the year for each climate zone. To construct these data, a slightly modified version of a procedure called the “Sandia method” (Hall et al. 1978) is applied to weather station records between 1998 and 2009. Each month of the typical meteorological year is selected from these records to be the most representative of the long-term distribution for that month. The historical month whose cumulative frequency distribution (CDF) has the smallest deviation from the long-term CDF is considered the most representative and the weather records from that month are used to describe the zone's climate. By request of the CEC, the same year-month was chosen for each climate zone when constructing its typical weather for each month. We convert these hourly temperature data to annual heating and cooling degree days by using the dry-bulb temperature to calculate heating and cooling degrees each hour, adding the results up by zone-year, and dividing by 24.

Summary statistics. Table 1 presents the summary statistics for the homes in our sample. The average home uses 188,600 BTUs per day, about 36% of which is electricity and 64% is natural gas. Even after eliminating homes built before 1947, the average year built in our sample is fairly low (1970), suggesting that we have many control homes that were not subject to building codes when they were constructed. The average home in our sample has almost 1,900 square feet of

¹⁷ An alternative would be to include billing cycle fixed effects rather than calendar month fixed effects. However, that would greatly increase the number of fixed effects and, because weather is autocorrelated, would not provide a clear advantage over the transformation we apply.

¹⁸ Available from http://www.energy.ca.gov/title24/2016standards/ACM_Supporting_Content/. Accessed November 2017. The data were created by White Box Technologies for the CEC. We thank Joe Huang for explaining the creation of these climate variables to us.

living space, 3.1 bedrooms, and 2.1 bathrooms. Finally, the average home in the sample is worth about \$490,000, and the average annual household income is about \$123,000.

Table 2 shows select summary statistics for each income quintile. In the bottom quintile, households in our sample earn about \$15,000 per year, live in houses with 1,313 square feet of living space, and use about 49,800 BTUs in electricity and 108,600 BTUs in natural gas daily. Household income in the second quintile averages almost \$35,600 per year and living space averages about 1,350 square feet. These houses use slightly more electricity than the bottom quintile (51,900 BTUs) but use less natural gas. Electricity and natural gas use increase monotonically for the remaining quintiles at a rate that is similar to or lower than increases in square footage. The top quintile's income averages 11.6 times more than the bottom quintile. Living space is about 70% larger than that of the bottom quintile, electricity use is 60% higher, and natural gas usage is about 25% higher.

The average household income in our sample is substantially higher than that of California as a whole. In the full RefUSA sample for California, median household income is about \$62,000, which is close to the 2010-2014 Census estimate of \$61,489 (U.S. Census 2016). The differences arise because we restrict the sample to single-family homes built between 1947 and 2006 that are located within 3 kilometers of a climate zone border. Each of these three restrictions has an independent and positive effect on the median and average incomes in the sample. Although our sample is thus richer than the typical California household, it still contains many households with fairly low incomes. Because we determine cutoffs for income quintiles using California microdata from the American Community Survey, the number of households in the lower quintiles is substantially smaller than the number of households in the upper quintiles (last column of Table 2).

Distribution of energy use per dollar by income decile. It is beyond the scope of this paper to fully analyze the distributional impacts of residential energy use taxation. In order to do that accurately, we would need to know how the price elasticity of energy demand varies by income, to make assumptions about how tax revenues would be utilized, and to specify how taxes would impact the nonlinear energy price schedules. However, our summary statistics suggest that, without explicit efforts to address distributional consequences, a simple tax on residential energy consumption would be regressive.

We illustrate this point further by calculating the ratios of average daily electricity and natural gas use to annual income for each household in our sample. Figure 4 shows the median, the 10th, and the 90th percentiles of these ratios by income decile (the exact point estimates can be found in Online Appendix Table A4). The bottom decile uses 12.9 BTUs per day per dollar of income in electricity and 5.34 BTUs per day per dollar of income in natural gas. These ratios decline monotonically with income, although the rate of decline is lower at higher deciles. The top income decile uses 15-19 times less energy on a per-dollar-of-income basis than the bottom decile: about 0.64 BTUs per day per dollar of income in electricity and 0.36 BTUs per day per dollar of income in natural gas.

Furthermore, there is a lot more heterogeneity in energy use per dollar of income in the lower deciles than in the top deciles. The 10th percentile of natural gas use declines from 2.27 BTUs per day per dollar (bottom decile) to 0.19 BTUs per day per dollar (top decile). By contrast, the 90th percentile declines from 12.01 to 0.67, a much larger drop in absolute terms. The pattern for electricity consumption is similar. These patterns further support the notion that, absent redistribution or significant adjustments by the bottom income deciles, residential energy consumption taxes would be highly regressive.

IV. Results

A. Testing the parallel trends assumption

Our identification strategy requires that climate zone boundaries be independent of geographically-varying time trends that influence energy usage. If that is not the case, we could mistakenly attribute post-1981 differences in energy usage to building codes rather than to pre-existing differential trends. Luckily, we have a large number of homes in our sample that were built prior to 1977, so we can directly test whether there are any trends differences that could not have been caused by building codes (absent differential sorting into older homes). We consider two of our main outcomes of interest: natural gas usage per square foot and electricity usage per square foot (in logs).

To probe our parallel trends assumption, we estimate a variant of equation (1) where we replace $\beta_1 Post_i * DiffCDD_i + \beta_2 Post_i * DiffHDD_i$ with $\sum_{t=1947, t \neq 1976}^{2006} \beta_{t1} \mathbf{1}[YB_i = t] * DiffCDD_i + \beta_{t2} \mathbf{1}[YB_i = t] * DiffHDD_i$. That is, we estimate an event study specification. The indicator $\mathbf{1}[YB_i = t]$ is equal to 1 if home i was built in year t . The reference category consists

of homes built in 1976, the year before California’s first statewide building code went into effect. Checking whether the $\beta_{t1}, \beta_{t2} = 0$ for $t < 1977$ assesses whether there are any pre-existing trends along climate zone borders that are correlated with how we measure building code strictness. Conditional on the parallel trends assumption holding, coefficients for years 1982 and later inform us of the causal effects of building codes differences across climate zones on energy use. As in our main specification, we omit homes built in 1977-1981.

Figure 5 provides a visual check of the parallel trends assumption and provides a preview of the treatment effect that we will discuss in the next section. Rather than plot β_{t1} and β_{t2} separately, we plot their linear combinations, using the average difference in heating and cooling degree days in our sample (see previous section for discussion). In both panels of Figure 5, we see little evidence of differential trends prior to 1977, although a few of the linear combinations are statistically different from zero. Combined with the fact that the older homes exhibit similar seasonal patterns in energy use (see Online Appendix Figure A1), this result increases confidence in our identification strategy. Immediately following the introduction of climate zones, we see both electricity and natural gas usage fall significantly, then rebound to zero in the mid-1980s. We also see significant and negative treatment effects in the mid-to-late 1990s and the 2000s. The absence of a significant treatment effect in other years could be indicative of heterogeneous treatment effects over time, either due to non-constant cross-border differences in code stringency (due to building code revisions), variation in enforcement, or in how binding building energy codes are. Because the focus of this paper is on distributional consequences, we do not examine this intertemporal heterogeneity further, leaving it for future research. Instead, we proceed by estimating these and other treatment effects using a more parsimonious regression specification.

B. Home characteristics

First, we consider the possibility that more stringent building codes affect the observable characteristics of homes constructed after the climate zones begin to be enforced, namely square footage and number of bedrooms. The results for our main 3-kilometer sample are shown in columns 3 and 4 of Table 3. The estimates indicate that, on average, stricter building energy codes do not significantly affect homes’ square footage but cause the number of bedrooms to decrease by 0.73%. We probe the robustness of these results by restricting the sample to homes

located within 1 kilometer of a climate zone border (columns 1 and 2). The estimated fall in square footage becomes larger (-1.8%) and statistically significant. The estimated change in the number of bedrooms in the 1 kilometer sample is very similar to the 3 kilometer sample. Finally, estimates are also similar when we do not take the log of square footage and bedrooms (see Online Appendix Table A5).

More interesting patterns emerge when we look at changes in living square footage for households in different income quintiles in Figure 6, again using a linear combination of the estimated coefficients.¹⁹ Specifically, the largest reductions in square footage occur in the bottom two income quintiles, with decreases of 4-6%. While estimates for the other quintiles are also negative, they are relatively small (less than 1%) and only the 4th quintile's estimate is significant at the 5% level. Because we have only about 5,500 homes in the first quintile compared to 165,000 homes in the fifth quintile, our standard errors for the poorer households are larger. However, even with large standard errors, we can rule out that households in the bottom income quintile do not experience any distortions to home attributes. Our findings for the number of bedrooms largely mirror our square footage results, with the exception that we see statistically significant reductions in all income quintiles.

The finding that the distortions are largest at the bottom of the income distribution is perhaps surprising given that lower-income individuals tend to live in smaller homes and that the geometric assumptions inherent in the compliance software penalize more rectangular homes, which are presumably larger. However, higher-income households may be willing to pay more for additional square footage, limiting the desirability of reducing their home size compared to lower-income households. Next, we see how these and other (unobservable to us) changes in building practices translate into changes in energy usage.

C. Energy usage

Table 4 shows the effect of a more stringent energy building code on total energy use (natural gas plus electricity), in BTUs. To test for both unconditional and conditional changes in energy use, we consider both aggregate energy use and energy use per square foot. Note that we conduct some of our analyses on a per-square-foot basis not because this is the correct way to estimate the treatment effect of building energy codes, but to see whether there are energy use

¹⁹ Point estimates and standard errors for Figures 6-8 can be found in Online Appendix Tables A6-A8.

reductions through channels other than square footage. If builders respond to energy codes by distorting attributes such as square footage, then the distortion and the subsequent energy use changes are part of the treatment effect and should not be ignored. Thus, the right estimate of a building energy code's net impact is the change in *total* energy use. In our 3-kilometer sample, we find no significant reduction in either total energy use or energy use per square foot (columns 3 and 4). In the 1-kilometer sample, we estimate a significant 1% reduction in total energy use, but again no change on a per-square-foot basis.

Tables 5 and 6 show the changes in natural gas and electricity consumption separately. Total natural gas use is unchanged, while total electricity consumption falls by about 1.0-1.3%. On a per-square-foot basis, we see no reductions in natural gas consumption, while electricity use per square foot falls by 0.6% in our preferred 3-kilometer sample. Overall, these results suggest that these (small) decreases in energy consumption are being achieved largely but not entirely through building smaller homes.

Figure 7 shows the 3-kilometer estimates from Tables 4-6 broken down by income quintile. On an aggregate basis, we see no reductions in energy use in any quintile except for the second one, where energy use decreases by 4.1%. However, our standard errors for the first quintile are large, so that we cannot rule out similarly sized decreases in energy use for the poorest households. For the top three quintiles, we can rule out decreases larger than 1.5% with 95% confidence. On a per-square-foot basis, we can rule out decreases of more than 1.6% for each of the five income quintiles, and we estimate that households at the bottom of the income distribution actually use *more* energy on a per-square-foot basis. Given that these households experience a decrease in home size, this finding could be mechanical if there are appliances whose energy use does not vary substantially with home size (e.g., dishwashers, refrigerators, washers/dryers, and water heaters).

When we consider natural gas and electricity separately, the largest decrease at the dwelling level is in the second lowest income quintile, which sees a decline of 4.7% and 4.5% in natural gas and electricity, respectively. No other income quintile shows a decrease in natural gas use, although it again should be noted that our confidence interval for the lowest quintile is very wide. By contrast, we see significant but small (1.1-1.2%) decreases in aggregate electricity use for the top quintiles. On a per-square-foot basis, we see no significant decreases in natural gas for any quintile, and electricity consumption falls only for the top income quintile. Again, the

decrease is small (0.8%). Because the households at the top of the income distribution make up a disproportionate share of our sample, it should not be surprising that our overall findings are driven by this group. Thus, the overall energy savings that building codes create in our context are small and can be explained almost entirely by changes in homes' square footage.

It is worth considering how these energy use reductions translate into reductions in energy expenditure. This is especially important because low-income households are eligible for the California Alternate Rates for Energy (CARE) program, which provides them with a 30-35% discount on electricity and a 20% discount on natural gas (California Public Utilities Commission, 2018). Households with slightly higher incomes are eligible for a 12% discount through the Family Electric Rate Assistance Program (FERA) program. Eligibility for each program is determined by a combination of household size and household income.

Table 7 uses the point estimates from Figure 7 combined with 2016 marginal rates in the first energy use tier to calculate average monetary savings for each quintile-utility combination.²⁰ Specifically, we multiply average annual usage in each quintile by the estimated percent saving and the average marginal energy rate, weighted by the proportion of CARE households in each income quintile. Unfortunately, our data from PG&E (about 16% of our sample) do not contain information about which households are enrolled in CARE, so we use the average shares from the other utilities to calculate the average marginal rate.

It is clear from Table 7 that the largest monetary savings accrue to lower-income households. For example, SDGE customers in the second income quintile save \$21 per year on natural gas and \$59 per year on electricity, while those in the fifth income quintile save an insignificant \$1.2 on natural gas and \$27 per year on electricity. Accounting for prices also reduces the differences between quintiles as compared to energy quantities, because many of the lower-income households face lower energy use rates.

D. Housing prices

Finally, we consider the effect of stricter building energy codes on housing prices. Building energy codes may differentially affect the prices of homes that do and do not face a more stringent code through two channels. First, Myers (2017) and Aydin et al. (2017) show

²⁰ In 2016, there were three pricing tiers for electricity and two pricing tiers for natural gas. The marginal rates used for this exercise as well as calculations for the higher tiers can be found in Online Appendix Tables A9-A10 and A11, respectively.

significant capitalization of energy costs into home prices. If building codes reduce energy use and the reduction is capitalized into the home price, then homes that were built under stricter codes will command a price premium. Second, building energy codes may affect other home attributes (e.g., square footage), some of which may be unobservable to us (e.g., how drafty a home is). Depending on how the attribute is perceived and valued by the buyer, the second channel could increase or decrease sale prices (e.g. Houde 2016).

Of course, homes facing more stringent building energy codes are also weakly more expensive to build (strictly if the codes are binding). Unfortunately, it is essentially impossible for us to calculate the additional costs imposed by the introduction and subsequent revisions of California's building energy codes, as we observe neither the attributes of pre-code homes (and thus cannot tell how out-of-compliance older homes were) nor the attributes of post-code ones. Intuitively, who bears the incidence of these costs depends on the relative elasticities of supply and demand and, in the case of imperfectly competitive markets, on the markup above marginal cost (Weyl and Fabinger 2013). Incidence can thus vary across the different areas in our sample: in areas where land available for building is relatively scarce (e.g., the Bay Area) and supply is therefore inelastic, buyers will bear a smaller share of the burden relative to areas where land for building is elastically supplied.

To understand what our estimates of the effect of building energy codes on home prices capture, we note that the assessed values we observe are largely based on existing homes being sold from one owner-occupant to another owner-occupant. Once a home is built, the construction costs become sunk costs and should not matter for resale. Consistent with this idea, Bruegge et al. (2016) find that the initial sale price premium for energy star homes disappears in the resale market. Because *new* homes are more expensive to build in the presence of a stricter building code, existing homes should command a higher price in stricter building code regimes. However, this effect should be the same for pre- and post-code homes on each side of the borders in our sample and thus not affect our estimates of the treatment effect.

Our difference-in-differences strategy is also helpful for purging the price estimates of other location-specific confounding factors. For example, Glaeser and Gyourko (2003) point to the importance of construction costs, land scarcity/building restrictions, and population growth for determining homes' sale prices. Intuitively, if an area is gaining population, the resale price of a house should be bounded below by construction costs. If there are no constraints on land and

no spatially differentiated amenities, the price of existing homes should also not *exceed* new construction costs (because a buyer can always buy a new house instead). Otherwise, prices can be above construction costs. By similar reasoning, if an area is losing population, home prices can fall below construction costs. However, both older (pre-code) and newer homes will be affected by these dynamics. Thus, our difference-in-differences strategy will isolate home value effects that are solely due to differences in the characteristics of existing homes (i.e., energy use and/or other attributes).

Finally, it is worth noting how energy prices would affect the capitalization of building energy codes into home values. At one extreme, if energy prices are zero, then so is the value of any reduction in energy consumption. Because building energy codes affect other attributes, however, we would not necessarily expect to find no effects on home value even if the price of energy were zero. In principle, we could estimate how the capitalization of building energy codes into home values varies with the local price of energy. However, electricity and natural gas rates depend on the utility and, within utility, on the geographic area in which the home is located. For both of these reasons, energy prices have a strong geographic component and are not as good as randomly assigned. Thus, any estimation involving energy prices would also pick up any heterogeneous treatment effects that are correlated with energy prices (e.g., a warmer climate). For this reason, we focus on the average treatment effect rather than introduce heterogeneity.

The effects of building energy codes on home prices across the entire income distribution are shown in Table 8. We detect a significant price increase of about 2.1% in the 3-kilometer sample, both on aggregate and on a per-square-foot basis (columns 3 and 4). The increase in home prices on a per-square-foot basis is similar in columns 1 and 2, where we restrict the sample to 1 kilometers around a border, while the aggregate price increase is smaller (about 0.8%).

Figure 8 shows the estimated changes in home prices by quintile. We see large and significant decreases in home prices in the bottom two quintiles (12% and 7.7%, respectively) and moderate significant increases in prices for the top two quintiles (2.8% and 1.6%, respectively). About half of the fall in prices in the bottom two quintiles can be explained by changes in square footage, as the percent change in price per-square-foot is about half the size of the total percent change in price. For the top two quintiles, the aggregate increases are similar to the per-square-foot increases, suggesting that building energy codes lead to changes in

(unobservable to us) home characteristics that are valued positively by households. Importantly, these price increases cannot be explained by the observed reductions in energy use. For the average home in these top quintiles, the observed estimates imply an increase in home value of over \$10,000. The largest reduction in electricity use we observe is 1.2%. Even assuming that the household is paying a high marginal price for electricity (see Online Appendix Table A11), the net present value of the reductions in energy use is about an order of magnitude smaller than the increase in home value. Of course, there may be a rebound effect, which we also cannot observe, but it would have to be very large to account for the entire difference in home value. Thus, it is more plausible that building energy codes cause changes in the characteristics of high-income homes that households value for reasons other than their impact on a home's energy use.

V. Discussion and conclusion

We adopt a novel approach to estimating the causal effect of building energy codes on energy usage by using temporal and spatial discontinuities in the code's strictness across California's 16 climate zones and a sample of homes that are located within 3 kilometers of climate zone boundaries. Comparing cross-border differences in the energy use of homes built before the introduction of a state-wide building code in 1977 to differences in the energy use and home characteristics of homes built after the introduction of modern climate zones in 1982, we find strong evidence that building energy codes led builders to change their building practices, resulting in smaller homes, especially for the lowest-income households. Higher-income households experience much smaller distortions in square footage and the number of bedrooms. These results suggest that such distortions are more costly to builders of high-income homes, perhaps because higher-income buyers are more price elastic with respect to square footage than lower-income buyers.

Homes built in more stringent building code regimes use less electricity, but the difference is small and largely driven by the smaller square footage. We see the largest percentage reductions in energy usage among lower-income households, but after accounting for the reduction in square footage, only a small electricity consumption decrease for the top income quintile remains. This result suggests that energy use reductions are achieved largely through reductions in home size. Nonetheless, because the home size reductions are part of the total treatment effect, building energy codes do save some energy, especially for lower income

households. However, low-income households on average consume 20-40% less energy than high-income households (depending on which quintiles we compare), making any proportional reduction smaller in absolute terms. Low-income households also pay less for energy, because California has an increasing block rate structure as well as a special low-income rate schedule. Thus, any reductions in energy use are worth less in absolute terms to these households.

We then examine the capitalization of building energy codes into home values. At the bottom of the income distribution, home prices fall, partly due to the decreased square footage of affected homes. At the top of the income distribution, prices increase for reasons that are unobservable to us. The energy use reductions for higher-income households are small, and the net present value of savings is an order of magnitude smaller than the increase in home prices, even if we assume a high marginal price of electricity. This fact suggests that building energy codes provide other benefits to these households that are difficult to measure directly, such as lower draftiness.

While our framework cannot identify the effect of the building code on consumer welfare across income quintiles, it is clear from our analysis that builders comply with building energy codes by changing secondary home attributes, i.e., those that are not directly targeted by the codes, such as square footage and the number of bedrooms. Even though the value of the energy savings is largest for households in the bottom half of the income distribution, these savings are brought about by reducing these households' home size. For them, the distortions in attributes lead to a reduction in home value, while for higher income households the distortions increase home value. If these changes do not correct for market failures such as asymmetric information or inattention, then building codes appear to be somewhat regressive.

The fact that energy building codes may lead to a distortion of secondary home attributes has two key implications, one for future research and one for policy. First, researchers should be wary of using solely intertemporal variation to estimate the causal effect of building codes on energy use, as bias could arise regardless of whether one controls for home attributes or not. Second, unless extremely well-designed, building codes can contain incentives not deliberately created by policymakers and subsequently have unintended consequences.

References

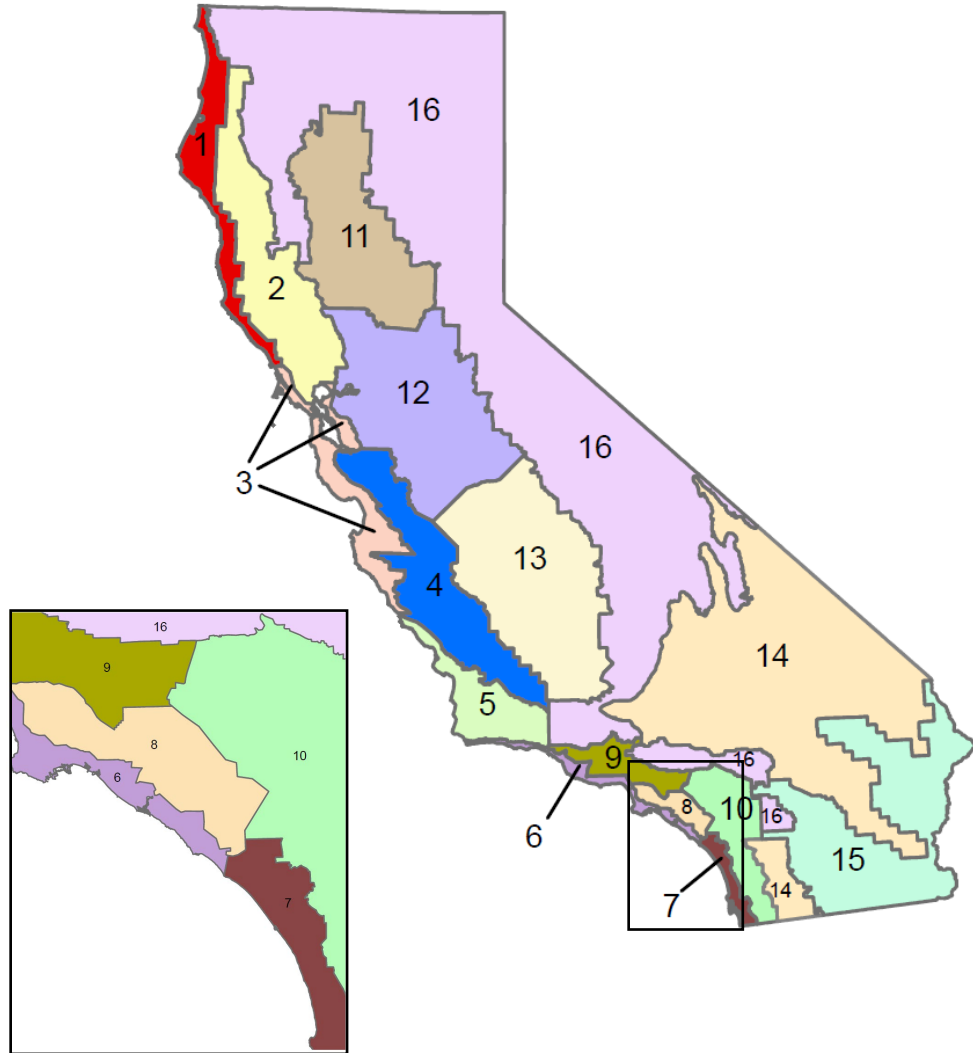
- Aroonruengsawat, Anin, Maximilian Auffhammer, and Alan H. Sanstad. 2012. The impact of state level building codes on residential electricity consumption. *The Energy Journal*: 31-52.
- Allcott, Hunt, and Michael Greenstone. 2017. Measuring the welfare effects of residential energy efficiency programs. NBER Working Paper 23386, National Bureau of Economic Research, Cambridge, MA.
- Angrist, Joshua D., and Jörn-Steffen Pischke. 2013. *Mostly harmless econometrics: an empiricists companion*. Princeton: Princeton University Press.
- Autor, David H. 2003. Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of labor economics* 21 (1): 1-42.
- Aydin, Erdal, Dirk Brounen, and Nils Kok. 2015. Capitalization of Energy Efficiency in the Housing Market. Working paper.
- Bento, Antonio M. 2013. Equity impacts of environmental policy. *Annual Review of Resource Economics* 5 (1): 181-196.
- Borenstein, Severin, and Lucas W. Davis. 2016. The distributional effects of US clean energy tax credits. *Tax Policy and the Economy* 30 (1): 191-234.
- Bruegge, Chris, Carmen Carrión-Flores, and Jaren C. Pope. 2016. Does the housing market value energy efficient homes? Evidence from the energy star program. *Regional Science and Urban Economics* 57: 63-76.
- California Energy Commission. 1978. Regulation Establishing Energy Conservation Standards for New Residential Buildings. R-31.
- California Energy Commission. 1995. California Climate Zone Descriptions for New Buildings. <http://www.energy.ca.gov/1995publications/P400-95-041.pdf>. Accessed September 2, 2016.
- California Public Utilities Commission. 2018. CARE/FERA Programs. <http://www.cpuc.ca.gov/General.aspx?id=976>. Accessed January 5, 2018.
- Davis, Lucas W., Alan Fuchs, and Paul Gertler. 2014. Cash for coolers: evaluating a large-scale appliance replacement program in Mexico. *American Economic Journal: Economic Policy* 6 (4): 207-38.
- Department of Energy (DOE). 2017. Status of State Energy Code Adoption. <http://www.energycodes.gov/status-state-energy-code-adoption>. Accessed October 27, 2017.

- Energy Information Agency (EIA). 2009. Household Energy Use in California. https://www.eia.gov/consumption/residential/reports/2009/state_briefs/pdf/ca.pdf. Accessed December 9, 2016.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2018. Do energy efficiency investments deliver? Evidence from the weatherization assistance program. *The Quarterly Journal of Economics* 133 (3): 1597–1644.
- Glaeser, Edward L., and Joseph Gyourko. 2003. The Impact of Building Restrictions on Housing Affordability. *FRB New York - Economic Policy Review* 9 (2): 21-39.
- Hall, Irving J., R. R. Prairie, H. E. Anderson, and E. C. Boes. 1978. *Generation of Typical Meteorological Years for 26 SOLMET Stations*. SAND78-1601. Sandia National Laboratories, Albuquerque, NM
- Hanna, Rema, Esther Duflo, and Michael Greenstone. 2016. Up in Smoke: The Influence of Household Behavior on the Long-run Impact of Improved Cooking Stoves. *American Economic Journal: Economic Policy* 8 (1): 80-114.
- Hassett, Kevin A., Aparna Mathur, and Gilbert E. Metcalf. 2007. The incidence of a US carbon tax: A lifetime and regional analysis. NBER Working Paper 13554, National Bureau of Economic Research, Cambridge, MA.
- Houde, Sébastien. 2014. How consumers respond to environmental certification and the value of energy information. NBER Working Paper 20019, National Bureau of Economic Research, Cambridge, MA.
- International Code Council. 2017. International Code Adoptions. <https://www.iccsafe.org/about-icc/overview/international-code-adoptions/>. Accessed December 8, 2017.
- Ito, Koichiro, and James M. Sallee. 2018. The Economics of Attribute-based Regulation: Theory and Evidence From Fuel-economy Standards. *The Review of Economics and Statistics* 100 (2): 319-336.
- Jacobsen, Grant D., and Matthew J. Kotchen. 2013. Are Building Codes Effective at Saving Energy? Evidence from Residential Billing Data in Florida. *The Review of Economics and Statistics* 95 (1): 34-49.
- Jacobsen, Mark R. 2013. Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity. *American Economic Journal: Economic Policy* 5 (2): 148-187.
- Kotchen, Matthew J. 2017. Longer-Run Evidence on Whether Building Energy Codes Reduce Residential Energy Consumption. *Journal of the Association of Environmental and Resource Economists* 4 (1): 135-153.

- Levinson, Arik. 2016. How Much Energy Do Building Energy Codes Save? Evidence from California. *American Economic Review* 106 (10): 2867–2894.
- Myers, Erica. 2017. Are Home Buyers Inattentive? Evidence From Capitalization of Energy Costs. E2e Working paper 024.
- Novan, Kevin, Aaron Smith, and Tianxia Zhou. 2017. Residential Building Codes Do Save Energy: Evidence From Hourly Smart-Meter Data. E2e Working Paper 031.
- Reiss, Peter C., and Matthew W. White. 2005. Household Electricity Demand, Revisited. *The Review of Economic Studies* 72 (3): 853-883.
- Weyl, E. Glen, and Michal Fabinger. 2013. Pass-through as an Economic Tool: Principles of Incidence Under Imperfect Competition. *Journal of Political Economy* 121 (3): 528-583.
- United States Census Bureau (U.S. Census). 2016. QuickFacts, California. <http://www.census.gov/quickfacts/table/PST045215/06>. Accessed December 5, 2016.
- West, Sarah E. 2004. Distributional Effects of Alternative Vehicle Pollution Control Policies. *Journal of Public Economics* 88 (3): 735-757.
- West, Sarah E., and Robertson C. Williams III. 2004. Estimates from a Consumer Demand System: Implications for the Incidence of Environmental Taxes. *Journal of Environmental Economics and Management* 47 (3): 535-558.
- Williams, Robertson, Hal Gordon, Dallas Burtraw, Jared Carbone, and Richard Morgenstern. 2014. The Initial Incidence of a Carbon Tax across Income Groups. *Resources for the Future Discussion Paper* 14-24.

FIGURES

Figure 1. Building climate zones in California



Source: California Energy Commission. Numbers inside or next to each area represent the climate zone number. Inset on the left represents area inside the black rectangle in the full map.

Figure 3. Locations of in-sample homes

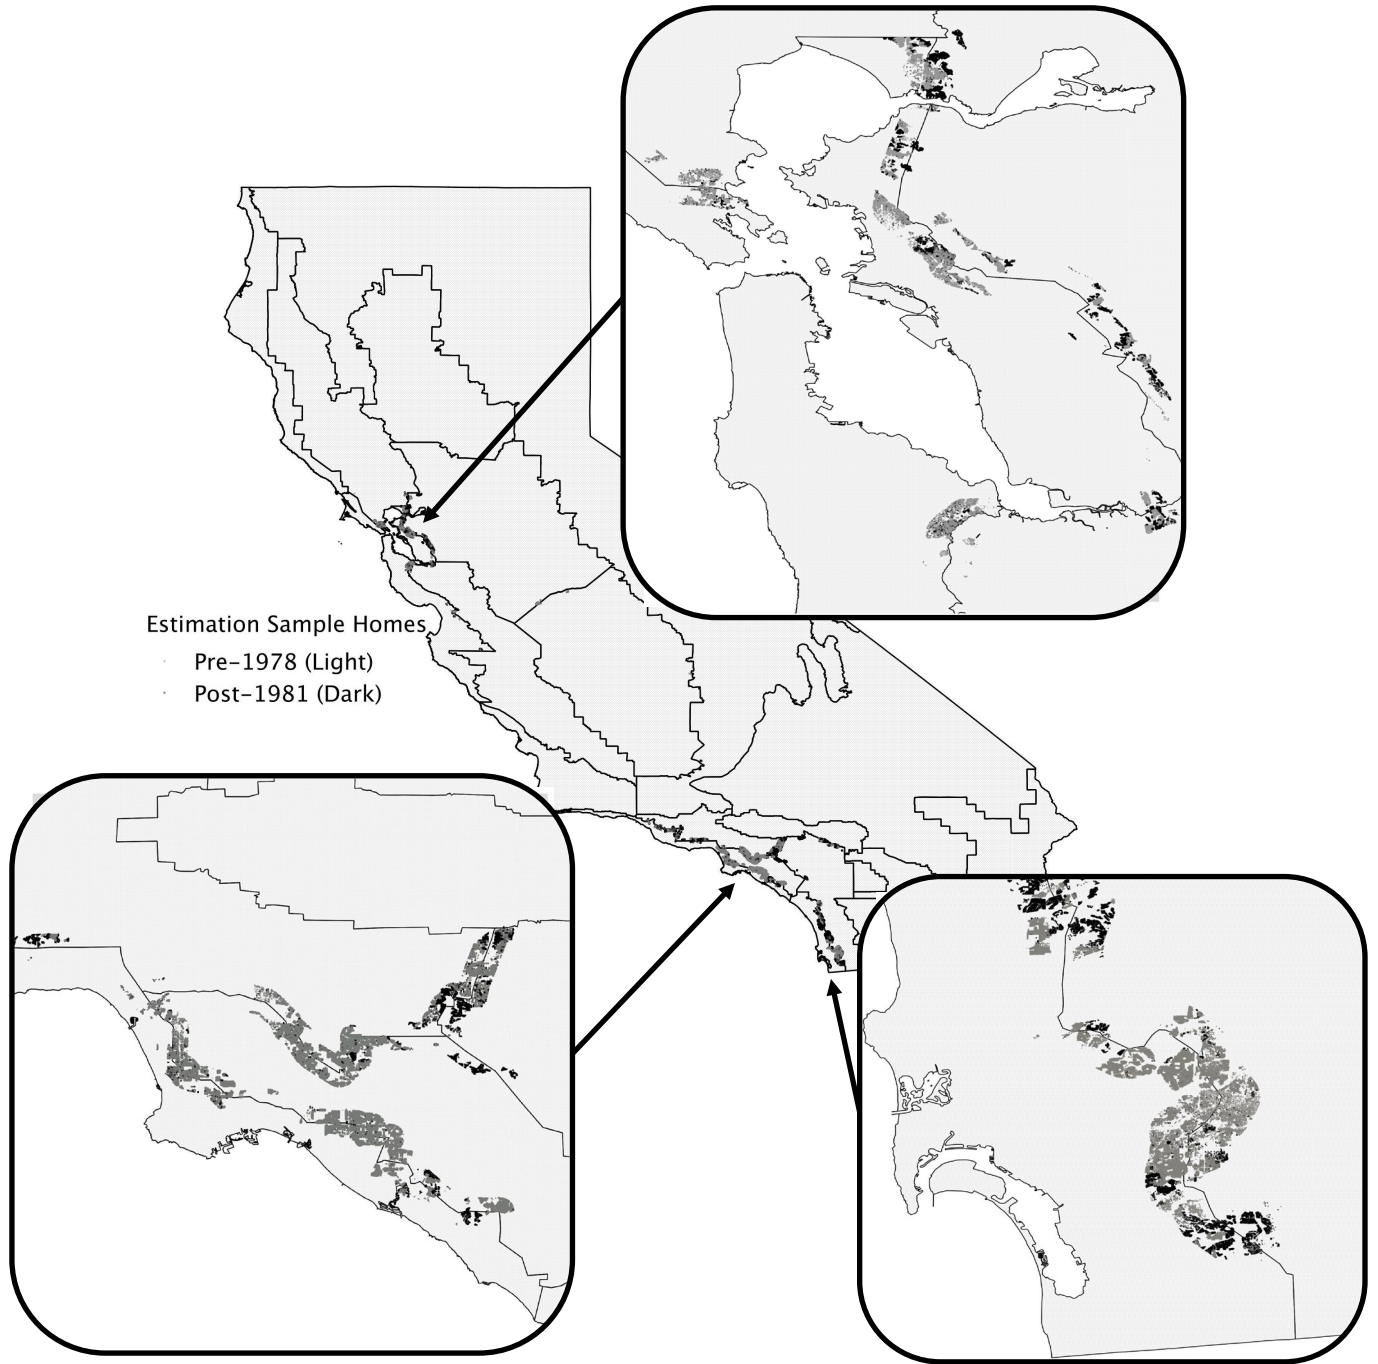

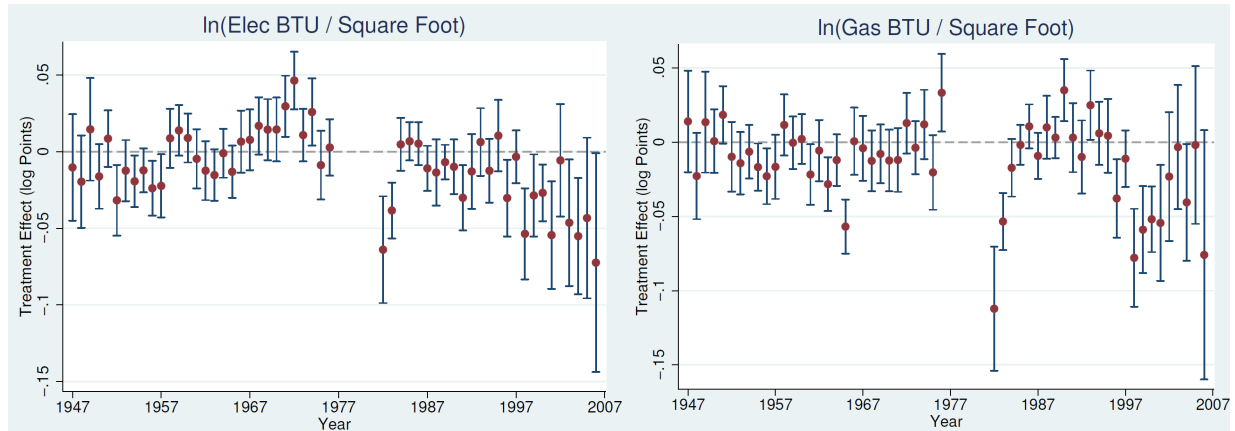


Figure 4. Energy use per dollar of income by income decile




Solid lines correspond to median daily electricity (left panel) or gas (right panel) use in BTUs per dollar of annual household income for each income decile. Dashed lines represent the 10th and 90th percentiles.

Figure 5. Trends in energy use per square foot before and after the introduction of climate zones



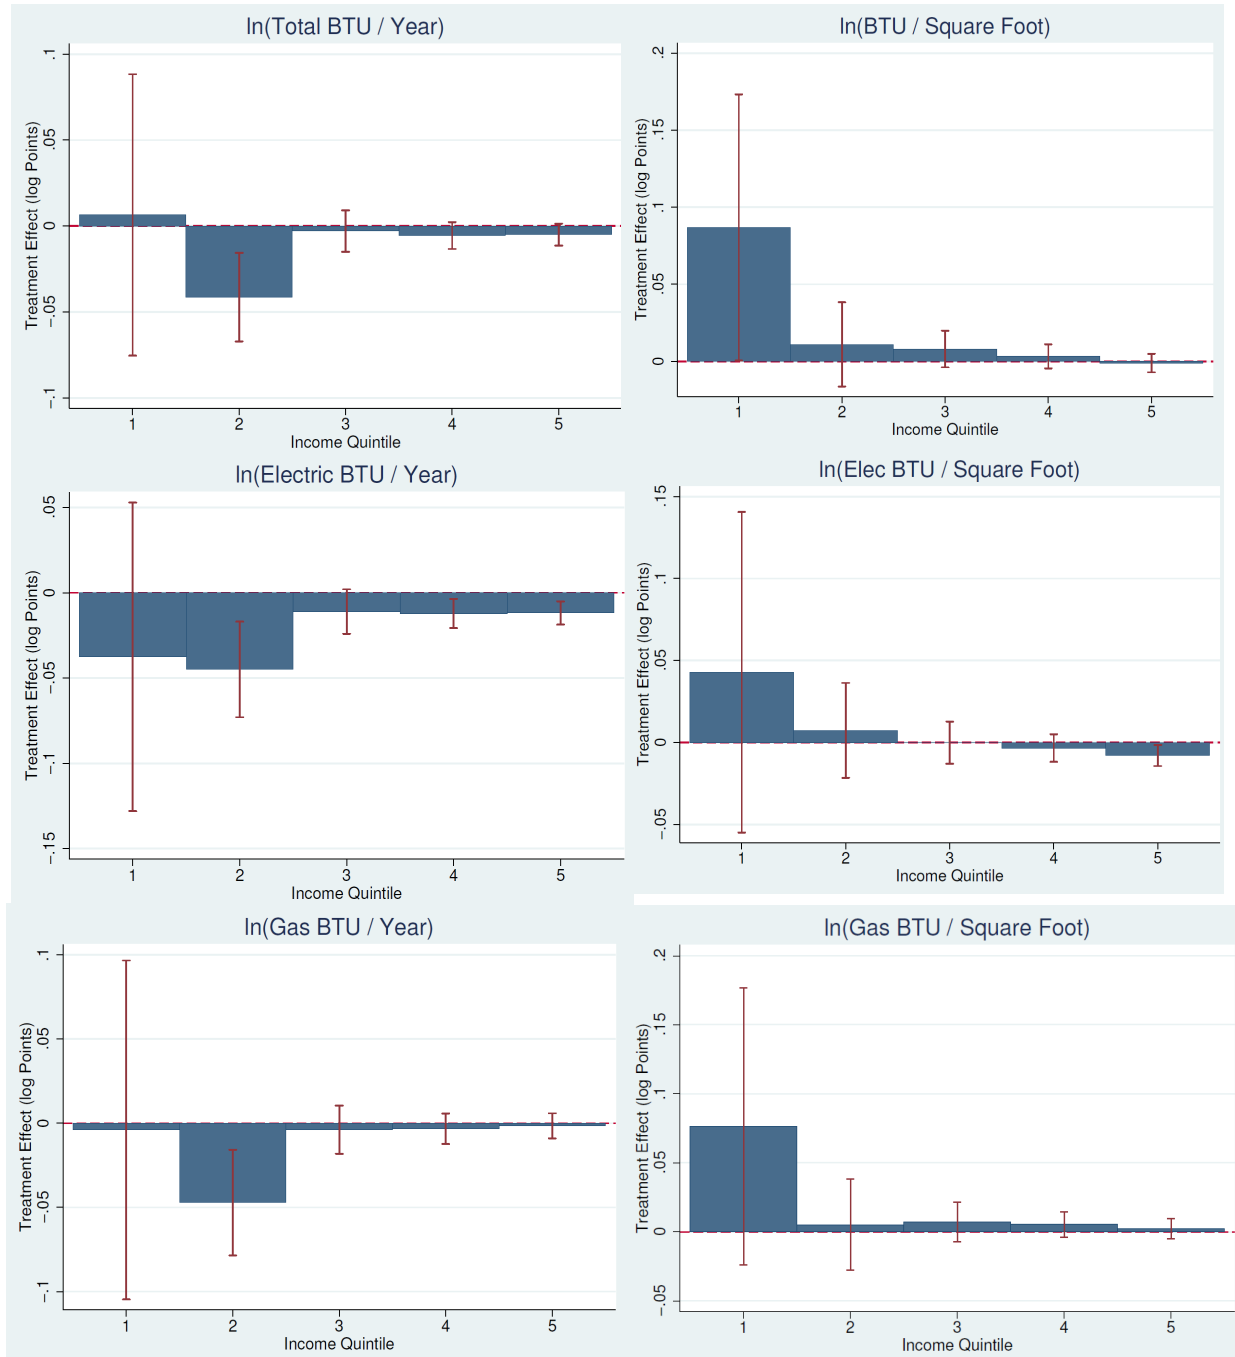
Panels show point estimates and 95 percent confidence intervals for the estimated treatment effect in each year, as specified by equation (1), relative to 1976 (the year before California’s building code went into effect). The energy metric is indicated above each panel.

Figure 6. Effect of building energy codes on home characteristics by income quintile




Panels show difference-in-differences coefficients and 95 percent confidence intervals for each income quintile, estimated using equation (2). Standard errors are clustered by zip code. Controls include income-quintile-by-border fixed effects, year-built-by-border fixed effects, and zip code fixed effects. The outcome variable is indicated above each panel.

Figure 7. Effect of building energy codes on energy use by income quintile



Panels show difference-in-differences coefficients and 95 percent confidence intervals for each income quintile, estimated using equation (2). Standard errors are two-way clustered by dwelling and month-of-sample. Controls include year-built-by-border fixed effects, border-by-income-quintile fixed effects, zip-code-by-month-of-sample fixed effects, and border-specific linear controls for latitude and longitude. The outcome variable is indicated above each panel.

Figure 8. Effect of building energy codes on home prices by income quintile



Panels show difference-in-differences coefficients and 95 percent confidence intervals for each income quintile, estimated using equation (2). Standard errors are clustered by zip code. Controls include income-quintile-by-border fixed effects, year-built-by-border fixed effects, and zip code fixed effects. The outcome variable is indicated above each panel.