Distributional Consequences of Policies for Electric Heat Conversion*

Soren Anderson† A. Justin Kirkpatrick‡

April 17, 2024

Abstract

We study the adoption of air-source heat pumps for home heating. We estimate that 5% of already-built homes converted to heat pumps during 2010–2020, while 22% of new homes built during this period chose heat pumps. New adoptions concentrate among rural households in the South. Conversions are more diffuse, covering urban and rural households in every region. Conversions are more prevalent in areas with mild winters and cheap electricity relative to other fuels but are less strongly associated with energy costs than adoptions in new homes. To better understand the distributional implications going forward, we calculate the annual energy-cost savings from switching to a heat pump for a large sample of U.S. households based on their current heating fuels, climate, and local energy prices. We find massive variation, with low-income households in the Northeast and Appalachians benefiting the most.

Key words: heat pumps, home heating, electrification

JEL classification numbers: Q41, Q48, Q52, Q55, D63

*Acknowledgements: For helpful comments and suggestions, we thank Catie Hausman, Erica Myers, Severin Borenstein, Rob Williams, Shanjun Li, three anonymous referees, and participants in the NBER Distributional Consequences of New Energy Policies Meeting.

†Michigan State University and NBER. Address: 486 West Circle Dr, East Lansing, MI 48824. Email: sta@msu.edu

‡Michigan State University. Address: 486 West Circle Dr, East Lansing, MI 48824. Email: jkirk@msu.edu

1 Introduction

Many states have adopted net-zero carbon emissions goals. Renewable electricity generation has grown substantially over the last 20 years, reducing average carbon emissions per kWh (Holland et al. 2020). As emissions rates have fallen, attention has turned to further reducing carbon emissions by electrifying home heating, hot water heating, and cooking. Currently, fossil-fuel combustion from these activities accounts for more than 10% of all U.S. carbon emissions (RMI 2021). To promote electrification, cities in California and elsewhere have banned new natural gas hookups, and New York has become the first state to ban natural gas in new homes under its state building code.

Efforts to electrify home heating have focused on electric heat pumps, and many states have introduced policies to encourage or subsidize their use (Nadel 2020); see Berg and Cooper (2020) for a complete list of state-level policies. Dating to the 1930s (Haldane 1930), heat pumps—which work like air conditioners in reverse, transferring heat from outside to inside the home—have made marked gains in efficiency over the last 10 years. These gains have expanded the range of cold climates for which heat pumps are now suitable for home heating. New innovations, potentially spurred by U.S. government "moonshot"-style initiatives, could further improve efficiency, making heat pumps viable in even colder climates (5 degrees Fahrenheit or below). Heat pumps have advanced to the point that they are now nearly synonymous with home-heating electrification, since electric resistance heating is no longer cost-effective in most new applications. Indeed, electric heat pumps now comprise more than 50% of all residential heaters shipped by manufacturers (see online appendix A).

Who is adopting these heat pumps? Previous research has studied electrification in new homes (Davis and Kilian 2011; Davis 2021). But new homes contribute less than 1% per year to the total housing stock and, once built, houses last for decades.[1] Thus, it is crucial to understand heat pump conversions in older homes. What fraction of existing homes have switched to heat pumps in recent years? To what extent do these retrofits differ by income, geography, and demographics? To what extent are they driven by energy prices and climate? How do retrofits compare to heat pump adoptions in new homes? Finally, what are the energy-cost savings from switching to a heat pump going forward and how do these benefits correlate with income, geography, and demographics at the household level? In particular, are low-income households and other disadvantaged groups likely to benefit from policies to promote electrification of home heating? Answering these questions will provide valuable insights into the efficiency, effectiveness, and distributional impacts of policies—such as adoption subsidies or carbon taxes or changes in electricity rate design—designed to relieve energy burdens, reduce carbon emissions, or both.

We address these questions by providing a range of evidence on heat pump conversions. Our main analysis

[1] See recent U.S. Census data here: https://data.census.gov/table/ACSDP1Y2022.DP04. The data show that 11.9% of housing units in 2022 were built in the twelve years 2011-2022, while more than 35% of homes were built before 1970.

relies on the 2009, 2015, and 2020 waves of the Residential Energy Consumption Survey (RECS). We measure net conversion rates via changes in the share of existing homes of the same vintage that use heat pumps, across different RECS waves. We then explore geographic and demographic correlates of heat pump conversions and offer suggestive evidence on price-responsiveness. The RECS does not identify household location below the state level or which heating technology individual households had before adopting a heat pump. Thus, to more accurately measure the benefits from converting to a heat pump, and how these benefits correlate with income, geography, and household demographics, we turn to detailed microdata from the American Community Survey (ACS). Using granular data on energy prices and climate at the U.S. Census Public Use Microdata Area (PUMA) level, we calculate the private energy-cost savings from converting to a heat pump for all 5.9 million ACS households during 2015–2019, based on their current heating fuels and other observable characteristics. We then explore the distributional implications of these savings both for modern in-service heat pumps, as well as for improved cold-climate heat pumps currently under development

Our analysis generates four main findings. First, for existing homes built before 2010 (92% of the housing stock), we estimate a 0.05 increase in the share using heat pumps as their primary heating technology over the last decade.[2] Meanwhile, the share of new homes built in 2010 or later that use heat pumps is 0.22. Thus, recent growth in the overall share of heat pumps is dominated by conversions in existing homes ($0.92 \cdot 0.05 \approx 0.046$) rather than installations in new homes ($0.08 \cdot 0.22 \approx 0.018$). These conversions coincide with decreases in the share of homes using other forms of electricity and heating oil.

Second, conversions are more common in the South, among black households, and in homes built in or around the 1980s. They are less common among higher-income and owner-occupied households. However, in regressions that control for energy prices and these other covariates simultaneously, we largely find null results. In particular, conversions are not strongly associated with income, race, home age, or region, although they are more common in rural areas. This contrasts with newer homes, where heat pumps are less common among Native American households and in the West and more common in the South. This also contrasts with other heating technologies, where both conversions and adoptions in new homes exhibit stronger correlations with income, home ownership, urban vs. rural, region, and home age.

Third, heat pump adoptions correlate strongly with energy costs, even after controlling for urban vs. rural, region, home vintage, and detailed household demographics. These estimates suggest that a doubling of relative energy costs (competing technologies vs. heat pumps) is associated with a 0.1 increase in the share of existing homes converting to heat pumps over the last decade. Meanwhile, we find that a doubling of relative energy costs is associated with a 0.3 increase in the share of new homes choosing heat pumps. These results are consistent with higher switching costs in older homes, for example, because the duct-work is

[2] We use "heating fuel" to refer to the fuel source, and "heating technology" to refer to the heating mechanism. We refer to electrification as a policy, but acknowledge that the primary heating technology used in electrification is the heat pump.

tailored to a particular heating technology, or because the existing furnace has many years to live, or because revised building codes apply mainly to new homes. Overall, our results imply that regional differences in heat pump adoption are largely driven by regional differences in energy prices and climate.

Fourth, we find that the average private energy-cost savings from converting to a heat pump ranges widely across U.S. localities (PUMAs) from negative $772 to positive $690 per year, though this average masks significant heterogeneity across households within PUMAs. This variation is driven by vast differences in heating demand (climate), energy prices, and baseline fuel types, all of which interact to determine the benefits from switching. Leveraging our household-level calculations, we find that the private energy-cost savings are highest for low- and medium-income households and people of Native American descent. These results are consistent with higher benefits for rural households, who disproportionately rely on costly propane and heating oil to heat their homes. Using proprietary data on shipments of heating equipment, we confirm that heat pump shares largely track the distribution of private energy-cost savings.

We contribute to an economics literature on heating technology choice, nearly all of which focuses on new homes or fails to differentiate new homes from old. Davis and Kilian (2011) model heating choice for new homes as a function of energy prices. Davis (2021) decomposes the 70-year trend in home electrification for new homes, emphasizing trends in energy prices and to a lesser extent geography, climate, home attributes, and income. Davis (2023), Shen (2023), and Edwards et al. (2023) all focus specifically on heat pumps, exploring demographic correlates of heat pump ownership in the cross section, while controlling for energy prices. Like us, Davis (2023) relies on nationally representative microdata from RECS, while Shen (2023) and Edwards et al. (2023) rely on zip-code and tract-level data, which may mask important correlations. We differ in three key ways. First, we differentiate heat pump conversions from adoptions in new homes, developing methods to infer conversions by comparing across multiple RECS waves.[3] We show that conversions in the last decade account for at least 32% of the installed base and 72% of the annual flow of adoptions. These conversions are not well-explained by income or other demographics, echoing what Davis (2023) finds for the installed base. Second, we construct detailed, household-level measures of energy demand and energy costs for heat pumps versus other technologies, which reflect local climate conditions and individual housing characteristics, in addition to energy prices. Like Davis (2023), we estimate price responses using cross-sectional variation in energy costs. We show that conversions respond strongly to energy costs but are less responsive to costs than heat pump adoptions in new homes. Third, we examine the net benefits of future adoption, as well as their correlations with income and race, showing that low-income households would

[3]Hlavinka et al. (2016) and Shen et al. (2022) also focus on adoption in existing homes, exploring the role of up-front incentives and advertising. But Hlavinka et al. (2016) is limited to monthly time-series data on adoption of ductless heat pumps among homes already using electric heating in Oregon, Washington, Idaho, and Montana. Only a small fraction of homes in our nationally representative data use ductless heat pumps. Meanwhile, Shen et al. (2022) study adoption in a portion of North Carolina during 2016–2020 using several waves of Zillow's assessor-reported property data. Their inclusion of property-level fixed effects implies that their estimates reflect conversions. But it is unclear how often or how accurately assessors update information on home heating systems over short time spans.

benefit most.[4]

We also contribute to a growing literature on long-run responses to energy prices and the purchase of energy-using durable goods. Buchsbaum (2023) shows that the long-run price elasticity of residential electricity demand exceeds the short-run elasticity by a wide margin, largely through a muted response to temperature fluctuations. Likewise, Alberini et al. (2011) show that the long-run price elasticity of energy demand is large for both electricity and natural gas but is weaker for low-income households. Our results suggest that fuel-switching is one potential mechanism for large long-run effects. Like us, previous research shows that households are sensitive to long-run differences in operating costs when purchasing energy-using durable goods, including air conditioners (Rapson 2014), gasoline cars (Allcott and Wozny 2014), electric vehicles (Bushnell et al. 2022), and home appliances (Houde and Myers 2021), with important implications for electrification (Rapson and Bushnell 2022). Houde and Myers (2021) show that households in high-income areas are more sensitive to energy costs for home appliances, while Bruegge et al. (2019) show that building codes have differential effects on energy-efficiency measures across income levels. Likewise, our work explores the correlation between income and the adoption of energy-saving durable goods.

The rest of this paper proceeds as follows. Section 2 describes our data sources. Section 3 describes our methods for inferring heat pump conversions and explores the relationship between heat pump conversions and energy costs, income, and other factors. Section 4 estimates the private benefits of heat pump conversion for a large, nationally-representative sample and explores the distributional implications by race, income, and geography. Finally, section 5 concludes.

2 Data sources

This section describes the various datasets we use to study heat pump conversions in recent years and to study the distributional implications of heat pump adoption going forward. These two analyses share some of the same datasets but use them in different ways and for different purposes, as we describe below.

2.1 Residential Energy Consumption Survey (RECS)

We construct our main sample for examining heat pump conversions using the RECS surveys from 2020, 2015, and 2009. The RECS is a regular survey fielded by the U.S. Energy Information Administration (EIA). It is designed to be a representative sample of U.S. households, measuring energy equipment and home energy consumption. We focus on the lower 48 states and the District of Columbia, omitting Alaska and Hawaii. We apply RECS household sampling weights throughout so that all statistics and regression

[4] Deetjen et al. (2021) also features a calculation-based analysis, studying heat pump adoption using simulated hourly consumption for houses in 55 cities across the United States. They do not examine past trends in heat pump conversions, and focus on the public (inclusive of reduced emissions) vs. private benefit payoff. In contrast, we use detailed, observed household consumption at the annual level, and focus on the private net benefits as these are most relevant to heat pump adoption.

coefficients reflect population values.

We focus on a household's reported main heating technology, which indicates the main fuel used for space heating, the type of space heating equipment (e.g., ductless heat pump or central gas furnace), and the age of this equipment. Age is given in ranges: less than 2 years, 2–4 years, 5–9 years, 10–19 years, and 20+ years. We use this information to identify when a home's heating equipment was most recently replaced. These questions are asked in each RECS wave. We combine all RECS waves from 2009–onward to form our main sample, resulting in repeated cross-section that includes over 30,000 households.

The RECS reports a variety of home characteristics, including a categorical variable for home vintage by decade. The vintage categories are consistent across RECS waves: pre-1950, 1950–1959, 1960–1969, and so on. Thus, we are able to track changes in heating technology over time for homes of the same vintage, by comparing across RECS waves. The RECS also includes information on each household's location, though the level of geographic detail varies. The 2020 RECS reports each household's state, making it the finest level of detail. The 2009 RECS reports large states or categories of 2+ smaller states; there are 27 such areas. The 2015 RECS only reports U.S. census division; there are 10 such divisions.[5]

We also extract additional data to estimate and predict energy demand as a function temperature, home attributes, and household characteristics. These data include: energy consumption used for space heating in millions of British thermal units (MMBTU); reported heating degree days (HDD) for the survey year and a 30-year average of annual HDD for 1981–2010; home size in number of bedrooms, total number of rooms, and square footage; the IECC climate zone; and household income reported in discrete intervals. Annual HDD is calculated as 65°F minus a day's mean temperature (with negative values set to zero), summed across days in the year. This variable captures heating needs for a given local weather and climate. Energy use for space heating is approximately linear in HDD assuming a constant energy-to-heat conversion efficiency.

2.2 American Community Survey (ACS) Public Use Microdata

To examine correlates between potential benefits from heat pump adoption and demographics, we use the American Community Survey (ACS) 5-year Public Use Microdata Sample (PUMS) for 2015-2019 from the U.S. Census. The ACS PUMS is an annual survey of approximately 1% of U.S. households. While ACS data tables and summaries are commonly used for tract- and block-level analyses, the U.S. Census publishes the individual household results as well, with geography limited to Public Use Microdata Areas (PUMAs). A PUMA is a census-defined spatial area that contains roughly 100,000 individuals or 40,000–45,000 households. Urban areas might contain 10–40 PUMAs, providing a high level of spatial detail, while sparsely populated areas might have PUMAs that cover large swathes of a state. Our 5-year 1% PUMS

[5] All three surveys additionally report the International Energy Conservation Code (IECC) climate zone, which we use to exclude Alaska and Hawaii from our analysis.

sample averages about 2,500 households per PUMA.[6]

We extract variables for annual household income, race of the head of household, total number of rooms, total bedrooms, year built, and heating fuel for each house and household in the sample. We drop any household that did not report on one or more variable. We also harmonize the household income, year built, total bedrooms, and total rooms variables into bins, vintages, and counts that match RECS reporting, allowing us to relate PUMS data for 2015–2019 to RECS data for 2015 and 2020. We also extract the census-reported household weights, and match each PUMA to its IECC climate zone and the total heating degree days for 2015–2019, calculated below.

While the ACS 5-year survey reports heating fuel, it does not report heating technology. Households that report using "electricity" could either be using heat pumps or inefficient electric resistance heating, and this distinction is important in our application. Thus, using the 2020 and 2015 RECS, we calculate the share of households using heat pumps vs. electric resistance heating for each IECC climate zone and household income bin. We then match each PUMS observation to its predicted heating technology share. Thus, we capture differences in the use of heat pumps vs. electric heat by climate zone and income bin for households that report using electric heat.

We drop 14% of observations that are missing income or home characteristics data, or that are located outside the lower 48 states and District of Columbia, yielding a full PUMS sample of 5.95 million observations in 2,336 PUMAs.

Census tracts have urban vs. rural designations; PUMAs do not. Thus, we assign urban vs. rural status to each PUMA based on the population-weighted majority of the constituent census tracts, which we obtain from the most recent decennial census in which the designations were published (2010). We relabel urban clusters as urban, simplifying the designation into a dichotomous variable.

2.3 Temperature data

For temperature, we use the Oregon State PRISM Climate Group 2015–2019 daily maximum and minimum temperatures for a 4km square grid spanning the continental United States (PRISM Climate Group 2023). We calculate all temperature variables from the full-resolution 4km grid, then aggregate to the relevant geography using population weights to avoid oversampling large but sparsely populated areas. Population is obtained from the Center for International Earth Science Information Network (CIESIN) census grid files for the 2010 decennial census (Warszawski et al. 2017).

[6]ACS Public Use Microdata can be found at https://data.census.gov/mdat

2.4 Heating efficiency

In all sections, we measure consumption of "usable" heat energy in MMBTU. Likewise, we measure prices in dollars per "usable" heat energy. While energy is billed per unit of volume, energy actually received for heating depends on the heat content of the energy source, as well as the efficiency of the heating technology. We convert volumes of input energy into MMBTU using standard conversion ratios from EIA.[7] We then convert input energy into "usable" heat using a fixed efficiency factor. This factor, known as the Average Fuel Use Efficiency (AFUE), depends on the heating unit and measures energy output per energy input. We obtain historic AFUE for natural gas, propane, and heating oil furnaces from EIA (2017). For each fuel, we use the upper limit of the mid-efficiency range for the year closest to 2004, as this represents the average vintage of furnaces likely to be replaced in 2015–2020. These AFUE values are 85% for propane, 93% for natural gas, and 90% for heating oil.[8] Electric resistance heating has an AFUE of 100%. In examining potential savings from heat pump conversions, we use usable heat consumption as a technology-agnostic measure of energy demand. In doing so, we assume that heating demand does not vary with heating technology, ruling out a rebound effect. We hold the efficiency factors fixed across all of our analyses. While furnace efficiency has improved substantially over the last 50 years, it has changed little over the last 20 years. For example, the modal efficiency of a natural gas furnace shipped between 2013–2020 was 95% AFUE (O'Brien and Vondra 2023), only 2 percentage points more efficient than the 2004 efficiency measure we use.

Unlike other heating technologies, a heat pump does not generate heat for use inside a home. Rather, it transfers heat from outside to inside the home, which allows it to achieve an efficiency rating greater than 100%. Further, a heat pump does not have a constant efficiency factor. Rather, its efficiency depends on temperature, becoming less efficient at temperatures below 45°F, as captured by the Coefficient of Performance (COP). We calculate the average COP separately for every 4km cell in the PRISM data based on the local distribution of daily temperature. In the end, the COP for heat pumps ranges from 1.91 (analogous to an AFUE of 191%) in North Dakota to over 3.5 (AFUE of 350%) in warmer locations. We then aggregate to the relevant geography, weighting by population. See section B of the online appendix for details, including a map of mean COP by PUMA.

2.5 Fuel prices

Fuel prices are used to better understand the role of energy costs in heat pump adoption, and to examine correlates between potential benefits from heat pump adoption and demographics. In all cases, we obtain fuel prices from publicly available Energy Information Administration (EIA) data. We use state-level annual or monthly prices when pairing with RECS (which at best reports the state of the responding households).

[7] See here: https://www.eia.gov/energyexplained/units-and-calculators/

[8] For example, a home that uses 100 MMBTU of propane for space heating (AFUE of 85%) has usable heat consumption of $0.85 \times 100 = 85$ MMBTU. To receive the same usable heat via a natural gas furnace (AFUE of 93%) would require $\frac{85}{0.93} = 91.4$ MMBTU of natural gas.

We use utility-level prices in contexts that allow for more spatial granularity. We only use data for residential customers and we adjust all fuel prices to January 2019 using the monthly Consumer Price Index (CPI-U) published by the Bureau of Labor Statistics. When adjusting annual prices, we use the monthly CPI averaged over each year.

2.5.1 Fuel prices by state

When matching to our RECS data, we use aggregate fuel prices by state and year. We further aggregate to the 27 states or groupings of 2+ states reported in the 2009 RECS to achieve compatibility with the 2020 RECS, using the household sampling weights from the 2020 RECS.

The EIA reports annual average residential natural gas prices by state, which they calculate as annual revenue divided by annual sales volume. We likewise calculate average residential electricity prices, dividing annual revenue by annual sales volume. Revenues for natural gas and electricity include fixed charges. Thus, the average prices we use will tend to overstate the variable prices that in theory should matter for heating technology choices. However, Ito (2014) shows that consumers in California respond to average electricity prices rather than marginal prices. Thus, we use average prices to best model household choices of heating technology.

The EIA reports weekly propane and heating oil prices by state or by Petroleum Administration for Defense Districts (PADD) region during the heating months (September to March) each year. For states that do not explicitly report, we use the price of the nearest PADD. Non-reporting states tend to have low usage of the respective fuel. For instance, southern states use little heating oil and do not report weekly heating oil prices. But they do use propane, which is reported. Annual prices are required for analysis; we calculate the straight average of these prices for the year. These prices are only reported for the winter heating months. Thus, the unweighted average will tend to reflect fuel costs most relevant for home heating.

Figure 1 shows the annual means of our energy prices for 2010–2020 weighted by the households in our 2020 RECS sample. Prices for natural gas, heating oil, propane, electricity (resistance), and heat pumps are all expressed in dollars per MMBTU of usable heat, i.e. accounting for different energy units and efficiency levels across these technologies (see above). Note that electric-resistance heating has the costliest energy on average, followed by propane and heating oil. Natural gas is the cheapest energy overall, but only by a hair; heat pumps are a very close second. Electricity is the most expensive energy in dollars per MMBTU of raw energy (top line). Yet the effective cost for a heat pump is only one-third as high due to the technology's superior efficiency. These averages mask substantial variation in prices across states. We use this price variation below to estimate price responses.

Figure 1: *Energy prices for different heating technologies*

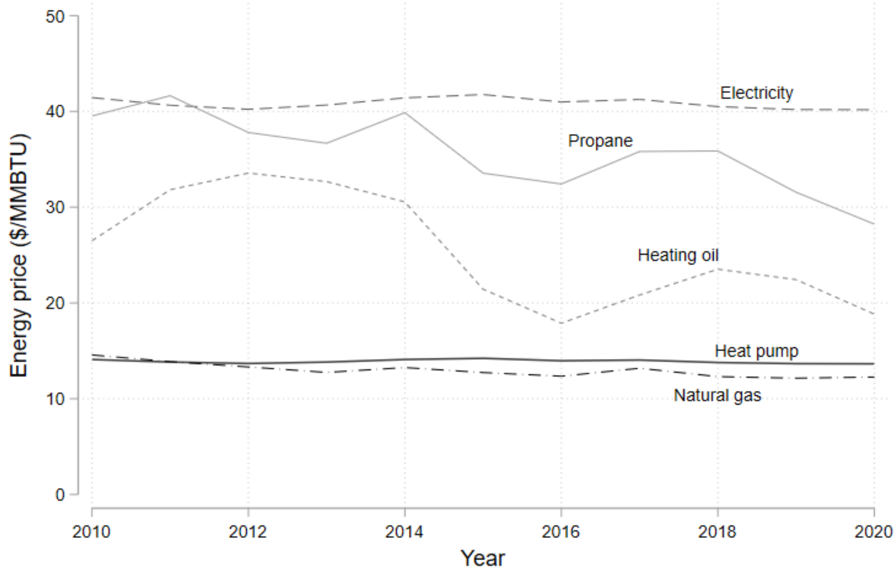

Note: This figure shows annual energy prices in dollars per usable MMBTU for natural gas, heating oil, propane, electricity (resistance), and heat pumps, accounting for the different efficiency levels of these technologies. Costs are annual means across all households in the 2020 RECS given their time-varying state-level energy prices and location-varying heat pump efficiency. Data source: Author calculations based on EIA and RECS data.

2.5.2 Fuel prices by utility

In our spatially disaggregate analyses of the potential benefits of heat pump conversion, we calculate the expected energy-cost savings from heat pump adoption by 2010 Public Use Microdata Area (PUMA) based on average fuel prices for 2015–2019. We continue to use state-level propane and heating oil prices from EIA as described above (or PADD when state is not available), but we use more spatially refined prices for natural gas and electricity available from EIA as described here.

We start by matching PUMAs to electricity and natural gas utility service territories using maps published in the Homeland Infrastructure Foundation-Level (HIFLD) Database, augmented with zipcode-level data compiled by the National Renewable Energy Laboratory (NREL).[9] These maps are compiled from EIA filings and maps published by utilities themselves. Utility territories for electricity are reported by zip code. For zip codes with multiple utilities, we select the largest utility (as measured by total number of customers). We then aggregate from zip code to PUMA. When a PUMA overlaps with more than one utility area, we use the utility of largest overlap.

[9]See here: https://hifld-dhs-gii.gov/HIFLD and https://catalog.data.gov/dataset/u-s-electric-utility-companies-and-rates-look-up-by-zipcode-2020

We obtain electricity prices for most utilities from EIA Form 861M, which reports monthly revenue and sales quantity by customer type for large utilities, plus a residual "state adjustment" that captures the non-reporting utilities. We calculate average electricity price as revenue divided by sales quantity for the October–April heating season, so that electricity prices reflect any seasonal pricing adjustments.[10] Not all PUMAs match to a reporting electric utility; in such cases we use average prices for the "state adjustment" areas. Prices are deflated to January 2019, and averaged over the period 2015–2020 when used for calculating potential benefits from heat pump conversion.

We calculate annual natural gas prices for all utilities that report on EIA Form 176. Filed annually, EIA Form 176 records total revenue and sales volume for each utility (with multi-state utilities reporting separately for each state). We calculate natural gas prices as revenue divided by sales volume, yielding a utility-specific average price. For natural gas utilities that do not report on EIA Form 176, we use the natural gas price for the nearest reporting utility. This approximates the likely price of natural gas should a utility expand to provide service to currently un-served areas. A total of 96% of all PUMAs map to a natural gas utility represented in EIA Form 176, with the plurality of un-matched PUMAs occurring in Florida, North Carolina, and Georgia.[11]

We augment the EIA Form 861M data for the state of Georgia with survey data from the Georgia Public Services Commission (GPSC).[12] The GPSC surveys Georgia's regulated utilities annually, recording bill totals for hypothetical consumption levels of 500, 1,000, 1,500, and 2,000 kWh. This survey is performed twice each year, generating a "winter" and a "summer" rate for each utility. We regress hypothetical revenues on consumption levels (four data points) for each utility, year, and season. We calculate average variable price as the mean of our slope estimates for 2014–2019. When available, we use these average prices to replace our similar estimates based on EIA Form 861M. The GPSC reports survey results for 92 utilities, of which 20 match to HIFLD-reported utility service areas.[13] As a robustness check, we also calculate prices without the GPSC data (see appendix D).

3 Trends in heat pump conversions

In this section we estimate trends in heat pump conversions over the last decade. We infer conversions from increases over time in the share of existing homes that rely on heat pumps. These approaches only allow us to infer net conversions in aggregate. We are unable to directly measure transitions from one heating technology to another, because RECS does not follow the same households over time.

[10]The heating season for heating oil and propane is reported as October to March. However, both heating oil and propane are storable and thus the purchase season is shorter than the use season.

[11]Unlike EIA Form 861M for electricity, EIA Form 176 for natural gas does not report total revenue and sales quantity for the "state adjustment" areas.

[12]See here: https://psc.ga.gov/utilities/electric/residential-rate-survey/

[13]Remaining unmatched utilities do not have unique, exclusive territories or do not appear in HIFLD maps.

Figure 2: *Heating technology by RECS wave and home vintage*

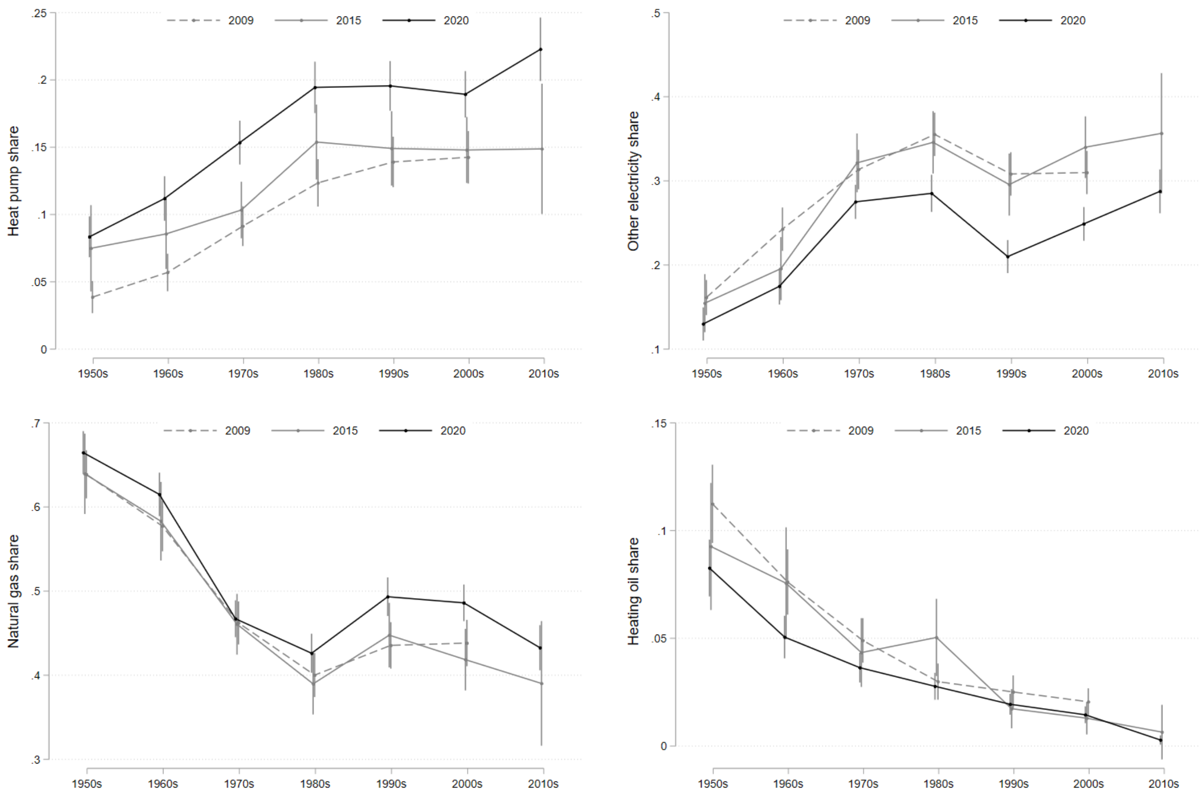

Note: This figure shows the share of households that rely on an air-source heat pump, other forms of electricity, natural gas, or heating oil conditional on home vintage for each RECS wave. In constructing the figure, we restrict the sample to households that heat their home with electricity, natural gas, heating oil, propane, or wood, and we account for RECS household sampling weights. Data source: RECS 2009, 2015, and 2020.

3.1 Inferring conversions across RECS waves

We identify net conversions over the last decade based on changes in the share of existing homes that use heat pumps as their main heating technology. We start with a pooled cross section of the RECS waves from 2009, 2015, and 2020. We restrict the sample to households whose homes were built in 2009 or earlier and who heat their home using air-source heat pumps, other forms of electricity (mainly baseboard and central electric resistance heating), natural gas, heating oil, propane, or wood. Figure 2 plots heating technology shares conditional on home vintage across the three RECS waves. The figure shows that homes are more likely to have heat pumps and natural gas and less likely to have other technologies in each subsequent wave. We interpret these changes as net conversions in heating technology. Heat pump shares increase across all vintage of homes. Meanwhile, the share of homes using other forms of electricity falls mainly among newer homes, while the share of homes using heating oil falls mainly among older homes. This finding foreshadows many of our later results: heat pump conversions are widespread but the replaced technologies vary by age

Table 1: *Regression results: heating technology by RECS wave for old and new homes*

Panel (a): Old homes

	(1) HP	(2) Elec	(3) Gas	(4) Oil	(5) Prop
Constant	0.089***	0.260***	0.507***	0.068***	0.051***
	(0.003)	(0.005)	(0.005)	(0.002)	(0.003)
RECS 2015	0.014*	-0.005	0.008	-0.015***	-0.008
	(0.006)	(0.008)	(0.009)	(0.004)	(0.004)
RECS 2020	0.050***	-0.050***	0.035***	-0.023***	-0.007*
	(0.004)	(0.006)	(0.007)	(0.003)	(0.003)
Observations	32506	32506	32506	32506	32506

Panel (b): New homes

	(1) HP	(2) Elec	(3) Gas	(4) Oil	(5) Prop
RECS 2009	0.143***	0.310***	0.438***	0.021***	0.072***
	(0.010)	(0.013)	(0.014)	(0.003)	(0.009)
RECS 2015	0.149***	0.356***	0.390***	0.006	0.071***
	(0.025)	(0.036)	(0.038)	(0.006)	(0.020)
RECS 2020	0.223***	0.288***	0.433***	0.003*	0.043***
	(0.012)	(0.013)	(0.014)	(0.001)	(0.005)
Observations	3500	3500	3500	3500	3500

Note: This table reports coefficient estimates for ten OLS regressions. The dependent variable in each column is an indicator for a given heating technology. Panel (a) reports coefficient estimates from a regression on a constant term and RECS wave dummies. This panel uses pooled data for homes built in 2009 or earlier from RECS waves 2009, 2015, and 2020. Panel (b) reports coefficient estimates from a regression on RECS wave dummies with the constant term suppressed. This panel uses pooled data for homes built in 2000 or later from the 2009 RECS and homes built in 2010 or later from the 2015 and 2020 RECS. Observations are weighted by RECS household sampling weights. Standard errors in parentheses are robust to heteroskedasticity. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. Data source: RECS 2009, 2015, and 2020.

and location, depending on the installed base of alternative technologies.

Formally, we measure conversions relative to 2009 via regression using a linear probability model:

$$y_i = \beta_0 + \beta_1 \text{RECS } 2015_i + \beta_2 \text{RECS } 2020_i + \epsilon_i, \tag{1}$$

where: y_i is a binary 0/1 variable indicating whether household i uses a given technology (heat pump, other electric, natural gas, heating oil, or propane) as their main source of heating; $RECS2015$ and $RECS2020$ are dummy variables indicating the 2015 and 2020 waves; ϵ_i is an idiosyncratic error; and the β's are parameters to be estimated. We estimate this model separately for every heating technology.

Table 1 panel (a) reports the OLS coefficient estimates. The constant term reports heating fuel shares in 2009, while the coefficients on 2015 and 2020 indicate changes relative to 2009.[14] The coefficient on the 2020 dummy in column (1) indicates a 0.050 increase in the share of homes that use heat pumps. Given the starting share of 0.089, this represents a net conversion rate of $0.050/(1 - 0.089) \approx 5.5\%$. Meanwhile, the other columns indicate net conversions toward natural gas and away from electricity, heating oil, and propane. The table omits a tiny fraction of households that heat with wood or other technologies.[15]

Note that the conversion rate is even higher if we consider that not all homes are equally at risk of replacing old or broken equipment. Figure 3 plots the distribution of heating equipment age in the 2020 RECS for homes of all vintages, and likewise for the 2009 and 2015 RECS. The age distribution has remained quite stable in recent years (compare 2020 to 2009 and 2015). The figure indicates that 46.5% of homes have equipment less than 10 years old, while one-quarter have equipment less than 5 years old. Thus, roughly half of heating systems will be replaced over the course of a decade. These values imply that the conversion rate for homes built in 2009 or earlier is closer to $5.5/0.465 \approx 12\%$ among homes at risk of replacing equipment.

For comparison, table 1 panel (b) reports choices of heating technology for recently built homes. The row labeled RECS 2009 reports estimated heating technology shares in 2009 for homes built in 2000–2009, while the rows labeled RECS 2015 and RECS 2020 report shares in 2015 and 2020 for homes built in 2010–2015 and 2010–2020. The share of new homes relying primarily on a heat pump is 0.22 in 2020 compared to just 0.14 in 2009. Meanwhile, the share relying on other forms of electricity, natural gas, heating oil, and propane are all lower in 2020 than in 2009. This evolution toward heat pumps and away from other technologies may reflect changes in where new homes are being built, along with changes in heating technology over time within a given location. Note that the 2015 shares are quite close to the 2009 shares. Thus, the table indicates a marked acceleration toward heat pumps in new homes in the latter half of the 2010s.

So which has contributed more to the recent growth in heat pumps: new or existing homes? Only 5% of existing homes switched to heat pumps during the 2010s, while 22% of new homes had heat pumps. Yet new homes accounted for just 8% of all homes in 2020, while old homes accounted for 92%. Thus, conversions in old homes contributed $0.05 \cdot 0.92 \approx 0.046$ to the overall share of homes with heat pumps in 2020, while new homes contributed just $0.08 \cdot 0.22 \approx 0.018$. Thus, a full $4.6/(4.6 + 1.8) \approx 72\%$ of the annual flow of heat pump adoptions over the last decade is due to heat pump conversions in existing homes. Meanwhile, these

[14] Table C.1 panel (a) in the online appendix shows the same information in levels: heating technology shares for each RECS wave.

[15] One potential concern is compositional changes in the location and vintage of existing homes across RECS waves, which could bias our estimates for conversion rates. Note that RECS purports to be a nationally representative survey, which should mitigate such concerns. To confirm, we replicate our analysis, including controls for urban vs. rural, census division, and home vintage. The coefficient on RECS 2020 for heat pumps in column (1) changes from 0.050 to 0.056, implying a slightly higher conversion rate. The coefficient for natural gas in column (3) changes the most, falling from 0.036 to 0.016. The other coefficients barely budge. Table C.1 panel (b) in the online appendix presents these results.

Figure 3: *Distribution of heating equipment age*

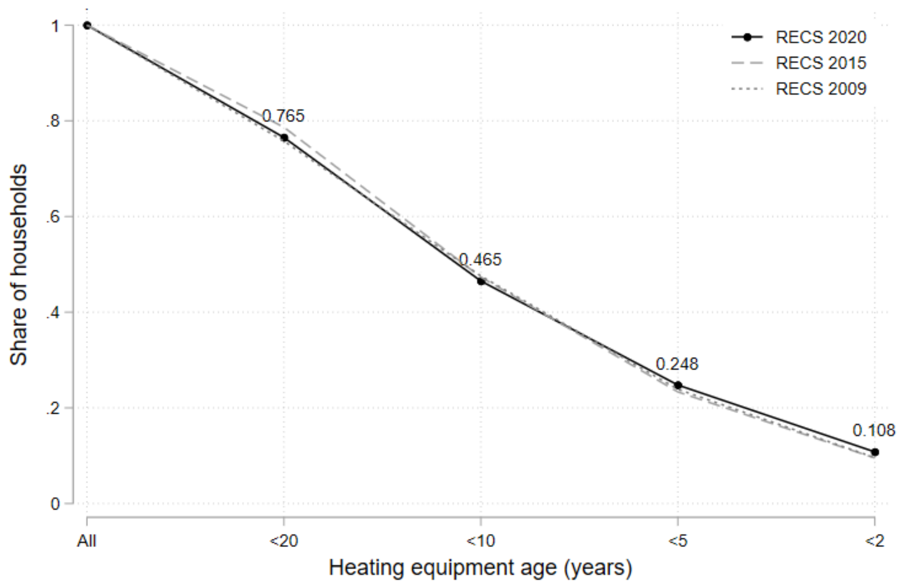

Note: This figure shows the cumulative distribution of heating equipment age in the 2009, 2015, and 2020 RECS among homes of all vintages. In constructing the figure, we account for RECS household sampling weights. Data source: RECS 2009, 2015, and 2020.

recent conversions accounted for $4.6/(0.92 \cdot 0.140 + 0.08 \cdot 0.22) \approx 31\%$ of the installed base in 2020.[16]

3.2 Correlates of heat pump conversions

In this section we explore how heat pump conversions correlate with energy costs, income, and other factors. We begin by describing how we measure energy costs for space heating for the homes in our RECS sample. We then illustrate our approach by presenting graphical evidence on the correlation between heat pump conversions and energy costs. Finally, we use a regression-based approach to reveal how heat pump conversions correlate with energy costs, geography, and demographics.

3.2.1 Measuring energy costs for space heating

We assume that households choose heating technology based in part on a comparison of relative energy costs. Thus, for every home in our 2009 and 2020 RECS sample, we calculate the expected energy cost for space heating using a heat pump vs. the expected cost using other technologies, accounting for differences in local energy prices, climate, and efficiency.

We first estimate expected heating demand (usable MMBTU) as a function of local climate and home characteristics. The RECS reports total MMBTU used for space heating. We convert this number into usable

[16]Note that 14.0% of homes built in 2009 or earlier had heat pumps in 2020. See table C.1 panel (a) in the online appendix.

MMBTU using efficiency conversion factors from EIA (2017). As MMBTU is non-negative, we estimate the following model via Poisson regression using pooled RECS data from 2009, 2015, and 2020:

$$E[MMBTU_i] = \exp(\beta_0 + \beta_1 \ln HDD_i + \beta_2 \ln sqft_i + vintage_i + zone_i), \tag{2}$$

where: HDD is annual heating degree days reported in the RECS; $sqft$ is home size in square feet; and $vintage$ and $climate$ capture home vintage and IECC climate zone fixed effects. We omit households that report using a heat pump and those reporting zero HDD.

We then use our estimated model to predict heating demand for all households in our 2009 and 2020 RECS samples (including those with heat pumps), replacing survey-year HDD with the 30-year annual mean HDD to yield a forward-looking expected value purged of annual temperature shocks. Note that predictions from a Poisson regression are unbiased in levels; in contrast, modeling log(MMBTU) with an additive error would require more complex adjustments to yield unbiased predictions in levels.

Finally, we multiply by the local price of electricity (adjusted for the local COP for heat pumps) to yield annual energy costs under a heat pump. Likewise, we multiply by the local prices of electricity, natural gas, heating oil, and propane (all similarly adjusted by their constant AFUE values) to yield expected annual energy costs under each of these technologies. Local energy prices are based on the 27 individual states or groups of 2+ states reported in the 2009 RECS. We are interested in modeling changes in heating technology that occur between 2009 and 2020. Thus, we use mean energy prices for the ten-year period 2010–2019.

In our regressions, we use logged heating costs for heat pumps. Meanwhile, we construct a composite logged "other fuel" heating cost for each home by taking a weighted average of logged costs for electricity, natural gas, heating oil, and propane, using as weights the local share of all 2009 and 2020 RECS households that choose these fuels (for the individual state or group of 2+ states).

3.2.2 Graphical evidence on heat pump conversions and heating costs

We begin with a graphical presentation to illustrate our approach, focusing on the role of energy costs. We again focus on households in the 2009 and 2020 RECS whose homes were built in 2009 or earlier and whose primary source of heating is a heat pump, other electricity, natural gas, heating oil, propane, or wood. To explore the relationship between heat pump conversions and energy costs, we estimate the following linear probability model using OLS:

$$hp_i = \beta_0 + \beta_1 \text{RECS } 2020_i + \epsilon_i, \tag{3}$$

where hp_i is an indicator for heat pump as the primary heating technology and the coefficient β_1 measures net conversions occurring between 2009 and 2020. We estimate this model separately for each state or group of 2+ states.

Figure 4 plots the resulting coefficient estimates vs. mean relative energy costs during 2010–2019 for each

Figure 4: *Heat pump conversions vs. relative energy costs by location*

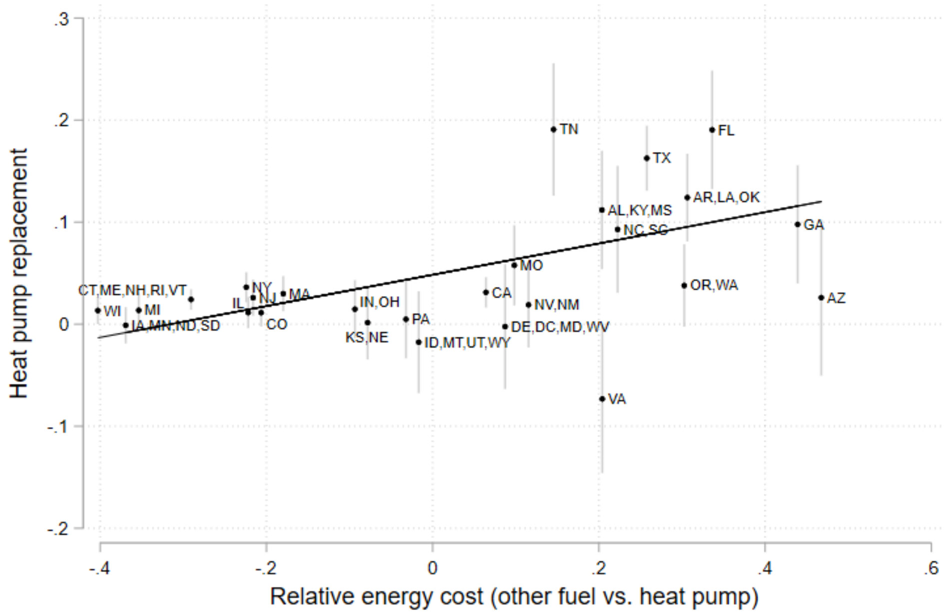

Note: This figure plots the estimated coefficient on the RECS 2020 dummy in equation (3) estimated separately for each state or group of 2+ states vs. mean relative energy costs in those same locations. The sample is all 2009 and 2020 RECS households whose homes were built in 2009 or earlier. The black dots represent the OLS point estimates, while the vertical bars represent the heteroskedasticity-robust 95% confidence intervals. The figure also shows an OLS fitted line through the point estimates, weighted by households in the 2020 RECS sample. We measure relative energy costs for each household as the weighted-average logged annual cost of heating via other fuels (electricity, natural gas, heating oil, and propane) minus the logged cost of heating via a heat pump; weights are based on heating technology shares in the 2009 and 2020 RECS for each location. We then calculate mean relative costs for each location using RECS household sampling weights.

location. There is a clear positive correlation: states with high energy costs for other fuels relative to a heat pump have seen more conversions to heat pumps in recent years. On average, a doubling of relative energy costs (moving from the far left to far right in the figure) is associated with a 0.15 increase in the share of homes converting to a heat pump.[17] Of course, this association does not control for other drivers of conversions. We therefore turn to a controlled regression below, asking whether the rate of heat pump conversions differs systematically with energy costs and other factors.

3.2.3 Interactions with energy costs and demographics

We now explore how heat pump conversions correlate with energy costs, demographics, and other factors in a more systematic way. To do this, we interact the RECS 2020 dummy with energy costs as well as

[17]The results are nearly identical when we omit households whose heating equipment is less than two years old and repeat the analysis (see online appendix C), alleviating concerns that the results are heavily influenced by the Covid pandemic.

indicators for different categories of income, race, renter vs. owner-occupied, urban vs. rural, census region, and home vintage. We continue to focus on homes built in 2009 or earlier in the 2009 and 2020 RECS. We estimate the following linear probability model using OLS:

$$hp_i = \beta_0 + \beta_1 \text{RECS 2020}_i + \gamma' X_i + \delta' X_i \cdot \text{RECS 2020}_i + \epsilon_i, \tag{4}$$

where hp_i is again an indicator for heat pump. Note that we have added controls for individual-level energy costs and other observables (X) along with interactions between the RECS 2020 dummy and these observables ($X_i \cdot$RECS 2020), where γ and δ are vectors of coefficients on the covariates and their interactions with the RECS 2020 dummy. We are mainly interested in the coefficients on the interactions (δ), which capture systematic differences in conversion rates by energy costs and other factors.

Table 2 reports the OLS coefficient estimates (δ) on the interactions between the RECS 2020 dummy and the observed covariates ($X_i \cdot$ RECS 2020). We suppress the coefficients on the new equipment dummy and the covariates themselves to focus on the interaction terms.[18]

Columns (1)–(7) explore one set of covariates at a time. Consistent with the graphical analysis above, conversions are positively associated with energy costs for other fuels relative to a heat pump (column 1). Note that the coefficient of 0.155 is nearly identical to the slope in figure 4. Conversions are more prevalent in the South (column 6). Meanwhile, they are less prevalent among homeowners (column 4) and high-income households (column 2).[19] But they are not strongly associated with major categories of race (column 3), rural vs. urban (column 5), or home vintage in the post-War period (column 7). Column (8) includes all of these variables simultaneously. The coefficient on rural turns positive. But the coefficients on income, race, owner, region, and vintage remain small or become even smaller. Thus, controlling for energy costs, heat pump conversions are distributed fairly evenly across households.[20] Note that these coefficients do not condition on a household's original heating equipment, which we do not observe. Meanwhile, heat pumps and natural gas are both positively associated with income and home-ownership in the 2009 base year, while other forms of electricity are negatively associated with these variables. Thus, the negative coefficients on high income and home-ownership might reflect weaker incentives to switch for individual households initially using low-cost natural gas and stronger incentives to switch for individual households using high-cost electricity.

The coefficient on ln(other cost) − ln(hp cost) in column (8) yields a precisely estimated coefficient of 0.101. Thus, a doubling of energy costs for other fuels relative to a heat pump is associated with an approximate 0.10 increase in the share of households converting to a heat pump over the last decade. Note that the baseline

[18] Note that all variables in this table are interactions between the RECS 2020 dummy and the indicated variable, even though the variable labels do not explicitly show the interaction with this dummy.

[19] Income categories differ across RECS waves due to shifting dollar-value cutoffs. We adjust the 2009 cutoffs for inflation and find that four of them map very closely to the $30,000, $60,000, and $100,000 cutoffs in 2020, which by coincidence correspond closely to income quartiles.

[20] The results are nearly identical when we omit households whose heating equipment is less than two years old and repeat the analysis (see online appendix C), alleviating concerns that the results are heavily influenced by the Covid pandemic.

Table 2: *Regression results: heat pump conversions in old homes*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(other cost) - ln(hp cost)	0.155***							0.101***
	(0.017)							(0.031)
$30-59k		-0.018						-0.011
		(0.012)						(0.012)
$60-99k		-0.037**						-0.015
		(0.013)						(0.012)
$100k+		-0.064***						-0.028*
		(0.012)						(0.013)
Black			0.029*					-0.007
			(0.014)					(0.014)
Nat.Am.			0.016					0.004
			(0.058)					(0.055)
Asian			0.009					0.009
			(0.018)					(0.017)
Other			0.051*					0.008
			(0.023)					(0.022)
Owner				-0.060***				-0.056***
				(0.009)				(0.009)
Rural					0.015			0.034**
					(0.013)			(0.012)
Northeast						0.010		0.009
						(0.008)		(0.008)
South						0.096***		0.050**
						(0.011)		(0.019)
West						0.010		-0.018
						(0.009)		(0.012)
1950s							0.021	0.015
							(0.012)	(0.012)
1960s							0.032*	0.018
							(0.013)	(0.012)
1970s							0.039**	0.018
							(0.013)	(0.012)
1980s							0.048**	0.025
							(0.015)	(0.014)
1990s							0.033*	0.024
							(0.015)	(0.014)
2000s							0.023	0.021
							(0.015)	(0.015)
Observations	27307	27307	27307	27307	27307	27307	27307	27307

Note: This table presents coefficient estimates from equation (4). The sample is 2009 and 2020 RECS households whose homes were built in 2009 or earlier. The dependent variable is an indicator for heat pump. The table only reports coefficients on the interactions between the RECS 2020 dummy and the variables indicated in the table; main effects are not reported here. Observations are weighted by RECS household sampling weights. Standard errors in parentheses are robust to heteroskedasticity. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. Data source: RECS 2009 and 2020.

conversion rate is 0.050 (see table 1). Thus, this association implies a conversion elasticity of $0.10/0.050 \approx 2$. This an economically meaningful effect.

While we do not find statistically large differences in the rate of heat pump conversion by demographics or location (controlling for energy costs), we do find large differences in conversions for other heating technologies, reflecting differences in baseline technology choices. For example, we find conversions away from electricity for homes built in the 1960s through the 2000s and in the South, conversions toward natural gas for homeowners and away from natural gas in rural areas, and conversions away from heating oil in the Northeast. These other technologies are generally less responsive to energy prices. See table C.3 in the online appendix, which reports results from linear probability models like column (8) estimated separately for these other heating technologies.

For comparison, we run a similar set of regressions to measure heating technology choices in new homes:

$$hp_i = \beta_0 + \delta' X + \epsilon_i, \tag{5}$$

where hp_i is an indicator for heat pump and δ measures the association between covariates X and heating technology choice. We limit our sample to 2020 RECS households whose homes were built in 2010 or later in an effort to capture original equipment choices. Nearly 90% of respondents whose homes were built in 2010 or later report heating equipment less than 10 years old.

Table 3 reports the OLS coefficient estimates. Among new homes, heat pumps are positively associated with energy costs for other fuels relative to heat pumps (column 1) and are less prevalent among high-income (column 2) and Native American households (column 3). They are more prevalent in rural areas (column 5) and much more prevalent in the South (column 6). These results contrast with those above for heat pump conversions, which are more geographically diffuse. Things again change when we include all covariates simultaneously. Heat pumps remain more prevalent in rural areas and less prevalent among among Native American households, even after controlling for energy costs. But the South shrinks in importance, while the West becomes negatively associated with adoption.

The coefficient on ln(other cost) − ln(hp cost) in column (7) implies that a doubling of energy costs for other fuels relative to a heat pump is associated with an approximate 0.33 increase in the share of new homes choosing heat pumps. Note that the baseline choice share is 0.22 (see table 1), implying an elasticity of $0.33/0.22 \approx 1.5$. Thus, while the absolute change in choice share is larger for new homes, the relative change in share (elasticity) is larger for heat pump conversions.

Meanwhile, table C.4 in the online appendix shows strong associations between the choices of other heating technologies and energy costs, owner-occupied vs. renter, urban vs. rural, and region, even controlling for energy costs. These associations may reflect large variation in the availability of natural gas across

Table 3: *Regression results: heat pump choices in new homes*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ln(other cost) - ln(hp cost)	0.391***						0.325***
	(0.045)						(0.075)
$30-59k		-0.014					-0.025
		(0.045)					(0.043)
$60-99k		-0.060					-0.075
		(0.042)					(0.041)
$100k+		-0.054					-0.050
		(0.039)					(0.040)
Black			0.036				-0.005
			(0.045)				(0.046)
Nat.Am.			-0.208***				-0.231***
			(0.024)				(0.047)
Asian			-0.049				-0.006
			(0.047)				(0.046)
Other			-0.101				-0.081
			(0.060)				(0.059)
Owner				0.011			-0.026
				(0.026)			(0.028)
Rural					0.084**		0.073**
					(0.027)		(0.027)
Northeast						-0.003	0.002
						(0.035)	(0.035)
South						0.198***	0.056
						(0.030)	(0.048)
West						-0.021	-0.110**
						(0.029)	(0.040)
Observations	1670	1670	1670	1670	1670	1670	1670

Note: This table presents coefficient estimates from equation (5). The sample is 2020 RECS households whose homes were built in 2010 or later. The dependent variable is an indicator for heat pump. The table reports coefficients on the variables indicated in the table. Observations are weighted by RECS household sampling weights. Standard errors in parentheses are robust to heteroskedasticity. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. Data source: RECS 2020.

locations.

3.3 Summary of findings for heat pump conversions

Overall, we estimate that 5% of existing homes converted to heat pumps from 2009 to 2020. We estimate that such conversions account for 72% of all heat pump installations since 2009. While 22% of new homes built since 2009 rely on heat pumps, such homes account for just 8% of the housing stock in 2020. These conversions occur in every region, among both urban and rural households, and among every income group. We infer that these conversions are replacing natural gas, heating oil, propane, and electricity, depending on which of these fuels is most prevalent locally. Finally, we find evidence that these conversions are strongly associated with the relative cost of heating via a heat pump vs. alternative fuels. In absolute terms, technology choices are more strongly associated with energy costs in new homes. But relative to baseline levels of adoption, conversions in old homes are equally or even more strongly associated with energy costs.

4 Private benefits of electrification via heat pumps

In the previous section, we newly document that the flow of heat pump adoptions is mainly driven by conversions in existing homes. We next examine the distributional implications of this trend. Our results above show that households of every race and income category have converted thus far. Yet we might expect uneven policy impacts going forward, as the energy cost savings from a heat pump vary widely depending on initial heating technology and across sub-state geographies—neither of which we observe in RECS data. Thus, to better understand the distributional implications of heat pump conversion policies going forward, we construct a highly detailed dataset of ACS Public Use Microdata Sample (PUMS) households from 2015 to 2019. This period overlaps the latter RECS waves from which we infer conversions. We give up detailed information on household energy consumption and home size from the RECS, but we gain substantial precision on location (PUMA) and a vastly larger dataset of 5.9 million observations.

We calculate the annual energy-cost savings from adopting a heat pump for each of the 5.9 million households in the PUMS dataset. For accuracy, we incorporate utility-level electricity and natural gas prices and state-level propane and heating oil prices. We impute each households's heating demand based on observed weather and home characteristics, given the relationship between these variables and heating demand estimated from 2015 and 2020 RECS data. We compute the annual cost of heating using the current fuel vs. a heat pump for 2015–2019 given the local efficiency (COP) for a heat pump and assuming a fixed demand for usable heat. We then estimate savings from heat pump adoption conditional on household income, race, and urban vs. rural status. Our main estimates do not include the fixed cost of switching, which can vary widely across baseline technologies, and even within a PUMA. But our calculations for the energy-cost savings, properly converted to a present value, can be compared to any hypothetical switching cost.[21] An advantage of using the PUMS is that it preserves correlations between heating fuel, home age, race, and income across individual households. Aggregate data are available at a finer spatial resolution (e.g. census tract) but lack this potentially important detail. In appendix D, we also study the savings from a hypothetical cold-climate heat pump that meets standards set by the Department of Energy's (DoE) Cold-Climate Heat Pump Challenge, which is a "moonshot" type effort by the DoE to accelerate the technology.

[21] As a robustness check, we incorporate rough approximations for switching costs in appendix D. Our results on distributional impacts are qualitatively unchanged.

4.1 Household usable heat consumption

We first estimate demand for usable heat as a function of local temperature and household characteristics. The RECS report total MMBTU used for space heating. We convert this number into usable MMBTU using efficiency conversion factors from EIA (2017). As MMBTU is non-negative, we fit the following Poisson regression model to RECS data from 2015 and 2020:

$$E[MMBTU_i] = \exp\left(\beta_0 + \sum_{p=0}^{3}\sum_{v=1}^{V}\beta_v^p I(v = vintage_i) HDD_i^p + \gamma^{'} X_i + \Gamma_{c(i)f(i)}\right), \tag{6}$$

where: HDD is heating degree days reported in the RECS survey data; X_i is a vector of household characteristics, including income, total rooms, and bedrooms; and Γ is a set of fixed effects for each climate zone $c(i)$ and fuel type $f(i)$. We allow a 3rd degree polynomial in HDD interacted with the home vintage to capture historic trends in home energy efficiency, some of which may vary across climates. We estimate the model omitting households that report using a heat pump and those reporting zero HDD.

Appendix table 8 presents the results from this estimation. We use this model to impute usable MMBTU for each household in the PUMS for which we observe each of the variables in equation 6. Notably, this imputation preserves correlations between income, home vintage, heating fuel, and home size.

Given imputed demand for usable heat and local energy prices, we can calculate heating cost under any initial heating fuel (propane, heating oil, natural gas, electricity) relative to a heat pump. A heat pump does not have a constant factor of efficiency. Rather, its efficiency depends on temperature, becoming less efficient at temperatures below 45°F. Thus, we calculate the average Coefficient of Performance (COP) separately for each PUMA. We discuss these methods in online appendix B.

4.2 Distributional impact of heat pump adoption

In this section we relate heat pumps savings to income, race, and urban vs. rural status. Many states have proposed or implemented programs to encourage heat pump adoption. Are these policies likely to be progressive or regressive? This depends on the income of marginal adopters (i.e., those induced to adopt by a subsidy or policy), who will tend to be the current non-adopters with the highest savings from switching to a heat pump. We measure these savings in three ways: absolute savings equal to heating costs under the household's current technology minus heating costs under a heat pump; absolute savings left-censored at zero, since the 33% of households with negative savings are unlikely to switch in any policy environment; and censored savings divided by income, to better capture relative incidence. We interpret this latter measure through an energy-justice lens.

For our main specification, we estimate the following regression using OLS:

$$savings_i = \beta_0 + \sum_{d \neq \$200k+} \beta_d^{inc} I(inc_i = d) + \sum_{e \neq white} \beta_r^{race} I(race_i = r) + \beta^{rural} I(rural_i = 1) + \theta_{p(i)} + \varepsilon_i, \tag{7}$$

Figure 5: *Correlates of annual energy-cost savings from heat pump adoption*

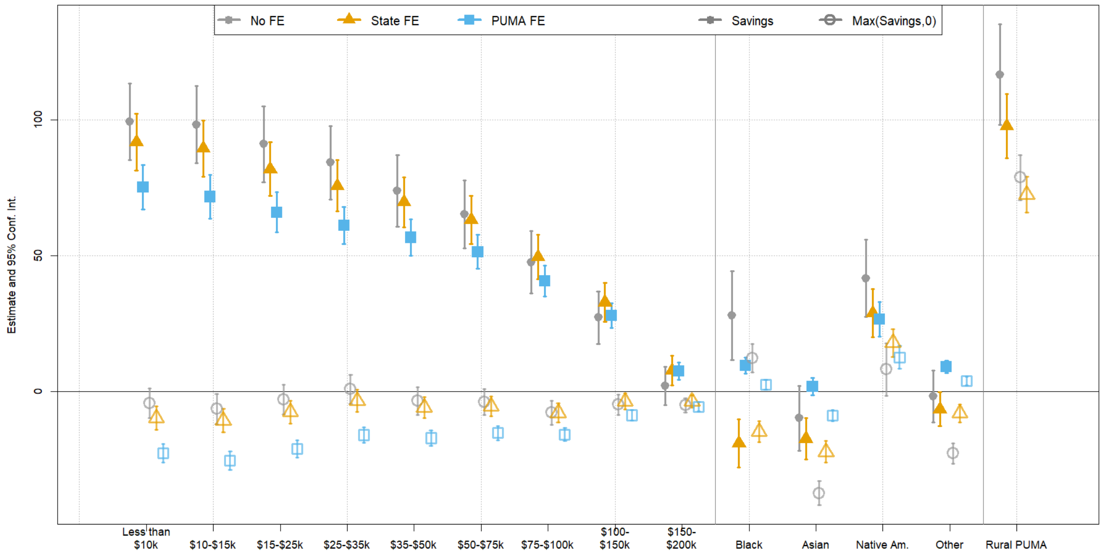

Note: This figure plots coefficient estimates from six OLS regressions of heat pump savings on income bins, race categories, and rural vs. urban status. All regressions use PUMS sampling weights. The base categories are income $> \$200k$, "White", and "urban". Different marker shapes correspond to different regressions. The outcome variable in these regressions is either absolute savings (solid markers) or savings left-censored at zero (hollow markers), and the regressions either control for no fixed effects (circles), state fixed effects (triangles), or PUMA fixed effects (squares). Positive signs indicate *greater* expected savings from heat pump adoption relative to the base categories. Vertical lines indicate 95% confidence intervals based on clustered standard errors at the PUMA level. Urban vs. rural status is assigned at the PUMA level. Thus, there is no estimate for "Rural" when using PUMA fixed effects.

where $savings_i$ is the predicted energy cost savings from adopting a heat pump for household i (or the censored or percentage measure); inc_i is income bin; $race_i$ is reported race; and θ_p is a state or PUMA fixed effect. We estimate equation (7) for absolute savings in levels, censored savings, and censored savings divided by income. If the savings from heat pump adoption are progressive, the β^{inc} coefficients will be positive and larger in magnitude for the lower income bins (note the omitted category is $200k+). In addition to income, we are interested in differences in savings for rural vs. urban households (β^{rural}) and minority vs. white households (β^{race}).

Figure 5 plots the results (see section D in the online appendix for corresponding table). Coefficients using absolute energy cost savings (solid markers) indicate that heat pump adoption is progressive, with the lowest-income households saving nearly $100 more per year than those at the high end; the pattern is smooth and monotonic across income bins. Higher-income households use more heat, but lower-income households save more from adopting a heat pump. This pattern holds conditional on state fixed effects (triangles) and

PUMA fixed effects (squares), which control for local weather, energy prices, and infrastructure. Thus, even within a PUMA, lower-income households have greater (though perhaps "less negative") savings. Coefficients on race are sensitive to the inclusion of fixed effects, except that Native American households save more than white households across all models. Native Americans are more likely to live in locations served by costly fuels like heating oil and propane, implying greater savings from a heat pump. This is reinforced by the large and positive coefficient on rural PUMA, even after controlling for state-level fixed effects.

Our finding above that heat pump adoption is progressive may only reflect *less negative* savings among low-income households, and households with negative savings are unlikely to adopt a heat pump in any case. Thus, we repeat our analysis using censored savings (hollow markers) and find that the progressive pattern reverses. In our preferred specification using PUMA-level fixed effects (square markers), the lowest-income households save about $25 less per year than those at the high end. Thus, among households that would see *positive* savings from adopting a heat pump, lower-income households save less. In addition to reversing sign, the magnitude of the effect is attenuated, since households with negative savings are assigned a value of $0 but are not dropped from estimation.

While low-income households save less in absolute terms, the effect could be different when expressed as a percentage of income. For example, a household earning $200,000 might save $100 per year, while a household earning $50,000 might save $50 per year. The low-income household saves less in dollars ($50) but more as a share of income (0.1% vs. 0.05%). Thus, we repeat our analysis using censored savings divided by income and scaled by 100 as the dependent variable. Households report a continuous measure of income, and 3% report earning less than $5,000 per year (including many zeros). To prevent unstable estimates for the bottom income bin, we reset these values to $5,000. Thus, our coefficient on this bin may be biased toward zero. Figure 6 plots the results (see section D in the online appendix for corresponding table).

We again find that the results reverse. In our preferred specification using PUMA-level fixed effects (square markers), we find that households in the second-lowest income bin save 1.1 percentage points more as a share of their income than those at the high end. Percent savings decreases monotonically as income increases. Thus, while censored savings in levels rises with income, there is a diminishing marginal effect, leading to a progressive trend when savings is measured as a percent of income. Households in rural areas also save more as a share of income.

A mass campaign to promote heat pumps via subsidies or non-pecuniary measures would primarily affect marginal adopters. Thus, regressions using censored savings are the most relevant measure for how adoption benefits are distributed. Households with positive savings that have not yet adopted (e.g., due to up-front costs or lack of information) tend to be lower-income, rural, and of Native American descent. In contrast, there is little evidence that households that have already adopted heat pumps are amassed at any particular

Figure 6: *Correlates of annual (censored) energy-cost savings as a percent of income from heat pump adoption*

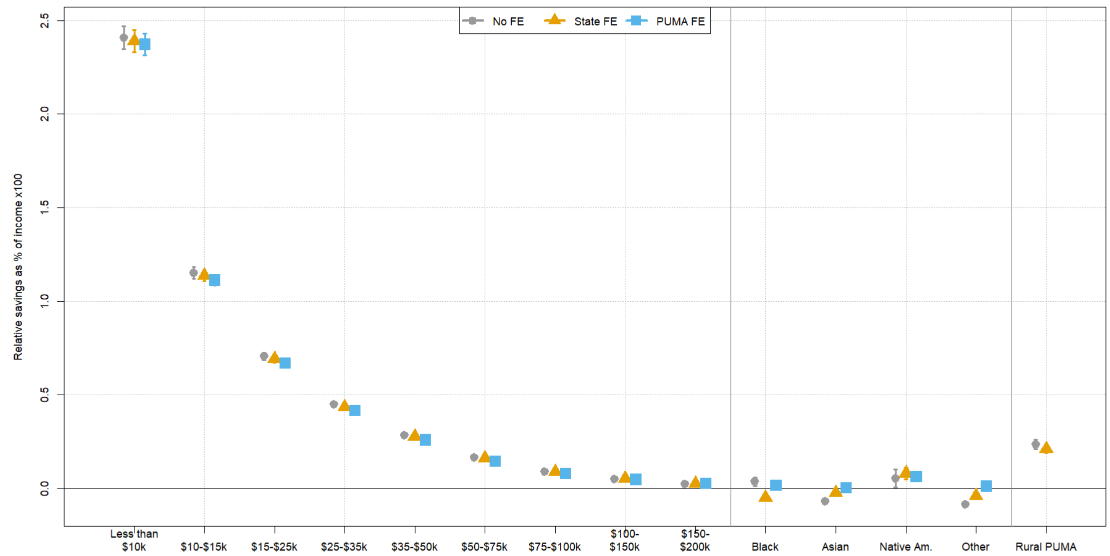

Note: This figure plots coefficient estimates from three OLS regressions of heat pump savings on income bins, race categories, and rural vs. urban status. All regressions use PUMS sampling weights. The base categories are income $> \$200k$, "White", and "urban". Different marker shapes correspond to different regressions. The outcome variable for all three regressions is censored savings ($\max\{0, savings\}$) divided by income and scaled by 100. The regressions either control for no fixed effects (circles), state fixed effects (triangles), or PUMA fixed effects (squares). Positive signs indicate *greater* expected savings from heat pump adoption relative to the base categories. Vertical lines indicate 95% confidence intervals based on clustered standard errors at the PUMA level.

part of the income distribution. While many households would not reduce energy costs by adopting a heat pump, those that would skew low-income.

4.3 Spatial variation in heat pump savings

In this section, we explore the spatial distribution of heat pump savings to illustrate how local climate, fuel prices, and fuel availability drive adoption benefits. We calculate the energy-cost savings for individual households as described above. These savings depend on which fuel the household currently uses, demand for usable heat, relative fuel prices, and the COP of a heat pump given local temperatures (see section B of the online appendix for a map of average COP by PUMA). We then calculate the average household-level savings for each PUMA, which is the finest spatial area available.[22] Note that some households within a PUMA can have positive savings while others can have negative savings, due to differences in their initial

[22]For households that report using electric heating (both heat pumps and electric resistance), we calculate savings conditional on using electric resistance, but weight each of these observations by the share of electricity-using households with a heat pump. We calculate these shares for each IECC climate zone and income level from the RECS. The resulting weighted-average savings are conditional on not already having a heat pump.

Figure 7: *Annual average savings from heat pump adoption by PUMA*

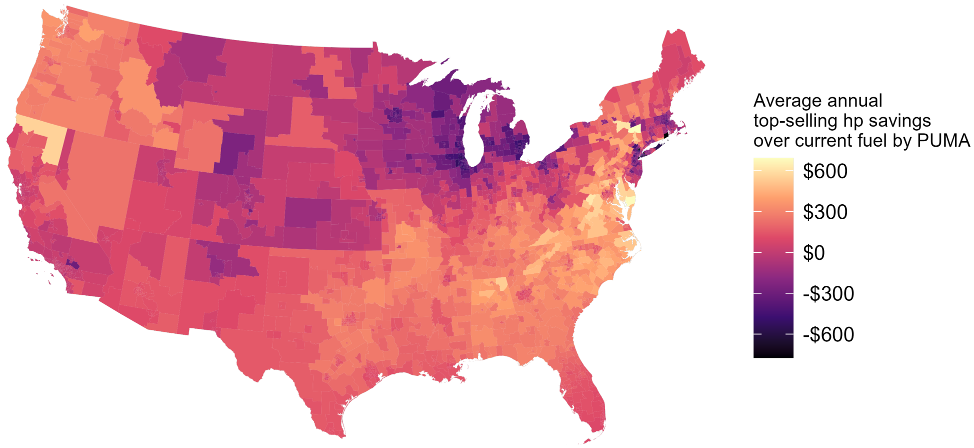

Note: This figure calculates the annual heating cost savings attributed to adoption of the top-selling heat pump. It incorporates local electricity and natural gas prices, state-level heating oil and propane prices, and weights incumbent technologies by their share in each PUMA. Usable MMBTU demand is based on the individual household PUMS data. Households that already have a heat pump contribute zero to the average savings. The negative extreme is the western Upper Peninsula of Michigan, which faces high electricity prices, low natural gas prices, and is primarily fueled by natural gas with a 67% share. The positive extreme is Sullivan and Ulster (West) Counties in New York, where nearly 67% use propane or heating oil for heating, and electricity is relative inexpensive.

heating fuel. Heat pumps are strictly more efficient than electric resistance heating in every location. Thus, the savings for households with electric resistance heating is always positive. Meanwhile, the savings for households with natural gas shows a much different pattern, with negative values in California and other places where natural gas is cheaper than electricity.[23] In California, the retail price of electricity even exceeds its the social marginal cost, working against efforts to decarbonize home heating (Borenstein and Bushnell 2018). We do not attempt to include the fixed costs of switching, which can vary widely across baseline technologies, and even within a PUMA, but our calculations for the energy-cost savings can be compared to any hypothetical adoption or switching cost.

Figure 7 shows our main map. Most PUMAs would see positive energy-cost savings from adopting a heat pump, but positive savings are far from guaranteed. Some areas around the Great Lakes, the eastern Great Plains, New York City, and Long Island would see negative savings, indicating that the current dominant fuels are less costly than electricity, even after accounting for a heat pump's efficiency. Meanwhile, areas near the Rockies, lower New England, and the Appalachian mountains would see the largest positive savings. But

[23]See section D in the online appendix, which shows average heat pump savings by PUMA separately for households using electricity, natural gas, heating oil, and propane. For households using electricity, savings exceed $2,500 in Fishers Island, NY and Block Island, RI due to the local cost of electricity delivered to the island. However, less than 15% of households in these PUMAs have electric heating. Middlesex County west of Boston, MA has the highest non-island savings of $1,844.

there is substantial within-state and within-region variation. Savings tend to be greatest in places with cold winters and high prices for baseline fuels relative to electricity, whether natural gas, propane, or heating oil. Here, the overall scale of heating demand, combined with large energy-price gaps, offsets somewhat lower COPs for heat pumps. Outside of these extremely cold areas, savings are greatest in the upper mid-Atlantic and along the Pacific Northwestern coast. In the latter, electricity prices are very low due to an abundance of federal hydropower projects, accentuating the potential savings (see section D of the online appendix for a map of electricity prices by PUMA).

Average savings range from negative \$772 per year on Block Island, Rhode Island to positive \$690 per year in Columbia and Greene Counties in New York. Savings are consistently positive in areas with high heat pump penetration, such as the Southeast (see Davis (2021)). In addition to providing heat, heat pumps also act as air conditioners, which helps explain why the Southeast has a higher adoption rate than New England, even though the savings are lower. Recent heat pump incentive programs in New York and New England coincide with positive savings. But this is not true in California, where electrification policy is at the forefront but savings are low due to high prices for electricity relative to natural gas (the state's most common heating technology by far). The same is true for New York City and Boston, where electricity prices are higher relative to the rest of New York and Massachussetts.

The Great Lakes and eastern Great Plains regions show negative savings due in part to lower efficiency for heat pumps below 45°F. However, the DoE has issued a "challenge" to manufacturers to meet higher efficiency standards in very-low temperatures. As of January 2023, no manufacturer has met the challenge's efficiency standards. But would it matter if they did? To answer this question, we calculate each household's savings from adopting a heat pump that meets the DoE's Cold Climate Heat Pump Challenge. As expected, we find that the average *gain* in heat pump savings from this technological change is highest in the coldest places, expanding the set of PUMAs with positive savings (see section D.2 of the online appendix).

To complement our calculations for heat pump savings, we leverage proprietary data on heat pump and other heating-technology shipments across much of the southern United States. We find that market shares for heat pumps are strongly negatively correlated with natural gas availability (see section D.3 of the online appendix).

5 Conclusion

We study heat pump conversions in the United States over the last decade. We estimate a 0.05 increase in the share of existing homes with heat pumps during 2009–2020. Our results imply that these conversions account for roughly three-quarters of all new heat pump adoptions during 2009–2020 and one third of the installed base in 2020. We show that these conversions are widespread, occurring throughout the income distribution, in both urban and rural areas, and in every region of the country. Thus, we build on previous studies that focus exclusively on new homes or that fail to differentiate new homes from conversions.

Our empirical analysis further shows that heat pump conversions are quite sensitive to relative energy costs, as are installations in new homes. Our measure of relative costs accounts for state variation energy prices, household-level variation in home heating demand based on climate and home characteristics, and the declining performance of heat pumps in cold climates. Thus, we build on previous studies that only consider the response to energy prices. One limitation of our approach is that we are only able to identify net conversions.

Finally, our calculations reveal large geographic dispersion in the private energy-cost savings from adopting a heat pump, driven by local energy prices and climate, ranging from negative \$772 to positive \$690. Cross-sectional correlations suggest that heat pump conversions would disproportionately benefit rural areas with large shares of low-income households, who tend to rely on expensive propane, heating oil, and electric-resistance heating. However, conditional on having positive annual savings, heat pumps savings are regressive, with greater benefits accruing to higher-income households. Expressed as a percent as income, low-income households (conditional on having positive savings) have larger savings from adopting a heat pump relative to high-income households. To explore the distributional implications of technical change, we consider a hypothetical heat pump that meets the Department of Energy's Cold-Climate heat pump challenge for a more efficient cold-climate heat pump. We find that this technology generates small but meaningful savings that are broadly shared across all demographics.

Our current analysis takes electricity and natural gas prices as given. Yet it is becoming increasingly clear that variable prices deviate substantially from private costs in many utility areas, due to inefficient two-part tariffs and steeply increasing block-rate schedules (Borenstein and Bushnell 2018). Thus, future research should consider distributional implications under the counterfactual assumption of efficient two-part tariffs.

References

Alberini, Anna, Will Gans, and Daniel Velez-Lopez. 2011. Residential consumption of gas and electricity in the us: The role of prices and income. *Energy Economics* 33 (5): 870–881.

Allcott, Hunt, and Nathan Wozny. 2014. Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics* 96 (5): 779–795.

Berg, Weston (ACEEE), and Emma (ACEEE) Cooper. 2020. State Policies and Rules to Enable Beneficial Electrification in Buildings through Fuel Switching — ACEEE. Technical report, ACEEE.

Borenstein, Severin, and James B Bushnell. 2018. Do two electricity pricing wrongs make a right? cost recovery, externalities, and efficiency. Technical report, National Bureau of Economic Research.

Bruegge, Chris, Tatyana Deryugina, and Erica Myers. 2019. The distributional effects of building energy codes. *Journal of the Association of Environmental and Resource Economists* 6 (S1): S95–S127.

Buchsbaum, Jesse. 2023. Are consumers more responsive to prices in the long run? evidence from electricity markets .

Bushnell, James B, Erich Muehlegger, and David S Rapson. 2022. Energy prices and electric vehicle adoption. Technical report, National Bureau of Economic Research.

Davis, Lucas. 2021. What matters for electrification? evidence from 70 years of us home heating choices. *NBER Working Paper* (w28324).

———. 2023. The economic determinants of heat pump adoption. *NBER Environmental and Energy Policy and the Economy* .

Davis, Lucas W, and Lutz Kilian. 2011. The allocative cost of price ceilings in the us residential market for natural gas. *Journal of Political Economy* 119 (2): 212–241.

Deetjen, Thomas A., Liam D Walsh, and Parth Vaishnav. 2021. Us residential heat pumps: the private economic potential and its emissions, health, and grid impacts. *Environmental Research Letters* 16.

Edwards, Morgan R., Jaime Garibay-Rodriguez, Jacob Erickson, Muhammad Shayan, Jingwen Tan, Xingchi Shen, Yueming (Lucy) Qiu, et al. 2023. Assessing inequities in electrification via heat pumps across the u.s.

EIA, US. 2017. Residential end uses: historical efficiency data and incremental installed costs for efficiency upgrades .

Haldane, T. G. N. 1930. The heat pump—an economical method of producing low-grade heat from electricity. *Journal of the Institution of Electrical Engineers* 68: 666–675.

Hlavinka, Alexander N, James W Mjelde, Senarath Dharmasena, and Christine Holland. 2016. Forecasting the adoption of residential ductless heat pumps. *Energy Economics* 54: 60–67.

Holland, Stephen P, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates. 2020. Decompositions and policy consequences of an extraordinary decline in air pollution from electricity generation. *American Economic Journal: Economic Policy* 12 (4): 244–74.

Houde, Sébastien, and Erica Myers. 2021. Are consumers attentive to local energy costs? evidence from the appliance market. *Journal of Public Economics* 201: 104480.

Ito, Koichiro. 2014. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review* 104 (2): 537–563.

Nadel, Steven (ACEEE). 2020. Programs to Electrify Space Heating in Homes and Buildings — ACEEE. Technical report, ACEEE.

O'Brien, Brendan, and Miles Vondra. 2023. Empowering consumer choices: Analyzing the impact of the energy star program on the adoption of high- efficiency gas appliances. Technical report, American Gas Association.

PRISM Climate Group. 2023. PRISM Gridded Climate Data.

Rapson, David. 2014. Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior. *Journal of Environmental Economics and Management* 68 (1): 141–160.

Rapson, David S, and James B Bushnell. 2022. The electric ceiling: Limits and costs of full electrification. Technical report, National Bureau of Economic Research.

RMI. 2021. It's Time to Incentivize Residential Heat Pumps - RMI.

Shen, Xingchi. 2023. What matters for the racial disparity in clean heating technology adoption? evidence from U.S. heat pumps .

Shen, Xingchi, Yueming Lucy Qiu, Pengfei Liu, and Anand Patwardhan. 2022. The effect of rebate and loan incentives on residential heat pump adoption: evidence from north carolina. *Environmental and Resource Economics* 82 (3): 741–789.

Warszawski, L, K Frieler, V Huber, F Piontek, O Serdeczny, X Zhang, Q Tang, et al. 2017. Center for international earth science information network—ciesin—columbia university.(2016). gridded population of the world, version 4 (gpwv4): Population density. palisades. ny: Nasa socioeconomic data and applications center (sedac). doi: 10. 7927/h4np22dq. *Atlas of Environmental Risks Facing China Under Climate Change* : 228.