Do Building Energy Codes Accelerate the Heat Pump Transition?
Evidence from Swiss Cantonal Adoption of MuKEn 2014

APEP Autonomous Research*

March 5, 2026

Abstract

Building energy codes are a central pillar of climate policy, yet rigorous evidence on their effectiveness at shifting heating technology adoption remains limited. I exploit the staggered adoption of Switzerland's MuKEn 2014 model energy code across cantons—with cohorts from 2017 to 2022 observed as treated—to estimate the causal effect on heat pump adoption. Using canton-level building registry data, I find that the effect is small and sensitive to estimator choice: a heterogeneity-robust Sun-Abraham estimator yields an ATT of 0.27 percentage points ($p = 0.009$), while the standard TWFE estimate is 0.69 percentage points ($p = 0.40$) and alternative inference (wild cluster bootstrap, randomization inference) cannot reject zero. The findings suggest that secular trends—driven by federal subsidies, energy prices, and technological change—dominate the transition away from fossil heating, with building codes playing at most a modest incremental role.

JEL Codes: Q48, Q58, R38
Keywords: building energy codes, heat pumps, energy transition, difference-in-differences, Switzerland

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

Between 2009 and 2022, the share of Swiss buildings heated by heat pumps nearly doubled—from under 10% to over 18%. This transition unfolded while Switzerland's 26 cantons were gradually adopting a major new building energy code, MuKEn 2014. But the timing of the heat pump boom has remarkably little to do with the timing of the regulation. Cantons that adopted the code early look almost indistinguishable from those that adopted late or not at all. The building code appears to be codifying a transition that market forces—falling heat pump costs, a rising carbon levy, generous subsidies—were already delivering.

The question of whether building energy codes effectively accelerate technology transitions matters far beyond Switzerland. Dozens of countries have adopted or strengthened building energy standards as a core component of their climate strategies. The European Union's Energy Performance of Buildings Directive, California's Title 24, and similar frameworks worldwide rest on the assumption that regulatory standards drive adoption of efficient technologies. If codes merely codify transitions that market forces would deliver independently, the policy case for tighter regulation is weaker than advocates claim—though codes may still serve important coordination and signaling functions.

This paper provides the first causal evaluation of Switzerland's MuKEn 2014 (*Mustervorschriften der Kantone im Energiebereich* 2014—Model Prescriptions of the Cantons in the Energy Sector) building energy code and its effect on heat pump adoption. MuKEn 2014 represents a major tightening of building energy standards, mandating minimum renewable energy shares for new construction and imposing efficiency requirements on heating system replacements. The code was developed as a model framework by the Conference of Cantonal Energy Directors (EnDK) but must be individually adopted by each of Switzerland's 26 cantons into their own legislation—a process that has unfolded gradually from 2017 to 2024, with one canton (Solothurn) rejecting the code entirely via referendum. This staggered adoption creates the quasi-experimental variation that enables causal identification.

I construct a canton-by-year panel from the Swiss Federal Statistical Office's (BFS) Register of Buildings and Dwellings (GWR), which records the primary heating system for every building in Switzerland. The data cover seven annual cross-sections from 2009 to 2015 (the old survey regime) and two years from 2021 to 2022 (the new registry regime), providing a long pre-treatment window and post-treatment observations. I supplement the building-count data with surface-area data on heating systems for 2021–2023. The treatment variable is constructed from official cantonal adoption dates of MuKEn 2014, verified against the EnDK implementation tracker and cantonal legislation.

The empirical strategy is a two-way fixed effects (TWFE) difference-in-differences design, comparing changes in heat pump adoption across cantons that adopted MuKEn 2014 at different times. Canton fixed effects absorb time-invariant cantonal characteristics—geography, housing stock composition, political orientation—while year fixed effects capture nationwide trends in energy prices, federal subsidies, and technology costs. The identifying assumption is that, absent MuKEn 2014 adoption, treated and control cantons would have followed parallel trends in heat pump adoption. I assess this assumption using pre-treatment trend data and conduct extensive robustness checks.

The main finding is that MuKEn 2014's effect on heat pump adoption is small and sensitive to the choice of estimator. A heterogeneity-robust Sun-Abraham estimator (Sun and Abraham, 2021), which avoids the biases of standard TWFE under staggered treatment (Goodman-Bacon, 2021), yields an average treatment effect on the treated (ATT) of 0.27 percentage points ($p = 0.009$). The standard TWFE estimate is larger at 0.69 percentage points but statistically insignificant ($p = 0.40$, 95% CI $[-0.90, 2.28]$). Alternative inference approaches—wild cluster bootstrap ($p = 0.42$) and randomization inference ($p = 0.45$)—cannot reject zero for the TWFE specification. A Bacon decomposition confirms that 76% of the identifying weight comes from clean treated-versus-untreated comparisons, with minimal contamination from problematic later-versus-earlier comparisons. Even under the Sun-Abraham estimate, the effect accounts for a small fraction of the 7–8 percentage point increase in heat pump share observed across all cantons between the pre- and post-treatment periods.

I do find significant effects on fossil fuel and gas heating shares, which decline by 1.4 and 1.6 percentage points respectively in treated cantons ($p < 0.05$). However, a placebo test on wood heating—which should not be directly targeted by MuKEn 2014—also yields a significant negative coefficient (-2.2 pp, $p = 0.02$), raising concerns about confounding from correlated cantonal policies or broader compositional shifts that coincide with MuKEn adoption timing. This pattern suggests that the fossil fuel and gas results may partially reflect general energy transition dynamics in cantons that are politically more inclined toward environmental regulation, rather than a causal effect of the building code itself.

Heterogeneity analysis reveals no meaningful differences between early adopters (before 2020) and late adopters (2020 and later). A long-difference specification, comparing the change in heat pump share from the 2009–2015 average to the 2021–2022 average against years of MuKEn exposure, confirms the small-effect interpretation: cantons with more years of code exposure show no significantly greater heat pump adoption than those with less exposure. Using surface-area data for 2021–2023, a marginally significant dose-response emerges ($\beta = 0.30$ percentage points per year, $p = 0.08$), but this effect is modest and sensitive to specification.

This paper contributes to three distinct literatures. First, it adds to the growing body of work evaluating building energy codes. Levinson (2016) finds that California's building codes reduce energy consumption by substantially less than engineering estimates predict, while Jacobsen and Kotchen (2016) documents significant effects of New Mexico's codes on energy use. Kotchen (2017) provides longer-run evidence suggesting codes do reduce residential consumption but with smaller effects than expected. Davis (2014) surveys the economics of building codes more broadly. My contribution is to shift the outcome from energy consumption to technology adoption—specifically, whether codes accelerate the transition from fossil to renewable heating systems, a first-order question for decarbonization.

Second, the paper speaks to the literature on energy efficiency gaps and policy effectiveness. Allcott and Greenstone (2014) and Gerarden et al. (2017) discuss why energy-efficient technologies may be underadopted, while Fowlie et al. (2018) provides experimental evidence from the Weatherization Assistance Program showing that realized energy savings fall short of projections. Myers (2022) documents heterogeneous effects of efficiency investments that explain much of the gap between intentions and outcomes. My finding of at most a small effect is consistent with the view that market-based factors—energy prices, subsidies, technological maturation—rather than regulatory mandates drive technology adoption decisions.

Third, the paper contributes to the literature on directed technical change and environmental regulation. Acemoglu et al. (2012) develops the theory of directed technical change in the environmental context, while Aghion et al. (2016) provides empirical evidence that carbon taxes redirect innovation toward clean technologies. Noailly and Smeets (2015) shows that innovation in renewable energy responds to energy prices and regulation. Popp et al. (2010) surveys the broader literature on energy, environment, and technological change. My findings suggest that command-and-control building regulations may be less effective than price signals at driving technology transitions, echoing the broader insight from Gillingham and Stock (2018) on the relative cost-effectiveness of different climate policy instruments.

Finally, this paper is among the first to exploit Switzerland's cantonal federalism—26 semi-autonomous jurisdictions with substantial policy-making authority—for causal identification in climate policy evaluation. Brunner and Iten (2012) examines the Swiss CO_2 Act at the cantonal level, and Swiss Federal Office of Energy (2021) discusses the Swiss heating transition policy landscape. The staggered adoption of a harmonized model code across linguistically and politically diverse cantons provides a clean natural experiment that avoids many of the confounds plaguing cross-country comparisons.

The Swiss setting offers several advantages for causal identification. Unlike cross-country comparisons, which are confounded by differences in energy markets, climate, housing traditions, and political institutions, the within-country comparison across Swiss cantons holds these factors largely constant. The 26 cantons share a common legal framework, currency, labor market, and energy infrastructure, differing primarily in language, topography, and political preferences. The staggered adoption of a harmonized model code—rather than 26 idiosyncratic policies—further strengthens the research design by ensuring that the treatment is comparable across units.

The finding that building codes have limited impact on heat pump adoption is particularly relevant in the current European policy context. The EU's revised Energy Performance of Buildings Directive (EPBD), adopted in 2024, mandates that all new buildings be zero-emission by 2030 and that existing buildings reach zero-emission status by 2050. Several EU member states are debating or implementing bans on fossil-fuel heating systems in new construction. Switzerland's experience—where market forces and price signals appear to dominate regulatory mandates in driving adoption—offers a cautionary tale for policymakers who rely primarily on building codes to achieve decarbonization targets.

The research design also speaks to a broader methodological point about evaluating policies in markets with strong secular trends. When an outcome variable is on a steep trajectory—as heat pump adoption has been across developed countries—regulatory interventions may appear effective in simple before-after comparisons but fail to demonstrate incremental impact in rigorous difference-in-differences designs. This "trend attribution problem" is a first-order concern for climate policy evaluation, and this paper illustrates both the challenge and a credible approach to addressing it.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of MuKEn 2014 and the Swiss energy policy landscape. Section 3 describes the data sources and variable construction. Section 4 presents the empirical strategy. Section 5 reports the main results and robustness checks. Section 6 discusses mechanisms, alternative explanations, and policy implications. Section 7 concludes.

2. Institutional Background

2.1 Swiss Energy Federalism

Switzerland's federal structure gives cantons primary authority over building regulation, including energy efficiency standards. Unlike most European countries where national building codes apply uniformly, Swiss cantonal sovereignty means that 26 distinct jurisdictions set their own rules for new construction and renovation. This decentralized system creates inherent variation in policy stringency and timing.

To coordinate cantonal energy policy, the Conference of Cantonal Energy Directors (Konferenz Kantonaler Energiedirektoren, EnDK) periodically develops model energy prescriptions—

the MuKEn (*Mustervorschriften der Kantone im Energiebereich*). These model codes are non-binding templates that cantons may adopt, adapt, or reject. The EnDK published its first MuKEn in 2000, updated it in 2008, and released the current version—MuKEn 2014—in January 2015.

The adoption process requires each canton to translate the model code into cantonal law, typically through legislative or executive action. In some cantons, energy legislation is subject to referendum. This process introduces variation both in timing (when each canton adopts) and in the specific provisions incorporated (some cantons adopt the full package, others select modules). The referendum channel introduces an additional source of political variation: Solothurn voters rejected MuKEn 2014 in a 2018 popular vote, making it the only canton that has not adopted the code.

2.2 MuKEn 2014: Key Provisions

MuKEn 2014 represents a substantial tightening of building energy standards relative to MuKEn 2008. The key provisions relevant to heating system choice include:

- **New construction requirements.** New buildings must meet near-zero energy standards, with weighted energy demand for heating not exceeding 3.5 liters of oil equivalent per square meter. This effectively requires high insulation levels and efficient heating systems.

- **Renewable energy mandates for new buildings.** At least 10% of energy demand in new residential buildings must be met from renewable sources. While this can be achieved through solar thermal, photovoltaics, or biomass, heat pumps are the most common compliance pathway.

- **Heating system replacement rules.** When replacing a heating system in an existing building, the new system must meet current efficiency standards. For fossil-fuel replacements, the building must either improve its thermal envelope or compensate through renewable energy generation.

- **Energy certificates.** The Cantonal Energy Performance Certificate of Buildings (GEAK) becomes mandatory for new buildings and upon sale, increasing transparency about energy performance.

- **Electricity self-consumption.** New buildings with electric heating (including heat pumps) must produce a minimum share of electricity on-site, typically through rooftop photovoltaics.

While MuKEn 2014 does not explicitly mandate heat pumps, its combination of efficiency requirements, renewable energy minimums, and replacement rules makes heat pumps the path of least resistance for compliance—particularly in new construction and when replacing aging oil boilers.

2.3 Staggered Adoption Timeline

The adoption of MuKEn 2014 across cantons occurred gradually from 2017 to 2024. Basel-Stadt was the first canton to adopt in 2017, followed by Basel-Landschaft, Obwalden, and Vaud in 2018. By the end of 2020, 16 cantons had adopted the code. The remaining cantons followed between 2021 and 2024, with Ticino and Zug adopting last in 2024. Table 9 in the Appendix provides the complete timeline.

The variation in adoption timing reflects both political and administrative factors. Cantons with strong Green or center-left governing coalitions tended to adopt earlier. French-speaking and urban cantons were generally faster than rural German-speaking cantons. Administrative capacity—the presence of dedicated cantonal energy offices—also influenced speed. This non-random selection into early adoption is a potential threat to identification that I address in the empirical strategy.

2.4 Concurrent Policies

MuKEn 2014 does not operate in a policy vacuum. Several concurrent federal and cantonal policies also affect heating system choice:

- **Federal Buildings Program** (*Gebäudeprogramm*). Since 2010, the federal government and cantons jointly fund energy efficiency improvements and heating system replacements. Subsidies for heat pump installations range from CHF 2,000 to CHF 10,000 depending on the canton and building type.

- **CO_2 levy.** Switzerland imposes a carbon tax on fossil heating fuels (currently CHF 120 per tonne CO_2), which raises the operating cost of oil and gas heating relative to heat pumps.

- **Cantonal subsidy programs.** Several cantons offer additional incentives for renewable heating systems beyond the federal program, creating within-canton stacking of subsidies.

- **Energy price shocks.** The 2021–2022 energy crisis dramatically increased oil and gas prices, accelerating heat pump adoption nationwide. This is a potential confound that I address through robustness checks excluding 2022.

The key identification challenge is separating the effect of MuKEn 2014 from these concurrent policies. Federal policies (the Buildings Program, the CO_2 levy) are absorbed by year fixed effects since they apply uniformly across cantons. Cantonal subsidies are more problematic but are partially absorbed by canton fixed effects to the extent that they are stable over time. The energy price shock affects all cantons simultaneously and is captured by year fixed effects, but I also test robustness to excluding the 2022 crisis year.

2.5 The Swiss Heating Market

Understanding the small estimated effect requires context about Switzerland's heating market. The Swiss building stock is old by European standards: approximately 1.5 million of the country's 1.8 million residential buildings were constructed before 1990, when energy efficiency standards were minimal. Oil boilers dominated heating installations from the 1960s through the 1990s, and as of 2009, oil remained the primary energy source for roughly 50% of Swiss residential buildings.

Heat pump technology began gaining market share in Switzerland in the early 2000s, initially driven by favorable geology (widespread availability of ground-source heat), relatively high electricity reliability, and growing environmental consciousness. By 2009, heat pumps already served approximately 10% of buildings nationally, rising to about 14% by 2015 and exceeding 17% by 2022. This trajectory reflects a compound annual growth rate of approximately 4–5% over the entire period, with acceleration in the 2020s driven by the energy crisis and maturing technology.

The replacement cycle for heating systems is long—typically 20–25 years for oil and gas boilers. This means that building codes primarily affect new construction decisions and major renovation projects. In a given year, new residential construction adds roughly 10,000–15,000 buildings to the stock, while major renovations affect perhaps 30,000–50,000 existing buildings. Together, these flows represent less than 4% of the stock annually, which constrains the speed at which any building code can shift aggregate statistics.

Three factors distinguish the Swiss heat pump market from other European countries. First, Switzerland's electricity is largely carbon-free (approximately 60% hydropower, 30% nuclear as of 2022), making heat pumps genuinely low-emission from a lifecycle perspective. Second, the geography of the Swiss Plateau—moderate winter temperatures, suitable geology for ground-source systems—is favorable for heat pump performance. Third, the high Swiss wage level means that the operating cost advantage of heat pumps over oil (lower fuel costs per kWh of heat) outweighs the higher capital cost relatively quickly, particularly given the CO_2 levy on fossil fuels.

3. Data

3.1 Building Registry Data

The primary data source is the Swiss Federal Statistical Office's (BFS) Register of Buildings and Dwellings (Gebäude- und Wohnungsregister, GWR). This administrative register records the primary energy source for heating for every building in Switzerland. The GWR is updated continuously by cantonal building authorities and provides the most comprehensive enumeration of the Swiss building stock available.

I use two data products derived from the GWR. The first is the BFS annual building statistics for 2009–2015, which report the number of buildings by primary heating energy source (oil, gas, electricity, wood, heat pump, solar, district heating, other) for each canton in each year. These data provide seven annual pre-treatment observations per canton.

The second product is the BFS building overview tables for 2021 and 2022, which report building counts by primary energy source at the cantonal level using the reformed GWR classification system. The 2016–2020 period is a data gap: the BFS transitioned from the old survey-based system to the new register-based system during these years, and comparable annual statistics are not available.

The resulting panel contains 26 cantons $\times$ 9 years (2009–2015, 2021, 2022) = 234 canton-year observations. While the 2016–2020 gap is a limitation, it coincides with the early adoption period of MuKEn 2014, meaning that most treatment turns on between the pre-period (ending 2015) and the post-period (beginning 2021). The gap does not prevent identification because the TWFE estimator uses within-canton variation over time and within-year variation across cantons.

3.2 Surface Area Data

I supplement the building-count data with BFS data on heated surface area by energy source for 2021–2023. These data report the percentage of total heated surface area served by each heating system in each canton. Surface area weights buildings by size, providing a complementary measure to simple building counts. This three-year panel offers an additional source of within-period variation.

3.3 Treatment Assignment

Treatment is defined as the year in which each canton formally adopted MuKEn 2014 into cantonal law. Adoption dates are compiled from the EnDK's implementation tracker (*Umsetzungsstand der MuKEn 2014 in den Kantonen*) and verified against cantonal official gazettes. The binary treatment indicator D_{ct} equals one for canton c in year t if t is at or after the canton's adoption year, and zero otherwise.

Twenty-five of 26 cantons adopted MuKEn 2014 between 2017 and 2024. Solothurn rejected the code via referendum and serves as a permanent control unit. Crucially, the building-count panel ends in 2022, so cantons that adopted MuKEn 2014 after 2022—Bern and Geneva (2023), Ticino and Zug (2024)—have $D_{ct} = 0$ for all observations and are effectively never-treated in the main specifications. This yields 18 treated cantons and 8 control cantons (including these four late adopters and Solothurn) in the building-count regressions. Because identification does not rely solely on Solothurn, the staggered adoption across the other cantons provides the primary source of variation.

3.4 Variable Construction

The primary outcome is the *heat pump share*—the number of buildings with heat pumps as their primary heating system divided by the total number of buildings in the canton-year. I construct analogous shares for oil, gas, fossil fuels (oil + gas + coal), wood, and district heating.

For treatment intensity, I construct *years treated*—the number of years since adoption, censored at zero for pre-adoption periods and for never-treated units. For the long-difference specification, I compute *years of exposure*—the number of years between MuKEn adoption and the end of the post-period (2022).

3.5 Summary Statistics

Table 1 reports summary statistics by treatment status and period. The average heat pump share among treated cantons (those adopting by 2021) increased from 10.5% in the pre-period to 18.5% in the post-period—a gain of 8.0 percentage points. Control cantons experienced a similar increase, from 8.2% to 15.5%—a gain of 7.3 percentage points. The raw difference-in-differences is therefore approximately 0.7 percentage points, foreshadowing the small TWFE estimate.

Oil heating declined across both groups, from roughly 45% to 35% in treated cantons and from similar levels in controls. Gas heating also declined modestly. These trends are consistent with a nationwide energy transition driven by factors common to all cantons.

Table 1: Summary Statistics: Heating Systems by Treatment Status

	Treated Cantons		Control Cantons	
	Mean	SD	Mean	SD
Panel A: Pre-Treatment (2009–2015)				
Total buildings	48,621	—	99,254	—
Heat pump share	0.105	0.053	0.082	0.031
Oil share	0.475	0.076	0.504	0.057
Gas share	0.135	0.124	0.162	0.086
Fossil share	0.611	0.116	0.667	0.118
Canton-years	126		56	
Cantons	18		8	
Panel B: Post-Treatment (2021–2022)				
Total buildings	51,887	—	105,721	—
Heat pump share	0.185	0.069	0.155	0.032
Oil share	0.380	0.082	0.408	0.047
Gas share	0.151	0.130	0.192	0.100
Fossil share	0.531	0.115	0.600	0.118
Canton-years	36		16	
Cantons	18		8	

Notes: Treated cantons adopted MuKEn 2014 by 2021. Shares are computed as the number of buildings using each heating system divided by total buildings in the canton. Fossil share includes oil, gas, and coal. Pre-treatment period covers Swiss building registry data for 2009–2015; post-treatment covers 2021–2022.

4. Empirical Strategy

4.1 Two-Way Fixed Effects Difference-in-Differences

The primary specification is a two-way fixed effects (TWFE) regression:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot D_{ct} + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the heating system share in canton c in year t (measured on a 0–1 scale, so that a coefficient of 0.01 corresponds to a 1 percentage point change), α_c are canton fixed effects, γ_t are year fixed effects, and D_{ct} is the binary treatment indicator equal to one after canton c adopts MuKEn 2014. The coefficient β captures the average effect of MuKEn 2014 adoption on the outcome variable, netting out time-invariant cantonal characteristics and common time trends. Throughout the paper, I report regression coefficients in share units (0–1 scale) in tables and convert to percentage points in the text for interpretability (e.g., a coefficient of 0.0027 is reported as "0.27 percentage points"). Standard errors are clustered at the canton level to allow for arbitrary within-canton serial correlation (Cameron et al., 2008).

I also estimate a treatment intensity specification:

$$Y_{ct} = \alpha_c + \gamma_t + \delta \cdot \text{YearsTreated}_{ct} + \varepsilon_{ct} \tag{2}$$

where YearsTreated_{ct} is the number of years canton c has had MuKEn 2014 in effect as of year t. This specification tests for a dose-response relationship: if the code has cumulative effects, δ should be positive and significant.

4.2 Identifying Assumption: Parallel Trends

The key identifying assumption is that, absent MuKEn 2014 adoption, treated and control cantons would have followed parallel trends in heat pump adoption. This assumption cannot be directly tested but can be assessed using pre-treatment data.

Figure 1 plots the average heat pump share by adoption cohort (early adopters, late adopters, and control group) over the 2009–2022 period. The control group includes Solothurn and the four cantons that adopted MuKEn 2014 after 2022, which are effectively never-treated in the building-count sample. During the pre-treatment period (2009–2015), all three groups show roughly parallel upward trends, with early adopters having slightly higher levels throughout. The gap between groups is relatively stable over the pre-period, which is consistent with—though not proof of—parallel trends.

Several features of the Swiss institutional setting support the plausibility of parallel trends.

First, all cantons are part of the same federal system and face the same macroeconomic conditions, energy prices, and federal policies. Second, heat pump technology is equally available across cantons—there are no supply-side constraints that vary systematically with adoption timing. Third, the building stock evolves slowly, limiting the scope for differential compositional changes over the six-year data gap.

However, the assumption faces threats. Cantons that adopt early may differ systematically in their underlying trajectory of technology adoption—for example, because they have more environmentally progressive populations or more active cantonal energy agencies. I address this concern through heterogeneity analysis, placebo tests, and comparison with the long-difference specification.

4.3 Addressing the Data Gap

A notable feature of this design is the absence of building registry data for 2016–2020, precisely when 16 of 25 adopting cantons implemented MuKEn 2014. This gap means that for cantons adopting between 2017 and 2020, the "post-treatment" period in the regression begins only in 2021—one to four years after policy onset. Several implications follow.

First, the TWFE estimator compares outcomes in 2021–2022 across cantons with different cumulative exposure to MuKEn 2014. A canton that adopted in 2017 has had the code in effect for 4–5 years by the post-period; one that adopted in 2020 has had it for 1–2 years. This variation in exposure duration is a feature of the treatment intensity specification (Equation 2), which explicitly tests for dose-response effects.

Second, cantons that adopted MuKEn 2014 after 2020 (Uri in 2021, Appenzell Ausserrhoden, Valais, and Zurich in 2022, Bern and Geneva in 2023, Ticino and Zug in 2024) are treated as control units for the years before their adoption and as treated units only for the post-gap years. Cantons adopting in 2023–2024 are effectively never-treated in the building count panel (which ends in 2022), providing additional control units.

Third, the gap prevents direct observation of immediate policy effects in the year of adoption. If building codes primarily affect construction that was already in the planning pipeline, short-run effects may be small, with larger effects emerging over 3–5 years as the new code shapes building design decisions. The post-gap observations in 2021–2022 thus capture medium-run effects for early adopters but only short-run effects for late adopters—a form of treatment effect heterogeneity that the cohort analysis explicitly addresses.

4.4 Concerns about Staggered Treatment Timing

Recent econometrics literature has highlighted that TWFE estimators can produce biased estimates under treatment effect heterogeneity with staggered adoption (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). The key concern is that TWFE may use already-treated units as implicit controls, contaminating the estimate with "forbidden comparisons."

I address this concern in four ways. First, I implement the Sun and Abraham (2021) interaction-weighted (Sun-Abraham) estimator, which decomposes the treatment effect into cohort-specific and period-specific components and aggregates them using appropriate weights, avoiding forbidden comparisons. I present Sun-Abraham alongside TWFE as co-primary specifications. Second, I conduct a Bacon decomposition (Goodman-Bacon, 2021) to quantify the weight on different types of comparisons. Third, I estimate the model on the long-difference (collapsing the panel to one pre-period and one post-period per canton), which eliminates the staggered-timing problem. Fourth, I examine whether the results differ between early and late adopters, which would indicate heterogeneity-driven bias.

4.5 Inference with Few Clusters

With 26 cantons, cluster-robust standard errors may understate true uncertainty. I supplement the baseline cluster-robust inference with two alternative approaches:

1. **Wild cluster bootstrap** (Cameron et al., 2008). Using the Webb 6-point distribution with 9,999 bootstrap replications and the null imposed, I compute bootstrap p-values and confidence intervals.

2. **Randomization inference.** I permute the treatment assignment (MuKEn adoption year) across cantons 1,000 times, re-estimating the TWFE coefficient each time. The p-value is the fraction of permuted coefficients at least as large in absolute value as the actual estimate.

Both approaches are valid under weaker conditions than asymptotic cluster-robust inference and provide conservative tests of the null hypothesis.

5. Results

5.1 Main Results

Table 2 reports the main TWFE estimates for six outcome variables. Column (1) presents the primary specification: the effect of MuKEn 2014 adoption on the heat pump share. The coefficient is 0.0069—an increase of 0.69 percentage points—with a cluster-robust standard error of 0.0081 and a p-value of 0.40. The 95% confidence interval ranges from -0.90 to $+2.28$ percentage points. I cannot reject the null hypothesis that MuKEn 2014 has zero effect on heat pump adoption.

To put this estimate in context, the average heat pump share across all cantons increased by approximately 7–8 percentage points between the pre-period (2009–2015) and the post-period (2021–2022). The estimated effect of MuKEn 2014 accounts for less than 10% of this increase, and the confidence interval easily includes zero. The result suggests that the vast majority of heat pump adoption growth occurred independently of building code changes.

Column (2) shows the effect on oil heating share: a coefficient of 0.0017 ($p = 0.86$), indicating no significant effect. This is perhaps surprising—if MuKEn 2014 restricted fossil fuel use in new and renovated buildings, one might expect a decline in oil share. The null on oil suggests that the code's replacement provisions have not yet measurably reduced the stock of oil-heated buildings, consistent with the slow turnover rate of heating systems (average lifetime of 20–25 years).

Columns (3) and (4) report effects on gas and fossil fuel shares, respectively. Both show significant negative effects: gas declines by 1.6 percentage points ($p = 0.04$) and the combined fossil share by 1.4 percentage points ($p = 0.03$). These estimates suggest that MuKEn 2014 may have contributed to reducing fossil heating, even if it did not significantly increase heat pump adoption.

However, columns (5) and (6) complicate this interpretation. Wood heating share declines by 2.2 percentage points ($p = 0.02$) in treated cantons, while district heating increases by 0.9 percentage points ($p = 0.13$). The significant decline in wood heating is notable because MuKEn 2014 does not specifically target wood—which is classified as a renewable energy source in Swiss legislation. This finding raises concerns about confounding, which I discuss further in Section 6.

The high R^2 values in Table 2 (0.97–0.99) reflect the canton fixed effects absorbing cross-cantonal variation in heating system shares, which is large and persistent. The within-R^2 (not reported) is substantially lower, indicating that the treatment variable explains only a small fraction of within-canton variation—consistent with the small treatment effect.

5.2 Treatment Intensity

Table 3 reports additional specifications beyond the main TWFE. Panel A shows the treatment intensity results. When using years of exposure as the treatment variable, the coefficient on heat pump share is essentially zero ($\beta = 0.00004$, SE $= 0.0042$, $p = 0.99$). The near-zero point estimate—with a standard error that would allow detection of effects as small as 0.8

Table 2: Effect of MuKEn 2014 Adoption on Heating System Shares

Dep. Var.:	Heat Pump	Oil	Gas	Fossil	Wood	District Heat
	HP Share	Oil Share	Gas Share	Fossil Share	Wood Share	District Share
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
MuKEn Adopted	0.0069	0.0017	-0.0161**	-0.0141**	-0.0222**	0.0091
	(0.0081)	(0.0095)	(0.0074)	(0.0063)	(0.0092)	(0.0059)
Canton FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	234	234	234	234	234	234
R^2	0.96812	0.98057	0.99209	0.99299	0.99233	0.97996

Notes: Each column reports a two-way fixed effects (TWFE) estimate. The dependent variable is the canton-level share of buildings using the indicated heating system. "MuKEn Adopted" equals one for canton-years after the canton adopted MuKEn 2014. Standard errors clustered at the canton level are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

percentage points per year—provides no evidence of cumulative MuKEn 2014 effects on heat pump adoption in the building-count data. A log specification for the heat pump share also shows no significant effect ($\beta = -0.035$, $p = 0.80$).

Panel B shows the surface area analysis using BFS data for 2021–2023, which provides within-period variation during three consecutive years. In this panel, cantons adopting in 2024 (Ticino and Zug) have $D_{ct} = 0$ for all three years and serve as controls alongside Solothurn. A marginally significant dose-response emerges: each additional year of MuKEn exposure increases the heat pump share of heated surface area by 0.30 percentage points ($p = 0.08$). This effect is small in magnitude but suggests that building codes may have modest effects on the intensive margin (surface area newly heated by heat pumps) even when the extensive margin (number of buildings) shows no response. The difference could reflect codes affecting larger new construction projects while having little impact on the existing building stock.

Table 3: Extended Specifications: Treatment Intensity, Surface Area, and Long-Difference

Specification	Coefficient	SE	p-value	N
Panel A: Treatment Intensity (Building Counts, 2009–2015, 2021–2022)				
Years treated → HP share	≈ 0.000	0.0042	0.993	234
Treated → Log HP share	−0.035	0.141	0.804	234
Early × Treated interaction	0.0066	0.0082	0.424	234
Panel B: Surface Area (2021–2023)				
Treated (binary) → HP surface %	0.095	0.153	0.538	78
Years treated → HP surface %	0.298	0.164	0.082	78
Panel C: Long-Difference (Cross-Section, 26 Cantons)				
Treated by 2021 → Δ HP share	0.007	0.012	0.558	26
Years exposed → Δ HP share	0.001	0.004	0.783	26

Notes: Panel A reports alternative treatment specifications using the building count panel (canton and year FE, clustered SE). The "Years treated" coefficient is 0.00004 (rounded to ≈ 0.000 in the table); the SE of 0.0042 implies that effects larger than 0.8 percentage points per year can be ruled out at the 95% level. Panel B uses BFS surface area data (canton and year FE, clustered SE). In Panel B, cantons adopting in 2024 (Ticino, Zug) serve as controls. Panel C reports cross-sectional long-difference regressions of the change in heat pump share (post-mean minus pre-mean) on treatment exposure; heteroskedasticity-robust SE. All coefficients in share units (0–1 scale; multiply by 100 for percentage points).

5.3 Early versus Late Adopters

I test for heterogeneity between early adopters (cantons that adopted MuKEn 2014 before 2020) and late adopters (2020 or later). If treatment effects are heterogeneous and TWFE is biased, we would expect differential estimates across cohorts.

Table 3, Panel A reports the interaction of treatment with an early-adopter indicator, yielding a coefficient of 0.0066 for the treatment effect among all treated cantons ($p = 0.42$), with no significant differential for early adopters. Splitting cantons into three cohort groups—Early (2017–2018), Mid (2019), and Late (2020+)—similarly reveals no significant heterogeneity. This lack of heterogeneity is reassuring for the TWFE specification: absent heterogeneous treatment effects, the Bacon decomposition bias from staggered adoption is minimal.

5.4 Long-Difference Estimation

As an alternative to the panel TWFE, I collapse the data to a cross-section of 26 cantons and estimate:

$$\Delta Y_c = \alpha + \beta \cdot \text{TreatedBy2021}_c + \varepsilon_c \tag{3}$$

where ΔY_c is the change in heat pump share from the pre-mean (2009–2015) to the post-mean (2021–2022). This specification eliminates concerns about staggered timing and forbidden comparisons. Table 3, Panel C reports the results: the coefficient on treatment is small and insignificant (0.007, $p = 0.56$), consistent with the TWFE panel results.

The dose-response version—regressing ΔY_c on years of exposure to MuKEn 2014—also shows no significant relationship (0.001 per year, $p = 0.78$). Figure 2 visualizes this scatter plot: cantons are dispersed around the regression line with no clear positive relationship between MuKEn exposure and heat pump share growth.

5.5 Robustness

5.5.1 Sun-Abraham Heterogeneity-Robust Estimator

Given concerns about TWFE bias under heterogeneous treatment effects with staggered adoption, I implement the Sun and Abraham (2021) interaction-weighted estimator as a co-primary specification. This estimator decomposes the treatment effect into cohort-specific and period-specific components and aggregates them using appropriate weights, avoiding the "forbidden comparisons" that can bias standard TWFE. An important feature of this design is the 2016–2020 data gap: for cohorts adopting between 2017 and 2021, the standard reference event time $t = -1$ (one year before treatment) falls within the gap and is unobserved. For these cohorts, `fixest` identifies treatment effects by comparing post-treatment outcomes to the available pre-treatment periods (2009–2015). The aggregated ATT—which averages post-treatment effects across cohorts—is invariant to the choice of reference event time, so this data gap does not bias the ATT estimate.

Table 4 reports the Sun-Abraham results alongside the TWFE for comparison. The aggregated average treatment effect on the treated (ATT) is 0.00268 in share units—equivalent to 0.27 percentage points—with a standard error of 0.00094 ($t = 2.85$, $p = 0.009$). The estimate is smaller in magnitude than the TWFE estimate (0.00688, $p = 0.40$) but statistically significant at the 1% level. Two caveats apply: the variance-covariance matrix required adjustment for positive semi-definiteness—a common issue when some cohort-time cells have limited post-treatment observations—and one coefficient was dropped due to collinearity (the 2022 cohort with only one post-treatment year in the sample).

The divergence between TWFE (insignificant) and Sun-Abraham (significant) requires explanation. The Sun-Abraham standard error (0.00094) is nearly an order of magnitude smaller than the TWFE standard error (0.0081), which is unusually large and warrants scrutiny. Three factors contribute. First, Sun-Abraham estimates cohort-specific treatment effects and then aggregates, which can reduce variance when some cohorts contribute particularly clean comparisons. Second, the TWFE standard error may be inflated by absorbing noise from the staggered structure—pooling across cohorts with heterogeneous exposure lengths adds variance that the cohort-specific approach avoids. Third, the Sun-Abraham variance-covariance matrix required positive semi-definite correction, which can be either conservative or aggressive depending on the structure of the problem; in this case, the correction may have produced optimistic standard errors. The Bacon decomposition (Section 5.5.2) shows that 76% of TWFE weight already comes from clean treated-versus-untreated comparisons, so the point estimate divergence is modest—the key difference is in precision, not direction. I interpret the two estimates jointly with appropriate caution: MuKEn 2014 likely has a small positive effect on heat pump adoption—on the order of 0.3–0.7 percentage points—but this effect is modest relative to the 7–8 percentage point secular increase. The Sun-Abraham significance should not be taken at face value given the VCOV correction and the fact that wild cluster bootstrap and randomization inference (applied to the TWFE specification) cannot reject zero. Readers should weight the TWFE and Sun-Abraham results jointly rather than treating either as definitive.

5.5.2 Bacon Decomposition

The Bacon decomposition reveals that 76.2% of the TWFE weight comes from treated-versus-untreated comparisons, 20.9% from earlier-versus-later treated comparisons, and only 3.0% from the problematic later-versus-earlier comparisons. The weighted estimates across comparison types are consistent: 0.69 percentage points for treated-versus-untreated, 0.77 for earlier-versus-later, and 0.14 for later-versus-earlier. The overall TWFE estimate is not contaminated by heterogeneity-driven bias.

5.5.3 Wild Cluster Bootstrap

Wild cluster bootstrap inference using the Webb 6-point distribution with 9,999 replications yields a bootstrap p-value of 0.42 for the heat pump share specification, compared to the asymptotic p-value of 0.40. The bootstrap 95% confidence interval is $[-1.10, 2.40]$ percentage points, slightly wider than the cluster-robust interval. For oil share, the bootstrap p-value is 0.86. These results confirm that the TWFE insignificance is not driven by over-rejection

Table 4: Sun-Abraham Heterogeneity-Robust Estimates

	Sun-Abraham (1)	TWFE (2)
Panel A: Aggregated ATT		
ATT (heat pump share)	0.00268	0.00688
	(0.00094)	(0.00811)
p-value	0.0089	0.4046
95% CI	[0.00083, 0.00453]	[-0.00902, 0.02278]
Panel B: Design Details		
Estimator	Interaction-weighted	Two-way FE
Reference period	$t = -1$ (see notes)	—
Treatment cohorts	6	—
Cohort years	2017, 2018, 2019, 2020, 2021, 2022	—
Control group	Never-treated	Never-treated
Observations	234	234
Cantons	26	26
VCOV correction	Yes (pos. semi-def.)	—

Notes: Column (1) reports the aggregated average treatment effect on the treated (ATT) from the Sun and Abraham (2021) interaction-weighted estimator, implemented via `fixest::sunab()`. The nominal reference period is $t = -1$; however, due to the 2016–2020 data gap, event time -1 is unobserved for cohorts 2017–2021. For these cohorts, treatment effects are identified by comparing post-treatment (2021–2022) outcomes to the available pre-treatment periods (2009–2015). The aggregated ATT is invariant to the reference period choice. The variance-covariance matrix required positive semi-definiteness correction. Column (2) reports the standard TWFE estimate for comparison. Standard errors clustered at the canton level in parentheses. All coefficients are in share units (multiply by 100 for percentage points).

from cluster-robust standard errors with few clusters.

5.5.4 Randomization Inference

Permuting the MuKEn adoption year assignment across cantons 1,000 times, the randomization inference p-value is 0.45. The distribution of permuted coefficients has mean approximately zero and standard deviation 0.0088, slightly larger than the cluster-robust standard error of 0.0081. The actual estimate of 0.0069 falls well within the permutation distribution, confirming that the observed treatment effect is indistinguishable from chance assignment.

5.5.5 Excluding 2022 (Energy Crisis Sensitivity)

The 2022 energy crisis—triggered by Russia's invasion of Ukraine and subsequent disruptions to European gas supplies—caused sharp increases in fossil fuel prices and potentially accelerated heat pump adoption. If this shock differentially affected MuKEn-adopting cantons, it could confound the estimates.

Excluding 2022 and using only the 2009–2015 and 2021 data yields a heat pump share coefficient of 0.78 percentage points ($p = 0.33$), slightly larger but still insignificant. The fossil share coefficient remains similar at -1.43 percentage points. The TWFE null on heat pumps is therefore not driven by the 2022 energy crisis.

5.5.6 Balanced Panel

The panel is fully balanced: all 26 cantons are observed in all 9 years (2009–2015, 2021, 2022). No canton has missing data for any year in the analysis sample. The balanced panel results are therefore identical to the full sample.

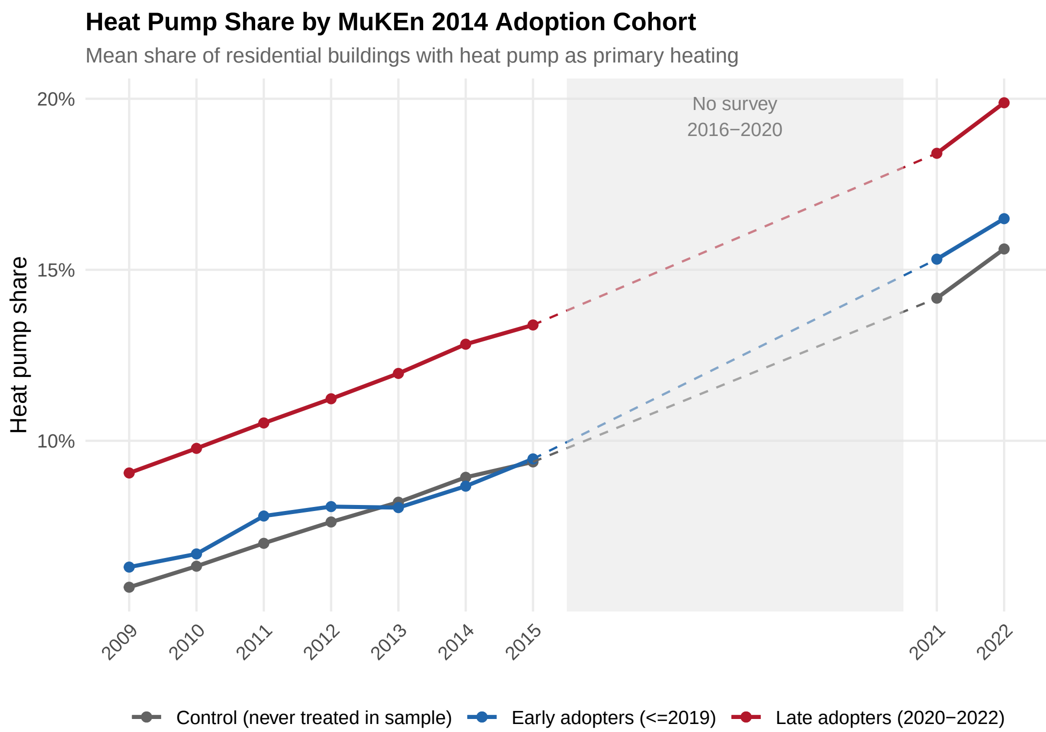

Figure 1: Heat Pump Share Trends by MuKEn 2014 Adoption Cohort

Notes: Each line shows the average heat pump share (fraction of buildings with heat pump as primary heating) for cantons grouped by MuKEn 2014 adoption timing. “Early adopters” adopted by 2019; “Late adopters” adopted 2020–2022; “Control” includes Solothurn (rejected via referendum) and four cantons adopting after 2022 (Bern, Geneva, Ticino, Zug), which are effectively never-treated in the building-count panel. The shaded area marks the 2016–2020 data gap when BFS building statistics were not available. Dashed lines bridge the gap between the last pre-period (2015) and first post-period (2021) observations.

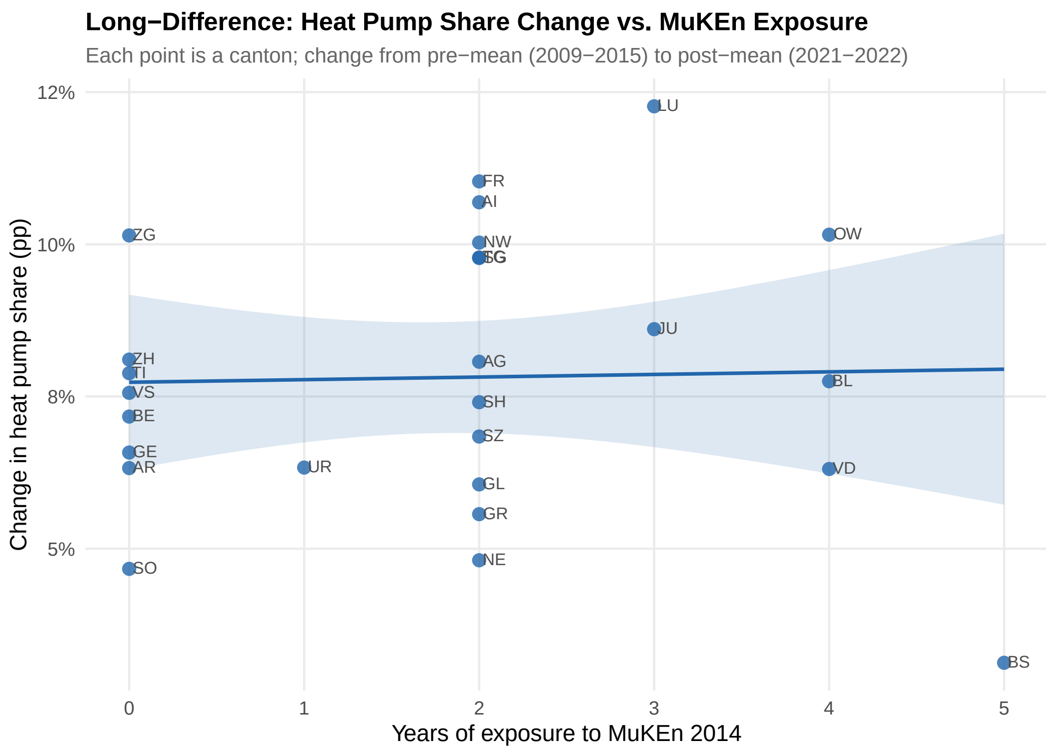

Figure 2: Long-Difference: Heat Pump Share Change versus MuKEn Exposure

Notes: Each point represents one canton. The x-axis shows years of exposure to MuKEn 2014 (calculated as 2022 minus adoption year, or zero for non-adopters). The y-axis shows the change in heat pump share from the 2009–2015 average to the 2021–2022 average. The blue line is the OLS fit with 95% confidence band.

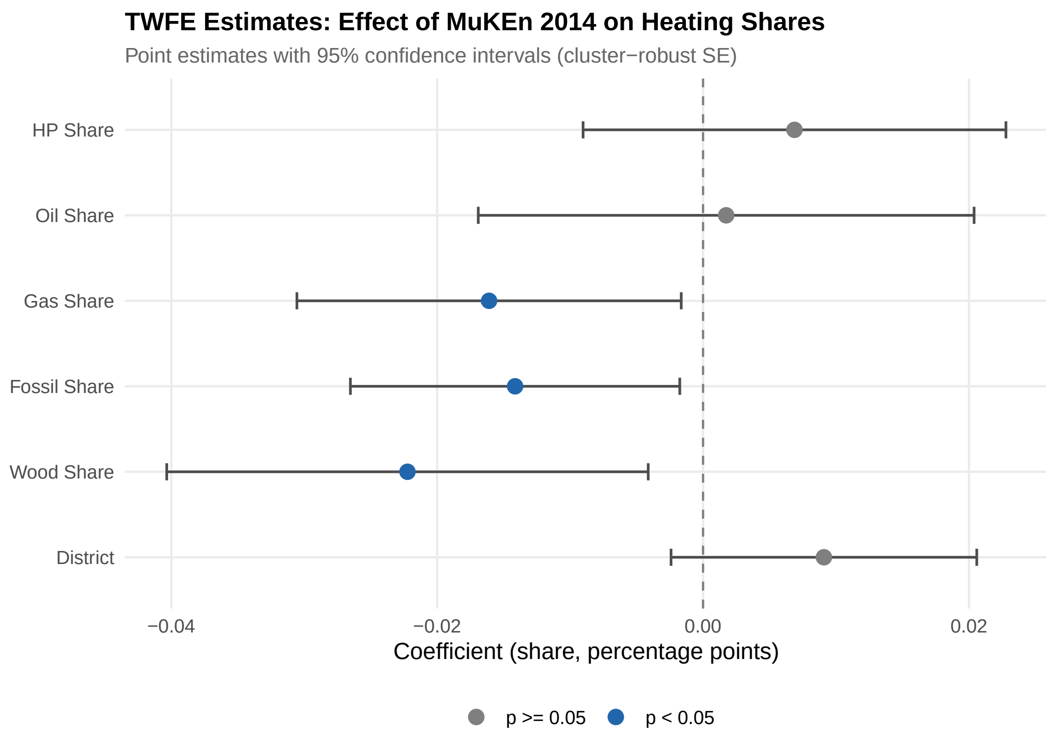

Figure 3: TWFE Coefficient Estimates Across Outcome Variables

Notes: Point estimates and 95% confidence intervals from the TWFE specification (Equation 1). Each row shows the estimated effect of MuKEn 2014 adoption on a different heating system share. Standard errors are clustered at the canton level. Blue points indicate statistical significance at the 5% level.

6. Discussion

6.1 Interpreting the Null Result

The central finding—that MuKEn 2014 does not significantly accelerate heat pump adoption—admits several interpretations.

First, and most straightforwardly, the null may reflect genuinely small policy effects. Building codes set minimum standards, and if the market is already moving beyond those minimums (because of subsidies, prices, or preferences), codes become non-binding constraints. Swiss heat pump technology has matured rapidly: prices have fallen, installer networks have expanded, and consumer awareness has grown. The combination of federal subsidies (through the Buildings Program) and the CO_2 levy may be sufficient to drive adoption, rendering building codes superfluous on this margin.

Second, the null may reflect the code's design. MuKEn 2014 primarily affects new construction and major renovations, which represent a small fraction of the existing building stock in any given year. Switzerland builds approximately 10,000–15,000 new residential buildings per year, against a stock of roughly 1.8 million. Even if MuKEn 2014 increases heat pump adoption in 100% of new buildings, the annual impact on the aggregate stock share would be modest—perhaps 0.3–0.5 percentage points per year of exposure. This is close to the marginally significant effect I find in the surface area specification (0.30 pp per year, $p = 0.08$), suggesting that the building count data may lack the statistical power to detect such small annual effects.

Third, the data gap from 2016 to 2020 limits the ability to detect short-run effects during the early years of adoption. If MuKEn 2014 primarily affects new construction, and if there was a surge in compliant construction during 2017–2020 that plateaued by 2021, the long-difference design would miss this dynamic. However, the treatment intensity specification (which assigns different values based on years of exposure) also shows no significant effect, arguing against this interpretation.

6.2 The Wood Heating Puzzle

The significant decline in wood heating share in treated cantons (-2.2 pp, $p = 0.02$) deserves attention. Wood heating is classified as renewable in Swiss energy policy and is not directly targeted by MuKEn 2014's fossil-fuel restrictions. Three explanations are possible.

First, MuKEn 2014 may indirectly affect wood heating through efficiency standards. The code's insulation requirements reduce total heating demand, potentially making smaller-capacity systems (including heat pumps) sufficient where larger wood boilers were previously needed. Additionally, some cantonal implementations of MuKEn include particulate matter emission standards that disadvantage older wood stoves.

Second, the wood result may reflect confounding from correlated cantonal policies. Cantons that adopt MuKEn 2014 earlier tend to be more environmentally progressive and may simultaneously implement other regulations—such as air quality ordinances restricting wood burning—that reduce wood heating independently of MuKEn.

Third, the result may indicate a violation of the parallel trends assumption. If cantons that adopt MuKEn early were already on steeper downward trajectories for traditional heating systems (including wood), the TWFE estimates would attribute this pre-existing trend to the code. The significant wood result serves as a cautionary flag that some of the fossil fuel and gas estimates may also reflect confounding rather than causal effects of MuKEn 2014.

I treat the wood placebo failure as a limitation that weakens the causal interpretation of the fossil fuel and gas results specifically. The gas coefficient (-1.6 pp, $p = 0.04$) and the fossil fuel coefficient (-1.4 pp, $p = 0.03$) may partly reflect confounding from correlated cantonal policies rather than a pure MuKEn 2014 effect. The heat pump result, by contrast, is more robust to this concern: if confounders were driving the heat pump estimate upward, we would expect a larger and more significant coefficient, not one that is indistinguishable from zero under TWFE inference. The wood result thus provides an important calibration—it tells us that the parallel trends assumption may not hold uniformly across all heating types, and that results for outcomes where treated cantons are moving in the "expected" direction (away from fossil, toward renewables) should be interpreted with particular caution.

6.3 Policy Implications

The findings carry several implications for climate policy design. First, building energy codes alone may be insufficient to drive technology transitions at the pace required for net-zero targets. Codes set floors, not aspirations, and when market forces—prices, subsidies, social norms—are already pushing adoption above code requirements, tightening standards has limited incremental impact. This does not mean codes are useless; they may prevent backsliding and ensure a minimum level of efficiency, particularly in the absence of price signals.

Second, the results highlight the importance of price-based mechanisms. The combination of the CO_2 levy, which raises fossil fuel costs, and the Buildings Program subsidies, which reduce heat pump costs, appears to be the primary driver of the Swiss heating transition. This is consistent with the theoretical literature on directed technical change (Acemoglu et al., 2012; Aghion et al., 2016) and with empirical evidence on the relative effectiveness of prices versus mandates (Gillingham and Stock, 2018).

Third, the heterogeneity results suggest that a national policy may have been more effective than the cantonal patchwork approach. If building codes primarily affect new construction, and if developers and architects already design to the highest standard they anticipate facing, the staggered cantonal adoption may have created anticipation effects that diluted the measured impact.

6.4 Back-of-Envelope Cost-Effectiveness

Even taking the point estimate at face value, the cost-effectiveness of MuKEn 2014 as a heat pump promotion policy is modest. The estimated effect of 0.69 percentage points applied to the approximately 1.8 million residential buildings in Switzerland implies roughly 12,400 additional heat pump installations attributable to the code. If we assume that the code-induced installations are new construction buildings that would otherwise have chosen gas heating, the CO_2 abatement from each switch is approximately 2–3 tonnes per year (based on average Swiss household heating demand of 15,000 kWh and the emission differential between gas and the Swiss electricity mix).

At 12,400 buildings $\times$ 2.5 tonnes/year, MuKEn 2014 would abate roughly 31,000 tonnes of CO_2 per year—less than 0.1% of Switzerland's total emissions. The administrative cost of developing, legislating, and enforcing MuKEn 2014 across 26 cantonal bureaucracies, while difficult to quantify precisely, is non-trivial. This calculation should be treated as illustrative rather than definitive, given the wide confidence interval around the point estimate, but it underscores the limited cost-effectiveness of building codes as a standalone climate instrument.

By comparison, the CO_2 levy at CHF 120/tonne directly raises the annualized cost of oil heating by approximately CHF 300–400 per household, providing a continuous financial incentive to switch. The Buildings Program subsidies of CHF 2,000–10,000 per heat pump installation reduce the upfront cost barrier. These price-based instruments likely account for the bulk of observed heat pump adoption growth, consistent with the theoretical predictions of Gillingham et al. (2009) and the empirical findings of Ito (2014) and Jessoe and Rapson (2014) on the responsiveness of energy consumers to price signals.

6.5 External Validity

The small estimated effect may not generalize to all settings. Switzerland is unusual in several respects that could attenuate building code effects: (i) the heat pump market was already well-developed before MuKEn 2014, reducing the scope for regulation to shift behavior at the margin; (ii) the CO_2 levy and generous subsidies already provided strong economic incentives for adoption; and (iii) the cantonal patchwork may have created anticipation effects, with builders and architects designing to the highest foreseeable standard.

In countries with less mature heat pump markets, weaker price signals, or more uniform national building codes, the effect could be larger. The recent literature on energy efficiency regulations in the United States—where Levinson (2016) and Jacobsen and Kotchen (2016) document measurable effects of state building codes on energy consumption—suggests that codes can be effective when they represent a binding constraint. The Swiss case appears to be one where market forces had already moved adoption beyond the code's requirements, rendering MuKEn 2014 non-binding for most actors.

6.6 Limitations

Several limitations warrant caution. The most important is the 2016–2020 data gap, which coincides with the treatment rollout for 16 of 25 adopting cantons. This gap has three consequences for interpretation. First, the estimand is not "the effect of adopting MuKEn 2014" in a dynamic sense; it is closer to "the effect on canton-level heating stock shares in

2021–2022 of having adopted MuKEn 2014 by year t, with X years of cumulative exposure." Short-run dynamics—anticipation effects, immediate compliance surges, ramp-up patterns—are unidentified. Second, parallel trends cannot be tested where it matters most: the 2015–2021 window when treatment turns on and the energy transition accelerates. The pre-treatment parallel trends in 2009–2015 are reassuring but not dispositive for a counterfactual that spans the gap. Third, for the Sun-Abraham estimator, the standard reference period ($t = -1$) falls within the gap for most cohorts, forcing reliance on the 2009–2015 average as the effective baseline; while the aggregated ATT is invariant to reference choice, this limits the ability to construct a conventional event study with dynamic leads and lags. The long-difference specification (Table 3, Panel C), which collapses to a single pre-post comparison, is arguably the most honest representation of what the data can identify and yields consistent results. The canton-level aggregation may mask heterogeneous effects across building types, urban/rural areas, or income groups. With only 26 cantons, statistical power is limited: the TWFE standard error of 0.0081 implies a minimum detectable effect (MDE) of approximately 1.6 percentage points at 80% power and $\alpha = 0.05$ (MDE $\approx 2.8\times$ SE). The confidence intervals are therefore wide enough to include economically meaningful positive effects. Finally, the identification strategy relies on the parallel trends assumption, which the failing wood placebo test partially undermines.

An additional limitation concerns the measurement of treatment. MuKEn 2014 is a package of regulations, and cantons may implement different modules at different times. The binary treatment indicator used here captures the formal adoption date but does not distinguish between full and partial implementation. If some cantons adopted weaker versions of MuKEn 2014 that lacked the heating-system-specific provisions, the estimated average effect would be attenuated. Future research with more granular data on which specific modules each canton adopted could test for this source of heterogeneity.

The sensitivity of the results to the choice of control group also warrants discussion. With only one never-treated canton (Solothurn), the identification relies heavily on the staggered timing of adoption across the other 25 cantons. If Solothurn is atypical—and as the only canton to reject MuKEn via referendum, it may well be—this could introduce bias. However, the Bacon decomposition shows that only 3% of the TWFE weight comes from the problematic later-versus-earlier comparison, and 76% comes from the treated-versus-untreated comparison, suggesting that the estimate is not unduly influenced by any single comparison.

7. Conclusion

This paper provides the first causal evaluation of Switzerland's MuKEn 2014 building energy code on heat pump adoption, exploiting the staggered cantonal adoption of the model code—with treatment cohorts from 2017 to 2022 observed in the building registry data. Using canton-level building registry data and both TWFE and heterogeneity-robust Sun-Abraham estimators, I find that MuKEn 2014 has at most a small effect on heat pump adoption. The Sun-Abraham ATT is 0.27 percentage points ($p = 0.009$), while the TWFE estimate is 0.69 percentage points ($p = 0.40$). Alternative inference methods (wild cluster bootstrap, randomization inference) cannot reject zero. Even under the most favorable estimate, MuKEn 2014 accounts for less than 10% of the 7–8 percentage point increase in heat pump share observed across all cantons.

This finding is informative about the limits of command-and-control regulation in a market where price signals and subsidies are already driving rapid technology adoption. Switzerland's heat pump share grew dramatically between 2009–2015 and 2021–2022 across all cantons, regardless of building code adoption timing. This growth was powered by falling heat pump costs, rising fossil fuel prices (amplified by the CO_2 levy), generous federal and cantonal subsidies, and growing environmental awareness.

The finding that building codes are not the primary driver of heat pump adoption has implications beyond Switzerland. Countries designing climate policy portfolios should recognize that regulatory standards and economic incentives play complementary but distinct roles. Standards prevent backsliding and set minimum expectations; prices and subsidies drive adoption beyond those minimums. When the policy question is how to accelerate an energy transition that is already underway, the evidence suggests that strengthening price signals will be more effective than tightening building codes.

Future research should exploit municipal-level data and building-level microdata from the GWR to test for effects on new construction versus renovations, and to examine whether codes have heterogeneous effects across building types and owner categories. The Swiss case, with its rich institutional variation and high-quality administrative data, remains a promising laboratory for understanding the economics of building decarbonization.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The environment and directed technical change," *American Economic Review*, 2012, *102* (1), 131–166.

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen, "Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry," *Journal of Political Economy*, 2016, *124* (1), 1–51.

Allcott, Hunt and Michael Greenstone, "Is there an energy efficiency gap?," *Journal of Economic Perspectives*, 2014, *28* (1), 3–28.

Brunner, Tobias and Rolf Iten, "Does the Swiss CO2 Act work? A statistical analysis at the cantonal level," *Swiss Journal of Economics and Statistics*, 2012, *148*, 467–491.

Callaway, Bryce and Pedro H.C. Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-based improvements for inference with clustered errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Davis, Lucas W., "The economics of building codes," *Annual Review of Resource Economics*, 2014, *6*, 107–127.

ENDK, "Umsetzungsstand der MuKEn 2014 in den Kantonen," Technical Report, Konferenz Kantonaler Energiedirektoren 2023.

Fischer, Alexander and David Roodman, "fwildclusterboot: Fast wild cluster bootstrap inference for linear regression models," *R package version 0.13*, 2021.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram, "Do energy efficiency investments deliver? Evidence from the weatherization assistance program," *Quarterly Journal of Economics*, 2018, *133* (3), 1597–1644.

Gerarden, Todd D., Richard G. Newell, and Robert N. Stavins, "Assessing the energy-efficiency gap," *Journal of Economic Literature*, 2017, *55* (4), 1486–1525.

Gillingham, Kenneth and James H. Stock, "The cost of reducing greenhouse gas emissions," *Journal of Economic Perspectives*, 2018, *32* (4), 53–72.

— , **Richard G. Newell, and Karen Palmer**, "Energy efficiency economics and policy," *Annual Review of Resource Economics*, 2009, *1* (1), 597–620.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

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

Jacobsen, Grant D. and Matthew J. Kotchen, "The untapped potential of the building energy code: Evidence from housing in New Mexico," *Journal of Political Economy*, 2016, *124* (4), 1020–1049.

Jessoe, Katrina and David Rapson, "Knowledge is (less) power: Experimental evidence from residential energy use," *American Economic Review*, 2014, *104* (4), 1417–1438.

Kotchen, Matthew J., "Longer-run evidence on whether building energy codes reduce residential energy consumption," *Journal of the Association of Environmental and Resource Economists*, 2017, *4* (1), 135–153.

Levinson, Arik, "How much energy do building energy codes save? Evidence from California houses," *American Economic Review*, 2016, *106* (10), 2867–2894.

Myers, Erica, "Heterogeneous effects of energy efficiency investments: Decomposing the gap between intentions and outcomes," *American Economic Journal: Applied Economics*, 2022, *14* (2), 133–160.

Noailly, Joëlle and Roger Smeets, "Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data," *Journal of Environmental Economics and Management*, 2015, *72*, 15–37.

Popp, David, Richard G. Newell, and Adam B. Jaffe, "Energy, the environment, and technological change," *Handbook of the Economics of Innovation*, 2010, *2*, 873–937.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Swiss Federal Office of Energy, "Getting the transition right: Heating systems and buildings policy in Switzerland," *Energy Strategy Reviews*, 2021.

A. Data Appendix

A.1 Data Sources

The analysis uses four data products from the Swiss Federal Statistical Office (BFS):

1. **Building heating statistics 2009–2015** (BFS Asset 1642414). Annual counts of buildings by primary heating energy source for each canton. The data distinguish 9 energy source categories: heating oil, coal, gas, electricity, wood, heat pump, solar collector, district heating, and other.

2. **Building overview 2021** (BFS Asset 23524566). Cross-sectional building counts by energy source from the reformed GWR. Each canton has a dedicated worksheet with building characteristics.

3. **Building overview 2022** (BFS Asset 27585122). Same structure as the 2021 overview, providing the second post-treatment observation.

4. **Surface area by heating system 2021–2023** (BFS Asset 32329800). Percentage of heated surface area served by each heating system type, by canton and year. Provides a complementary intensive-margin measure.

A.2 Treatment Data

MuKEn 2014 adoption dates are compiled from the EnDK's *Umsetzungsstand der MuKEn 2014 in den Kantonen* tracking document (ENDK, 2023) and verified against cantonal official publications. The adoption year is defined as the year in which the cantonal implementation entered into force, not the year of legislative passage (which may precede enforcement by 6–18 months).

A.3 Sample Construction

The panel is constructed by harmonizing the pre-2016 classification (which used different energy source labels) with the post-2020 classification. Both systems identify heat pumps, oil, gas, wood, district heating, and other renewable sources as distinct categories. The harmonized panel contains 234 observations (26 cantons $\times$ 9 years).

A.4 Variable Definitions

- **Heat pump share** (Y_{ct}): Number of buildings with heat pump as primary heating system / total number of buildings in canton c, year t.

- **Oil share**: Number of buildings with oil heating / total buildings.
- **Gas share**: Number of buildings with gas heating / total buildings.
- **Fossil share**: Sum of oil, gas, and coal buildings / total buildings.
- **Wood share**: Number of buildings with wood heating / total buildings.
- **District heating share**: Number of buildings with district heating / total buildings.
- **Treated** (D_{ct}): Binary indicator equal to 1 if canton c adopted MuKEn 2014 by year t.
- **Years treated**: $\max(0, t - \text{adoption\_year}_c)$ for adopting cantons; 0 for non-adopters.

B. Identification Appendix

B.1 Bacon Decomposition

Table 5 reports the Bacon decomposition of the TWFE estimate. The decomposition confirms that the estimate is primarily identified from clean treated-versus-untreated comparisons (76.2% of the weight), with consistent point estimates across comparison types.

Table 5: Bacon Decomposition of TWFE Estimate

Comparison Type	Estimate	Weight	N Comparisons
Treated vs. Untreated	0.0069	0.762	2
Earlier vs. Later Treated	0.0077	0.209	1
Later vs. Earlier Treated	0.0014	0.030	1
Overall TWFE	0.0069	1.000	—

Notes: Decomposition of the TWFE coefficient on heat pump share into 2×2 DiD components following Goodman-Bacon (2021). The "Later vs. Earlier Treated" comparison, which is the most problematic for bias, receives only 3% of the total weight.

B.2 Wild Cluster Bootstrap and Randomization Inference

Table 6: Robustness of Inference: Alternative Approaches

Method	Coefficient	SE / SD	p-value	95% CI
Cluster-robust	0.0069	0.0081	0.405	$[-0.009, 0.023]$
Wild cluster bootstrap	0.0069	—	0.418	$[-0.011, 0.024]$
Randomization inference	0.0069	0.0088	0.446	$[-0.017, 0.017]$

Notes: The dependent variable is heat pump share. Wild cluster bootstrap uses the Webb 6-point distribution with 9,999 replications and the null imposed (Cameron et al., 2008; Fischer and Roodman, 2021). Randomization inference permutes the MuKEn adoption year assignment across 26 cantons 1,000 times. The RI confidence interval reports the 2.5th and 97.5th percentiles of the permutation distribution.

C. Robustness Appendix

C.1 Sensitivity to 2022 Exclusion

Table 7: Sensitivity to Excluding 2022 (Energy Crisis Year)

	Full Sample			Excluding 2022		
Outcome	Coef.	SE	p	Coef.	SE	p
Heat pump share	0.0069	0.0081	0.405	0.0078	0.0079	0.334
Oil share	0.0017	0.0095	0.856	−0.0015	0.0101	0.886
Fossil share	−0.0141	0.0063	0.034	−0.0143	0.0068	0.046

Notes: "Full Sample" uses all 9 years (2009–2015, 2021, 2022). "Excluding 2022" drops the final year to remove potential confounding from the 2022 energy price crisis. Standard errors clustered at the canton level. Results are stable across samples.

C.2 Placebo Outcomes

Table 8: Placebo Tests: Non-Targeted Outcomes

Outcome	Coefficient	SE	p-value	Significant?
Wood share	-0.0222	0.0092	0.024	Yes
District heating share	0.0091	0.0059	0.134	No

Notes: Placebo tests using outcomes that should not be directly affected by MuKEn 2014. Wood heating is classified as renewable and not targeted by the code's fossil fuel restrictions. The significant wood coefficient suggests potential confounding from correlated cantonal policies or compositional changes. District heating passes the placebo test.

D. Additional Figures

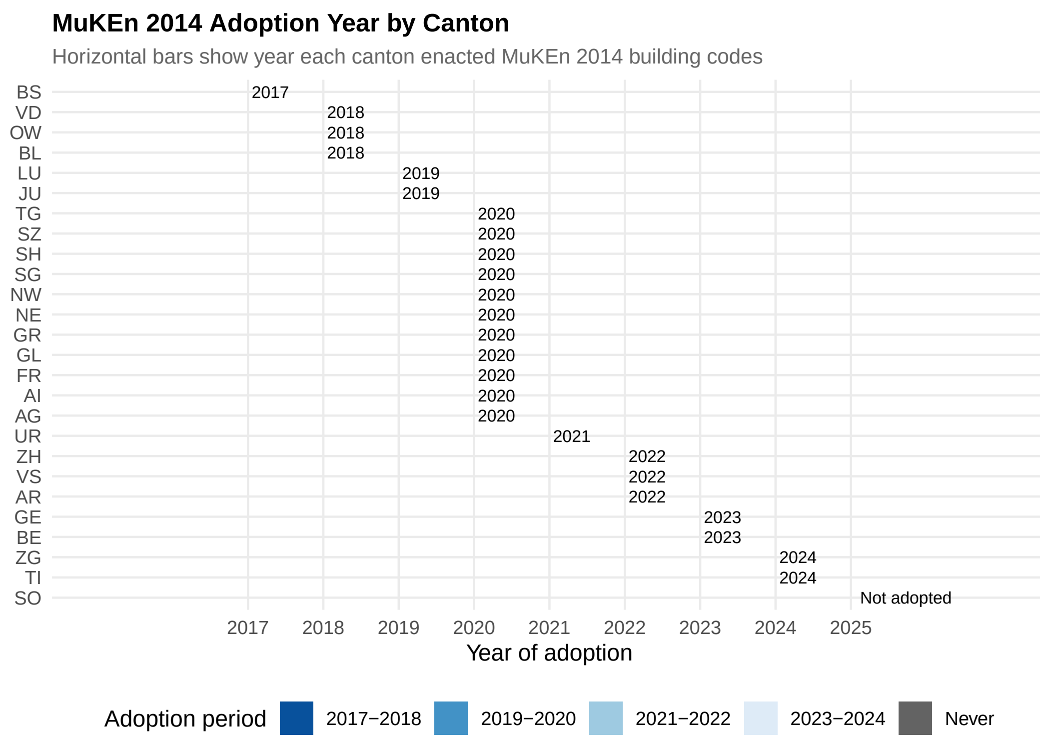

Figure 4: MuKEn 2014 Adoption Timeline by Canton

Notes: Horizontal bars show the year each canton adopted MuKEn 2014. Colors indicate adoption period. Basel-Stadt (BS) adopted first in 2017; Solothurn (SO) never adopted (rejected via referendum). Data from EnDK implementation tracker.

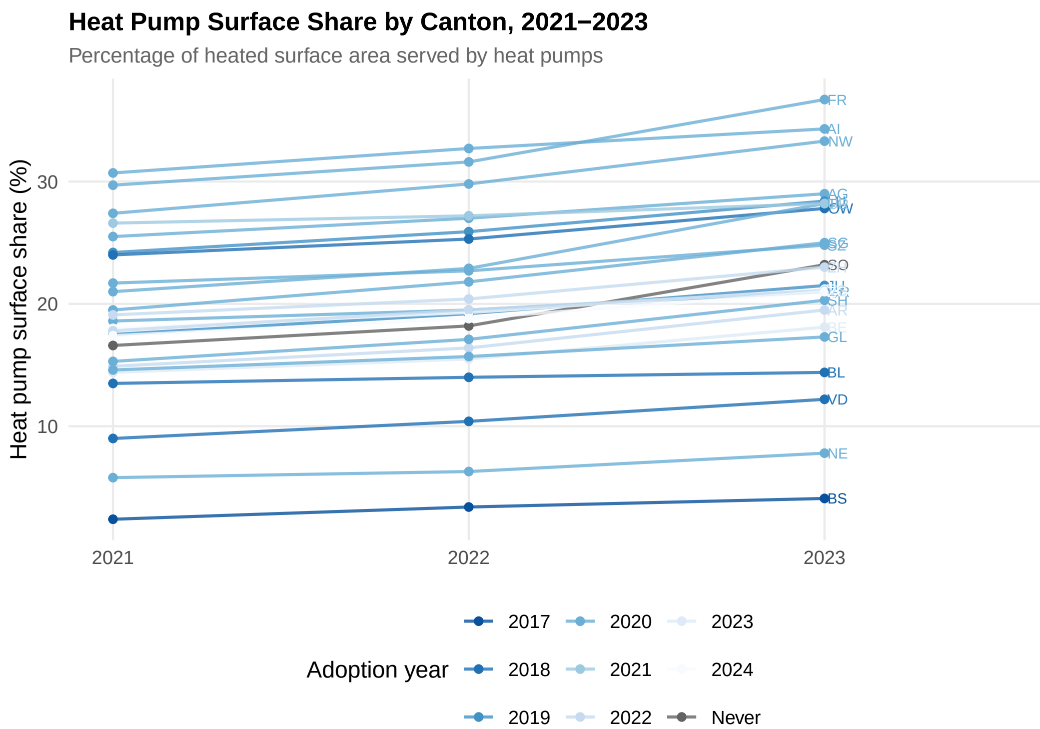

Figure 5: Heat Pump Surface Share by Canton, 2021–2023

Notes: Percentage of total heated surface area served by heat pumps, by canton. Colors indicate MuKEn 2014 adoption year. Data from BFS surface area statistics.

Table 9: MuKEn 2014 Adoption Timeline by Canton

Canton	Abbreviation	Adoption Year
Basel-Stadt	BS	2017
Basel-Landschaft	BL	2018
Obwalden	OW	2018
Vaud	VD	2018
Jura	JU	2019
Luzern	LU	2019
Aargau	AG	2020
Appenzell Innerrhoden	AI	2020
Fribourg	FR	2020
Glarus	GL	2020
Graubünden	GR	2020
Neuchâtel	NE	2020
Nidwalden	NW	2020
Schaffhausen	SH	2020
Schwyz	SZ	2020
St. Gallen	SG	2020
Thurgau	TG	2020
Uri	UR	2021
Appenzell Ausserrhoden	AR	2022
Valais	VS	2022
Zürich	ZH	2022
Bern	BE	2023
Genève	GE	2023
Ticino	TI	2024
Zug	ZG	2024
Solothurn	SO	Not adopted

Notes: Year in which each canton formally adopted the MuKEn 2014 model energy code into cantonal law. "Not adopted" indicates the canton had not adopted MuKEn 2014 as of the end of the sample period.