Do Consumers Distinguish Fixed Cost from Variable Cost? "Schmeduling" in Two-Part Tariffs in Energy*

Koichiro Ito
University of Chicago and NBER

Shuang Zhang
Imperial College London and NBER

This version: August 28, 2023

Abstract

A central assumption in economics is that consumers properly distinguish fixed cost from variable cost. This assumption is fundamental to various economic theories, including optimal taxation, redistribution, and price discrimination. Using a quasi-experiment in heating price reform in China, we find empirical evidence that is inconsistent with this conventional assumption and more consistent with the "schmeduling" model in Liebman and Zeckhauser (2004). As we demonstrate the policy implications for two-part energy tariffs, this consumer behavior makes fixed cost directly relevant to the perceived relative prices of goods, and therefore alters the welfare implications of price, tax, and subsidy designs.

*Ito: Harris School of Public Policy, University of Chicago, and NBER (e-mail: ito@uchicago.edu). Zhang: Department of Economics and Public Policy, Imperial College London, and NBER (e-mail: szhang7@ic.ac.uk). For helpful comments, we thank Hunt Allcott, Kjetil Storesletten, Douglas Almond, Severin Borenstein, Fiona Burlig, Steve Cicala, Lucas Davis, Tatyana Deryugina, Larry Goulder, Michael Greenstone, Kelsey Jack, Ryan Kellogg, Catherine Wolfram, and seminar participants at Stanford, Colorado Boulder, UC Berkeley, University of Chicago, Resources for the Future, the Federal Reserve Board of Governors, NBER Chinese Economy Meeting, and NBER Environment and Energy Program Meeting. We thank Theodor Kulczycki, Jing Qian, Chenyu Qiu, Andrew Smith, Max Snyder, and Yixin Zhou for excellent research assistance.

1 Introduction

A central assumption in economics is that individuals properly distinguish fixed cost from variable cost. In public finance, a lump-sum tax or subsidy is considered to be non-distortionary because it does not distort the relative prices of goods as long as taxpayers distinguish variable cost from fixed cost (Stiglitz, 1988). In industrial organization, a two-part tariff—a price schedule with a fixed charge and a variable charge—allows profit-maximizing firms to price-discriminate and natural monopolists to achieve allocative efficiency under the assumption that consumers distinguish variable cost from fixed cost (Tirole, 1988).

In this paper, we provide the first empirical evidence that this assumption may not be consistent with consumer behavior in reality. Despite the fact that this assumption is fundamental to many theoretical models and empirical works in economics, there is little to no direct empirical evidence on this question. The closest literature is the studies on tiered marginal price schedules in which individuals face multiple marginal prices or taxes for the same good. In this context, many studies find evidence that consumers and taxpayers tend to respond to average price rather than marginal price (de Bartolome, 1995; Borenstein, 2009; Kahn and Wolak, 2013; Ito, 2014; Rees-Jones and Taubinsky, 2020). However, in studying a two-part tariff, Borenstein and Davis (2012) clarify that evidence from this literature cannot demonstrate whether consumers distinguish variable cost from fixed cost because differentiating between a fixed cost and a single variable cost can be much less complex than identifying a correct marginal price from tiered marginal pricing that involves multiple variable prices in a price schedule.

To empirically test this assumption, we use a quasi-experiment in a recent heating price reform in China. Until recently, the majority of Chinese households paid only fixed charges for their winter heating consumption. Starting in 2005, the Chinese Ministry of Housing and Urban-Rural Development (MOHURD), in collaboration with the World Bank, introduced new pricing called consumption-based billing (CBB)—a two-part tariff with a fixed charge and a single variable price.

We exploit three unique features of this reform to examine our research question. First, the policy change induced an increase in variable cost but a decrease in fixed cost. With this price variation, many consumers experienced an *increase* in marginal price but a *decrease* in average price. Standard theory predicts that these consumers would reduce heating usage because the marginal price of heating increased.[1] However, an alternative theory, originating to the theory of "schmeduling" by Liebman and Zeckhauser (2004), predicts that they may perceive their average price as the true marginal price. In this case, these consumers may *increase* heating usage even though their marginal price increased. We exploit this price variation to develop a simple non-parametric test on our research question.

Second, in collaboration with a regulated heating provider, the World Bank, and the MOHURD, we were able to obtain newly-available administrative data on household-level daily heating usage from 2007 to 2019 in the city of Tianjin. Our data address a key empirical challenge that has been common in the literature. Usually, individually-metered usage data are available only after the introduction of metered pricing because firms tend to install meters at the same time as they introduce metered pricing. This makes empirical analysis challenging because individual-level usage data are unobserved before the policy change. Our data overcome this challenge because regulators in our context required household-level metered data to be collected at least one year before the introduction of metered pricing. This allows us to access daily household-level usage data *before and after* the reform started.

Third, the reform was introduced by staggered rollouts. Using this quasi-experimental variation in treatment timings, we can estimate the causal effects of the reform using the staggered difference-in-differences (DID) method. Although the validity of the identification assumptions is untestable, we show that the rollout timings are uncorrelated with observables and that our event-study figures support parallel trends between treated and untreated households in the pre-treatment periods. Following the recent econometric literature on the estimation of the staggered DID, we implement an estimation method that allows heterogeneous treatment effects across households using the estimation method developed by de Chaisemartin and D'Haultfœuille (2020).

We begin by estimating the overall impact of the reform on heating usage. The intention-to-treat (ITT) estimates indicate that the reform decreased heating usage on average by 10.1% in the first year, 10.7% in the second year, and 8.7% in the third year. These impacts are economically

[1] The income effect of the CBB policy, if any, was likely to be very small. The CBB policy reduced the annual fixed charge by about \$226 per household. The average household income in Tianjin in our sample period was \$15,041. Therefore, \$200 was about 1.5 percent of household income. In the literature on residential energy demand, short-run income elasticity estimates are found to be inelastic and estimated around 0.239 on average based on a recent meta analysis (Zhu and Yang, 2018). This implies that the income effect of the CBB on heating usage is 0.36 percent. We have a detailed discussion in Section 6.2.

substantial and long-lasting compared to a variety of policies on residential energy usage studied in the literature (Wolak, 2011; Ito, 2014, 2015; Ito, Ida and Tanaka, 2018; Deryugina, MacKay and Reif, 2020; Shaffer, 2020).

We show that this overall reduction in usage can be interpreted as an improvement in social welfare in the standard theoretical framework of two-part tariffs. The marginal price was zero before the reform, and it was set equal to the private marginal cost of heating production after the reform. Therefore, if we assume that consumers distinguish fixed cost from variable cost, the reduction in usage produces an improvement in allocative efficiency. If we also consider an environmental externality from coal-based heating production, the social marginal cost is above the private marginal cost. In this case, the overall welfare gain of the policy is even larger.

However, this standard framework cannot be applicable if consumers do not respond to marginal price by properly distinguishing fixed cost from variable cost. We show that the welfare impact of the reform is ambiguous if consumers respond to the average price of the bill rather than their marginal price as predicted by the theory of "schmeduling" (Liebman and Zeckhauser, 2004). Both of the gains from allocative efficiency and environmental externalities are likely to be smaller than those calculated in the standard framework, and the overall social welfare gain from the reform can be ambiguous in theory.

To empirically test this question, we investigate whether consumers respond to marginal price by properly distinguishing fixed cost from variable cost. As described above, many consumers in our data experienced an *increase* in marginal price but a *decrease* in average price. For these consumers, we use the staggered DID method to identify the causal impact of the reform on heating usage. Our results show that the reform caused a statistically significant *increase* in heating usage even though these consumers had an increase in marginal price. In addition, the reform made some consumers in our data experience an *increase* in marginal price but nearly zero change in average price. We find that the policy-induced change in heating usage for these consumers was near zero.

The set of our empirical findings suggest that consumer behavior in our data is more consistent with the schmeduling model than the standard economic theory. However, there can be at least three alternative mechanisms that might be able to explain this consumer behavior—category budgeting, income effects, and spurious correlations. In Section 6, we provide additional empirical evidence from our data that are inconsistent with these alternative explanations. First, we exploit the fact that consumers could opt out from the new pricing and show that the opt-out decision is inconsistent with the prediction from the theory of category budgeting (Hastings and Shapiro, 2013). Second, we show that the income effect in our setting is very small and unlikely to explain our empirical findings. Third, we explore whether our findings are driven by spurious correlations between household characteristics. Our analysis suggests that spurious correlations are unlikely to explain our key empirical findings.

Finally, we calculate the welfare impact of the reform based on our empirical findings. We obtain the estimated environmental externality by using ambient air pollution data and the estimated willingness to pay for clean air in Ito and Zhang (2020). We first calculate welfare impacts based on the standard theoretical framework of two-part tariffs, which assumes that consumers distinguish fixed cost from variable cost. In this case, the total social welfare gain is 25.5 USD per year per household and 109.3 million USD per year for Tianjin. The one-time administrative cost of the reform was 99 USD per household. This implies that if we conduct the cost-benefit analysis assuming standard economic theory, the policy's net present value of the benefit would exceeded its cost in 4 years with the discount rate of 3% in China.

We show that this standard framework is likely to substantially overstate the benefit of the reform once we incorporate the schmeduling behavior into the framework. In our second welfare calculation, we incorporate the empirical finding that consumers may not properly distinguish fixed cost from variable cost. In this case, the total social welfare gain is 6.3 USD per year per household and 26.8 million USD per year for Tianjin. This implies that the policy requires 21 years (as opposed to 4 years) to make its benefit exceed its cost. That is, the CBB is still likely to be a cost-effective policy in the long run, but it takes much more time to become cost-effective than one would consider with the standard approach. These results imply that the schmeduling behavior could alter the welfare implications of price, tax, and subsidy designs substantially to possibly change the key policy implications.

Related literature and our contributions—This paper provides three primary contributions to the economics literature and the design of economic policy. First, our findings provide a new insight to the literature on consumer inattention under complex pricing (Busse, Silva-Risso and Zettelmeyer, 2006; Gabaix and Laibson, 2006; Chetty, Looney and Kroft, 2009; Finkelstein, 2009; Brown, Hossain and Morgan, 2010; Hastings and Shapiro, 2013; Ito, 2014; Rees-Jones and Taubinsky, 2020). As noted above, our study provides the first empirical evidence that consumers may not distinguish fixed costs from variable costs as the standard economic theory predicts. We show that this behavior critically changes the welfare implication of two-part tariffs. Likewise, this behavior could alter key conclusions of many fundamental economic models, including optimal taxation, redistribution, natural monopoly, and price discrimination (Stiglitz, 1988; Tirole, 1988). This is because if consumers do not properly distinguish fixed costs from variable costs, a fixed payment or subsidy could affect the relative prices of goods, and therefore, have a direct impact on the efficiency of price, tax, or subsidy design. Our welfare analysis shows that this implication is empirically substantial and policy-relevant as taking into account for the schmeduling behavior significantly changes the cost-benefit of the CBB policy in China.

Second, our results have important policy implications for energy and climate policy across the globe because many energy policies in practice involve combinations of fixed and variable incentives. As we show in Section 7, the introduction of metered energy pricing could have different welfare implications if consumers do not properly distinguish fixed costs from variable costs. Another important relevant policy is the compensation scheme of carbon pricing in climate change policy. When considering the introduction of carbon pricing, many governments, including the US federal government, propose monetary compensations to citizens who would be negatively impacted by carbon pricing. Policymakers usually propose a lump-sum fixed credit to energy bills, hoping that a fixed credit would not distort the marginal incentive to conserve energy. However, if customers do not distinguish fixed costs from variable costs, the fixed credit on energy bills may still discourage conservation and defeat the purpose of carbon pricing.[2]

Finally, we provide one of the first pieces of empirical evidence on the long-run responses to energy prices in developing countries. In the coming decades, most of the global increase in energy demand will come from developing countries (Wolfram, Shelef and Gertler, 2012). Understanding how to design energy pricing in these countries is, therefore, a first-order priority for addressing climate change and global scarcity in natural resources. However, the energy demand literature has focused on developed nations because of the availability of administrative billing data.[3] More-

[2]An example includes the compensation scheme proposed in the American Clean Energy and Security Act of 2009, described on page 901 of Congress (2009). Burtraw (2009) and Burtraw, Walls and Blonz (2010) also note that distributing a fixed credit may not work in the desired way if residential customers do not pay attention to the difference between their marginal price of electricity and their electricity bill.

[3]For example, see Borenstein (2012), Aroonruengsawat and Auffhammer (2011), Wolak (2011), Ito (2014), Jessoe

over, nearly all existing studies focus on estimating short-run demand elasticity because long-run exogenous variation in energy prices is rarely available.[4] We collected administrative billing data in China, and our quasi-experimental design allows us to estimate three-year responses to long-run price variation. Our findings on the sizable and long-lasting welfare gains highlight that inefficiency in energy pricing is likely to be large in developing countries, and, therefore, it is important to conduct rigorous studies in these countries.[5]

2 Key Features of the Heating Price Reform

This section describes the key features of the heating price reform in Tianjin for our research design. First, the city of Tianjin required household-level metered usage data to be collected at least one year before the introduction of consumption-based billing (CBB), allowing us to access daily household-level usage data *before and after* the price reform. Second, the reform was introduced with staggered rollouts, which creates quasi-experimental variation in treatment. Third, the policy-induced price variation allows us to test if consumers properly distinguished between fixed and variable costs.

2.1 Metered Data

Since 1958, the Chinese government has provided centralized, coal-fired heating to cities north to the Huai River. These cities constitute roughly half of China's urban population. The urban heating sector accounts for about 25 percent of total commercial energy use north of the river. This coal-based heating system has been inefficient for two major reasons. First, the heating facilities were mostly built in the 1950s and 1960s based on standards of Soviet technology that do not allow

and Rapson (2014), Ito (2015), Ito, Ida and Tanaka (2018), Deryugina, MacKay and Reif (2020), and Shaffer (2020) for studies based on administrative energy billing data in the United States, Japan, and Canada. Recently, researchers have started to collect such data in developing countries: Mexico (Davis, Fuchs and Gertler, 2014), South Africa (Jack and Smith, 2015, 2020), Colombia (McRae, 2015*b*,*a*), Brazil (Costa and Gerard, 2018), and Kenya (Lee, Miguel and Wolfram, 2020).

[4] Deryugina, MacKay and Reif (2020) emphasize this point and estimate two-year responses to electricity prices in Illinois. They find that Illinois households gradually responded to changes in electricity prices, which is consistent with our findings for Chinese households. Another related study is Costa and Gerard (2018), but their study focuses on persistent responses to a temporal policy shock and is therefore different from Deryugina, MacKay and Reif (2020) and our study.

[5] For example, Wolak (2011) and Ito, Ida and Tanaka (2018) find that the introduction of residential dynamic electricity pricing in the United States and Japan, which created increases in peak-hour prices by 100–300%, induced 10–15% reductions in electricity usage. Another policy that has been extensively studied in many developed countries is information provision with peer comparison in energy usage, which typically induced 1–2% reductions in energy use (Allcott and Rogers, 2014).

households to control their heating and do not have meters to record household usage. Second, without metered usage data, it has been practically impossible to provide incentives for households to respond to market-based energy costs. Billing is based on a flat per square meter price for an entire heating season, regardless of actual heating usage.

China's Ministry of Housing and Urban-Rural Development (MOHURD), in collaboration with the World Bank, piloted a reform in 2005 in seven cities to improve energy efficiency in the heating sector. The reform created a market mechanism so that consumers pay for their actual heating consumption. Individual manual or thermostatic valves are installed to enable households to control indoor temperature, and household meters are installed to establish metering consumption and then introduce consumption-based billing (CBB).

A common challenge to evaluate a metered price reform is that utility companies usually start metered pricing right after meters are installed. Household metered data are therefore only observable after the price reform, which prevents the comparison of household-level usage before and after the reform. The city of Tianjin provides a unique opportunity to overcome this challenge. The city government requires that the CBB starts at least one year after household meters are installed and household-level metered data are collected. Household daily heating usage under the flat-rate billing is metered for at least one heating season, making it possible to compare household usage *before and after* the price reform.

2.2 Staggered Rollouts

The CBB was introduced in Tianjin in staggered rollouts between 2008 and 2016. The long window of staggered rollouts enables us to estimate the policy's long-run effect using an event-study design. The vast majority of households in Tianjin live in condominiums, and, therefore, the rollouts were done at the condominium building level. Thus, in our estimation, we cluster the standard errors at the building level. There are 429 apartment buildings that had introduced consumption-based billing by 2016, and there are 16,425 households in these buildings. Figure A.1 shows substantial variation in the rollout by the number of households.

The city's annual operating budget for the reform was constrained, which was the reason the staggered rollouts spanned nine years. According to the city officials, the rollouts were done in an unsystematic order, though the timing was not randomly assigned. Hence, our estimation relies on the standard identification assumptions for quasi-experimental event study design, as we describe in Section 4. For example, we test if the rollout timings are correlated with building characteristics. We do not find statistically significant relationships between the rollout timings and the observable building characteristics, including year of building, square meters of residence, and house value, as we discuss in Section 4.2.

Households are fully informed about the start of the new billing scheme. Before heating is turned on in mid-November of the first season of CBB, the HOA office sends every household a letter in October to announce the change in the billing method. At the same time, every household also receives a user handbook from utility companies. The handbook explains the new billing policy in detail, including how households can adjust indoor temperature, how household usage is metered, how a consumption-based bill is calculated, etc.

Once a building was assigned to start CBB, all households were defaulted into CBB. However, there was an option to opt out from CBB and stay with the pre-reform fixed payment. To do so, households had to opt out before the first winter of CBB. In our data, we observe that about 70 percent of households complied with CBB and 30 percent of households opted out. For this reason, we estimate the intention-to-treat (ITT) and the average treatment effect on the treated (ATET) in Section 4.[6]

2.3 Price Variation Created by the Reform

Before the policy change, households paid an annual fixed charge equal to 3.97 dollars times their residence's square meters. For example, a household with 100 square meters of residence paid 397 dollars for every winter, regardless of how much heating was used.

After the policy change, a heating bill became a two-part tariff: 1) a new annual fixed charge equal to 1.895 dollars times square meters, and 2) a variable charge equal to 1.4 cents per kWh of heating usage.[7] This policy change provides useful variation for our empirical analysis because many consumers experienced an increase in marginal price but a decrease in average price. For example, consider a household with 100 square meters of residence whose typical usage was 10,000 kWh per winter (usage per square meter is 100 kWh). The pre-reform payment was 397 dollars

[6]For both CBB compliers and non-compliers, we observe their daily metered heating usage.

[7]The regulator set the marginal price equal to the marginal cost based on information on heating production.

with zero marginal price. With the same amount of usage, the post-reform payment would be 338.5 ($= 198.5 + 0.014 \cdot 10,000$) dollars with marginal price equal to 1.4 cents. Thus, given the same amount of usage, this household experienced an increase in marginal price but a decrease in average price.

Figure 1 visualizes how the CBB policy changed the marginal price and average price of heating for a given level of usage per square meter. The change in marginal price was common to all households—it changed from 0 to 1.4 cents per kWh. However, the change in average price depended on heating usage per square meter. Given the same usage level, households whose usage per square meter was less than 142 kWh were likely to experience a decrease in average price and the rest of households were likely to experience an increase in average price. This implies that many consumers were likely to experience an *increase* in marginal price but a *decrease* in average price, which is the key variation we use in Section 5.

3 Data

3.1 Household Heating Usage Data

We obtained administrative data on household-level daily heating usage from a regulated heating provider in a district of Tianjin. The data cover usage for all of the provider's residential customers from December 2007 to February 2019. Heating usage is automatically recorded once a day and uploaded to the provider's database. With a confidentiality agreement, we obtained direct access to the database. To our knowledge, our study is among the first to use such high-frequent administrative data on energy usage in developing countries. Most previous studies in developing countries rely on survey data on energy usage, which could suffer from issues of self-selection and measurement errors.

Every year, a winter heating season starts in mid-November and ends in mid-March. Each year's exact start date in November and end date in March depend on that year's temperature. To make our analysis consistent between years, we focus on daily usage in three full months of a heating season—December, January, and February. In these three months, heating is on every day of every year so that we do not have missing days for any year. In our empirical analyses, we aggregate the household daily data to the household-by-month level, and we analyze average daily heating usage by household and month.

All households in our data have at least one year of metered heating data (three winter months) prior to the start of the CBB. For about 40% of households, we observe at least two years of metered heating data in the pre-reform period (six winter months). For the post-reform period, all households have at least three years of metered heating data.

We observe daily heating usage data from 16,425 households in 429 buildings. We also observe each household's address, apartment number, square meters, and house value. Table 1 reports summary statistics. The average heating usage is 95 kWh per day and 11,600 kWh per winter. The average heating bill per winter is 413.6 dollars before the reform and 370.6 dollars after the reform. The average size of residence is 105.3 square meters. The take-up rate of the consumption-based billing is 70%.

3.2 Air pollution data

To examine the impact of the heating price reform on environmental externalities, we ask two questions: 1) does the change in household heating usage at home, induced by the reform, affect pollution emissions at the heating plant of the utility company? and 2) how does local ambient air quality change?

Our research site, a district of Tianjin, provides an interesting setting to answer these questions. First of all, the heating plant is located close to the residential area where most households in our data live, and a monitor of ambient air pollution is located in this residential area. The heating plant is a major local emitting source in winter and is about 8 kilometers away from the pollution monitor. We would expect that, if changes in household heating usage affect pollution emissions of the plant, the local ambient air quality might also be affected. Second, the district of our study is located in a relatively isolated part of Tianjin, about 55 kilometers away from the Tianjin metro area. The remoteness of its location is useful for our analysis because other emitting sources in the metro area are less likely to affect local air quality in this district.

We obtained pollution data from two sources. To measure pollution emissions of the heating plant, we collected hourly emission concentrations for SO_2, NO_x and PM from the Continuous Emission Monitoring System (CEMS) monitor placed at the heating plant. To measure local ambient air quality, we compiled daily readings of SO_2, PM_{10} and $PM_{2.5}$ from the pollution monitor in this district. The PM_{10} readings are particularly useful for the welfare analysis on externalities, because we can combine the changes in PM_{10} with the measure of marginal willingness-to-pay for PM_{10} in Ito and Zhang (2020) to measure the changes in welfare gain from externalities.

4 Impacts of the Consumption-Based Billing on Heating Usage

In this section, we estimate the causal effect of the CBB on heating usage. As described in Section 2.2, the CBB was implemented with an opt-out option, and about 30% of households opted out. This created one-sided incomplete compliance because all households in the control group were untreated, while there was incomplete compliance in the treatment group. For this reason, we estimate both the intention-to-treat effect (ITT) and the average treatment effect on the treated (ATET).

4.1 Overall Policy Impacts

We begin by estimating the ITT. Our identification strategy is based on the staggered rollout of the CBB as described in Section 2.2. Our estimating equation is,

$$y_{it} = \alpha_i + \gamma_t + \sum_{k=a}^{b} \phi_k D_{it}^k + u_{it}, \tag{1}$$

where y_{it} is the natural log of average daily heating usage for household i in year-month t, α_i indicates household-level fixed effects, and γ_t indicates year-month fixed effects. We use $k = [a, b]$ to denote the *event-time* relative to the first month of treatment. For example, $k = 0$ is the last month of the pre-treatment period and $k = 1$ is the first month of treatment. Note that we use data from three winter months, between the first day of December and the last day of February. Therefore, if we consider a household whose first treatment month was December 2010, k equals 0 in February 2009, 1 in December 2010, 2 in January 2011, 3 in February 2011, 4 in December 2012, and etc. Dummy variables $D_{it}^k = 1$ if year-month t falls within the event-time k for household i.

Recent developments in the econometrics literature highlight that conventional OLS could produce biased estimates for two-way fixed effect models such as equation (1) if treatment effects are heterogeneous across households and/or time (de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2021). To address this problem, we use a method developed by de Chaisemartin and D'Haultfœuille (2020) to estimate equation (1) so that we do not impose the assumption of homogeneity in the treatment effects. For comparison, we also show our results based on the conventional OLS in Appendix A. We indeed find that the results are different when we impose the assumption of homogeneity in the treatment effects. de Chaisemartin and D'Haultfœuille (2020) show that the reason why the conventional OLS may not produce the correct average treatment effects is because it produces an incorrectly weighted average of treatment effects across cohorts and time. Moreover, some of these incorrect weights can be negative in theory, which could make the OLS estimates significantly different from the correctly weighted average of the cohort-by-time treatment effects. In Figure A.4, we show that this is indeed the case in our data. We find that 46% of cohort-by-time weights are negative if we use the conventional OLS, which suggests that it is important to incorporate heterogeneity in the treatment effects in our setting.[8]

The primary variables of interest are ϕ_k. These coefficients provide the ITT estimates of mean log average daily usage for event time k, controlling for household fixed effects and time fixed effects. The excluded group is $k = -1$, the last month of the pre-treatment period. Thus, we can interpret ϕ_k as the difference in mean log average daily usage between event month k and the last month of the pre-treatment period. We need identification assumptions that are standard in the difference-in-differences method. In the absence of the treatment, there would be no difference in mean log usage between the event months—that is, k would be zero. The validity of this identification assumption is untestable, but we can assess if the parallel trend assumption is consistent with data in the pre-treatment period.

In Figure 2, we show the estimates of ϕ_k in $k = [-5, 9]$—between two years before and three years after the reform. This figure provides three key results. First, there is no statistically significant pre-existing trend before the treatment. Second, the ITT estimate for the first year is about a 10 percent reduction in heating usage. Third, the policy impact is similar in the second and third years after the policy implementation.

[Figure 2 about here]

[8] Results in an earlier version of our working paper were produced by the conventional OLS with two-way fixed effects. In Appendix A, we compare estimation results based on de Chaisemartin and D'Haultfœuille (2020) and those based on the conventional OLS and provide a discussion in detail.

In Table 2, we provide the ITT estimates of the CBB for each of the three post-reform years in column 1. These results suggest that the CBB resulted in reductions in heating usage by 10.1% in the first year, 10.7% in the second year, and 8.7% in the third year (-0.107, -0.113, and -0.091 log points, respectively).[9] These results are consistent with the findings in Figure 2.

[Table 2 about here]

In column 2 of Table 2, we provide the ATET of the CBB for each of the three post-reform years. We estimate equation (1) by replacing the treatment assignment D_{it}^k by household i's actual treatment status T_{it}^k in event time k. We use D_{it}^k as instruments for T_{it}^k to obtain the IV estimates. With the standard assumptions required for the local average treatment effects (Imbens and Angrist, 1994), the IV estimates can be interpreted as the ATET because the incomplete compliance due to the opt-out decision was only in the treatment group. The ATET of the CBB are 14% in the first year, 15% in the second year, and 12% in the third year (-0.153, -0.161, and -0.13 log points, respectively).

4.2 Validity of the Identification Assumptions

The validity of our quasi-experimental approach is subject to a standard set of identification assumptions for the difference-in-differences method with staggered rollouts. A key assumption is the parallel trend in the counterfactual outcome—in the absence of the treatment, the trajectory of the outcome variable (i.e., heating usage in our context) has to be parallel between treatment and control groups. Although this is an untestable assumption, we provide two pieces of supporting evidence. The first evidence is the absence of pre-existing trends in Figure 2. The trajectory of heating usage was not statistically different between the treatment and control groups prior to the treatment periods. Second, we test if building characteristics are associated with the staggered timings of the policy implementation. We do not find statistically significant relationships between the timings of the CBB rollout and the vintage of building, size of residence, house value, and annual heating usage prior to the CBB in Table A.1.

To interpret our IV estimates as the ATET, we also need the standard set of assumptions for the LATE (Imbens and Angrist, 1994). A potential concern is that the Stable Unit Treatment

[9] y_{it} is heating usage in log. Therefore, ϕ_k (-0.107, -0.113, -0.091 in column 1 of Table 2) provide the ITT estimates in log points. The ITT estimates in percentage changes can be obtained by $\exp(\phi_k) - 1$.

Value Assumption (SUTVA) can be violated if a household's usage is affected by other households' compliance decisions. To examine this concern, we test if the change in household heating usage is correlated with the compliance rate of its neighbors. In Table A.3, we do not find statistically significant correlation between one's response to the CBB and the compliance status of one's neighbors either next door or in upper and lower floors.

4.3 Interpretation of the Overall Impacts of the CBB

Overall, the findings in Table 2 indicate that the CBB resulted in statistically and economically significant changes in heating usage. Reductions in residential energy usage by 10% in the ITT and 15% in the ATET are sizable relative to the estimates of other policies evaluated in the literature. For example, the effects of non-price energy conservation programs such as home energy reports with peer comparisons usually produced 1-2% reductions in residential energy usage (Allcott, 2011). The short-run effects of dynamic electricity pricing are found to be between 10% and 15% (Wolak, 2010; Ito, Ida and Tanaka, 2018, forthcoming).

In addition, the staggered rollout of the CBB allows us to estimate long-run effects, which are challenging to estimate in most existing studies because obtaining long-run exogenous variation in energy price is generally difficult.[10] Our results suggest that the CBB produced long-lasting reductions in heating usage for at least three years after the introduction of the policy.

Hence, if a policymaker's objective is to reduce residential heating usage, our analysis in this section indicates that the CBB is effective to produce long-lasting sizable impacts. However, as we show in detail in Section 7, the welfare implications of the CBB depends on whether consumers distinguish between fixed and variable costs when they are faced with a two-part tariff such as the CBB. For this reason, we investigate this question in the next section before we consider the welfare implications of the CBB in Section 7.

[10] Most studies in the literature on residential energy demand estimate short-run effects based on price variation that lasted for a few months to a year (Wolak, 2010; Ito, Ida and Tanaka, 2018). Few papers are able to estimate long-run effects, including Allcott and Rogers (2014) and Deryugina, MacKay and Reif (2020).

5 Do Consumers Distinguish Fixed Cost from Variable Cost?

In Section 4, we find that the CBB induced overall reductions in heating usage *on average*. However, this finding by itself does not reveal whether consumers distinguished between fixed and variable cost in responding to the introduction of the two-part tariff. In this section, we exploit the price variation created by the CBB to test this question. As we show in section 7, testing this hypothesis is key to the welfare implications of the CBB and two-part tariffs in general.

5.1 Conceptual Framework

Consider a utility maximization problem for heating demand y. A consumer has income I, marginal price of heating p, and fixed charge f. We consider a quasi-linear utility function $u = v(y) - py - f + I$.[11] A standard utility maximization problem simply solves the first order condition for the utility function with respect to y, leading to $v'(y^*) = p$. Therefore, the optimal usage y^* can be obtained when the marginal utility from heating usage equals marginal price. In particular, this standard model predicts that an increase in marginal price should result in a decrease or no change in heating usage.

In contrast, Liebman and Zeckhauser (2004) suggest the "schmeduling" model as an alternative model of consumer behavior. A consumer who is faced with nonlinear pricing may misperceive a change in average price (p_a) as a change in marginal price. Two-part tariffs could create this misperception if a consumer does not properly distinguish between changes in fixed costs and changes in variable costs. In this case, the optimal usage can be characterized by $v'(y^{**}) = p_a$. Importantly, the shmeduling model predicts that an increase in marginal price could result in a *increase* in heating usage when a consumer has an increase in marginal price but a decrease in average price.

5.2 Empirical Tests for Schmeduling

We propose a simple non-parametric test for the schmeduling with a two-part tariff. Our approach exploits the unique price variation created by the introduction of the CBB. As we described in Figure 1, some consumers in our data experienced a policy-induced increase in marginal price and

[11] A quasi-utility function assumes that there is no income effect. This assumption is likely to be valid in our empirical context because the income effect of the CBB policy was likely to be very small. We explore this point in Section 6.2.

a *decrease* in average price. This is because the CBB increased the marginal price and lowered the fixed charge. We use β to denote the impact of the CBB on heating usage for these consumers. As we described in the conceptual framework in Section 5.1, the standard model predicts that $\beta \leq 0$, but the schmeduling model predicts that $\beta > 0$. Therefore, we can apply the estimation method described in Section 5 to the subgroup of consumers to conduct this simple statistical test. An advantage of this test is that we do not need to impose functional form assumptions on the demand curve. We simply test how the marginal changes in price affected heating consumption.

A naive way to identify this subset of consumers is to look at the actual average price paid by each consumer. This approach would, however, create an endogeneity concern because the actual average price is a function of contemporaneous heating usage. To address this issue, we follow the literature on nonlinear income taxation and pricing and identify the *policy-induced change* in average price (Saez, Slemrod and Giertz, 2012; Ito, 2014). For each customer, we construct the *predicted* changed in average price by using their heating usage two years prior to the introduction of the CBB. This predicted change in average price does not depend on each customer's heating usage after the introduction of the CBB, and therefore is purely driven by the changes in the price schedule induced by the policy change. Consistent with the previous studies in this literature, we find that the predicted change and actual change are highly correlated, as we show below.

We begin by testing for visual evidence in Figures 3. We divide households into a decile based on the predicted change in average price. Separately for each decile group, we estimate the ITT effect of the CBB using the estimation method described in Section 5. Recall that all consumers had the same change in marginal price—an increase in marginal price by $0.014 per kWh. However, the changes in average price were different across groups. The decile groups 1 to 3 had decreases, the decile group 4 had nearly no change, and the decile groups 5 to 10 had increases in average price.

[Figure 3 about here]

The standard model predicts that all groups would reduce usage because everyone had an increase in marginal price. However, the changes in heating usage in Figure 3 are inconsistent with this prediction. The decile groups 1 to 2 have increases, the decile groups 3-4 has nearly no change, and the decile groups 5 to 10 have decreases in heating usage. The relationship between the changes in average price and the changes in usage suggest that the empirical evidence is more consistent with the schmeduling model.

With the insight from Figure 3, we conduct statistical tests for schmeduling in Tables 3 and 4. Our procedure is similar to the approach used in Figure 3 except that we use a quartile rather than a decile to increase the precision of the statistical tests. We divide households into a quartile based on the predicted change in average price. Separately for each quartile group, we estimate the ITT effect of the CBB using the estimation method described in Section 5.

[Table 3 about here]

Table 3 shows that while all groups had the same increase in marginal price, quartile group 1 had a decrease, the group 2 had nearly no change, and groups 3 and 4 had increases in average price. For quartile group 1, the CBB increased heating usage by 0.216 log points (a 24.1 percent increase).[12] For quartile group 2, the impact of the CBB on heating usage is statistically insignificant from zero.

In Table 4, we provide further analyses by focusing on quartile group 1, who experienced policy-induced increase in marginal price and decrease in average price. The first row includes all customers in this quartile to reproduce the results shown in column 1 of Table 3. The null hypothesis that consumer behavior is consistent with the standard economic model (against the schmeduling model) is $H_0 : \beta \leq 0$, in which β is the ITT of the CBB on heating usage. We provide the p-value of this test in the last column. We reject this null at the 1% statistical significance level for the full sample, indicating that consumer behavior is inconsistent with the standard model.

[Table 4 about here]

We also examine if there are subgroups of customers who might behave more consistently with the standard model. We further divide quartile group 1 by home value in rows 2 and 3 and by home sizes in rows 4 and 5. Although the magnitudes of the ITTs are heterogeneous across these groups, we reject the null of the standard model for all of these subgroups.

In Tables A.4 and A.5, we provide the same analyses based on the ATET instead of the ITT. We find that qualitatively, our findings do not change if we use the ATET and we reject the null

[12] As we explained above, the treatment effect in log points (β) can be converted to the the percent change by $\exp(\beta) - 1$.

for all subgroups based on the ATET as well. In fact, the ATET rejects the standard theory more strongly than the ITT because the ATET on consumption is larger than the ITT on consumption in absolute value as we have one-sided incomplete compliance.[13]

5.3 Mean Reversion and Validity of Identification Assumptions

The empirical evidence in the previous section suggests that consumer behavior is inconsistent with the standard economic model and more consistent with the schmeduling model. An important factor that we want to be careful about is mean reversion in the outcome data and its possible threat to identification in consumer behavior with nonlinear price schedules (Saez, Slemrod and Giertz, 2012; Ito, 2014). Many types of economic panel data, including household-level heating consumption, tend to have mean reversion. For example, consider a household who had a negative (positive) transitory shock to their heating consumption in one period. Then, this household's consumption tends to be higher (lower) in other periods because of mean reversion. This natural mean reversion is important to be considered when researchers analyze consumer behavior in nonlinear price schedules.

In our context, we calculate each household's policy-induced change in average price using their consumption in the pre-CBB baseline period. Consider a household whose usage was low in the pre-CBB baseline period. Naively comparing their usage in the baseline period against their usage in other periods could lead to a misleading conclusion if this household had a transitory negative shock in the baseline period. Their usage will be higher in other period because of mean reversion, not necessarily because of the policy impact.

Our empirical analysis controls for this mean reversion by using the staggered DID design presented in the previous sections. Instead of naively comparing a treated household's usage before and after the CBB, our staggered DID design uses data from untreated households to control for time-varying changes in usage, including mean reversion. To make it clear, consider example households A and B who had an identical level of consumption in the pre-CBB baseline period. Suppose that in the staggered rollout of the CBB, household A started the CBB several years earlier than household B. In this example, our staggered DID design controls for the mean reversion for household A by using data from household B.

[13]In our context, consumers could have opted out from the CBB once they are assigned to treatment, but those who were not assigned to treatment were not able to have the CBB. Therefore, the ITT eqals to the products of the ATET and take-up rate, and thus, the ATET is larger than

A key assumption in this approach is the standard parallel trend assumption—time-varying unobserved factors that affect heating usage, including the mean reversion, should not be systematically different between households who had an earlier rollout of the CBB and those who had a later rollout.

This identification assumption is unstable, as is the case for any DID design with quasi-experimental data, but we can assess the validity of the assumption by examining the pre-trend data in the event study figure in Figure 4. The pre-trends suggest that time-varying unobservable factors that affect heating usage are not systematically different between treated and untreated households in the pre-CBB period. In particular, if mean reversion is the reason for an increase in usage for quartile group 1 and decreases in usage for quartiles 3 and 4, we should observe these changes in usage in all years, including the years before the CBB.[14] However, Figure 4 suggests that we do not observe such changes until the beginning of the CBB. This provides supporting evidence that time-varying unobservable factors including mean reversion are unlikely to be systematically different between households with different rollout timings.[15]

[Figure 4 about here]

6 Alternative Explanations

In the previous sections, we find empirical evidence that is consistent with the prediction from the schmeduling model. We explore three alternative explanations in this section.

6.1 Category Budgeting

The category budgeting model suggests that households may keep track of category-specific budgets and try to maintain category spending at a target level (Heath and Soll, 1996; Antonides and Raaij, 2011; Hastings and Shapiro, 2013). Hastings and Shapiro (2013) show that one way to model this

[14] Recall that we use the first-available year's data in the pre-CBB period to define the quartile groups. Therefore, if mean reversion drives the increase or decrease in usage, we should observe it before the start of the CBB including one and two years before the CBB. However, the pre-trend in these figures do not show such evidence.

[15] In addition, we provide further robustness checks in Table A.2 and Figure A.3. In these analyses, we include further flexible controls for mean reversion by interacting the time fixed effects with the calendar year of the pre-CBB baseline period. In this way, we can allow mean reversion and other time-varying unobservable factors to be different between households whose pre-CBB baseline period fell in different calendar years. We find that the results are robust to these controls.

behavior is to consider that households experience disutility from spending an atypical amount on a certain good. In our context, this model implies that consumers experience disutility from spending a different amount of money on heating in the new price schedule (the two-part tariff) as compared to what they used to pay in the old price schedule (the fixed payment).

We test if consumer behavior in our data is consistent with the categorical budgeting model. To do so, we exploit the opt-out structure of the reform. Recall that households could opt-out from the new pricing if they wanted to stay with the old pricing. If households wanted to maintain the pre-reform level of heating expenditure, the most obvious way to do so would be to opt out from the new pricing. Importantly, the category budgeting model implies that the incentive to opt out was equally large for those who expected a large increase in payment and those who expected a large decrease in payment with the new pricing.

We test this prediction in Figure 5. For each household, we calculate the predicted change in annual billing based on their pre-CBB heating usage. That is, we calculate how much a household would pay more (or less) with the CBB compared to the old pricing if they use the same amount of heating. The horizontal line in the figure is the decile of the predicted change in annual billing. The dashed line shows the average change in annual billing for each decile group, which ranges from a decrease of about 150 dollars to an increase of about 150 dollars. The solid line shows the opt-out rate relative to the opt-out rate for the decile group 1. We find that the opt-out rate is monotonically increasing—-that is, those who expected an increase in payment were more likely to opt out, and those who expected a decrease in payment were less likely to opt out. This empirical finding is inconsistent with the prediction from the category budgeting model. This evidence is consistent with the response to average price because the signs and directions of the predicted changes in annual billing are the same as those of the predicted changes in average price.

[Figure 5 about here]

6.2 Income Effect

In Section 5.1, we consider a quasi-utility function, which abstracts from the income effect. In theory, a change in fixed cost could create an income effect that affects heating usage. However, in our empirical setting, the income effect of the CBB is likely to be very small.

The CBB reduced the fixed cost of heating by $226 per year. The average household income in Tianjin in our sample period was $15,041. Therefore, the change in the fixed cost was about 1.5 percent of household income. In the literature on energy demand, the short-run residential income elasticity estimates are found to be moderately inelastic and estimated around 0.239 on average based on a recent meta analysis (Zhu and Yang, 2018). This implies that the income effect of the CBB would be an increase in heating usage by 0.36 percent.

This effect is too small to explain our findings in Table 4. We find that customers who experienced an increase in marginal price and a decrease in average price increased usage by 24.1 percent (a 0.216 log-point increase), which is a lot larger than 0.36 percent. To explain this finding by the income effect, the short-run income elasticity of heating demand would have to be 16.07, which is far larger than upper bounds of empirical findings in the literature.

6.3 Spurious Correlations

Recall that customers who experienced an increase in marginal price and a decrease in average price were those who had relatively low usage per square meter before the introduction of the CBB (see Figure 1). One potential concern is that this customer type is correlated with other important household characteristics, and this type of customer increased usage in response to the CBB for a reason unrelated to their average price.

An example of potential spurious correlations is the location of housing units. For instance, although we showed in Section 4.2 that one's response to the CBB is unlikely to be affected by their neighbors in our setting, we could consider the following hypothesis. Consider households who were in the middle floors or non-corner units and suppose that they received heating spillovers from their neighbors prior to the reform. For this reason, their pre-reform usage per square feet can be lower than others. After the reform, many neighbors reduced usage, and therefore, these particular types of households needed to increase usage.

We test this hypothesis in Table A.6 in the appendix. In Panel A, we estimate the ITT of the CBB policy separately for households living on middle floors vs. top/bottom floors. Similarly, in Panel B, we estimate the ITT of the CBB policy separately for households living in corner vs. non-corner units. In each panel, we find that both types of customers reduced usage. If anything, the reduction in usage was larger for those who were on the middle floor and those how were in non-corner units in the point estimates. These findings are inconsistent with the hypothesis above.[16]

7 Welfare Implications

In Figure 6, we describe the social welfare gains from the CBB. Panel A considers the standard theory of consumer behavior on a two-part tariff. This model assumes that consumers distinguish between fixed and variable cost, and therefore, their consumption is determined by the intersection of their demand curves and marginal price. In this case, both consumers A and B in the figure would reduce usage because the CBB increased the marginal price of heating from 0 to $0.014 per kWh. Then, the social welfare gains from the CBB would be the shaded areas under the social marginal cost.

[Figure 6 about here]

In contrast, the welfare implications are different with the schmeduling model in Panel B. In this model, consumers do not distinguish between fixed and variable costs, and they respond to their average price. This implies that for consumer B in the figure, the reduction in usage is smaller than the one in Panel A because the change in average price is smaller than the change in marginal price. Moreover, consumer A would increase usage because the CBB lowered the average price. This implies that the change in social welfare is negative for consumer A in this case.[17]

In Table 5, we use our data and empirical findings to calculate the policy's social welfare gains based on the two approaches shown in Figure 6.[18] In the first two columns, we calculate the social welfare gain per household per year. Our estimate suggests that the standard approach overstates the welfare gain. By incorporating the schmeduling, we estimate that the social welfare gain from the CBB is 6.3 USD per household per year, which is 26.8 million USD per year for the city of Tianjin.

[16] To locate households living in corner units and non-corner units, we collected information on building structure and apartment locations. In buildings that host three households per floor, four households per floor and eight households per floor, one can clearly define corner units vs. non-corner units. We therefore focus on this sub-sample in Panel B.

[17] Note that in theory if the social marginal cost is low enough, the social welfare gain for consumer A can be positive because the socially optimal level of consumption can be closer to y_1^A than y_0^A in that case. However, empirically we find a substantial level of environmental externalities in Appendix B. We draw the social marginal cost in the figure consistent with our empirical setting.

[18] In Appendix B, we use ambient air pollution data to calculate the negative environmental externalities. We add these externalities and the private marginal cost of heating to compute the social marginal cost.

[Table 5 about here]

Our analysis suggests that incorporating the schmeduling behavior could substantially change the welfare implications of the two-part tariff. In the implementation of the CBB, the city of Tianjin reported that the one-time cost of introducing CBB—including the cost of installing meters—was about 99 USD per household. With China's 3% discount rate in 2015, the conventional approach based on the standard theory suggests that the policy's net present value of the benefit would exceed its cost in about 4 years, whereas the schmeduling model implies that it takes about 21 years to be cost-effective. This result suggests that the CBB is still likely to be a cost-effective policy in terms of the social welfare gain in the long run, but the cost-benefit analysis would substantially overestimate the benefit if the shmeduling behavior is not took into account.

Finally, we also evaluate if the CBB increased or decreased the consumer surplus. The consumer's surplus is another important welfare measure, as it may be challenging to obtain political support for a policy if many consumers would have decreases in consumer surplus even if the CBB results in an overall increase in social welfare.

We find that the CBB resulted in an increase in consumer surplus by 172 USD per household per year on average. We also calculate the change in consumer surplus by home value, which is a proxy for wealth. We find that the CBB increased the consumer surplus by 119 USD per household per year for households in quartile 1 of home value, 202 USD in quartile 2, 191 USD in quartile 3, and 249 in quartile 4. These findings suggest that the CBB resulted in an increase in consumer surplus for a broad set of the population.

8 Conclusion

In this paper, we examine long-run effects of a recent heating pricing reform in China that replaced fixed payment with a two-part tariff. Using staggered policy rollouts and administrative data on household-level daily heating consumption, we find that the reform induced significant and persistent reductions in heating usage. We find that consumer behavior in our data is inconsistent with the standard economic theory that assumes that consumers properly distinguish fixed cost from variable cost and more consistent with the "schmeduling" model in Liebman and Zeckhauser (2004). As we show its policy implications for two-part energy tariffs, this consumer behavior makes fixed cost directly relevant to the perceived relative prices of goods, and therefore, alters the welfare implications of price, tax, and subsidy designs.

References

Allcott, Hunt. 2011. "Social norms and energy conservation." *Journal of public Economics*, 95(9-10): 1082–1095.

Allcott, Hunt, and Todd Rogers. 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review*, 104(10): 3003–37.

Antonides, Gerrit, I Manon De Groot, and W Fred Van Raaij. 2011. "Mental budgeting and the management of household finance." *Journal of Economic Psychology*, 32(4): 546–555.

Aroonruengsawat, Anin, and Maximilian Auffhammer. 2011. "Impacts of Climate Change on Residential Electricity Consumption: Evidence from Billing Data." National Bureau of Economic Research, Inc NBER Chapters.

Borenstein, Severin. 2009. "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing." *Center for the Study of Energy Markets Working Paper*, 195.

Borenstein, Severin. 2012. "The Redistributional Impact of Nonlinear Electricity Pricing." *American Economic Journal: Economic Policy*, 4(3): 56–90.

Borenstein, Severin, and Lucas Davis. 2012. "The equity and efficiency of two-part tariffs in US natural gas markets." *The Journal of Law and Economics*, 55(1): 75–128.

Brown, Jennifer, Tanjim Hossain, and John Morgan. 2010. "Shrouded Attributes and Information Suppression: Evidence from the Field." *Quarterly Journal of Economics*, 125(2): 859–876.

Burtraw, Dallas. 2009. "Hearing on Climate Change Legislation: Allowance and Revenue Distribution." *US Senate Committee on Finance.*

Burtraw, Dallas, Margaret Walls, and Joshua A. Blonz. 2010. "Distributional Impacts of Carbon Pricing Policies in the Electricity Sector." *U.S. Energy Taxes, Gilbert E.Metcalf eds. American Tax Policy Institute.*

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer. 2006. "1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions." *The American Economic Review*, 96(4): 1253–1270.

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

Chetty, R., A. Looney, and K. Kroft. 2009. "Salience and taxation: Theory and evidence." *The American Economic Review*, 99(4): 1145–1177.

Congress. 2009. "The American Clean Energy and Security Act of 2009 (HR 2454)."

Costa, Francisco, and François Gerard. 2018. "Hysteresis and the welfare effect of corrective policies: Theory and evidence from an energy-saving program." National Bureau of Economic Research.

Davis, Lucas W., Alan Fuchs, and Paul J. Gertler. 2014. "Cash for Coolers: Evaluating a Large-Scale Appliance Replacement Program in Mexico." *American Economic Journal: Economic Policy*, 6(4).

de Bartolome, Charles A. M. 1995. "Which tax rate do people use: Average or marginal?" *Journal of Public Economics*, 56(1): 79–96.

de Chaisemartin, Clément, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–96.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2020. "The long-run dynamics of electricity demand: Evidence from municipal aggregation." *American Economic Journal: Applied Economics*, 12(1): 86–114.

Finkelstein, Amy. 2009. "E-ZTAX: Tax Salience and Tax Rates." *The Quarterly Journal of Economics*, 124(3): 969–1010.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *The Quarterly Journal of Economics*, 121(2): 505–540.

Hastings, Justine S., and J. M. Shapiro. 2013. "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics*, 128(4): 1449–1498.

Heath, Chip, and Jack B Soll. 1996. "Mental budgeting and consumer decisions." *ournal of consumer research*, 23(1): 40–52.

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica*, 62(2): 467–475.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–63.

Ito, Koichiro. 2015. "Asymmetric Incentives in Subsidies: Evidence from a Large-Scale Electricity Rebate Program." *American Economic Journal: Economic Policy*, 7(3): 209–37.

Ito, Koichiro, and Shuang Zhang. 2020. "Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China." *Journal of Political Economy*, 128(5): 1627–1672.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. "Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand." *American Economic Journal: Economic Policy*, 10(1): 240–67.

Ito, Koichiro, Takanori Ida, and Mokoto Tanaka. forthcoming. "Selection on Welfare Gains: Experimental Evidence from Electricity Plan Choice." *American Economic Review.*

Jack, B. Kelsey, and Grant Smith. 2015. "Pay as you go: Prepaid metering and electricity expenditures in South Africa." *American Economic Review*, 105(5): 237–41.

Jack, B. Kelsey, and Grant Smith. 2020. "Charging ahead: Prepaid electricity metering in South Africa." *American Economic Journal: Applied Economics.*

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review*, 104(4): 1417–38.

Kahn, Matthew E., and Frank A. Wolak. 2013. "Using Information to Improve the Effectiveness of Nonlinear Pricing: Evidence from a Field Experiment : Final Report." Stanford University Working Paper.

Lee, Kenneth, Edward Miguel, and Catherine Wolfram. 2020. "Experimental Evidence on the Economics of Rural Electrification." *Journal of Political Economy.*

Liebman, Jeffrey, and Richard Zeckhauser. 2004. "Schmeduling." *www.hks.harvard.edu/jeffreyliebman/schmeduling.pdf.*

McRae, Shaun. 2015*a*. "Efficiency and Equity Effects of Electricity Metering: Evidence from Colombia." *Working Paper.*

McRae, Shaun. 2015*b*. "Infrastructure quality and the subsidy trap." *American Economic Review*, 105(1): 35–66.

Rees-Jones, Alex, and Dmitry Taubinsky. 2020. "Measuring "Schmeduling"." *Review of Economic Studies*, 87(5): 2399–2438.

Saez, Emmanuel, Joel B. Slemrod, and Seth H. Giertz. 2012. "The elasticity of taxable income with respect to marginal tax rates: A critical review." *Journal of Economic Literature*, 50(1): 3–50.

Shaffer, Blake. 2020. "Misunderstanding Nonlinear Prices: Evidence from a Natural Experiment on Residential Electricity Demand." *American Economic Journal: Economic Policy*, 12(3).

Stiglitz, Joseph E. 1988. *Economics of the public sector.* WW Norton.

Tirole, Jean. 1988. *The theory of industrial organization.* MIT press.

Wolak, Frank A. 2010. "An Experimental Comparison of Critical Peak and Hourly Pricing: The PowerCentsDC Program." *Stanford Univeristy Working Paper.*

Wolak, Frank A. 2011. "Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment." *The American Economic Review*, 101(3): 83–87.

Wolfram, Catherine, Orie Shelef, and Paul Gertler. 2012. "How will energy demand develop in the developing world?" *Journal of Economic Perspectives*, 26(1): 119–38.

Zhu, Xing, Lanlan Li Kaile Zhou Xiaoling Zhang, and Shanlin Yang. 2018. "A meta-analysis on the price elasticity and income elasticity of residential electricity demand." *Journal of Cleaner Production*, 201: 169–177.

Figures

Figure 1: Policy-Induced Changes in Marginal and Average Prices

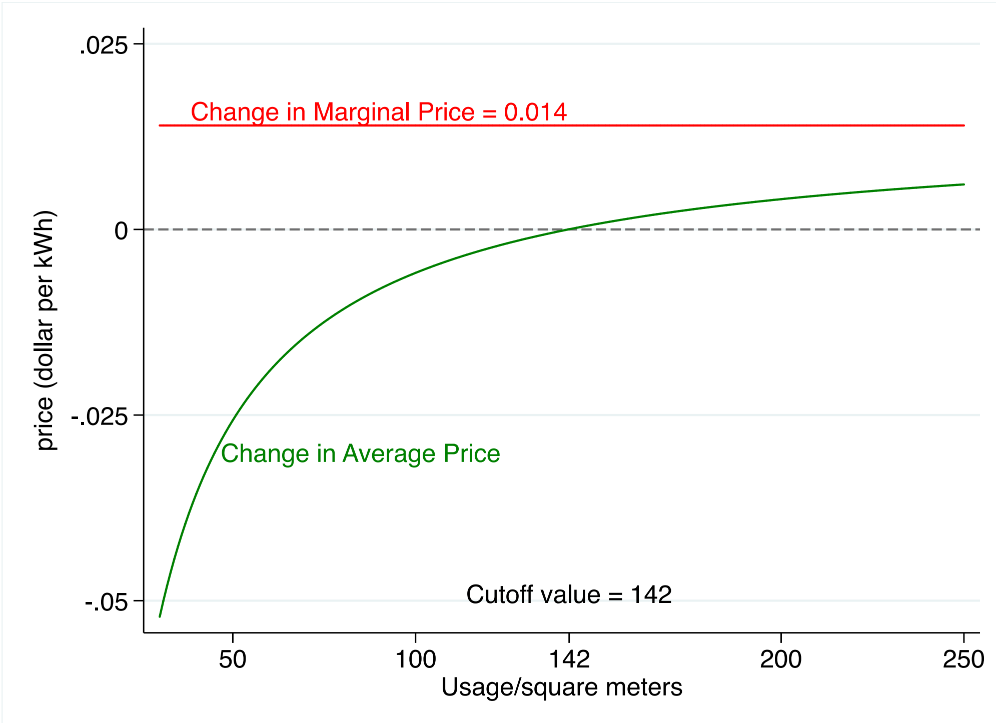

Notes: This figure shows the changes in marginal price and average price induced by the introduction of the consumption-based billing policy. All consumers had the same change in marginal price, but the policy-induced changes in average price depended on heating usage per square meter, and some consumers experienced an increase in marginal price and a *decrease* in average price.

Figure 2: Staggered Difference-in-Differences Analysis: Intention-to-Treat (ITT)

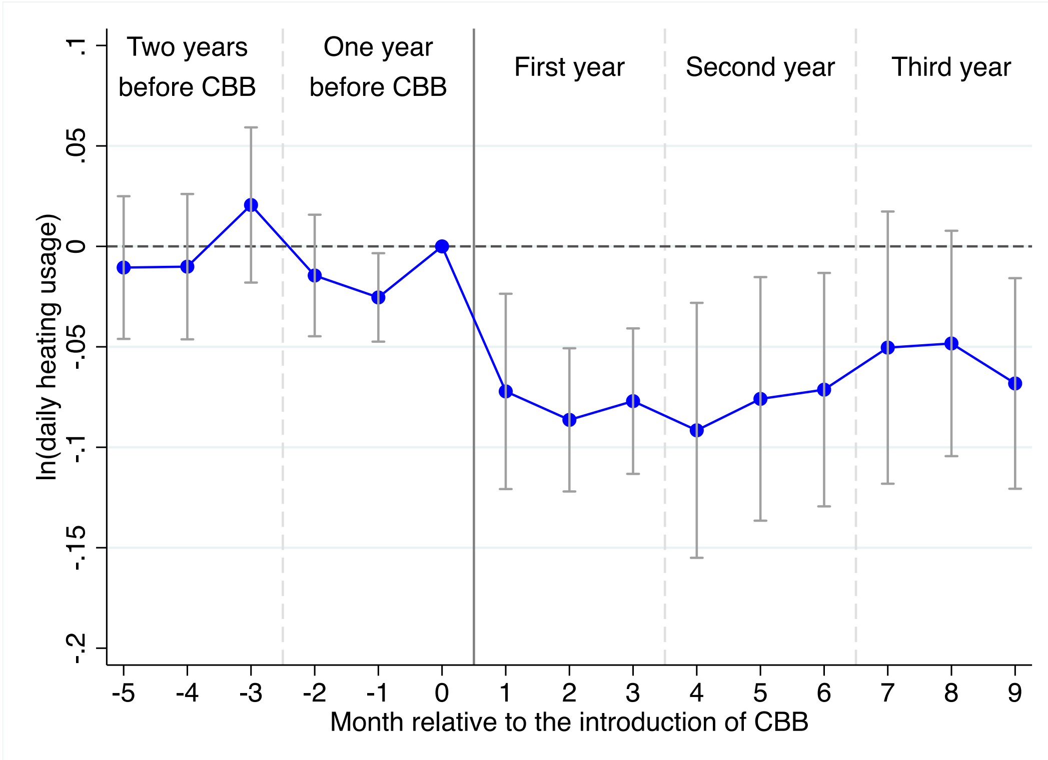

Notes: This figure shows the ITT estimates of the staggered difference-in-differences analysis described in equation (1) based on the estimation method developed by de Chaisemartin and D'Haultfœuille (2020). There are three heating months in each year because the heating season is December, January, and February. The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the building level.

Figure 3: Policy-Induced Changes in Average Price and Usage

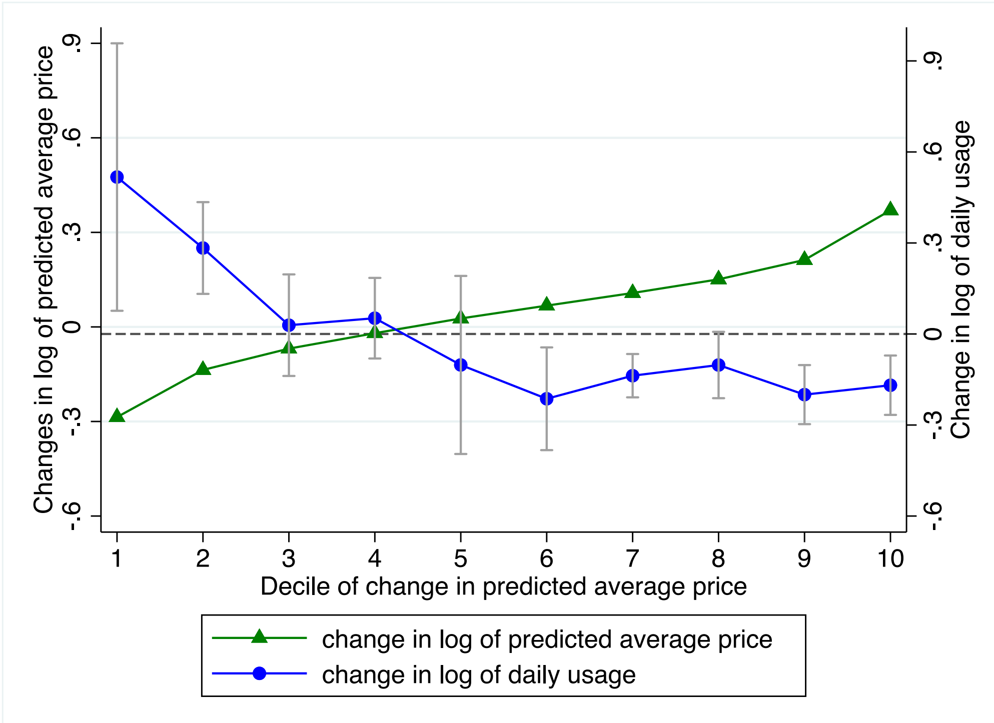

Notes: We divide customers by decile based on their policy-induced changes in average price. For each decile, we estimate the ITT of the CBB on the log of heating usage based on the difference-in-differences estimation method developed by de Chaisemartin and D'Haultfœuille (2020). We also apply the same method to estimate the ITT on the log of the policy-induced change in average price. The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the building level.

Figure 4: Staggered Difference-in-Differences by Quartile of the Policy-Induced Change in Average Price

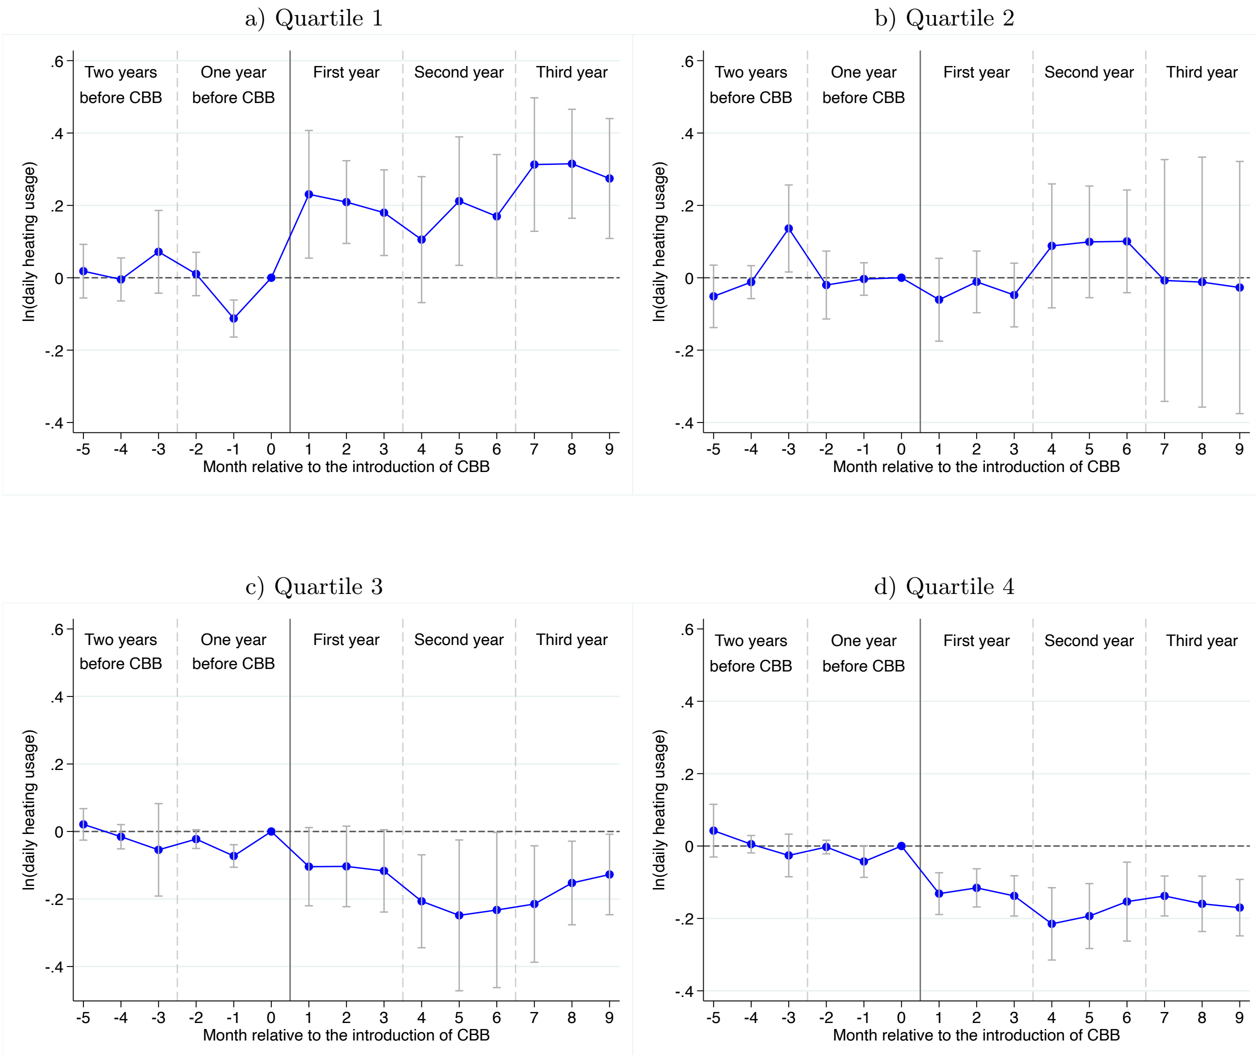

Notes: This figure shows the ITT estimates of the staggered difference-in-differences analysis described in equation (1) based on the estimation method developed by de Chaisemartin and D'Haultfœuille (2020). There are three heating months in each year because the heating season is December, January, and February. The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the building level. We divide customers by quartile based on their policy-induced changes in average price. We then estimate equation (1) for each quartile group separately to make these event study figures.

Figure 5: Household Opt-out Decision

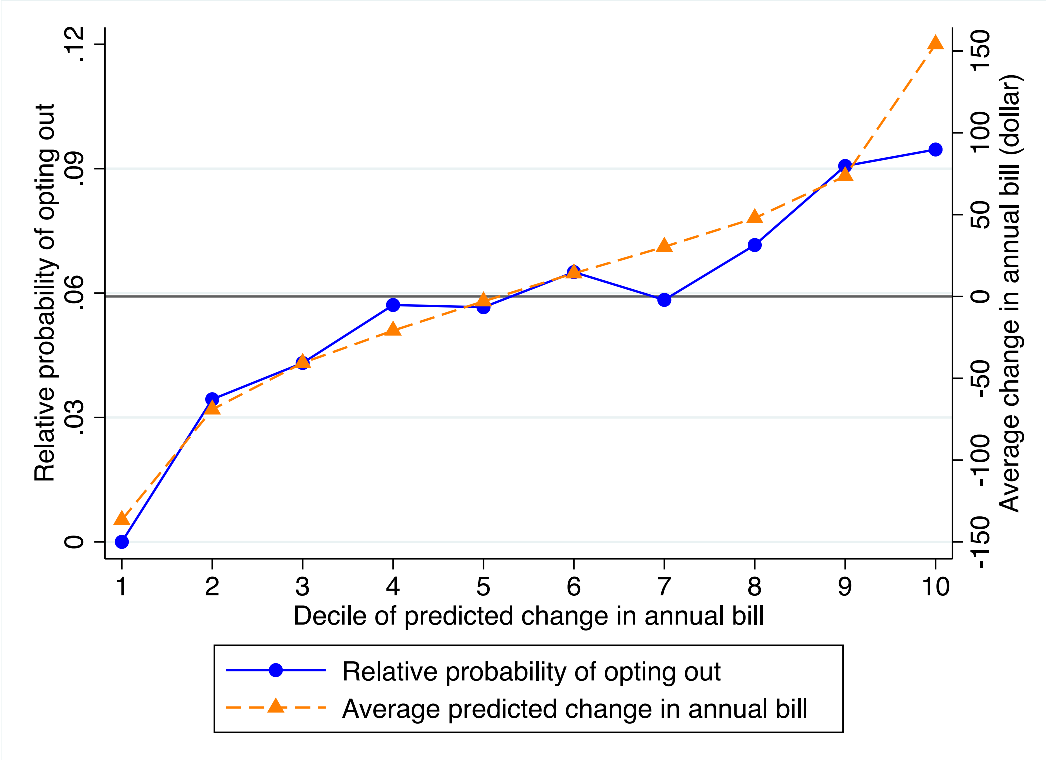

Notes: The consumption-based billing was introduced with an option to opt-out, and about 30% of customers in our data opted out. In this figure, we examine if the opt-out decision was related to the predicted change in annual payment at each customer's average heating usage in the pre-reform period. The blue solid line shows opt-out rates relative to the 1st decile of predicted change in annual billing. The orange dashed line shows the average predicted change in the annual bill. This suggests that selection was positively related to the expected gain from the policy, which is consistent with selection on the level in **?** and **?**.

Figure 6: Welfare Implications

Panel A: Standard theory

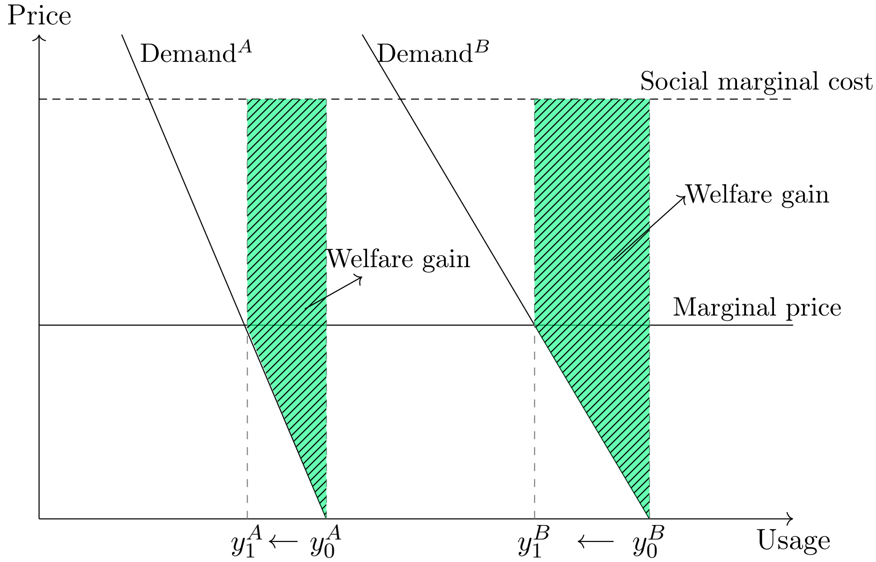

Panel B: Schmeduling

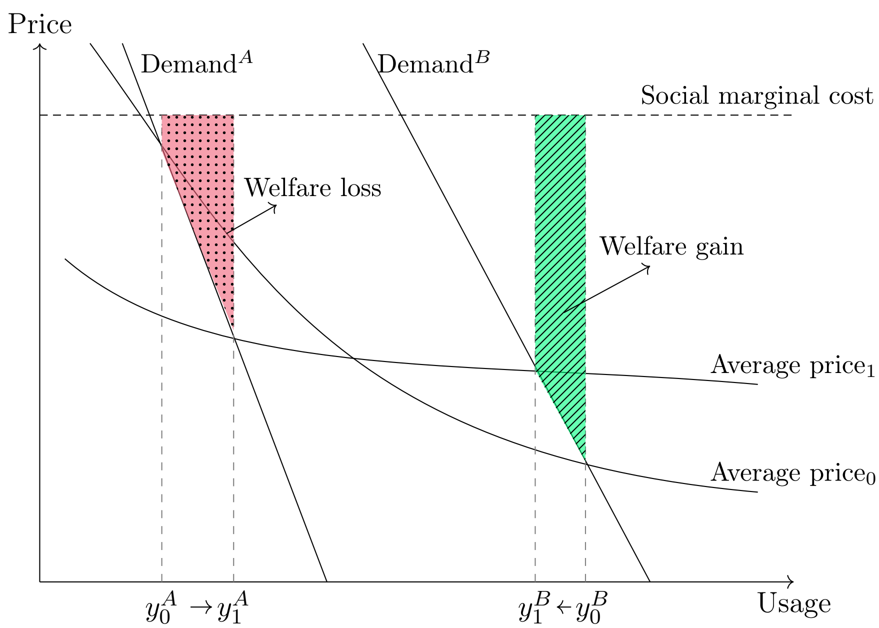

Notes: These figures show the social welfare gains from the CBB based on two alternative assumptions regarding consumer behavior. Panel A considers the standard theory of consumer behavior on a two-part tariff. Panel B considers the schmeduling model, in which consumers do not distinguish fixed cost from variable cost, and therefore respond to average price. These figures suggest that the social welfare gains depend on whether consumers respond to the two-part tariff as predicted by the standard model or the schmeduling model.

Tables

Table 1: Summary Statistics

Variable	Mean (Standard Deviation)
Average daily heating usage (kWh)	95.27
	(49.87)
Total heating usage per heating season (kWh)	11600
	(6261)
Heating bill per heating season before CBB (dollar)	413.6
	(172.0)
Heating bill per heating season after CBB (dollar)	370.6
	(146.0)
Square meter of the residence	105.3
	(41.01)
Take-up rate of the CBB policy	0.70
	(0.46)
Number of households	16,425
Observations (household-by-month)	278,041

Table 2: Impacts of the Consumption-Based Billing on Heating Usage

Dependent variable: Log of daily heating usage		
	ITT	ATET
First year of CBB	-0.107	-0.151
	(0.014)	(0.024)
Second year of CBB	-0.113	-0.163
	(0.031)	(0.035)
Third year of CBB	-0.091	-0.134
	(0.030)	(0.043)
Observations	278,041	278,041

Notes: This table shows the estimation results of equation (1). The estimation includes household fixed effects and year-by-month fixed effects. Standard errors in parentheses are clustered at the building level.

Table 3: Impacts of the CBB by the Quartiles of the Predicted Changes in Average Price

Dependent variable: Log of daily heating usage

	ITT			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
CBB	0.216	0.019	-0.159	-0.154
	(0.060)	(0.052)	(0.060)	(0.028)
Observations	61,495	70,262	52,715	61,473
Change in Marginal Price	0.014	0.014	0.014	0.014
ITT on ln(Average Price)	-0.221	0.044	0.190	0.270
ITT on ln(Predicted Average Price)	-0.116	0.055	0.057	0.230

Notes: We divide customers by quartile based on their policy-induced changes in average price and estimate equation (1) for each quartile group separately. The estimation includes household fixed effects and year-by-month fixed effects. Standard errors in parentheses are clustered at the building level.

Table 4: Testing for Schmeduling

	ITT				H_0: $\beta \leq 0$
	Marginal price	ln(Predicted average price)	ln(Actual average price)	ln(Usage)	(p-value)
Full sample	0.014	-0.116	-0.221	0.216	0.0002
		(0.005)	(0.047)	(0.060)	
Households with home value > median	0.014	-0.117	-0.166	0.170	0.038
		(0.009)	(0.065)	(0.096)	
Households with home value $\leq$ median	0.014	-0.113	-0.182	0.153	0.007
		(0.008)	(0.054)	(0.063)	
Households with home size > median	0.014	-0.115	-0.155	0.175	0.07
		(0.009)	(0.072)	(0.120)	
Households with home size $\leq$ median	0.014	-0.117	-0.229	0.194	0.002
		(0.008)	(0.056)	(0.068)	

Notes: As described in the text, we use households who are in the first quartile of the policy-induced change in average price to provide a test for schmeduling. In this table, we report the ITT estimates for this group. Standard errors in parentheses are clustered at the building level.

Table 5: Welfare Implications

	Welfare gain per household (USD/year)		Welfare gain for Tianjin (million USD/year)	
	Standard	Schmeduling	Standard	Schmeduling
Total social welfare gain	25.5	6.3	109.3	26.8

Notes: This table shows the social welfare gain from the consumption-based billing based on two different assumptions regarding consumer behavior. Columns 1 and 3 consider the standard theory of consumer behavior on a two-part tariff. Columns 2 and 4 consider the schmeduling model, in which consumers do not distinguish fixed cost from variable cost, and therefore respond to average price.

Online Appendices: Not For Publication

Appendix A: Results using the OLS with Two-Way Fixed Effects

Recent developments in the econometrics literature highlight that the OLS could produce biased estimates for two-way fixed effects models such as equation (1) if treatment effects are heterogeneous across households and/or time (de Chaisemartin and D'Haultfœuille, 2020; ?). To address this problem, in our main analysis in the paper, we use a method developed by de Chaisemartin and D'Haultfœuille (2020) to estimate equation (1) so that we do not impose the assumption of homogeneity in the treatment effects.

In this appendix, we compare our estimation results based on the method developed by de Chaisemartin and D'Haultfœuille (2020), which we report in Table 3, and our estimation results based on the conventional OLS with two-way fixed effects, which we report in Table A.7 in this appendix. We also would like to note that an earlier version of our working paper also used the conventional OLS with two-way fixed effects.

We find that the results are indeed different between the two estimation methods. de Chaisemartin and D'Haultfœuille (2020) show that the conventional OLS does not produce the correct average treatment effects when the treatment effects are heterogeneous across individuals and/or time. This is because the conventional OLS produces an incorrect weighted average of treatment effects across cohorts and time. In addition, some of these wrong weights can be negative.

To be more precise, de Chaisemartin and D'Haultfœuille (2020) define cohort g and time t for a staggered difference-in-differences method. Cohort g is the group of units who share the timing of the start of treatment. t is the time period of the data. de Chaisemartin and D'Haultfœuille (2020) show that $E\left[\hat{\beta}_{OLS}\right] = E\left[\sum_{g,t} W_{g,t}\Delta_{g,t}\right]$, where $\hat{\beta}_{OLS}$ to be an estimate from the OLS with two-way fixed effects, $\Delta_{g,t}$ is the ATE for group g and time t, and $W_{g,t}$ is weights summing to one.

If $W_{g,t}$ are the relative sample size in (g, t), $E\left[\hat{\beta}_{OLS}\right]$ is equal to the ATE across (g, t). However, de Chaisemartin and D'Haultfœuille (2020) shows that $W_{g,t}$ in the OLS are not necessarily equal to the relative sample size in (g, t) when the treatment effects are heterogeneous across g and/or t. Moreover, many of $W_{g,t}$ can be negative. If many of $W_{g,t}$ are negative, $E\left[\hat{\beta}_{OLS}\right]$ could have an opposite sign of the correct ATE over (g, t).

To explore this point in our data, we use the approach developed by de Chaisemartin and D'Haultfœuille (2020) to compute the weights $W_{g,t}$ in the conventional OLS with two-way fixed effects. In our application, the cohorts (g) are defined by the staggered timings of the introduction of the CBB, and the time (t) is year-by-month.

In Figure A.4, we show $W_{g,t}$ against the "correct" cohort-by-time weights, which are the relative sample size across the cohort-by-time cells. If the OLS uses weights that are equivalent to the correct weights (i.e., the OLS weights and correct weights line up at the 45-degree line in the figure), we can obtain the correct average treatment effect using the OLS.

However, the figure shows that many weights are not on the 45-degree line. Furthermore, 46% of the weights used by the OLS are negative. These results imply that the OLS estimates could be substantially different from the ATE and even could have a wrong sign. Indeed, we find that the sign of the estimates are different between the two methods for the quartile group 1 (column 1 in each table).

de Chaisemartin and D'Haultfœuille (2020) and **?** describe that the OLS with two-way fixed effects could produce biased estimates because it effectively uses all units, including already-treated units, as control units. The estimation methods developed by de Chaisemartin and D'Haultfœuille (2020), which we use in our main analysis, address this problem. Another alternative approach is to use OLS to estimate a cohort-specific treatment effect. We divide our sample into cohorts based on the staggered timing of the introduction of the CBB. For each cohort, we use households who were not-yet-treated as a clean control group, which allows the cohort-specific OLS estimation to be a standard difference-in-differences without a staggered roll-out. To check the robustness of our main approach, we also estimate this cohort-specific OLS regression and take the weighted average of these estimates using the relative sample size as weight. We estimate that this approach produces results consistent with our main results.

Appendix B: Calculation of the Externalities from Air Pollution

To calculate the additional welfare gain from reduced environmental externalities, we need to understand how household energy conservation affects local ambient pollution. First, we examine whether pollution emissions of the heating plant correlate with household heating usage. We use emission concentration data from the CEMS (Continuous Emission Monitoring System) monitor placed at the heating plant. Figure A.5a shows that as the daily total heating usage of households in this district increases, daily average SO_2 concentration also increases. This positive association is also observed for NO_x and PM in Figure A.5b and Figure A.5c. As a major polluting source in winter, the heating plant's emissions likely affect the local ambient air quality.

Second, we estimate the correlation between household heating usage and ambient air pollution, using air pollution data from a pollution monitor located in the residential area. In Table A.8, the ambient pollution data and the total heating usage of all households are both at the daily level, and we control for weather conditions, year-by-week fixed effects, and day of week fixed effects. We find that 1 percent increase in heating consumption is associated with 0.88 percent increase in ambient PM_{10} concentrations, where the baseline PM_{10} concentrations before the reform was 131. For each of the decile groups, we combine these estimates with the ITT estimate on heating usage to calculate the amount of reductions in PM_{10} concentrations following the reform. Ito and Zhang (2020) and the average household income in Tianjin suggests that a Tianjin household's marginal willingness to pay for a reduction in PM_{10} is 1.43 dollars per ug/m^3 of PM_{10} per year. We then multiply these two estimates to measure the WTP for the policy-induced reduction in PM_{10}. We find that the marginal cost of the environmental externality is 0.0153 USD per kWh of heating usage.[19]

[19]Note that this estimate is likely to be a lower bound estimate for environmental externalities because this calculation does not include other potential environmental externalities than PM_{10} and the MWTP for reductions in PM_{10} in Ito and Zhang (2020) is a lower bound estimate for reasons described in that study.

Appendix C: Additional Tables

Table A.1: Timing of the CBB policy

	Dependent variable: Rollout year of usage-based pricing
Year of build	0.042
	(0.095)
Average condo size (square meter)	-0.005
	(0.006)
Average home price per square meter (1,000 dollars)	0.032
	(0.035)
Annual heating usage prior to CBB (1,000 kWh)	0.019
	(0.051)
Number of Buildings	429
R^2	0.98

Notes: In this table, we test if observed building characteristics are associated with the staggered rollout timings of policy implementation. The observations are at the building level. The dependent variable is the rollout year of consumption-based billing. The estimation includes the meter installation year fixed effects.

Table A.2: Robustness of the Impacts of the CBB by the Quartiles of the Predicted Changes in Average Price (non-parametric controls of cohort trends are included)

Dependent variable: Log of daily heating usage

	ITT			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
CBB	0.174	-0.030	-0.105	-0.140
	(0.050)	(0.048)	(0.031)	(0.033)
Observations	61,495	70,262	52,715	61,473

Notes: We divide customers by quartile based on their policy-induced changes in average price and estimate equation (1) for each quartile group separately. The estimation includes household fixed effects and year-by-month fixed effects. Non-parametric controls of cohort trends are included. Standard errors in parentheses are clustered at the building level.

Table A.3: Does Heating Usage Depend on Neighbors' Compliance Status?

	ln(daily heating usage)
CBB*complied	-0.100
	(0.020)
CBB*next door neighbor complied	-0.001
	(0.020)
CBB*upper or lower level neighbor complied	0.003
	(0.018)
Observations	201,540
R^2	0.59

Notes: In this table, we test if changes in heating usage are correlated with neighbors' compliance status. The regression includes household fixed effects and year-by-month fixed effects. The first coefficient implies that changes in heating usage are negatively correlated with households own compliance status, which is consistent with our main findings on the policy's treatment effects. The rest of the coefficients indicate that there is little statistical evidence that changes in heating usage are correlated with neighbors' compliance status.

Table A.4: ATET: Impacts of the CBB by the Quartiles of the Predicted Changes in Average Price

Dependent variable: Log of daily heating usage

	ATET			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
CBB	0.306	0.037	-0.248	-0.217
	(0.060)	(0.082)	(0.082)	(0.025)
Observations	61,495	70,262	52,715	61,473
Change in Marginal Price	0.014	0.014	0.014	0.014
ATET on ln(Average Price)	-0.316	0.053	0.292	0.383
ATET on ln(Predicted Average Price)	-0.166	0.079	0.082	0.329

Notes: The regression includes household fixed effects and year-by-month fixed effects. This table reports the DID estimates (ATET) of the overall reform effect by quartile group of predicted average price.

Table A.5: ATET: Testing for Schmeduling

	ATET				H_0: $\beta \leq 0$
	Marginal price	ln(Predicted average price)	ln(Actual average price)	ln(Usage)	(p-value)
Full sample	0.014	-0.166 (0.007)	-0.316 (0.052)	0.306 (0.060)	0.00005
Households with home value > median	0.014	-0.168 (0.010)	-0.252 (0.076)	0.237 (0.084)	0.002
Households with home value ≤ median	0.014	-0.162 (0.011)	-0.256 (0.072)	0.208 (0.089)	0.001
Households with home size > median	0.014	-0.164 (0.011)	-0.217 (0.061)	0.229 (0.089)	0.005
Households with home size ≤ median	0.014	-0.167 (0.013)	-0.332 (0.072)	0.284 (0.080)	0.0002

Notes: This table reports the DID estimates (ATET) of the reform effect in quartile 1 of predicted average price.

Table A.6: Middle vs. Top and bottom floors, and Corner vs. Non-corner unites

	ln(daily heating usage)		
	Panel A: Middle vs. Top and bottom floors		
	(1) Middle floors	(2) Top and bottom floors	(3) Difference P-value
CBB	-0.083 (0.025)	-0.062 (0.026)	0.28
Observations	205,432	41,100	
	Panel B: Non-corner vs. Corner units		
	(1) Non-corner units	(2) Corner units	(3) Difference P-value
CBB	-0.081 (0.027)	-0.043 (0.019)	0.12
Observations	54,971	78,746	

Notes: The regression includes household fixed effects and year-by-month fixed effects. In Panel B, we use a subsample of buildings where we can identify corner vs. non-corner units: buildings with three, four or eight households on the same floor.

Table A.7: Results Based on the Conventional OLS with Two-way Fixed Effects

	ln(daily heating usage)			
	(1) Quartile 1	(2) Quartile 2	(3) Quartile 3	(4) Quartile 4
First year of CBB	-0.056	-0.111	-0.055	-0.048
	(0.036)	(0.023)	(0.022)	(0.023)
Second year of CBB	-0.074	-0.139	-0.114	-0.052
	(0.055)	(0.036)	(0.040)	(0.047)
Third year of CBB	-0.146	-0.234	-0.229	-0.152
	(0.072)	(0.049)	(0.051)	(0.065)
Observations	61,495	70,262	52,715	61,473
R^2	0.53	0.64	0.67	0.69

Notes: This table shows the ITT estimates using the conventional OLS with two-way fixed effects. The estimation includes household fixed effects and year-by-month fixed effects. Standard errors in parentheses are clustered at the building level.

Table A.8: Household Heating Usage and Ambient Pollution

	(1) lnSO2	(2) lnPM2.5	(3) lnPM10
ln(Daily total heating usage)	1.562	1.414	0.877
	(0.434)	(0.476)	(0.414)
Observations	461	459	444
R^2	0.74	0.65	0.58
Weather controls	Y	Y	Y
Year-week FE	Y	Y	Y
Day-of-week FE	Y	Y	Y

Notes: In this table, we estimate the relationship between heating usage and ambient pollution. Ambient pollution data on the concentration of SO_2, $PM_{2.5}$ and PM_{10} are from a pollution monitor located in the district of Tianjin where this study is conducted. This pollution monitor is the only one located in the district of our study, and this district is relatively isolated from other districts of Tianjin, with a road distance of about 55 kilometers from Tianjin's city center. Weather controls include temperature, precipitation and wind speed.

Appendix D: Additional Figures

Figure A.1: Rollout of the reform

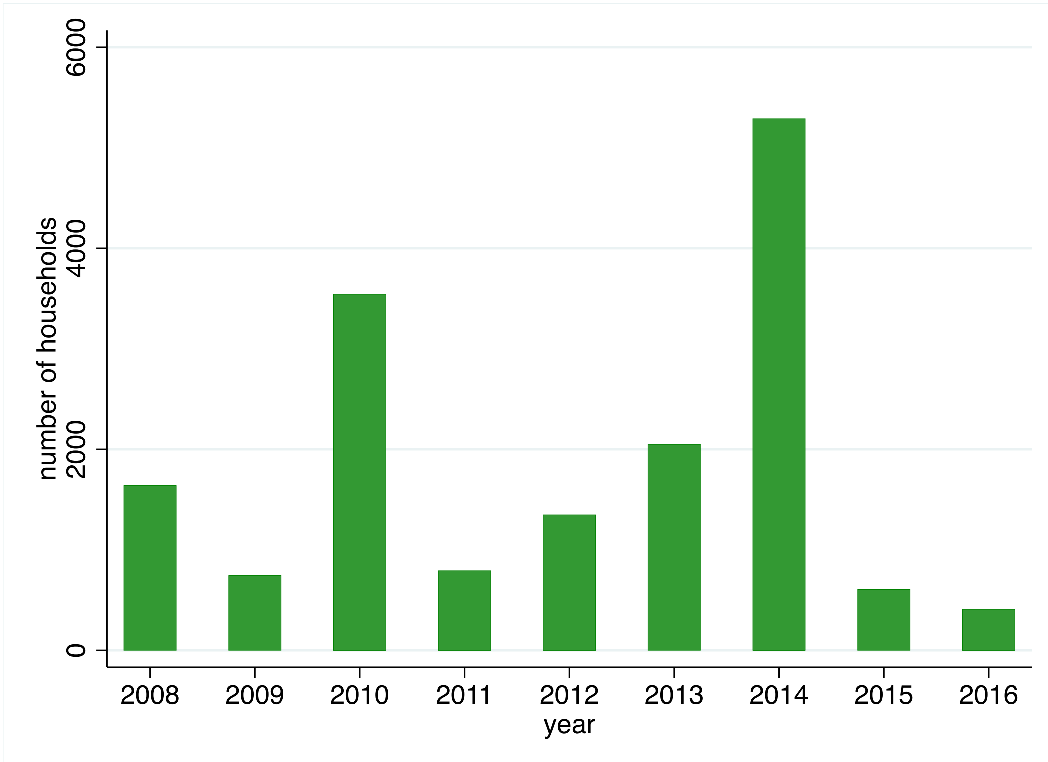

Notes: This figure shows the rollout of the consumption-based billing policy.

Figure A.2: Robustness of Policy-Induced Changes in Average Price and Usage (non-parametric controls of cohort trends are included)

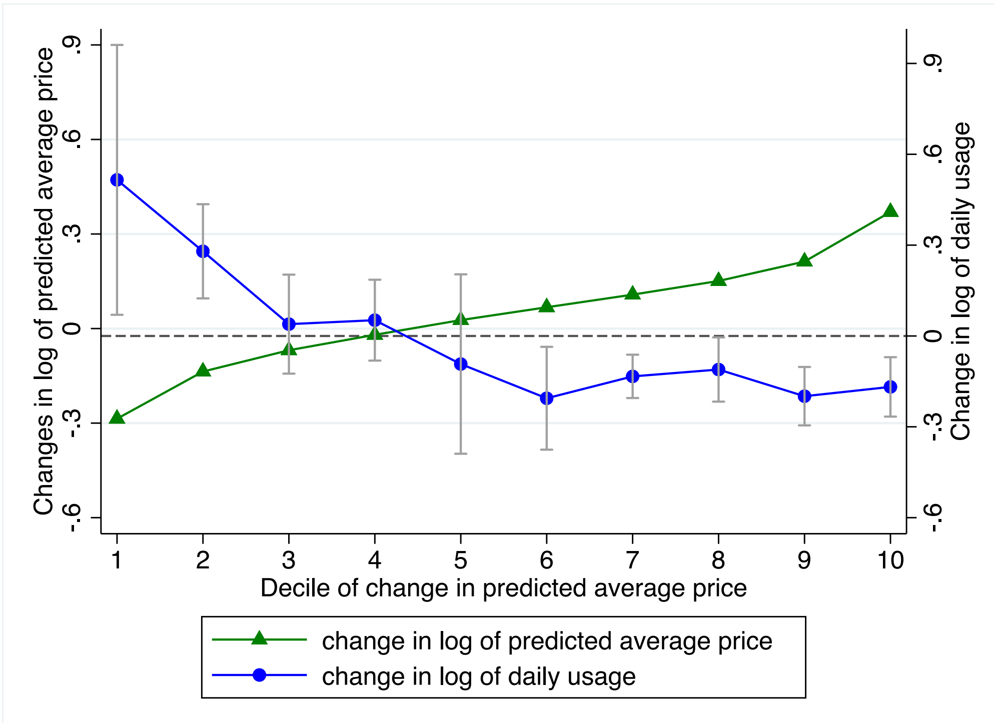

Notes: We divide customers by decile based on their policy-induced changes in average price. For each decile, we estimate the ITT of the CBB on the log of heating usage based on the difference-in-differences estimation method developed by de Chaisemartin and D'Haultfœuille (2020). We also apply the same method to estimate the ITT on the log of the policy-induced change in average price. Non-parametric controls of cohort trends are included. The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the building level.

Figure A.3: Robustness of the Event-Study by Quartile (non-parametric controls of cohort trends are included)

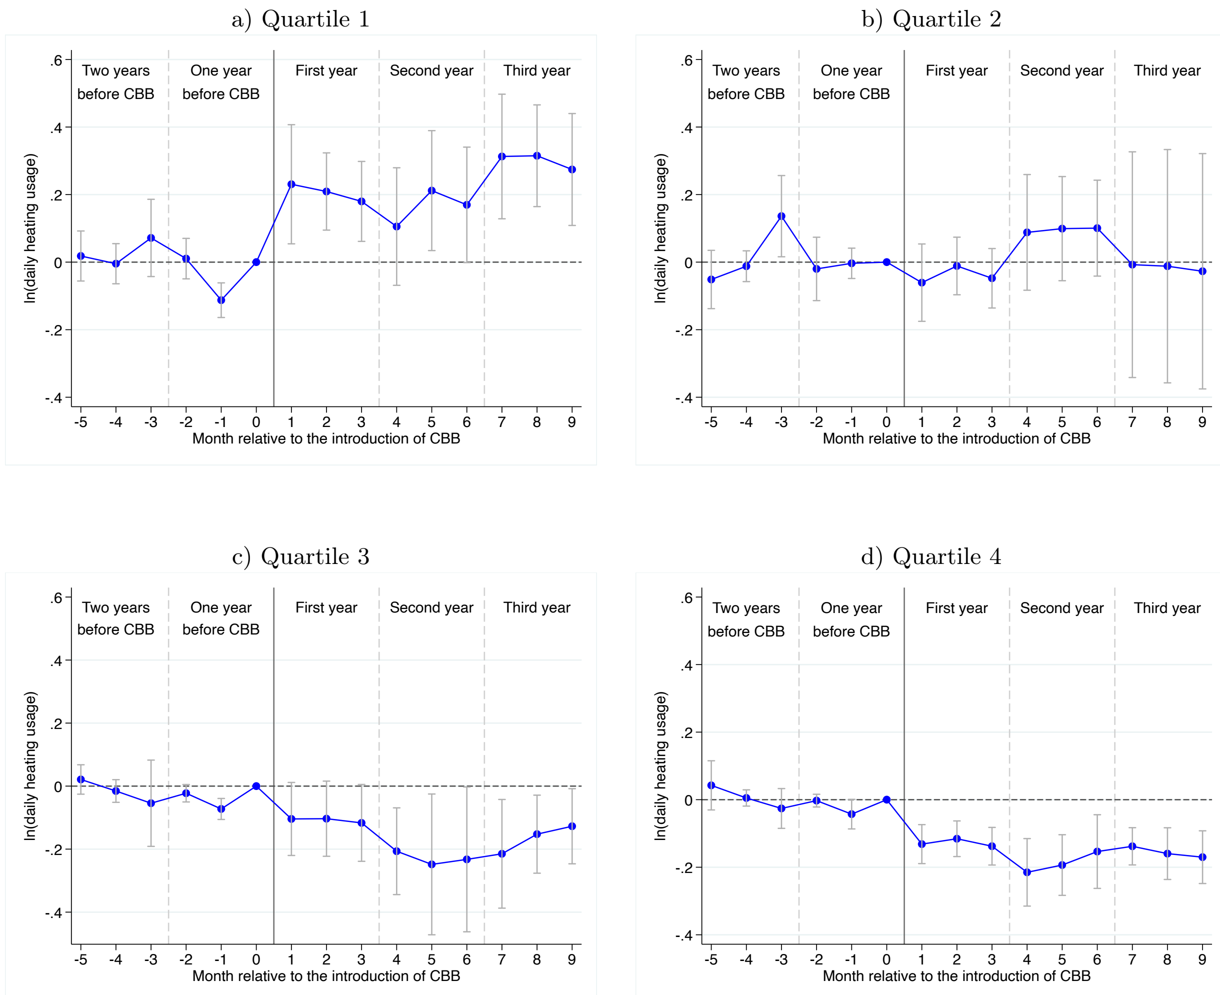

Notes: This figure shows the ITT estimates of the staggered difference-in-differences analysis described in equation (1) based on the estimation method developed by de Chaisemartin and D'Haultfœuille (2020). There are three heating months in each year because the heating season is December, January, and February. The bars indicate the 95 percent confidence intervals. Standard errors are clustered at the building level. We divide customers by quartile based on their policy-induced changes in average price. We then estimate equation (1) for each quartile group separately to make these event study figures. In this analysis, we interact time fixed effects with cohorts.

Figure A.4: Cohort-by-Time Weights Imposed by the OLS with Two-Way Fixed Effects

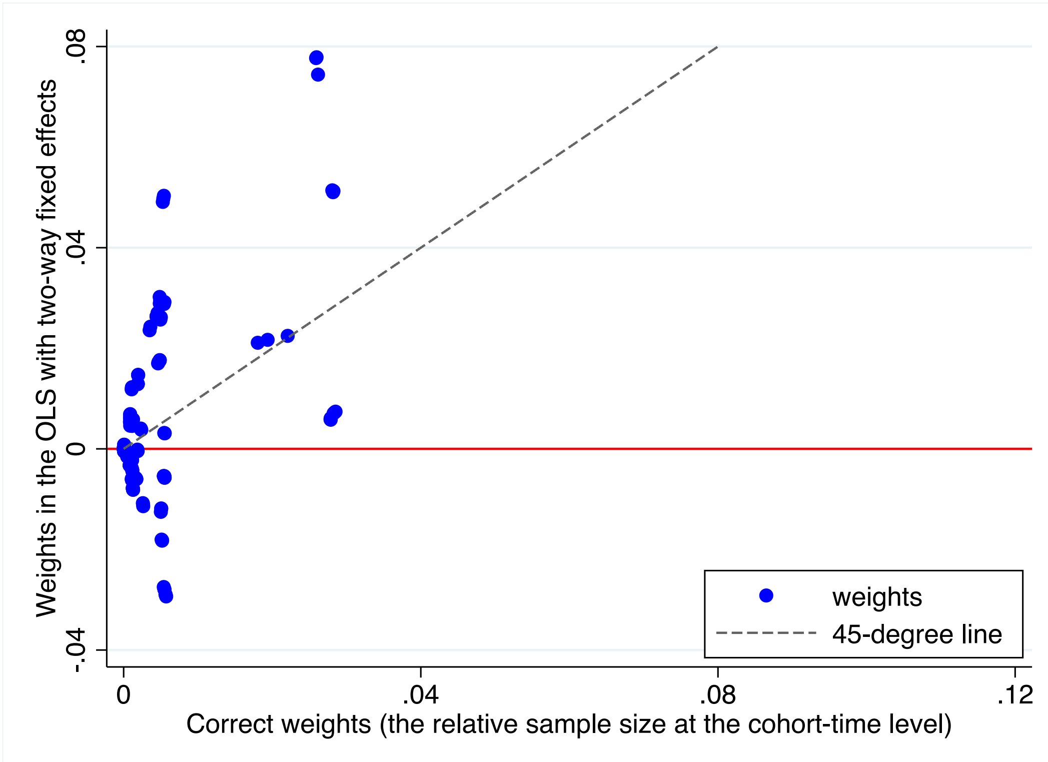

Notes: We compute the cohort-by-time weights that are imposed by the conventional OLS using the method developed by de Chaisemartin and D'Haultfœuille (2020). We then plot them against the "correct" cohort-by-time weights, which are the relative sample size across the cohort-by-time cells. If the OLS uses weights that are equivalent to the correct weights (i.e., the OLS weights and correct weights line up at the 45-degree line in the figure), we can obtain the correct average treatment effect using the OLS. However, the figure shows that many weights are not on the 45-degree line. Furthermore, 46% of the weights imposed by the OLS are negative.

Figure A.5: Household Heating Usage and Pollution Emissions of the Heating Plant

(a) SO_2 concentration (mg/m^3)

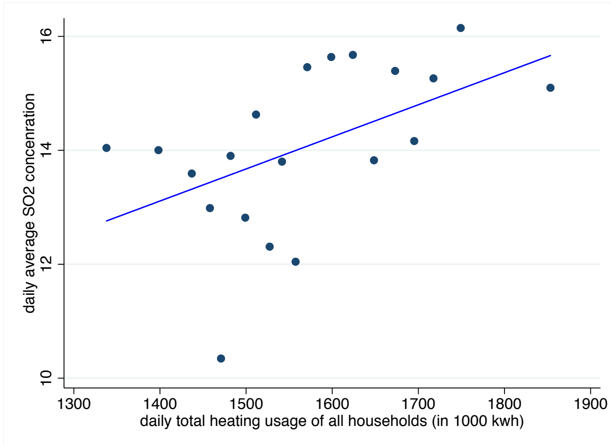

(b) NO_x concentration (mg/m^3)

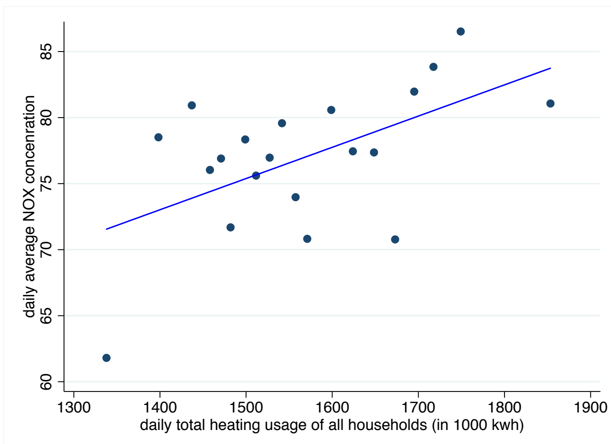

(c) PM concentration (mg/m^3)

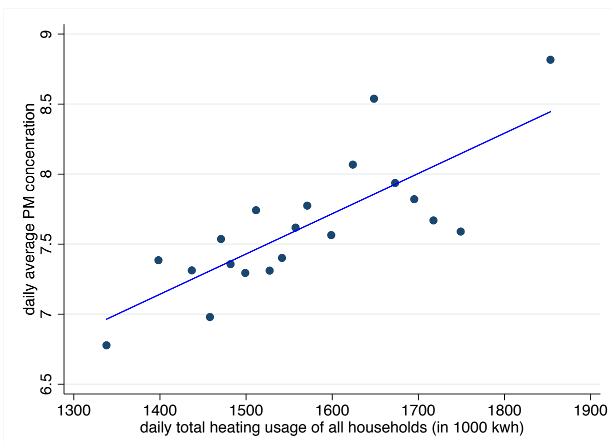

Notes: Data on emission concentrations are from the CEMS monitor placed at the heating plant.