

Surcharges versus Rebates: Evaluating a Municipal Electricity Conservation Program

Brandon Schaufele

bschaufele@ivey.uwo.ca

June 22, 2018

Abstract

Using monthly account level data for over 27,000 households between 2007 and 2014, this study evaluates a revenue-neutral municipal electricity conservation program. Rebates for the purchase of energy efficient appliances were financed via a small surcharge on high consuming households. The results demonstrate that the program mainly transferred money between residents with almost no effect on electricity consumption. Using variation in the timing of the rebate checks, none of the energy efficiency incentives yielded a statistically or economically meaningful reduction in electricity consumption compared with a counterfactual where no rebate was offered. Using a bunching estimator and exploiting changes in behavior around the high consumption threshold, a small reduction in electricity consumption is attributable to the surcharge, suggesting that prices are better than subsidies at reducing electricity consumption. Overall, the change in behavior attributable to the electricity conservation program is small, supporting recent evidence that many energy efficiency programs underperform in real-world settings.

Keywords: Electricity demand, energy conservation, feebate

JEL codes: H23, Q48, Q58

Introduction

Energy efficiency programs are popular with utilities and governments across North America. Between 1994 and 2012, the US spent more than \$34 billion on energy conservation (Boomhower and Davis, 2014), including \$17 billion allocated in the 2009 Recovery and Reinvestment Act (Allcott and Greenstone, 2017). The experience is similar in Canada with the federal government committed to “ramping up its effort to encourage building owners to invest in energy retrofits” (McCarthy, 2017). Growing evidence suggests that these programs deliver substantially less savings than initially promised. Yet, most research focuses on large scale, tax-financed initiatives launched at the national or state level. Fewer studies investigate targeted city-specific programs. It is plausible that local officials are better able to tailor critical parameters to specific local or regional characteristics, enabling them to achieve improved outcomes. Likewise, self-financed initiatives provide two instruments to achieve conservation – for example, when rebates for the purchase of energy efficient appliances are funded via surcharges on households who have high electricity consumption, both the subsidy and higher price can reduce energy consumption. This study uses monthly account-level data for over ten years to study the savings generated by a municipal electricity conservation program in Canada. The results contribute to and accord with the growing literature on incentivizing investments in energy efficiency but are among the handful focused on small-scale program design.

Despite the popularity of energy efficiency programs, a puzzle known as the “energy efficiency gap” has been identified (Allcott and Greenstone, 2012). The energy efficiency gap states that we should observe substantially more investment in energy efficiency than we actually do. The basis for this claim rests on the difference between the projected cost savings from energy efficient investments and the observed investments in the market: households should be more willing to invest in energy efficiency than they are. The basic story is as follows. Households should invest in energy efficiency, for example, by purchasing more efficient appliances. Of course, these investments are costly for the household – high energy efficiency appliances cost more than low efficiency alternatives – yet the premium paid for energy efficiency purportedly reduces lifetime energy consumption by an amount that is greater than the initial outlay (i.e., price differential). Because total household energy consumption declines, utility bills decline and private investments in energy efficiency pay for themselves. But this behavioral response is typically not observed. Households appear to underinvest in energy efficiency and so the energy efficiency gap refers to this underinvestment.

The missing investment in energy efficiency has implications for social welfare. Generation of electricity produces environmental externalities such as emissions of carbon dioxide equivalent (CO₂e) and other local pollutants. Climate change and local air pollution have real economic costs that are borne by citizens and governments. As total energy consumption declines, environmental quality and human health improves. Stated differently, as household energy efficiency improves, less total energy is needed. So, when household’s fail to invest in energy efficiency, not only do they forego the private benefits, but the social benefits also fail to materialize.

These unrealized social benefits have prompted governments to intervene in the energy efficiency market in an attempt to promote greater investment. These initiatives take many forms.

The program studied in this paper is known as Hat Smart and was created by the City of Medicine Hat, a municipality of approximately 60,000 people, located in the Canadian province of Alberta. Originally launched in 2008, Hat Smart was viewed as among Canada's "most successful municipally offered program of its kind" (Row and Welk, 2011, pg.7) and, as of December 2015, had distributed over \$4 million in energy efficiency incentives via 14,000 rebates aimed primarily at reducing electricity and natural gas consumption (Hat Smart, 2017).^{1,2} Several features make Medicine Hat a unique context to study. First, the local utility is wholly owned and managed by the municipal government. This is atypical in Canada, especially in Alberta, a province with a market focused electricity sector. Second, the climatic conditions of Medicine Hat are uncommon in Canada. Medicine Hat is among the hottest and driest cities in the country and air conditioning is pervasive in summer, while forced-air natural gas furnaces are almost the exclusive source of heating in winter. This means that there is a large peak in summer electricity demand. Finally, Hat Smart was a revenue neutral program – it is a feebate (Rivers and Schaufele, 2017; Grant, 2017). Subsidies for energy efficiency rebates were entirely financed through a surcharge levied on high consumers of electricity.

Most analyses of energy efficiency programs use engineering estimates to calculate energy savings. Projected energy savings are derived from simulation models or tests run in laboratory settings. Unfortunately, engineering estimates frequently fail to account for important real-world features. Technologies may be installed incorrectly and households' behaviour often changes as a result of incentives (Fowlie et al., 2015). A common example of an unintended behavioural change is colloquially known as the "beer fridge problem": offering a rebate for energy efficient refrigerators often increases, rather than reduces, electricity usage because households continue to operate their old unit – i.e., households buy a new primary refrigerator, but keep their old unit as a secondary, "beer-fridge" (thus, the net effect is simply adding a new refrigerator to the grid). Upgrading is another means through which behaviour and incentives interact. Consumers may purchase larger or feature-enhanced appliances because the incentive makes these cheaper to acquire. Appropriately evaluating programs such as Hat Smart requires measuring combined technological plus behaviour changes.

Until recently, surprisingly little was known about the actual effectiveness of utility-based conservation programs in the real-world (Allcott and Greenstone, 2012). Research has emerged over the past decade suggesting that it is challenging to obtain many of the promised benefits of energy efficiency. Fowlie et al. (2015), for instance, evaluate a large weatherization incentive program in Michigan. They find that engineering models over-estimate actual energy savings by more than 2.5 times and that these over-estimates cannot be attributed to rebound effects or upgrading. Davis et al. (2014) look at appliances. They evaluate a large-scale appliance replacement program that helped 1.5 million Mexican households purchase new energy efficient refrigerators and air conditioners (colloquially, referred to as "Cash for Coolers"). Using household electricity billing records, similar to those used in this study, Davis et al. find replacing a household's refrigerator reduced electricity consumption by 11 kWh per month. In contrast, the air conditioner incentives led to an increase in electricity consumption of 6 kWh per month, with

¹ This \$4 million includes funds from both the Province of Alberta as well as incentives allocated to natural gas efficiency.

² In addition to the electricity policy, Hat Smart had a separate program for natural gas conservation. Only the electricity portion of the program is evaluated in this paper.

even larger increases during the summer (up to 20 kWh).³ Moreover, they explicitly state that their estimates are “considerably less than what was predicted ex ante by the World Bank and McKinsey based on engineering models that ignore behavioral responses. The World Bank study, for example, predicted savings for refrigerators that were about four times larger” (p. 208). Examining the same Cash for Coolers program Boomhower and Davis (2014) find that between 69 and 84% of Mexican households were inframarginal, meaning that they would have purchased a new, energy efficient fridge even without the subsidy. The subsidy was, in other words, unnecessary to achieve improved energy conservation. Rivers and Shiell (2016) provide one of the few studies of a Canadian energy efficiency program. Studying incentives to replace forced-air natural gas furnaces between 2007 and 2011, they find that more than 70% of replacements would have occurred without any subsidy or tax credit and that middle and high-income households were more likely to receive benefits compared with lower income families. Finally, following the financial crisis, the US Government helped state governments subsidize households’ purchases of energy efficient appliances through the Energy Efficient Appliance Rebate Program. Houde and Aldy (2017) evaluate this program and demonstrate that approximately 90% of consumers who claimed a rebate did not contribute to an improvement in energy efficiency. New refrigerator, clothes washer and dishwasher purchases led to an expected improvement in energy efficiency of 2 kWh per year at most. Rebates mainly contributed to appliance upgrading, where households purchased a larger appliance or one with additional features.⁴

The emerging consensus on the efficacy of energy conservation programs appears pessimistic. As stated, a unique feature of Hat Smart is that it is revenue neutral: all funds allocated towards energy efficiency were collected from a small per kwh surcharge on high electricity demanders. Indeed, it turns out that this surcharge generated energy savings whereas the rebates did not. Nonetheless, the response to the surcharge is best characterized as trifling: in a city of 60,000, roughly 536MWh were conserved over 9 years. The reason that the surcharge generated more electricity conservation is almost entirely because there are virtually no statistically measurable conservation benefits from the rebates. While several point estimates suggest minor energy savings, the confidence intervals are wide. Higher prices appear to be a more effective conservation

³ Several features differentiate Mexico’s Cash for Coolers program from Hat Smart. First, it was a nationwide program, which meant that fixed administrative costs could be spread over a large number of participants. Second, sellers needed to verify that the existing appliances met certain requirements. In order to qualify for rebates, for example, the old refrigerator or air conditioner needed to be operational and at least 10 years old. Further, the retailer needed to remove the old appliance at the time of replacement (old appliances were permanently destroyed). Size restrictions were also imposed and households could only redeem one rebate – i.e., for either a fridge or an air conditioner. Nonetheless, despite these restrictions, Davis et al. emphasize that “increases in appliance size and appliance features (e.g., through-the-door ice) worked to substantially offset the potential reductions in electricity consumption” (p. 208).

⁴ Often these larger fridges, dishwashers or clothes washers had a better efficiency rating per unit of appliance services (e.g., per cubic meter of fridge space), but actually required more total electricity when compared with the counterfactual purchase (i.e., the most likely appliance that would have been purchased if there was no subsidy).

instrument than rebates, even in markets, such as electricity, where demand is extremely inelastic. Ultimately, the results in this study show that the revenue neutral Hat Smart served to transfer money between households without generating any consequential costs or benefits.

Methodology

This section presents an overview of the methodology used to evaluate Hat Smart. It is separated into three parts. First, a description of the program and data are discussed. Next, the method used to infer benefits of the rebate payments is discussed. Finally, the measurement of costs and benefits of the surcharge is reviewed. It is important to emphasize that I am seeking to measure the effect of the surcharge and incentive payments on electricity consumption behaviour. As in Aldy and Houde (2017), I do not quantify the welfare from new or upgraded appliances.

Program Structure and Data

Hat Smart was launched in 2008. The initial program was designed in conjunction with a similar scheme offered by the Canadian province of Alberta. In fact, the first wave of rebate recipients obtained funding from both the city and province. With only minor tweaks, the basic structure of Hat Smart has remained constant over the seven years studied in this research.

Hat Smart is a revenue neutral energy efficiency program. It offers rebates to rate payers for the purchase of a pre-defined set of efficiency investments. Specifically, it helps households “to make better choices regarding upgrades to their homes” (Hat Smart, 2017). Predominantly, this involves rebating a fixed amount of the purchase of new air conditioners, refrigerators, dishwashers and clothes washers.⁵ These rebates were financed via an “Environmental Efficiency Charge” (ECC). The ECC is a per kilowatt-hour (kWh) surcharge levied on billable electricity consumption above a 950kWh threshold. That is, if an account holder consumed, say, 1100kWh within a billing period, they would pay the monthly rate for the first 950kWh of consumption and then the monthly rate plus the ECC on the remaining 150kWh. The ECC did not vary during the sample period, equalling \$0.0074/kWh throughout.

Several comments on the rebates are needed. First, the funds collected from the ECC were placed into a pool and paid out according to a fixed budget. Once the annual rebate budget was exhausted, residents could no longer claim any money; thus, there was an advantage to trying to obtain a rebate early in the calendar year. Second, residents were not required to verify that they either disposed of their old energy inefficient appliance or purchased a model with enhanced efficiency. Rebates were given as long as the newly purchased model had an Energy Star rating. Third, rebates were promptly paid, usually within the month. Fourth, the city advertised the rebate scheme in both household electricity bills and in the local newspaper, so residents were largely aware of the plan. Finally, not all rebates were available in all years. For example, incentives for efficient clothes washer were available during the initial phase of Hat Smart but not in subsequent years.

⁵ A separate set of rebates were offered in natural gas conservation, but only the electricity component is considered in this paper.

The account-level data used in the study cover every household in the city from 2007 through 2014.⁶ Essentially, information was provided on the billed electricity consumption for all addresses in the city. All residential accounts pay the same per kWh rate in each month, with the exception of the ECC, so there is no cross-sectional variation in prices. Rates do vary intertemporally however. Households are billed ten times a year; so while the billing cycle does not precisely correspond to months, the period of observation will be referred to as a month for convenience. During this period all households in the city were also converted from analogue metering to digital metering. These conversions occurred over several years and it is unknown when a specific household switched. This conversion has implications for the analysis, as prior to the digital meters, meter-readings were completely twice a year and monthly bills were based on estimated electricity consumption in the given month. Information is not available for the month in which the meters were read.

Table 1 provides several summary statistics. During any given month, there are roughly 27,000 accounts billed by the city. The sample used in the regression analysis varies, but there are over 2.2M observations in the data. The average monthly consumption equals 663.24kWh, and after trimming the top and bottom one percent, had a minimum of 36 kWh and a maximum of 2,216 kWh. The ECC surcharge was paid by 20% of households in any given month. Four types of rebates are examined. The table shows the conditional summary statistics (i.e., conditional on receiving a rebate). An average rebate of \$198 was given for air conditioners, of which the vast majority of cheques were for \$200. Only a small set of households received \$50 rebates for the purchase of a window air conditioner unit. All recipients of dish washer cheques received an identical \$100. There is no variation in this amount. Like with air conditioners, most recipients of refrigerator cheques received \$200, with a small group getting \$100. Thus, the mean refrigerator subsidy equals \$198. The most variation in rebates is for clothes washers as this program coincided with the provincial program. The average clothes washer rebate is \$178, with a minimum of \$75 and a maximum of \$775.

Table 1: Summary Statistics

	Mean	Std. Dev	Min.	Max.
Electricity consumption (kWh/month)	663.24	392.87	36.00	2216.00
Share of households paying ECC ^a	0.20	0.40	0.18	0.24
Rebates (\$)				
Air conditioners	198.01	17.16	50	200
Dishwashers	100	0	100	100
Refrigerators	198.25	13.11	100	200
Clothes washers	178.18	26.5	75	775

a - minimum and maximum refer to monthly values

⁶ All data were provided under a strict confidentiality agreement with the City of Medicine Hat.

Measuring the Benefits of Hat Smart Rebates

Conceptually, the benefits of conservation programs are easy to understand. The objective of Hat Smart is to reduce electricity consumption. Private benefits are therefore the dollar-valued amount of energy conserved. However, this only captures one part of the benefit calculation. Market failures are pervasive in electricity generation and therefore may justify public intervention into energy conservation. A wide range of market failures have been highlighted (Fowlie et al., 2015). Examples include imperfect information (e.g., consumers are unaware of the benefits of energy efficiency), capital market failures (e.g., consumers cannot obtain financing for profitable investments in efficiency), split incentive problems (Papineau, 2017) (e.g., the individual paying the utility bill may be different than the individual consuming energy) as well as a series of behavioural economic explanations such as myopia and inattentiveness (Allcott and Greenstone, 2017). Market failures also entail that the public or social benefit of energy efficiency, from, say, reduced CO₂e emissions, does not factor into private decisions to spend on more efficient clothes washers.

Reduced emissions and the associated environmental and health improvements imply that programs such as Hat Smart really have two benefits that must be quantified. The first is the private savings from lower electricity bills. Private benefits are calculated from a reduced form regression as the amount of energy saved multiplied by the rate per kWh multiplied by the number of rebate receiving households. The second benefit arises from the reduction in harmful emissions. This includes CO₂e abated and lower ambient concentrations of local pollutants. Medicine Hat has virtually none of the air quality issues that are prevalent in larger urban centres. As a result, the social benefits of energy efficiency can be limited to tonnes of CO₂e abated.

Both the private and social benefits are due to changes in energy consumption. These benefits are therefore directly measured as the incentive-induced reduction in electricity consumed. This is estimated via:

$$y_{it} = \alpha \cdot \text{Hat Smart Incentive}_{it} + \gamma_i + \tau_t + \varepsilon_{it}$$

where y_{it} is energy consumption by household i in period t . Energy consumption is measured as kWh of electricity per month. This represents the energy for which a household is billed in a given month. The number of households, i , included in any specific econometric model changes based on the source of identifying variation. In the broadest model, the sample includes all households in Medicine Hat. Hat Smart is a voluntary program however. This means that households self-select into it. In more restricted specifications, therefore, the sample is limited to only those households that received an incentive for a particular category of purchase (e.g., refrigerators). The rationale underlying the different samples is that selection bias poses a problem if those households that received a rebate for, say, a new dishwasher are fundamentally different than the control group (i.e., those that did not obtain a rebate). If they are fundamentally different, it may be the case that the parameter of interest, α , will over- or under-estimate the true effect of Hat Smart. γ_i is an address fixed effect. Including γ_i captures a wide range of variables, such as a house's square-footage and location, that are time invariant but fundamentally unobservable. γ_i alleviates many concerns over potential omitted variable bias. Time is measured as months-of-sample and common time-specific shocks such as weather are captured by τ_t , the

time fixed effect. ε_{it} is the error term that captures everything that varies at the household-by-time level.

The main source of identification exploits differences in the timing of rebate cheques conditional on the time and address fixed effects. For example, one household may have received their rebate in January of 2012, while another obtained theirs in May of that year. An unbiased evaluation of Hat Smart requires that, conditional on address and time fixed effects, ε_{it} is uncorrelated with incentive payments. This assumption is reasonable especially for the restricted samples that exploit variation in timing of Hat Smart cheques paid for the identical types of investment (e.g., rebates on insulation).

α is the coefficient of interest, which represents the change in energy consumption per \$100 of incentive.

Measuring the Economic Benefits and Costs of Hat Smart's Surcharge

A common misperception is that whatever money is paid to households via programs such as Hat Smart is a cost of the program. This is incorrect as transfers are not costs. Given Hat Smart's financing structure, only features that introduce distortions in decision-making are costs. As Hat Smart is completely funded via a surcharge on high consumption households, economic costs only arise from the deadweight loss due to reduced demand for electricity. As with rebates, reduced consumption generates social benefits. Thus, both the deadweight loss of the ECC as well as the social benefits from reduced electricity generation must be measured.

Deadweight loss from surcharge

Figure 2 illustrates the economic costs from Hat Smart using the standard supply and demand graph. The downward sloping blue curve is the demand curve. This represents a household's demand for electricity. The red curves then are the within month supply functions for this household. A household's supply function depends on their total monthly consumption and the threshold at which the ECC kicks in. If a household consumes less than 950kWh/month. The standard constant rate supply curve applies to all consumption. For those households that exceed 950kWh per month, the supply curve jumps to $Supply^{ECC}$ for all additional consumption. After the threshold, these high energy consuming households must pay the additional ECC fee. The blue triangle represents the extent to which households change their behavior – reduce demand – because of the higher price for electricity. Without the fee, they would consume Q^* . With the fee, they consume Q^{ECC} . The triangle is the deadweight loss due to the energy conservation surcharge; it is the economic cost of Hat Smart. Of course, this triangle only exists for consumption in excess of the ECC threshold.

The size of this triangle depends on the responsiveness of demand with respect to the ECC. This is encapsulated in the price elasticity of demand.




Figure 1: Deadweight Loss due to ECC

The deadweight loss (DWL), or costs, of Hat Smart, in single market’s price-quantity space, is calculated as:

$$DWL = \frac{1}{2} \eta_Q p Q \left(\frac{ECC}{p} \right)^2$$

where p is the per kWh price, and Q is demand above the 950 kWh/month threshold. η_Q is the elasticity of demand. This deadweight loss calculation assumes that the marginal utility of income is constant, a reasonable assumption given the budget share of electricity.

The deadweight loss formula hinges on the elasticity of demand and measures the gross economic costs from Hat Smart – i.e., the costs without factoring in the social value from the reduction in electricity consumption. The elasticity of demand, whose estimation is discussed next, also summarizes the reduction in electricity consumption caused by the ECC. Fewer kWh consumed implies fewer tonnes of CO₂e emitted. Thus the elasticity of demand permits the calculation of social benefits too. The net economic costs of Hat Smart subtract social benefits from the surcharge’s deadweight loss.

Estimating the Elasticity of Electricity Demand

Two empirical methodologies are applied to estimate the elasticity of demand for electricity. First, models similar to the ones estimated to evaluate the benefits of Hat Smart are formulated.

An important difference from those models is that the regressions required to estimate the elasticity of demand have no cross-sectional variation in prices across households. This means that a household whose average consumption is, say, 700 kWh per month pays an identical per kWh price as another household with consumption of 700 kWh per month. When evaluating the benefit side of the equation, it is possible for otherwise similar households to receive rebate cheques at different points in time and this idiosyncratic variation allows for clear identification of the parameter of interest. Restricting the analysis to time series variation limits the reliability in the elasticity estimates as it is possible for time-varying shocks that are correlated with price (e.g., an economic downturn) to bias the coefficients.

As a result, an alternative “bunching estimator” is used to infer the elasticity of demand in the cross-section of the immediate vicinity around the ECC threshold. Bunching estimators are applied in Sallee and Slemrod (2012), Bastani and Selin (2014) and Kleven (2016). The idea is that the discrete jump in prices at 950 kWh, attributable to the ECC, can be exploited to infer elasticity of electricity demand in the region around the surcharge. Specifically, if there is excess mass on the low price side of the threshold, this mass can be interpreted as a behavioural response to the surcharge. The elasticity takes the form:

$$\hat{\varepsilon} = \frac{\hat{B}/\hat{h}_0(z^*)}{z^* \ln\left(\frac{p_1}{p_2}\right)}$$

where z^* is the price threshold, \hat{B} is the measured excess mass to the left of the surcharge threshold and $\hat{h}_0(z^*)$ is the estimated mass that we would expect to see in a counterfactual “no surcharge” scenario. p_1 is the price of electricity before the surcharge is levied and p_2 is the post-surcharge price. Calculating this elasticity requires estimating several regions of the electricity demand distribution. $\hat{h}_0(z^*)$, in particular, is the key. This counterfactual is estimated in an interval around z^* : $[z^* - \delta_b, z^* + \delta_b]$, where δ represents the width around the threshold z^* , b indexes the actual region considered and c will index the counterfactual region. The region around z^* is an area where the density of electricity demand is expected to be smooth but where there is, in fact, bunching. Start by defining three regions: \hat{h}_-^* : $[z^* - \delta_b - \delta_c, z^* - \delta_b]$, \hat{h}^* : $[z^* - \delta_b, z^* + \delta_b]$ and \hat{h}_+^* : $[z^* + \delta_b, z^* + \delta_b + \delta_c]$. It is possible to use the densities in each of these three regions to calculate the following cumulative densities: $\hat{H}_-^* = \delta_c \hat{h}_-^*$, $\hat{H}^* = 2\delta_b \hat{h}^*$ and $\hat{H}_+^* = \delta_c \hat{h}_+^*$. Given these quantities it is possible to define actual excess mass as:

$$\hat{B} = \hat{H}^* - \frac{\delta_b}{\delta_c} (\hat{H}_-^* + \hat{H}_+^*)$$

And the counterfactual mass as:

$$\hat{h}_0 = \frac{1}{2} (\hat{h}_-^* + \hat{h}_+^*)$$

With \hat{B} and \hat{h}_0 in hand it is possible to calculate, the elasticity of demand. The masses in the three regions around the threshold z^* - \hat{h}_-^* , \hat{h}^* and \hat{h}_+^* - are estimated using Epanechnikov kernels. A width of 45 kWh per month is used for δ_b .

Results

Energy Savings Attributable to Hat Smart Rebates

The change in electricity consumption for each category of rebate is shown in Table 2. For each class of incentive, three separate econometric models are estimated. These models are distinguished by the underlying source of variation that is used to statistically identify the parameter of interest. Column (1) uses all households in the city as a baseline. Column (2) restricts the sample to just those that received at least one electricity incentive at some point in time. The incentive could have been for an appliance category that is unrelated to the received rebate. The logic underlying this sample restriction is that there may be some fundamentally unobservable difference between households that received a rebate and those that did not. This unobserved difference may bias the estimates and thus needs to be adjusted for. Column (3) take this one step further. It provides the most credible econometric identification. Column (3) focuses exclusively on households that receive identical rebates, but exploits differences in the timing at which those rebates were received. The idea is that two households that received, say, an incentive to purchase a new dishwasher – but where one received her cheque in January, while the other received her rebate in June – are more similar than households who did not receive a dishwasher rebate.

All econometric specifications contain household and month-of-sample fixed effects. Throughout, all standard errors are clustered on individual addresses (i.e., at the household level). All coefficients should be interpreted as the reduction in kWh per \$100 rebate.

Regression results

Table 2 presents the results. Four panels are included, one for each air conditioners, clothes washers, refrigerators and dish washers.

Air conditioners. Rebates for air conditioners led to the largest reduction in electricity consumption, but none of the point estimates are statistically distinguishable from zero. The baseline model, column (1), shows that a \$100 incentive reduces monthly electricity usage by 6.8 kWh per month. This decreases to a statistically insignificant 5.5 kWh per month in column (2). Column (3), providing the most credible identification, shows the largest reduction in electricity consumption at 12 kWh per month or 144 kWh per year. Still the confidence interval is wide and the true value could be notably larger or smaller.

Nonetheless, relative to the other categories of rebates, air conditioners appear to yield the largest reductions in electricity use. The US Department of Energy projects the typical lifespan of an air conditioner to be 15 to 20 years (DOE, 2017). Assuming an air conditioner lasts for 15 years, the total electricity savings per \$100 incentive is approximately 2,160 kWh. At a rate per kWh of \$0.08 and with a 4% discount, the private return from this \$100 rebate for an energy efficient air conditioner is -31.9%. If the air conditioner's lifespan is 20 years, then electricity savings total 1,814 kWh and the private return equals 60.7%. This suggests that investing in energy efficient air conditioning may be privately beneficial. Including social benefits from abated CO₂e makes investing in air conditioner efficiency more attractive, but, to repeat, these estimates must be interpreted with caution due to the imprecision of the coefficients.

Table 2: Energy Savings Attributable to Hat Smart Rebates

	(1)	(2)	(3)
<i>Panel A: Air conditioner rebates</i>			
kWh per \$100 incentive	-6.797 (8.105)	-5.512 (8.171)	-12.0005 (8.675)
Number of households	27,921	3,925	228
Number of observations	2,200,266	308,069	17,295
<i>Panel B: Clothes washer rebates</i>			
kWh per \$100 incentive	-0.035 (2.104)	1.251 (2.386)	2.068 (2.520)
Number of households	27,921	3,925	2,435
Number of observations	2,200,266	308,069	199,462
<i>Panel C: Refrigerator rebates</i>			
kWh per \$100 incentive	-1.847 (3.402)	-3.699 (3.562)	-2.764 (4.813)
Number of households	27,921	3,925	833
Number of observations	2,200,266	308,069	65,579
<i>Panel D: Dishwasher rebates</i>			
kWh per \$100 incentive	0.139 (0.085)	0.187* (0.088)	0.163* (0.073)
Number of households	27,921	3,925	675
	2,200,266	308,069	53,049

Clothes washers. Panel B in Table 2 presents the results for clothes washers. Column (1) shows that a \$100 rebate decreased electricity consumption by 0.04 kWh per month. This value *increases* to 1.3 and 2.1 kWh per month in columns (2) and (3). None of the specifications have coefficients that are statistically distinguishable from zero. Further, not only do the confidence intervals include zero, the standard errors are large. The imprecision of these estimates means that it is difficult to claim that rebates on clothes washers had any effect on household electricity consumption. And while no evidence of an effect is different than finding no effect, the positive point estimate suggests that it is unlikely much energy savings was obtained via clothes washer rebates.

Washing machines have seen some of the largest efficiency gains over the past two decades. Thus, this result may seem out-of-place. It is important to re-emphasize that these models are measuring the impact of the incentives and not the effect of the underlying technologies. Washer efficiency has improved, but these regressions demonstrate that the incentives did not induce any incremental, or marginal, improvement in efficiency. Further, while it is not possible

to test explicit mechanisms with the data available, it is plausible that households engaged in upgrading behaviour. Larger and feature-enhanced models likely replaced smaller and more basic appliances. This upgrading likely offset any rebate-induced improvements in energy efficiency.

Refrigerators. As with clothes washers, Panel C shows that refrigerator incentives have no statistically significant effect on electricity consumption. Again, wide standard errors make it difficult to make definitive claims. Column (1) shows a point estimate of -1.9 kWh per month from a \$100 rebate. This increases slightly to -3.7 and -2.8 kWh per month in columns (2) and (3). While statistically indistinguishable from zero, these point estimates are larger than those found for refrigerators in Houde and Aldy (2017), but smaller than those in Davis, Fuchs and Gertler (2014). Ultimately, as with clothes washers, these models suggest that little energy savings are gained by incentivizing the purchase of energy efficient refrigerators (at least, given the existing structure of Hat Smart, where households were not required to remove their old fridges).

Dishwasher. Finally, Panel D of Table 2 displays the results from the dishwasher regressions. Column (1) where all other households in Medicine Hat act as a control group shows that a \$100 dishwasher incentive increases electricity consumption by 0.1 kWh per month. This estimate is not statistically distinguishable from zero. Restricting the sample to households that received any rebate changes the estimate to a 0.2 kWh per month in columns (2) and (3). These two models do show a statistically significant increase, but the magnitudes are trivial. Model C, for instance, suggests that a dishwasher incentive increased electricity consumption by 0.2 kWh per month. In essence, given the comparatively precise standard errors, it is safe to claim that dishwasher incentives have no meaningful effect on electricity consumption and, hence, Hat Smart produced no benefit from providing these rebates.

Costs and Benefits of Hat Smart Surcharge

Few benefits from Hat Smart rebates are identifiable in Table 2. The surcharge is investigated next. As described, the gross economic costs equal the deadweight loss attributable to the surcharge, which is a function of the elasticity of electricity demand. The elasticity estimates and deadweight loss are discussed first. Table 3 presents three estimates for the elasticity of electricity demand with respect to price. Electricity demand is normally viewed as extremely inelastic with limited response to changing prices. The gross costs are not the full costs of program however. The net costs (or benefits) of Hat Smart require adjusting for the social value of reduced CO₂e emissions. These emissions are valued using Canada's social cost of carbon, which equals \$40.70/tCO₂e and are discussed second.

Table 3 shows that, using time series variation, neither the short- nor long-run elasticities of electricity demand are statistically distinguishable from zero. In fact, the point estimates for both elasticities suggest that quantity demanded increases as prices increase. The point estimate on the short-run demand for electricity is 0.3, implying that a 1% increase in price leads to a 0.3% increase in quantity demanded. The corresponding long-run estimate is also 0.3. The estimates suggest two things. First, electricity demand may be extremely inelastic and it may not be possible to distinguish the true response from zero. In other words, the true demand response is very small (virtually a vertical line). If electricity demand is indeed perfectly inelastic, it implies

that Hat Smart effectively has no economic cost beyond its administration costs. Second, it is possible that other time-varying factors such as the state of the economy are correlated with both price and electricity demand and, as a result, the time-series coefficients are biased. Given the size of the standard errors, it is difficult to infer anything meaningful about consumer responses to electricity prices.

The bottom part of Table 3 uses the cross-sectional bunching estimator to infer the elasticity of demand. It exploits the discontinuity in electricity pricing near the threshold for the ECC, by comparing the behaviour of households slightly below and slightly above the 950 kWh per month cut-off. As Table 3 illustrates, the elasticity of electricity demand with respect to price, in the cross-section, equals -0.05. This estimate is statistically significant at the 0.1% level. This implies that households do, in fact, respond to prices by reducing their demand and that there is a cost to financing Hat Smart (as well as benefits from less electricity consumption).

Table 3: Elasticity of Electricity Demand

<i>Time series variation</i>		
Short-run elasticity	0.273	
	(0.309)	
Long-run elasticity		0.330
		(0.372)
Month fixed effects	Y	Y
Location-year fixed effects	Y	Y
Number of observations	2,200,260	2,200,260
<i>Cross-sectional variation</i>		
Elasticity	-0.052***	
	(0.001)	
Number of observations	43,194	

This bunching elasticity is used to calculate both the deadweight loss and the reduction in electricity consumption attributable to the ECC. Given the paucity of the elasticity estimate, these values are small. (Also, as the time series models did not yield statistically significant elasticities of demand, they should be used with caution and interpreted as an upper bound on the true costs.) The gross deadweight loss from the ECC surcharge equals a paltry \$1,984.56 over the entire 2008 to March 2014 period. This deadweight loss does not include any fixed or variable administrative costs involved in managing the program, but is best labeled as tiny. The main conclusion is clear: the ECC surcharge on electricity generates tiny market distortions and hence financing Hat Smart involves negligible economic costs.

The excess burden of Hat Smart represents the gross of environmental benefits cost of the program. Using the cross-sectional elasticity, the program also reduced electricity consumption by

536.37 MWh. All of Medicine Hat's electricity is generated using natural gas. Applying NRCan (2018) conversion factors, the ECC reduced CO₂e emissions by 108.1tCO₂e. At a social cost of carbon of \$40.70, this means that the ECC produced gross environmental benefits of \$4,400.98. The *net benefits* – environmental benefits less the deadweight loss – from the surcharge then equal \$2,416.42. Thus, while the surcharge did improve economic welfare, its chief outcome appears to involve transferring money between households in the city.

Conclusion and Policy Recommendations

Economists have long argued that the savings from energy conservation programs tend to be overstated (Joskow and Miron, 1992). These results support this. Few Hat Smart rebates has any statistically measurable effect on electricity consumption, but, more importantly, point estimates are not economically meaningful. The surcharge levied on high consumers did produce net benefits, but the magnitude of these is also negligible. Indeed, Hat Smart served primarily to transfer money between households in the city. While perhaps disappointing, this paper fits within a growing swath of research emerging on the economics of energy efficiency programs. Unfortunately, from an efficiency perspective, this research paints a cynical picture of energy conservation initiatives. Utility-based and government-funded energy efficiency programs simply have not delivered their promised electricity reductions. This begs the question: are the policy tweaks that support better results? Three options are discussed: targeting and verification, rebates conditional on energy conservation and higher prices.

Targeting and verification. A common recommendation for energy efficiency programs is more precise targeting and verification (e.g., Allcott and Greenstone, 2017). Targeting is easy to understand but hard to do well. Targeting is actually is catch-all term that encompasses several themes. Targeting may mean that funds are directed to low income households or towards “energy hogs” – i.e., houses with unusually high consumption for their profile with the hope that these households have greater scope for improvement per dollar incentive. Regardless of which targets are selected, targeting relies on some underlying heterogeneity in the population where rebates induce a particular set of households to invest in energy efficiency and reduce their electricity consumption.

Similar to targeting, verification may also improve program performance, especially when beer-fridge-type problems are a concern. Verification means that program administrators require evidence that old appliances are removed prior to issuing rebates. Eligibility for rebates could mimic the Mexican Cash for Coolers program, where recipients must demonstrate that they were replacing appliances that were at least 10 years old and opting for models of approximately the same size. These verification steps may mitigate energy consuming upgrading behaviour.

Despite the appeal of targeting and verification, caution is warranted before pursuing these strategies. Simply, the payoff may not materialize. Both targeting and verification introduce administration costs and can be unpopular with residents who are familiar with a “no questions asked” program. Indeed, administrators of the Hat Smart program voiced precisely this concern. Moreover, the additional energy savings from targeting and verification may be small. Fowlie et al. (2015), for example, demonstrated that a large-scale encouragement program, one targeted

at low income households, yielded only small gains but the costs of this encouragement were approximately \$1000 per household.

Payments conditional on energy conservation. Rather than directly targeting appliances or home heating investments, Hat Smart rebates could be directly tied to energy consumption. Cheques could be issued if households reduce their energy usage versus some benchmark (e.g., by 5% of previous year's electricity consumption). This design could be similar to a program devised by the Canadian province of British Columbia known as Team Power Smart. Team Power Smart is a voluntary program that offers households the opportunity to undertake annual conservation "challenges" (Fraser, 2017). Households that are able to reduce their annual, weather-adjusted electricity use by 10%, relative to the previous 12 month period, receive a payment of \$75 (Fraser, 2017).

The advantage of this style of program is that households can choose the best method to reduce energy consumption, rather than being restricted to a finite set of rebates. The activities that would potentially earn a reward could also include behavioral change. For instance, a family that actively reduces its energy consumption by, say, reducing air conditioning in summer would not currently be eligible for a Hat Smart payment. Under a redesigned scheme, this family may be able to make a large contribution to conservation goals and should become eligible for rebates.

Higher prices. Finally, the main conclusion of this research is that pricing works. If policy-makers' primary concern is improving energy conservation and reducing emissions, electricity prices could be increased substantially. Indeed, in the case studied in this paper, the City of Medicine Hat appears to have significant scope to increase the price of electricity before substantial consumer behavioural changes are undertaken. Higher prices mean that substantial additional revenue would be collected by the utility or municipality, funds could be recycled, used to offset other taxes or to fund community projects. Ultimately, the experience of Hat Smart shows that pricing appears to work while rebates disappoint.

Works Cited

- Allcott, H. and M. Greenstone, 2012. Is there an Energy Efficiency Gap? *Journal of Economic Perspectives*, 26(1): 3-28.
- Allcott, H. and M. Greenstone, 2017. Measuring the Welfare Effects of Residential Energy Efficiency Programs. NBER working paper, no. 23386.
- Bastani, S. and H. Selin, 2014. Bunching and Non-bunching at Kink Points of the Swedish Tax Schedule. *Journal of Public Economics*, 109: 36–49.
- Davis, L.W., A. Fuchs and P.J. Gertler, 2014. Cash for Coolers: Evaluating a Large-Scale Appliance Replacement Program in Mexico. *American Economic Journal: Economic Policy*, 6(4): 207-238.
- Energy Efficiency Alberta (EEA), 2017. Home Efficiency Tool. <https://www.energycanada.ca/residential/home-efficiency-tool/>.
- Fowlie, M., M. Greenstone and C. Wolfram, 2015. Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program. NBER working paper, no. 21331.
- Fraser, A., 2017. Financial Incentives or Persistent Habits? How Households Respond to Financial Rewards for Electricity Conservation. Working paper, <http://www.alastairfraser.ca/>.
- Grant, L., 2017. Carrot and stick: Using a regression kink design to identify impacts of energy fees.
- Hat Smart, 2017. About Hat Smart. <https://www.medicinehat.ca/government/departments/utility-sustainability/hat-smart/about-hat-smart> (retrieved July 2017).
- Heyes, A., N. Rivers and B. Schaufele, 2016. Pollution and Politicians: The Effect of PM on MPs. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2852886.
- Houde, S. and J.E. Aldy, 2017. Consumers' Response to State Energy Efficient Appliance Rebate Programs. *American Economic Journal: Economic Policy*, 9(4): 227-255.
- Joskow, P.L. and D.B. Marron, 1992. What Does a Negawatt Really Cost? Evidence from Utility Conservation Programs. *Energy Journal*, 13(4): 41-74.
- Kleven, H.J., 2016. Bunching. *Annual Review of Economics*, 8: 435-464.
- McCarthy, S., 2017. Ottawa Eyes Energy Retrofits in Buildings. *The Globe and Mail* (September 20, 2017).
- Natural Resources Canada (NRCan), 2017. ENERGY STAR in Canada. <http://www.nrcan.gc.ca/energy/products/energystar/18953>.
- Natural Resources Canada (NRCan), 2018. APPENDIX B – CO2 Emissions Factors. <http://www.nrcan.gc.ca/energy/efficiency/industry/technical-info/benchmarking/canadian-steel-industry/5193>.
- Papineau, M., 2017. Energy Efficiency Premiums in Unlabeled Office Buildings. *The Energy Journal*, 38(4): 195-212.

Rivers, N. and L. Shiell, 2016. Free-Riding on Energy Efficiency Subsidies: the Case of Natural Gas Furnaces in Canada. *The Energy Journal*, 37(4): 239-266.

Rivers, N. and B. Schaufele, 2017. New Vehicle Feebates. *Canadian Journal of Economics*, 50(1): 201-232.

Row, J. and E. Welk, 2011. *The Future of Energy Use in Medicine Hat: Engaging Individuals and Businesses*. Pembina Institute.

Sallee J.M. and J. Slemrod, 2012. Car Notches: Strategic Automaker Responses to Fuel Economy Policy. *Journal of Public Economics*, 96:981–99.

US Department of Energy (DOE), 2017. Central Air Conditioning. <https://energy.gov/energysaver/central-air-conditioning>.

US Environmental Protection Agency (EPA), 2014. *Emission Factors for Greenhouse Gas Inventories*. Washington, DC.

US Environmental Protection Agency (EPA), 2017. AVERT, U.S. national weighted average CO₂ marginal emission rate, year 2016 data. Washington, DC.

Ward, C., 2015. It's an ill wind: The effect of fine particulate air pollution on respiratory hospitalizations. *Canadian Journal of Economics*, 48(5): 1694–1732.