Energy Policy 150 (2021) 112157

Contents lists available at ScienceDirect

Energy Policy

journal homepage: http://www.elsevier.com/locate/enpol

Lessons from a utility-sponsored revenue neutral electricity conservation program[☆]

Brandon Schaufele

ARTICLE INFO

JEL classification:
H23
Q48
Q58

Keywords:
Electricity demand
Energy conservation
Program design
Revenue neutral

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.

1. Introduction

Most research on energy conservation programs focus on large scale, tax-financed initiatives launched at the national or state level (Boomhower and Davis, 2014; Allcott and Greenstone, 2017).[1] Yet, electricity regulators often require local utilities to pursue similar efficiency programs, with programming targeted at local customers.[2] Regulated utilities have different constraints than governments: utilities often have small footprints and cannot fund programming through the tax base. Utility-sponsored conservation must be underwritten through higher rates or surcharges levied on ratepayers. Further, because utilities cannot fund conservation programs via taxes, they often design self-financed or revenue neutral schemes. Revenue neutrality means that there are two instruments influencing electricity consumption: rebates are offered for, say, the purchase of new energy efficient appliances, with the subsidies funded via fees added to ratepayers' bills. Subsidies tend to be coarse and are limited to the purchase of a pre-determined class of durables that indirectly influence electricity consumption behaviour. Fees, in contrast, directly target electricity consumption, thus resemble prices and Pigouvian regulation. Of course, while surcharges fund the sought-after investment in energy efficiency (e.g., the purchase of new appliances), these fees may inadvertently introduce a wedge in the electricity market, creating deadweight loss.

This study uses monthly account-level data from over ten years to investigate the savings generated by a revenue neutral 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 with a revenue neutrality constraint. The main conclusions are that surcharges reduce household electricity demand, while subsidies for the purchase of energy efficient appliances have few economically meaningful effects on consumption. Using the Government of Canada's current benchmark social cost of carbon (SCC) suggests that the program is socially welfare-enhancing, but overall gains

☆ Funding and data for this paper were generously provided by the City of Medicine Hat, Alberta, Canada.

E-mail address: bschaufele@ivey.ca.

[1] Energy conservation 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).

[2] This may be because governments are downloading responsibility for energy conservation to utilities or, alternatively, utilities may possess an informational advantage and are better able to target programs according to local or regional characteristics, achieving improved outcomes. Further, the stated objectives of these programs are usually two-fold: (i) to improve environmental outcomes and (ii) to lower the future costs of energy infrastructure.

https://doi.org/10.1016/j.enpol.2021.112157
Received 4 June 2020; Received in revised form 11 December 2020; Accepted 9 January 2021
Available online 2 February 2021

are tiny and hinge on the assumed value for this SCC parameter. In other words, even though the subsidies did not change household electricity consumption relative to a counterfactual scenario where the program did not exist, the second instrument, the one used to finance the appliance rebates – i.e., the surcharge on high consuming households – did ensure that total benefits outweighed total costs.

Despite the popularity of energy efficiency programs, a puzzle known as the "energy efficiency gap" persists (Allcott and Greenstone, 2012).[3] According to the energy efficiency gap literature, 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 with a less efficient model). 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, hence, there is a "gap" in energy efficiency outlays.

The ostensible missing investment in energy efficiency has implications for social welfare. Generation of electricity produces environmental externalities such as emissions of carbon dioxide (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 households fail to invest in energy efficiency, not only do they forego private benefits of lower energy bills, but social benefits such as improved health and environmental quality also fail to materialize.

These unrealized social benefits have prompted utilities (and 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).[4] 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. 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., 2018). 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 (Houde and Aldy, 2017). Appropriately evaluating utility-managed energy conservations 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. (2018), 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 (informally, 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 even larger increases during the summer (up to 20 kWh).[5] 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, mostly 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

[3] There is debate, including in Allcott and Greenstone (2012), as to whether the energy efficiency gap exists. An emerging literature, including several papers discussed below, suggest that it may be smaller than originally believed.

[4] This $4 million is not exclusively for the electricity version of Hat Smart studied in this paper. In addition to the electricity policy, a separate program for natural gas conservation existed. The $4 million sums total spending on energy efficiency programs in the city. This includes funds allocated to conserving natural gas consumption and other programs initiated by the Province of Alberta. Only the electricity portion of the program is evaluated in this paper.

[5] Several features differentiate Mexico's Cash for Coolers program from Hat Smart. First, it was a nation-wide 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).

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.[6] Further recent evidence includes Qui and Kahn (2019), research that examining building retrofits in Phoenix, and Zha et al. (2020), a study evaluating appliance energy label program in China.

This emerging consensus on the efficacy of energy conservation programs appears pessimistic. Still, 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. That is, the fee led to a reduction in the quantity of electricity demanded, but spending on more efficient appliances had no discernable effect. This arises because the demand for electricity is responsive to price, while appliance subsidies are non-additional. Notwithstanding these results, the response to the surcharge is best characterized as trifling: in a city of 60,000, roughly 536 MWh were conserved over 9 years. This is approximately 1 kWh per person-year.[7] And to be clear, 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 instrument than rebates, even in markets, such as electricity, where demand is extremely inelastic.

Ultimately, the results of this study show two things. First, the revenue neutral Hat Smart served to transfer money between households without generating many consequential costs or benefits. Second, conclusion offers several potential modifications that can be applied to this class of programs. Yet, the main lesson from this analysis is that prices appear to be more effective than subsidies for influencing energy conservation.

2. Program structure and data

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

Hat Smart is a revenue neutral energy efficiency program. It offers rebates to ratepayers 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 price of new air conditioners, refrigerators, dishwashers and clothes washers. Hat Smart was an energy conservation program, whose objective – and the purported benefits of the program – was to reduce electricity consumption.

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 950 kWh threshold. That is, if an account holder consumed, say, 1100 kWh within a billing period, they would pay the monthly rate for the first 950 kWh of consumption and then the monthly rate plus the ECC on the remaining 150 kWh. The ECC did not vary during the sample period, equalling \$0.0074/kWh throughout.[8]

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. Funds on average ran out in September and citizens were informed of this via notices in the local newspaper and in their bills. Second, residents were not required to verify that they either disposed of their old energy inefficient appliance or acquired 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 washers were available during the initial phase of Hat Smart but not in subsequent years.

The account-level data used in the study cover every household in the city from 2007 through 2014.[9] Information was provided on the billed electricity consumption for all addresses in the city. Except for the ECC, 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 is referred to as a month for convenience. Likewise, the terms household, account and consumer are used interchangeably. During this period all households in the city were also converted from analogue 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 completed 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

[6] 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).

[7] Only a subset of households received rebates or paid the surcharge, so this estimate under-estimates the treatment effect on treated units.

[8] ECC charges were clearly described on all customer bills and paper bills were the only type of billing used during this period. Inserts described the program and its funding structure. Because bills continued to be mailed to all customers, households needed to open and examine the envelopes in order to pay the correct amounts. This suggests that residents were aware of the surcharge and HatSmart program. Further, city administrative staff strongly believed that the program was well-known and understood. Despite these reassurances however, it is difficult to definitively claim or disclaim whether households paid attention to or knew about the ECC and HatSmart. Residual lack of awareness is, on its own, a lesson for program designers.

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

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	–	100	100
Refrigerators	198.25	13.11	100	200
Clothes washers	178.18	26.50	75	775

[a] Minimum and maximum refer to monthly values.

meters were read.[10]

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 663.24 kWh, 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.[11] 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 largest variation in rebate amounts 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.

3. Conceptual framework and empirical methodology

The conceptual framework and empirical methodology used to evaluate Hat Smart is discussed in two parts. First, the method used to infer benefits from rebate payments is discussed. The economic implications of the surcharge are then reviewed. It is important to emphasize that I am seeking to measure the effect of incentive payments and surcharges on electricity consumption behaviour. As in Houde and Aldy (2017), I do not quantify welfare from new or upgraded appliances. Consumers likely obtain additional utility from the purchase of new appliances: there are benefits from owning a new and improved refrigerator or dishwasher. While these benefits are critical for the cost-benefit analyses of tax-financed energy efficiency programs, it is less obvious that they should be included in the evaluation of self-financed programs initiated by regulated utilities. Utilities are urged to develop energy efficiency programs to *conserve electricity consumption and improve environmental outcomes*. Their focus is electricity consumption and the associated environmental spillovers. The program's focus is not general household well-being. As a result, for this analysis, the prospective benefits from the mere purchase of new appliances is set aside and attention is exclusively on derived electricity demand.

3.1. Evaluating the rebates

The stated goal of utility-sponsored efficiency programs is electricity conservation. Conceptually, therefore, the benefits of conservation programs are straightforward. The objective of programs such as Hat Smart is to reduce electricity consumption by shifting a household's monthly (derived) electricity demand curve to the left and reducing how much electricity a household consumes. This is how the benefits are ultimately measured: as the effect of the rebate on electricity consumption. Still, it is worth framing how a subsidy influences household decision-making, by facilitating this reduced demand for electricity.

This conceptual discussion proceeds in two steps. First, rebates make appliances cheaper, so this influences how purchasers trade-off appliance attributes. The discussion follows Houde and Aldy (2017), illustrating, in Fig. 1 via a conventional indifference curve-budget line graph, the range of potential outcomes. Second, how rebates yield gains in surplus is outlined from the perspective of the electric utility. The marginal benefits of energy efficiency from both a private and social perspective are considered, holding constant other appliance attributes. Fig. 2, shown below, isolates how and why the benefits from a subsidy for the purchase of an energy efficient appliance arise and shows how a lack of targeting can lead to sizable inframarginal transfers.

Households trade-off attributes when purchasing new appliances. This trade-off is shown in Fig. 1. For simplicity, assume appliances are only described by two characteristics, size and efficiency. Households choose an appliance with a mix of size and efficiency such that their indifference curve is tangent to their budget constraint. In panel (a) of Fig. 1, a household's pre-rebate optimal appliance choice is shown as point A. Panel (a) also shows how rebates change the household's appliance decision. An appliance-level rebate, like those used in Hat Smart, shifts the household's budget constraint rightward. The household reoptimizes across appliance attributes and selects a different appliance, one that places them on a higher indifference curve. In panel (a), the post-rebate optimal appliance choice is shown by point B. At B, the appliance obtained by the household is larger as B sits above A. It is also more energy efficient. In this graph, the improvement in efficiency is measured as the horizontal distance between A and B. A sufficiently large horizontal shift means that the flow of services provided by the appliance at B is likely to use less electricity than appliance at point A, even though it is larger.

Panel (a) can be contrasted with the situation in panel (b). Panel (b) replicates the subsidy on energy efficiency but with a markedly different outcome. Households in panel (b) start at a pre-subsidy point C. The rebate, as in panel (a), shifts the budget line, moves households to a higher indifference and makes them better off. However, in panel (b), the shape of the indifference curves yields a different prediction. In panel (b), the rebate has a much smaller influence on energy efficiency. As households move from C to D, they largely maintain the same level of efficiency but convert the additional savings, the rebate, into larger appliances, a phenomenon known as upgrading. At point D, households purchase a larger appliance with roughly the same level of energy efficiency as the option at point C. The larger appliance makes the household better off but does not meaningfully reduce electricity demand or contribute to the utility's objective. Of course, whether the situation in

[10] This has the potential to introduce measurement error as estimated electricity consumption of household i in period t, y_{it}, is a noisy measure of true electricity consumption, $\hat{y}_{it}$. Using estimated meter readings implies that $y_{it} = \hat{y}_{it} + v_{it}$, where the components of v_{it} include a time fixed effect, τ_t, and measurement error, e_{it}: $v_{it} = \tau_t + e_{it}$. HatSmart incentives, the treatment variable, are measured without error, so imprecisely estimated standard errors due to classical measurement error would be the first-order effect. A second concern relates to unbiased estimation of the parameter of interest. The potential for measurement error requires me to assume $E[u_{it}|HatSmartIncentive_{it}] = 0$, where (1) u_{it} is the error term in the main regression model, presented in section 3.2 and (2) $u_{it} = e_{it} + \psi_{it}$ where ψ_{it}, and the error associated with main regression (see section 3.2 for further discussion) and is e_{it} the error from the potentially mismeasured electricity consumption (follows from linearity of the expectation operator). In other words, I must assume is that both errors are uncorrelated with HatSmart incentives. Any violations of this assumption are likely to attenuate the estimated parameters of interest.

[11] As mentioned, a parallel natural gas conservation program, which was similar but not identical to the electricity conservation program, was offered at the same time as Hat Smart. Rebates for windows, insulation and furnaces were offered. Even though reducing electricity consumption was not their main objective, these rebates, those offered within the natural gas program, were also examined in the context of electricity conservation. No meaningful results were found.

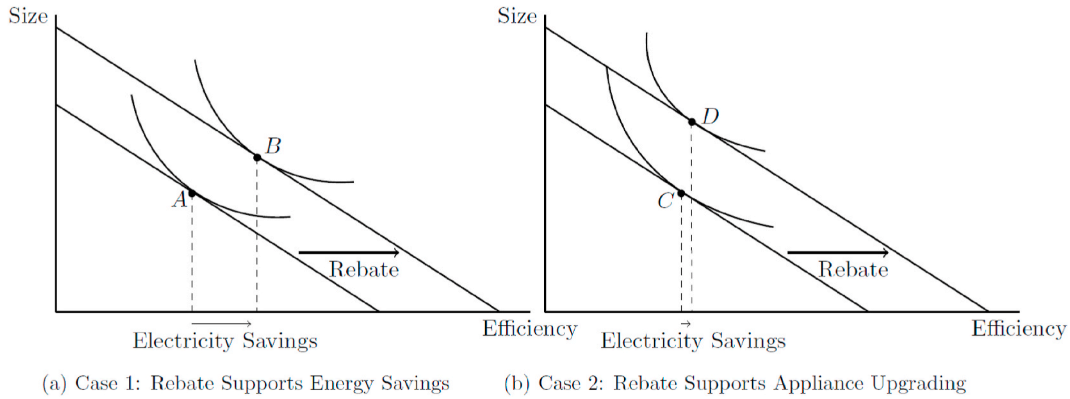

Fig. 1. Influence of rebate on Household's private decision to invest in energy efficiency.
Note: Panel (a) represents a scenario where a rebate supports energy efficiency. The pre-rebate optimal appliance choice is given by point A. The post-rebate optimal appliance choice is shown by point B. At B, the appliance purchased by the household is larger and uses less electricity per unit services as B sits above and the right of A. The improvement in efficiency is measured as the horizontal distance between A and B. A sufficiently large horizontal shift means that the flow of services provided by the appliance at B is likely to use less electricity than appliance at point A, even though it is larger. The same shifts occur in Panel (b); however, in this panel, while there is an improvement per unit services when a household shifts from point C to D, these savings in electricity are insufficient to compensate for the increase in appliance size.

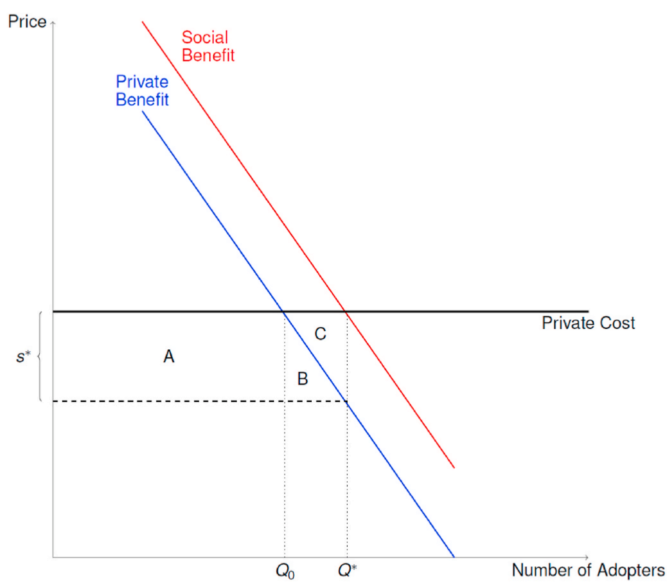

Fig. 2. Demand of energy efficiency in the appliance market (source: Boomhower and Davis, 2014).

panel (a) or panel (b) is likely to prevail is an empirical question.

Fig. 1 helps convey how rebates for appliances influence purchase decisions through a consumer choice mechanism, yet this study is interested in evaluating the impacts of changes in actual electricity consumption at the utility level. To this end, I investigate the direct costs and benefits of subsidies on electricity consumption. Utility-sponsored programs such as Hat Smart generate both private and social benefits. Private benefits are the dollar-valued amount of electricity conserved directly attributable to the program. These benefits differ from those illustrated in Fig. 1. They are program specific and accrue at the program-level (i.e., they exclude transfers). Moreover, they only materialize, from the utility's perspective, if the program addresses a pre-existing market failure. Without a market failure, private actions in the market would be optimal. For example, Hat Smart sought to improve environmental outcomes because the price of electricity failed to internalize the full damages from its generation. While CO2e emissions are the focus in this paper, a wide range of alternative market failures have been highlighted (Fowlie et al., 2018). 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; Jessoe et al., 2020) (e.g., the agent paying the utility bill may be different than the agent consuming energy) as well as a series of behavioural economic explanations such as myopia and inattentiveness (Allcott and Greenstone, 2017). Market failures imply that, absent an incentive from the utility, 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.

Fig. 2, illustrating the market for energy efficiency and based on Boomhower and Davis (2014), depicts the primary economic elements of this class of programs. The blue downward sloping line plots the demand for energy conservation investment, the number of households who adopt an efficiency-enhancing product. This reflects the willingness of households to pay for electricity savings when purchasing a new appliance. As in Fig. 1, the underlying idea is that appliances are differentiated products and the efficiency of clothes washers and dryers, for instance, represent a core attribute of these goods. The blue line reflects the demand for the characteristic of "efficiency", holding all other appliance characteristics constant. Also, drawn in Fig. 2 is a horizontal private cost curve. This is the (unsubsidized) "total price" that consumers must pay for the energy efficiency attribute, holding all other characteristics constant. This total price is comprised of any premium paid for the appliance plus the lifetime discounted operating costs of the appliance. (Operating costs primarily depend on the price of electricity.) A second, dashed horizontal line is also drawn. This is the price of efficiency net of the subsidy rate, where s* is the rebate provided by the utility. The asterisk (*) on s* is used as in Boomhower and Davis (2014) to illustrate the socially optimal subsidy level. This curve sits below the initial private cost curve because the subsidy for energy efficiency is designed to reduce both the cost of purchasing an energy efficient appliance and the lifetime cost of energy consumption.

Fig. 2 highlights several important dimensions of the program. The subsidy level is represented by s*. s* is the amount of money provided by the utility to the household for the purchase of a new energy efficient appliance. The marginal consumer is willing to pay for efficiency until the private benefits equal her private costs. Without a subsidy, the demand for efficiency equals Q_0. s* lowers the private cost of efficiency and thus increases the number of adopters from Q_0 to Q*. Given these equilibria, three regions are apparent in Fig. 2. First, utilities are not able to discriminate between those that place high and low values on efficiency (e.g., the utility cannot distinguish between the household in panel (a) or (b) of Fig. 1). Households with high willingnesses to pay for electricity conservation would have invested in efficient appliances without a subsidy. Area *A*, therefore, represents a transfer to infra-marginal households, those consumers who intended to purchase an energy efficient appliance irrespective of the subsidy (or those, like in panel (b), who obtain an appliance so large that it negates any electricity reduction). In other words, even in a counterfactual scenario where $s = 0$, these households would make identical decisions. (A positive subsidy, of course, lowers the cost of their investment.) As these households would have invested in energy efficiency even in counterfactual scenario where there is no subsidy, any reduction in monthly electricity consumption from these accounts is not a conservation benefit attributable to Hat Smart; the program did not change outcomes and area *A* represents an economic transfer. Yet, while reduced electricity consumption from households in area *A* is not a benefit attributable to Hat Smart, it is equally important emphasize that transfers are not economic costs.[12] They are purely distributional, reflecting a shifting of economic surpluses between groups.

Area *B* in Fig. 2 does represent the private benefits attributable to Hat Smart. This triangle captures the additional conservation investment that is directly induced by the subsidy (e.g., this reflects the horizontal increase in energy efficiency shown in panel (a) of Fig. 1).

For any given subsidy, s*, the economic benefits of the program vary with the slope of the private benefit curve. More elastic demand for efficiency implies larger benefits from utility-provided subsidies. Inelastic demand for efficiency, in turn, implies that subsidy programs may struggle to induce additional conservation, because households are unresponsive along the efficiency margin. As benefits in this context are limited to electricity conservation, this means that the elasticity of electricity consumption with respect to the subsidy is the key parameter needed to measure private benefits from the program.

Fig. 2 also illustrates the social benefit attributable to the conservation program. This is shown by area *C*. The red social benefit line represents the sum of the private benefits from electricity conservation plus any additional social benefits coming via spillover effects. Reduced emissions and the associated environmental and health improvements imply that private investments in efficiency have positive spillover effects and that subsidies can increase these social benefits in conjunction with the private gains. Spillovers include CO2e abated and lower ambient concentrations of local pollutants. The size of area *C* depends on both the slope of the social benefit curve and marginal value of the externality. Fig. 2 shows a scenario where the subsidy just so happens to equal the marginal damage from CO2e emissions. In practice, these are often not equal. Further, in this analysis, social benefits equal a constant marginal damage multiplied by the tonnes of emissions abated. Medicine Hat has virtually none of the air quality issues that are prevalent in larger urban centres. As a result, the social benefits of efficiency can be limited to tonnes of CO2e abated. The constant therefore reflects the social cost of carbon (SCC).

3.2. Estimating the subsidy elasticity of electricity consumption

As mentioned, no information is available on which appliances households purchased. Thus, measurement of benefits is constrained to the evaluating the reduced-form effect of subsidies on electricity consumption, the derived demand from the horizontal axis in Fig. 1. The utility's goal is to reduce electricity consumption, however, so this is the relevant statistic from their perspective. Simply, rebates induce the purchase of efficient appliances, efficient appliances reduce demand for electricity, reduced electricity demand entails that the utility achieves its objective.

A reduced-form equation captures the effect of incentives of electricity consumption and hence measures the benefits attributable to this narrow objective. Specifically, the model estimated is:

$$y_{it} = \alpha \cdot HatSmart\ Incentive_{it} + \gamma_i + \tau_t + u_{it}$$

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 electricity 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. Households self-select into it. In more restricted specifications, the sample is therefore 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. u_{it} is the error term that captures everything that varies at the household-by-time level.

α is the coefficient of interest, representing the change in energy consumption per \$100 of incentive. If $\alpha = 0$, then subsidies are infra-marginal as shown by region A in Fig. 2. In contrast, $\alpha < 0$, suggests that the subsidy did decrease electricity consumption. Identification of this parameter in the most restrictive models exploits differences in the timing of rebate cheques conditional on the time and address fixed effects. Fig. 3 illustrates this timing. Household 1, for example, receives a rebate in, say, January, while household 2 receives their cheque in March. Unbiased evaluation of Hat Smart requires that, conditional on address and time fixed effects, u_{it} is uncorrelated with incentive payments. The identifying assumption is that households 1 and 2 are conditionally identical but for the timing of their Hat Smart rebates. This means it is possible to use household 2 to formulate a counterfactual for household 1's electricity consumption in the absence of the rebate cheque. This assumption is viewed as 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 refrigerators).

3.3. Economic implications of Hat Smart's surcharge

A common misperception is that whatever money paid to households via programs such as Hat Smart is a cost of the program. But transfers are not economic 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 (excluding the provincial funding used in the first year), economic costs only arise from the deadweight loss due to reduced demand for electricity. Moreover, surcharges have an additional feature: as with rebates, reduced consumption attributable to a fee or surcharge generates spillovers and social benefits. Thus, both the deadweight loss of the ECC as well as the social benefits from reduced electricity generation must be measured.

[12] Transfers are predominantly funded by ratepayers, but the province did provide money in first wave of Hat Smart. As taxpayers contributed to the program, there would also be inefficiencies associated with the collection of tax revenue. Because these tax distortions are believed to be small (because the funds from taxpayers are small), they are excluded from the analysis and ignored in Fig. 3.

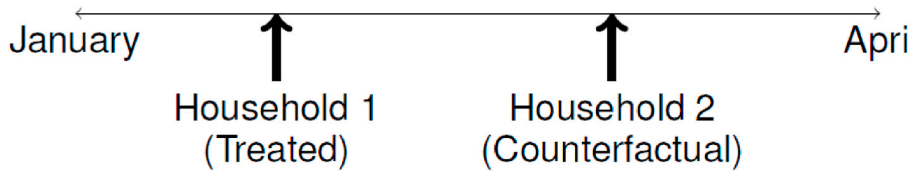

Fig. 3. Illustration of Timing used to Identify Rebate Elasticity.

3.4. Deadweight loss from surcharge

Fig. 4 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 are the within month supply functions facing 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 950 kWh/month, the standard constant rate supply curve applies to all consumption. After the threshold, these high energy consuming households must pay the additional ECC fee. For those households that exceed 950 kWh per month, the supply curve jumps to $Supply^{ECC}$ for all consumption exceeding that threshold. 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 and is the economic cost of Hat Smart. Of course, this triangle only exists for consumption in excess of the ECC threshold.[13]

The size of this triangle critically hinges on the responsiveness of demand with respect to the ECC. This is encapsulated in the price elasticity of demand. A smaller elasticity of demand (in absolute value) suggests that the demand curve is steep and households do not notably alter their behaviour in response to the surcharge. The deadweight loss from the ECC is small in this case. This can be contrasted with a flatter demand curve where the elasticity of demand is larger (in absolute value) and the economic costs are potentially large.

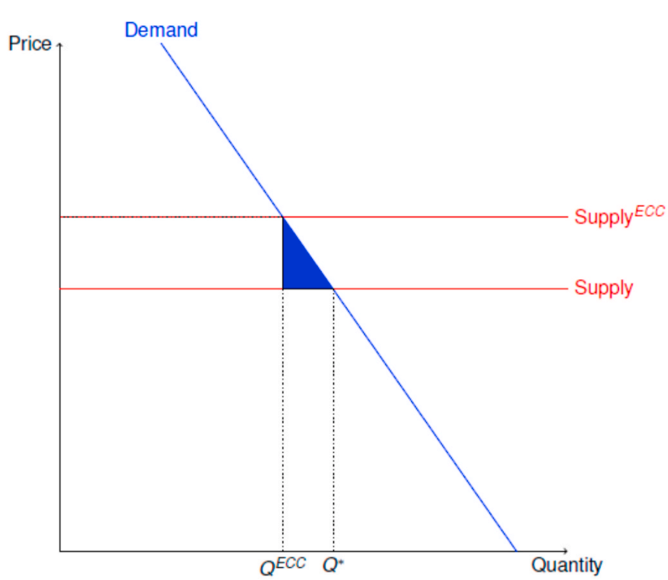

Fig. 4. Deadweight loss due to ECC

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

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

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.

As stated, the deadweight loss formula varies with 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 buy 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. The social benefits are converted to a dollar-value by multiplying the number of tonnes of CO2e abated – i.e., the reduction in electricity demand attributable to the ECC – by the constant SCC.

3.5. 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, however, is that the regressions required to estimate the elasticity of demand have no cross-sectional variation in prices across households. This means that the source of identifying variation used to pin down α cannot be used to determine the elasticity of demand, because, 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.[14] The bunching estimator applied here builds on Shaffer (2020) (see also, 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 the 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:

[13] This characterization of the deadweight loss triangle potentially overlooks other distortions that are common to electricity pricing. For example, electricity rates are frequently set equal to average, rather than marginal, cost. Likewise, the triangle is also gross of the costs of environmental externalities, a point considered below.

[14] Observations are pooled and treated as cross-sectional.

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

where z^* is the price threshold, $\widehat{B}$ is the measured excess mass to the left of the surcharge threshold and $\widehat{h}_0(z^*)$ 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. $\widehat{h}_0(z^*)$, in particular, is critical. This counterfactual is estimated in a region around z^*: $[z^* - \delta_b, z^* + \delta_b]$, where δ represents the width of an interval 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: $\widehat{h}^*_-$: $[z^* - \delta_b - \delta_c, z^* - \delta_b]$, $\widehat{h}$: $[z^* - \delta_b, z^* + \delta_b]$ and $\widehat{h}^*_+$: $[z^* + \delta_b, z^* + \delta_b + \delta_c]$. It is possible to use the densities in each of these three regions to calculate the following cumulative densities: $\widehat{H}^*_- = \delta_c \widehat{h}^*_-$, $\widehat{H}^* = 2\delta_b \widehat{h}^*$ and $\widehat{H}^*_+ = \delta_c \widehat{h}^*_+$. Given these quantities it is possible to define actual excess mass as:

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

And the counterfactual mass as:

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

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

4. Results

4.1. Energy savings attributable to Hat Smart rebates

The change in electricity consumption for each category of rebate is shown in Table 2. Three separate econometric models are estimated, one for each air conditioners, clothes washers, refrigerators and dish washers. These models are distinguished by the underlying source of variation that statistically identifies the parameter of interest. Column (1) uses all households in the city as a baseline. Column (2) restricts the sample to households that received any Hat Smart incentive any point in time. For example, when I evaluate the effect of, say, the clothes washer incentive, the counterfactual is formulated by using households who received dishwasher, fridge and air conditioner rebates as well as the differential timing of clothes washer rebates. The reason for 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 as illustrated in Fig. 3. 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. Defining the samples this way implies that the sample sizes for columns (1) and (2) will be identical across rebate types. The sample in column (3) varies across rebate types, so contains different numbers of observations.

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.

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.[15] These coefficients are unstable and the confidence intervals are wide, implying that the true value could be notably larger or smaller. As such, it is difficult to draw meaningful inferences.

Despite the imprecision of the point estimates, relative to the other categories of rebates, air conditioners appear to yield the largest reductions in electricity use. It is possible to use these point estimates to provide a sense of the underlying factors that drive the energy efficiency gap. Expectations about asset durability turn out to be a major contributing factor. The US Department of Energy projects the typical lifespan of an air conditioner to be 15–20 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 $0.08/kWh rate for electricity consumption and with an 8% discount, the private return from this $100 rebate for an energy efficient air conditioner is −1.4%. Extending the assumed air conditioner lifespan to 20 years yields electricity savings totaling 2,880 kWh with a private return from this $100 equaling 13.1%, suggesting that investing in energy efficient air conditioning is privately beneficial. Including social benefits from abated CO_2e, of course, makes investing in air conditioner efficiency more attractive as does a smaller discount rate. Yet, while these values seem promising (i.e., investing in air conditioning yields net benefits for

Table 2
Energy savings attributable to hat smart rebates.

	(1)	(2)	(3)
Panel A: Air conditioner rebates			
kWh per $100 incentive	−6.797	−5.512	−12.005
	(8.105)	(8.171)	(8.675)
Number of households	27,921	3,925	228
Number of observations	2,200,266	308,069	17,295
Panel B: Clothes washer rebates			
kWh per $100 incentive	−0.035	1.251	2.068
	(2.104)	(2.386)	(2.520)
Number of households	27,921	3,925	2,435
Number of observations	2,200,266	308,069	199,462
Panel C: Refrigerator rebates			
kWh per $100 incentive	−1.847	−3.699	−2.764
	(3.402)	(3.562)	(4.813)
Number of households	27,921	3,925	833
Number of observations	2,200,266	308,069	65,579
Panel D: Dishwasher rebates			
kWh per $100 incentive	0.139	0.187	0.163
	(0.085)	(0.088)	(0.073)
Number of households	27,921	3,925	675
	2,200,266	308,069	53,049

Notes: Value in the parenthesis report standard errors clustered at the account level.

[15] Unlike clothes washers, refrigerators and dishwashers, air conditioners are only used in during warm months. If the model in column (3) is re-estimate, but just for summer months, the point estimate decreases by 2 kWh/month to −10.211 kWh per $100 rebate. Similar to the model in Table 2, this standard error on this point estimate is large, equaling 9.784.

reasonable parameter values), the empirical support for strong claims is shaky. It is likely too shaky, in fact, to be useful for policy analysis or to guide program design. Indeed, the estimates should be interpreted with healthy caution given the imprecision of the coefficients: it is difficult to draw meaningful conclusions from these models.

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 argue 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, at first glance, this result seems odd. 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 through the adoption of more efficient units. 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 may have 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, increasing 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 et al. (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 at conventional levels, but the magnitudes are trivial. Model (3), 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.

4.2. Costs and benefits of Hat Smart surcharge

Few concrete benefits from Hat Smart rebates are identifiable in Table 2. Hat Smart's energy conservation incentives failed to translate into reduced electricity consumption. These results echo others that are found in the literature. But a distinguishing feature of Hat Smart is it is a revenue neutral program with two instruments: rebates and surcharges. The surcharge is investigated next.

As described, the gross economic costs of Hat Smart equal the deadweight loss attributable to the surcharge, which is a function of the elasticity of electricity demand. So, the first step in the evaluating the surcharge involves estimating the elasticity of electricity demand with respect to price. Table 3 presents three estimates for this statistic.

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, both point estimates 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. These coefficients suggest two things. First, electricity demand may be extremely inelastic. Electricity demand is ordinarily viewed as extremely inelastic with limited response to changing prices and it may not be statistically possible to distinguish the true response from zero. In other words, the true demand response is very small (virtually a vertical line in Fig. 4). Interestingly, if electricity demand is indeed perfectly inelastic, Hat Smart effectively has no economic cost beyond its administration expenses. Because Hat Smart is revenue neutral and funded by the surcharge on electricity consumption, the only economic costs arise via quantity contraction in the electricity market. A perfectly inelastic demand response implies that consumers do not reduce electricity consumption due to the ECC and, hence, there is no deadweight loss. Instead, with perfectly inelastic demand, the ECC is distributional, transferring surplus from households to the utility (who then recycle it back to households).

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 the resulting time-series coefficients are biased. Indeed, given that only time series variation is used to recover these coefficients, this is likely the most plausible explanation for the positive price elasticities of demand. Further, given the size of the standard errors, it is difficult to infer anything meaningful about consumer responses to electricity prices.

While the top of Table 3 uses time series variation to infer the elasticity of demand, the bottom panel applies the cross-sectional bunching estimator. This bunching estimator 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. The identifying assumption in this analysis is that, even though the variation is cross-sectional, households that are within a small bandwidth of the threshold are exchangeable or sufficiently similar to infer a credible estimate of the behavioural response to higher electricity prices.

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 implying 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). Importantly, the −0.05 elasticity of electricity demand is virtually identical to the one obtained for British Columbia by Shaffer (2020).

Table 3
Elasticity of electricity demand.

Time series variation		
Short-run elasticity	0.273	
	(0.309)	
Long-run elasticity		0.330
		(0.372)
Month-of-year fixed effects	Y	Y
Location-year fixed effects	Y	Y
Number of observations	2,200,260	2,200,260
Cross-sectional variation		
Elasticity	−0.052	
	(0.001)	
Number of observations	40,299	

Notes: The top panel contains clustered standard errors with clustering at the account level. Standard errors for the bunching estimator are obtained via bootstrapping with 50 replications.

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 $1,985 over the entire 2008 to March 2014 period. Granted, this deadweight loss does not include any fixed or variable administrative costs involved in managing the program but is best labeled as miniscule. Simply, the combination of three factors, (i) a tiny ECC, (ii) a small elasticity of demand and (iii) a modest share of households paying the ECC, leads to a clear conclusion: the ECC surcharge on electricity generates trivial market distortions and hence financing Hat Smart involves negligible economic costs.

The $2,000 excess burden of Hat Smart represents the "gross of environmental benefits" cost of the program. This "cost" doesn't factor in prospective social benefits from reduced electricity production as these gross costs do not factor in environmental improvements. 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. Using the cross-sectional elasticity, the program cumulatively reduced electricity consumption by 536.37 MWh. All of Medicine Hat's electricity is generated via natural gas. Thus, applying NRCan (2018) conversion factors, the ECC reduced CO2e emissions by 108.1tCO2e over the sample period. At a social cost of carbon of $40.70, this means that the ECC produced gross environmental benefits of $4,401. The *net benefits* – environmental benefits less the deadweight loss – from the surcharge then equal $2,416. Of course, this net benefit calculation hinges on the assumed value of the SCC. A smaller SCC – or a larger elasticity of demand – could easily make the net benefits of ECC become net costs. Thus, while the surcharge did improve economic welfare, this conclusion is sensitive to key assumptions.

Based on these results, the chief outcome of Hat Smart involves transferring money between households in the city with few meaningful economic costs or benefits – and little energy conservation.

5. Conclusion and policy implications

Economists have long argued that the savings from energy conservation programs tend to be overstated (Joskow and Marron, 1992). The results in this paper support this. Few Hat Smart rebates had any statistically measurable effect on electricity consumption, and, more importantly, most point estimates are not economically meaningful. Using reasonable assumptions for the SCC and a locally inferred elasticity of electricity demand, the surcharge levied on high consumers did produce net benefits, but the magnitude is negligible. Indeed, Hat Smart served primarily to transfer money between households in the city. Interestingly, while the subsidies appear to be dogged by potential non-additionality or upgrading behaviour, the mandated revenue neutrality – and the requirement to fund the subsidies via a surcharge on ratepayers – did yield benefits. Typically, revenue neutrality constraints are viewed as second-best restrictions on program design. In this scenario, this constraint on an appliance rebate program compensated for households' behavioural responses. Of course, this is not a feature of revenue neutrality but a consequence of levying prices on externalities. In essence, prices work.

This paper fits within an emerging literature on the economics of energy efficiency programs. Unfortunately, from both an energy and economic efficiency perspective, this research paints a cynical picture of energy conservation initiatives. Utility-based and government-funded energy efficiency programs are under-performing by not delivering their promised electricity reductions. This begs the question: are there policy tweaks that might support better results? Three options are discussed: targeting and verification, rebates conditional on energy conservation and higher prices.

Targeting and verification. A commonly advocated 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 due to information asymmetry problems. 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 subset of households to invest in energy efficiency and reduce their electricity consumption. This heterogeneity also needs to be correlated with some observable variable (e.g., consumption level or address) and the utility needs to be able to sort based on this observable variable. Practically, solving this imperfect information problem that may be intractable or politically impossible for many utilities.

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 studied by Davis et al. (2014). Recipients in this program demonstrated 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.

Hat Smart's absence of a verification policy also opens it to prospective moral hazard problems (Giraudet et al., 2018). The program may induce households change their behavior and undertake unobservable actions that undermine the program's effectiveness. In fact, the "beer-fridge" problem, described in the introduction, can be interpreted an example of moral hazard: households are provided incentives whose objective is to reduce electricity consumption. Yet, the program leads these households to adjust their behavior (i.e., adding an appliance), making adjustments that offset the prospective gains from the program. Verification could limit these moral hazard challenges.

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. (2018), 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. Finally, targeting and verification fail to address how intensive household behaviour may change due to the subsidy (e.g., there may be rebound effects).

Payments conditional on energy conservation. Rather than directly targeting appliances, 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 style of program was implemented by the provincial utility in the Canadian province of British Columbia via a scheme known as Team Power Smart. Team Power Smart is a voluntary program that offers households the opportunity to undertake annual conservation "challenges" (Fraser, 2020). 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, 2020).

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. That is, it is aimed at behaviour rather than at technology. For instance, a family that actively reduces its energy consumption by, say, reducing air conditioning in summer would not currently eligible for a Hat Smart payment. Under Team Power

Smart-style program, this family may be able to make a large contribution to conservation goals and hence would be eligible for payments.

Payments conditional on energy conservation suffer from several of the disadvantages afflicting targeting and verification. Administration costs, for example, may increase and they may be unpopular with households familiar with the "no questions asked" program. Further, Fraser (2020) demonstrates that continued conservation depends on repeated conservation challenges (i.e., interventions). Once participation lapses, Fraser shows that households backslide, forgoing a share of their energy savings.

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 here, 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 that 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, at least in this context, disappoint.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The City of Medicine Hat, Alberta, Canada graciously provided data and funding to support this research.Link grant sponsor: City of Medicine Hat, Alberta, Canada

References

Allcott, H., Greenstone, M., 2012. Is there an energy efficiency gap? J. Econ. Perspect. 26 (1), 3–28.

Allcott, H., Greenstone, M., 2017. Measuring the Welfare Effects of Residential Energy Efficiency Programs. NBER working paper no. 23386.

Bastani, S., Selin, H., 2014. Bunching and non-bunching at kink points of the Swedish tax schedule. J. Publ. Econ. 109, 36–49.

Davis, L.W., Fuchs, A., Gertler, P.J., 2014. Cash for Coolers: evaluating a large-scale Appliance replacement program in Mexico. Am. Econ. J. Econ. Pol. 6 (4), 207–238.

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

Fowlie, M., Greenstone, M., Wolfram, C., 2018. Do energy efficiency investments deliver? Evidence from the weatherization assistance program. Q. J. Econ. 133 (3), 1597–1644.

Fraser, A., 2020. Success, failure, and information: how households respond to energy conservation goals. Working paper. http://www.alastairfraser.ca/research/papers/Fraser_BCH.pdf.

Giraudet, L.-G., Houde, S., Maher, J., 2018. Moral hazard and the energy efficiency gap: theory and evidence. J. Assoc. Environ. Resour. Econ. 5 (4), 755–790.

Houde, S., Aldy, J.E., 2017. Consumers' response to state energy efficient appliance rebate programs. Am. Econ. J. Econ. Pol. 9 (4), 227–255.

Jessoe, K., Papineau, M., Rapson, D., 2020. Utilities included: split incentives in commercial electricity contracts. Energy J. 41 (5), 271–303.

Joskow, P.L., Marron, D.B., 1992. What does a negawatt really cost? Evidence from utility conservation programs. Energy J. 13 (4), 41–74.

Kleven, H.J., 2016. Bunching. Ann. Rev. Econ. 8, 435–464.

McCarthy, S., 2017. Ottawa eyes energy retrofits in buildings. The Globe and Mail (September 20,2017).

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. Energy J. 38 (4), 195–212.

Qiu, Y., Kahn, M.E., 2019. Impact of voluntary green certification on building energy performance. Energy Econ. 80, 461–475.

Rivers, N., Shiell, L., 2016. Free-riding on energy efficiency subsidies: the case of natural gas furnaces in Canada. Energy J. 37 (4), 239–266.

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

Sallee, J.M., Slemrod, J., 2012. Car notches: strategic automaker responses to fuel economy policy. J. Publ. Econ. 96, 981–999.

Shaffer, B., 2020. Misunderstanding nonlinear prices: evidence from a natural experiment on residential electricity demand. Am. Econ. J. Econ. Pol. 12 (3), 433–461.

Smart, Hat, 2017. About Hat Smart. https://www.medicinehat.ca/government/departments/utility-sustainability/hat-smart/about-hat-smart retrieved July 2017.

Zha, D., Yang, G., Wang, W., Wang, Q., Zhou, D., 2020. Appliance energy labels and consumer heterogeneity: a latent class approach based on a discrete choice experiment in China. Energy Econ. 90, 104839.