Ivey Business School
WESTERN UNIVERSITY · CANADA
Energy Policy and Management Centre

POLICY BRIEF
SEPTEMBER 2019

What Do We Know About Household Energy Conservation Programs? Evidence from Medicine Hat

By Brandon Schaufele

EXECUTIVE SUMMARY

- Programs designed to promote investment in energy conservation are popular with governments and utilities across North America as these schemes promise win-win outcomes: consumers invest in energy efficiency and lower their utility bills, while less overall energy is combusted reducing harmful environmental and health externalities.

- Until recently, little was known about the effectiveness of many of these programs. Growing evidence suggests that many initiatives designed to promote private investment in energy efficiency deliver fewer benefits than initially promised.

- This policy brief evaluates a program designed to reduce electricity and natural gas consumption in the City of Medicine Hat. Results demonstrate that this program, known as HAT Smart, generated very little energy conservation. Program funding crowded out private investment at a rate of one-to-one. Despite little meaningful energy conservation, the analysis also demonstrates that the program imposed virtually no economic costs on the city's residents.

INTRODUCTION

In 2008, the City of Medicine Hat, Alberta, launched a large-scale energy conservation program, HAT Smart, an initiative targeted at reducing residential per capita energy consumption. As of December 2015, HAT Smart had distributed over $4 million in energy efficiency incentives via 14,000 household rebates (HAT Smart, 2017). Rebates were offered for a wide range of energy conservation investments, including for the purchase of Energy Star refrigerators, dishwashers, air conditioners or for the installation of insulation, efficient furnaces, or new windows and doors. The program was entirely self-financed via an 'Environmental Conservation Charge' (ECC), a small levy per kilowatt-hour (kWh) of electricity or per gigajoule (GJ) of natural gas consumed added to the utility bills of high energy-consuming households.[1] After nine years, it is time to evaluate the actual energy savings and net economic benefits – both in electricity and natural gas – achieved by the program. The main focus of this brief is to quantitatively evaluate HAT Smart; this analysis also offers an opportunity to learn from this program and situate the results of HAT Smart within a growing body of research studying similar programs offered in other jurisdictions.

[1] Households that consume more than 950 kWh of electricity per month pay an additional $0.0075/kWh on their above-threshold electricity consumption. Similarly, accounts that consume more than 22 GJ of natural gas per month pay an additional $1.01/GJ on above-threshold consumption. (The natural gas threshold was decreased in 2016 and 2017, and now equals 19 GJ.)

Analysis of energy conservation interventions requires defining and then estimating an appropriate counterfactual scenario. Counterfactual scenarios are "but-for" situations. They are hypothetical states of the world, where everything is identical *but-for* the existence of an energy conservation program such as HAT Smart. Estimating the appropriate counterfactual scenario isolates the true effect of energy conservation programs, disentangling changes in energy consumption attributable to a program such as HAT Smart from the myriad of other changes that happen simultaneously, such as technological improvements. Evaluating HAT Smart – or any other energy conservation program – therefore involves comparing rebate-receiving households' energy consumption to the consumption of those same households, but in the absence of a rebate or subsidy. In order to do this, sophisticated econometric methods are applied to perform these counterfactual exercises.[2]

Formulating these counterfactual scenarios involves incorporating both technological improvements in energy efficiency and behavioural changes of households. Most analyses of energy efficiency programs use engineering estimates to calculate energy savings, limiting themselves to the technological dimensions of, say, the appliance. Projected energy savings are derived from simulation models and tests run in laboratory settings. However, engineering estimates frequently fail to account for important real-world features. Insulation, for example, may be installed incorrectly, or household behaviour may change as a result of incentives like the ones provided by HAT Smart (Fowlie et al., 2018). One significant example of an unintended behavioural change has been referred to as the "beer fridge problem," in which offering a rebate for energy-efficient refrigerators may increase, rather than reduce, electricity usage because households continue to operate their old unit – i.e., a household buys a new primary refrigerator, but keeps its 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, therefore requires accounting for these behavioural changes while controlling for improvements in technology..

The main conclusion from the analysis of HAT Smart is that the program yielded modest reductions in energy consumption. Between 2008 and 2014, total net benefits to rate payers from HAT Smart equalled $50,000, with approximately 8800 MWh of electricity and 3300 GJ of natural gas conserved, equal to less than 0.01% of the city's energy consumption. Indeed, for the majority of rebate payments, there is no statistically discernable evidence of any effect on electricity or natural gas consumption. Notable exceptions are rebates for air conditioners, natural gas dryers, and hot water heaters. For air conditioners and hot water heaters, the reductions in energy consumption, though small, are attributable to the joint effect of improved technology alongside behaviour – and are directly attributable to HAT Smart incentives. In contrast, natural gas dryers involved fuel switching, whereby electricity consumption declined, as households removed their old electric dryers, but natural gas consumption increased as demand for natural gas accordingly increased.

While the benefits from HAT Smart are modest, the *economic* costs are also virtually zero. This is a consequence of the program's revenue-neutral design, where fees levied on inelastic electricity and gas demand are used to fund the energy conservation incentives. Hence, the analysis highlights that the main outcome of HAT Smart involved what economists refer to as 'transferring surplus' between high-consuming ratepayers and rebate recipients within the city.

[2] These methods are described in the companion academic paper, posted on the Ivey Energy Policy and Management Centre website.

PUTTING HAT SMART INTO CONTEXT:
Understanding Recent Academic Evidence on the Effectiveness of Energy Conservation Programs

A growing body of evidence on the efficacy of energy conservation initiatives, across a spectrum of jurisdictions and designs, has accumulated in recent years. Analyses of these programs have found that much of the promise of energy conservation programming appears to be elusive or underwhelming.

Energy efficiency initiatives such as HAT Smart are often pitched as "win-win" policies. They are popular with utilities and governments across North America. In fact, the US spent more than $34 billion on conservation and efficiency programs between 1994 and 2012 (Boomhower and Davis, 2014), including $17 billion allocated in the 2009 Recovery and Reinvestment Act (the stimulus bill) (Allcott and Greenstone, 2017). In Canada, the federal government recently announced that it "is ramping up its effort to encourage building owners to invest in energy retrofits" (McCarthy, 2017)., Further to that point, HAT Smart itself was initially developed in conjunction with a program offered by the Government of Alberta.

The first "win" in the win-win characterization of energy efficiency programs refers to the private benefits from investing in energy efficiency. The reasoning is as follows. Households invest in energy efficiency, for example, by purchasing more efficient appliances or adding insulation to their attics. These investments are costly for the household – e.g., high energy-efficient appliances cost more than low energy-efficient appliances. The payoff for the household is that these investments in energy efficiency should reduce lifetime energy consumption by an amount that is greater than the initial outlay or efficiency-based price differential.

The private benefits from energy efficiency have become conventional wisdom and an oft-repeated refrain. As an example, the Government of Canada's 2017 building energy efficiency program claims that the payback periods on building retrofits are as short as 18 months (McCarthy, 2017), while Energy Efficiency Alberta launched an online calculator highlighting the potential savings from conservation activities (EEA, 2017). There is a catch to these calculations, however. The vast majority of projected cost-savings are based on engineering forecasts and not the deployment of efficiency upgrades in the real world. Behavioural responses, such as the beer fridge problem, and improper use, such as imperfectly installed insulation, can meaningfully affect the realized benefits from investments in energy efficiency. Indeed, as the recent research shows, increasing evidence suggests that many of the purported benefits from energy conservation are overstated, and that savings are difficult to obtain.

The second "win" from the win-win framing is the social benefit from energy efficiency. Generation of electricity and combustion of natural gas produces environmental externalities such as carbon dioxide (CO_2e) emissions 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. Of course, if the benefits from the first win are smaller than expected, those from the second win are also reduced.

The win-win characterization of energy efficiency programs is persuasive. However, until recently, surprisingly little was known about the actual effectiveness of utility-based conservation and efficiency programs in the real world (Allcott and Greenstone, 2012). This includes programs that are similar to HAT Smart. Evidence has accumulated over the last nine years, suggesting it is quite challenging to obtain many of the promised benefits of energy conservation, irrespective of how well the programs are designed and administered.

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. Moreover, they demonstrate that these over-estimates cannot be attributed to rebound effects or upgrading. Their estimates suggest that the private benefits from investing in improved weatherization yield an internal rate of return of -2%. By including the social benefits, the program appears more attractive, yielding an internal rate of return of -0.8%. The cost per tonne of CO_2e abated from the specific Weatherization Assistance Program that they study equalled $329. Finally, they also found that the transaction costs of increasing participation in the program were large; a 5 percentage point increase in participation would cost more than $1,000 per household.

The implications of this research are stark. Home retrofit programs are substantially less effective than forecast. Low adoption rates, consumers' reluctance to participate, and the energy efficiency gap may in fact be due to over-estimated private benefits. Even when program staff visit homes, make direct phone calls, and schedule follow-up appointments – all activities that have administrative costs – many households do not experience any benefit from energy efficiency investments. Common arguments supporting the allocation of public funds towards energy conservation programming is that consumers may be myopic, underweighting future energy savings, or that they may misperceive potential benefits. Fowlie et al.'s results, in contrast, suggest that consumers are accurately evaluating the costs and benefits of energy efficiency investments.

Davis et al. (2014) study a large-scale appliance replacement program that helped 1.5 million Mexican households purchase new energy-efficient refrigerators and air conditioners. Colloquially, this program is referred to as Cash for Coolers. Using household electricity billing records, similar to those used in this study, the authors found that 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).

There are several features that 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 replaced appliances met certain requirements. In order to qualify for rebates, the old refrigerator or air conditioner had 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).

Davis et al. also place their estimates in context. 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). Even given the strict replacement and removal requirements – requirements that are not part of the HAT Smart program – Davis et al. demonstrate that it is difficult to obtain the promised efficiency benefits from conservation programs.

Next, studying 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. Along similar lines, Rivers and Shiell (2016) provide one of the few studies of a Canadian energy efficiency program. Rivers and Shiell studied incentives to replace forced-air natural gas furnaces in Canada between 2007 and 2011, finding 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. In the end, the cost per tonne of CO_2e abated from these furnace subsidies is potentially as high as $800. These results underscore a key trend in energy consumption: improvements in technology and changes in behaviour are slowing the increase in electricity and natural gas demand. Accurate evaluation of programs such as HAT Smart, must acknowledge and control for the underlying trend in energy consumption. It cannot merely examine what occurred before and after the program's launch, because the underlying technological and behavioural trend will attribute too much of the change in consumption to the program. Instead, it is important to estimate what would have occurred in the absence of the program (i.e., the counterfactual scenario), because electricity demand may, for example, be lower in 2012 than 2008, irrespective of whether an incentive for clothes washers was provided.

Finally, following the financial crisis, the US Government passed the Recovery and Reinvestment Act. One element of this Act was the State Energy Efficient Appliance Rebate Program. State governments were provided with funding to subsidize households' purchases of energy-efficient appliances. Houde and Aldy (2017) evaluate this program and demonstrate underwhelming results. 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 improvement in energy efficiency of 2 kWh per year at most, compared with a counterfactual scenario where the State Energy Efficient Appliance Rebate Program was not implemented. Rebates mainly contributed to appliance upgrading, where households purchased a larger appliance or one with additional features. Often these larger refrigerators, dishwashers and 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).

An emerging consensus on the efficacy of energy conservation programs appears pessimistic. Yet, this background provides crucial context for Canada's continued commitment to funding investments in energy efficiency, and helps situate the results from the HAT Smart experience. HAT Smart yielded very little energy savings, but this outcome is similar for many other programs. "Missing" energy conservation does not appear to be due to poorly designed programs or mismanagement. Simply, across a wide range of jurisdictions, these types of schemes have not performed as well as anticipated because, until recently, many of the behavioural responses were poorly understood.

BACKGROUND ON HAT SMART

HAT Smart was launched in 2008. Originally designed in conjunction with a similar scheme offered by the 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 in this research.

HAT Smart was a revenue-neutral energy efficiency program. It offered rebates to ratepayers for the purchase of a pre-defined set of efficiency investments. Specifically, it was marketed as a way to help households "make better choices regarding upgrades to their homes" (HAT Smart, 2017). Predominantly, this involved rebating a fixed amount of the purchase of new air conditioners, refrigerators, windows, and furnaces. HAT Smart was an energy conservation program, whose objective – and the purported benefits of the program – was to improve household welfare through a reduction in electricity and natural gas 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, and if gas consumption was greater than 22 GJ. For example, if an account holder consumed, say, 1100 kWh within a billing period, they would pay the monthly rate for the first 950 kWh of electricity 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 and $1.01/GJ 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. Funds, on average, ran out in September, and citizens were informed of this via notices in the local newspaper. 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 proof of purchase was provided. Third, rebates were promptly paid, usually within the month. Fourth, the city advertised the rebate scheme in both household utility 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. Information was provided on the billed electricity and natural gas consumption for all addresses in the city. All residential accounts pay the same rate in each month, except for the ECC. 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 below regression analysis varies, but there are over 2.2 million observations in the data. The average monthly consumption equals 663.24 kWh, and after trimming the top and bottom 1%, had a minimum of 36 kWh and a maximum of 2216 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 dishwasher 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 was $178, with a minimum of $75 and a maximum of $775. For natural gas-focused rebates, recipients of natural gas stoves and dryers rebates recovered similar amounts of $195 and $200, respectively. Households getting a home energy audit received an average of $1,049, but the range was wide, spanning $210 to $4,250. Likewise, there was a wide range on insulation rebates, covering $125 to $2,250, with an average of $442. The average furnace rebate was $381 and the average water heater cheque equaled $285.

Table 1 | Summary Statistics

	Mean	Std. Dev	Min.	Max.
Electricity consumption (kWh/month)	663.24	392.87	36.00	2216.00
Natural gas consumption (GJ/month)	5.11	5.29	0.03	28.07
Share of households paying ECC	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
Furnace	381.41	136.14	100	500
Nat. gas stove	195.07	21.73	100	200
Nat. gas dryer	200	0	200	200
Insulation	442.72	219.49	125	2250
Water heater	285	23.51	250	300
Home energy audit	1049.11	563.95	210.25	4250

Minimum and Maximum refer to monthly values

ECONOMICS OF CONSERVATION PROGRAMS

Market failure-type arguments usually justify public intervention into energy conservation with a wide range of market failures highlighted (Fowlie et al., 2018). These 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 (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 CO_2e emissions, does not factor into private decisions to spend on insulation or efficient clothes washers.

Conceptually, the benefits of conservation programs appear easy to understand, but there are subtleties. The objective of HAT Smart was to reduce electricity and natural gas consumption. Initiation of the program was part of the city's commitment to reduce CO_2e emissions. Private benefits, those accruing to ratepayers, are therefore the dollar-valued amount of energy conserved. Yet, because one goal of the program was to reduce CO_2e emissions, the private benefits only capture some of the benefit calculation. 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 and gas bills. This is straightforward to calculate, as the amount of induced energy savings multiplied by the current rate per kWh or GJ.[3] The second benefit arises from the reduction in harmful emissions. This includes CO_2e abated or lower ambient concentrations of local pollutants. Medicine Hat does not suffer from 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 CO_2e abated. The magnitude of benefits from reduced CO_2e emissions critically hinge on the value of a tonne of CO_2e

[3] In the cost-benefit calculation, a rate of $0.08 per kWh is applied for electricity savings. This is the average variable charge for Medicine Hat in 2016. Likewise, the average price per GJ of natural gas in 2016 was $3.38.

abated. Tonnes of CO_2e abated are valued at the Government of Canada's 2016 social cost of carbon, which equals \$40.70/tCO_2e. There is substantial uncertainty in this estimate, however.[4]

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 energy consumed. Incentive-induced reductions in energy consumption are measured against a counterfactual scenario, a state of the world where everything else is held constant but HAT Smart did not exist. Figure 1 uses a standard supply and demand graph to outline what is meant by this.

Figure 1, illustrating the market for energy efficiency, depicts the main economic elements of this class of programs. The blue downward sloping line plots the demand for energy conservation investment, or, alternatively, the number of households who adopt an efficiency-enhancing product. This reflects the willingness of households to pay for, say, electricity savings when purchasing a new appliance. The underlying idea is that appliances are differentiated, and the energy efficiency of clothes washers and dryers, for instance, represent a core attribute of these products. The purple line reflects the demand for the characteristic of "energy efficiency," holding all other appliance characteristics constant. Also drawn in Figure 1 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 or natural gas. A second, dashed horizontal line is also drawn. This is the price of energy efficiency net of the subsidy rate, where s^* is the subsidy provided by the utility. 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.

Figure 1 | Demand of Energy Efficiency in the Appliance Market

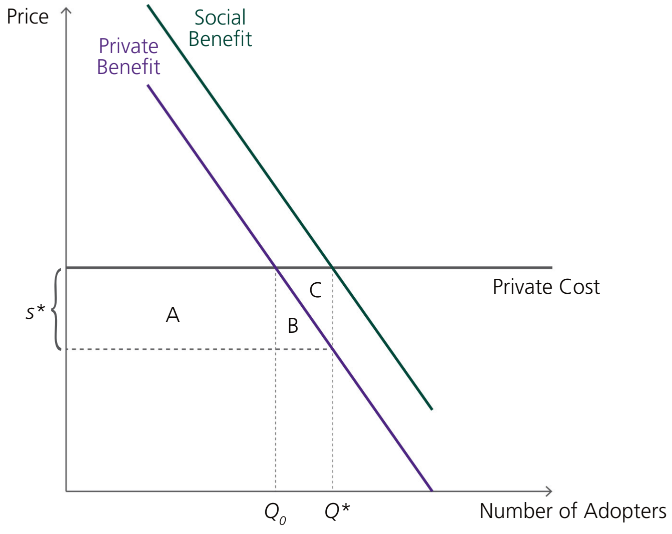

Source: Author created based on a similar figure in Boomhower and Davis (2014)

[4] Electricity will be converted to CO_2e at a rate of 0.596kg/kWh (EPA, 2017), while natural gas will be converted at a rate of 56kg/GJ (EPA, 2014).

In Figure 1, 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 will pay for energy efficiency until the private benefits equal her private costs. Without a subsidy, the demand for energy efficiency equals Q_0. s^* lowers the private cost of energy efficiency and thus increases the number of adopters from Q_0 *to* Q^*. Three regions are apparent in Figure 1. First, utilities are not able to discriminate between those that place high and low values on efficiency. Households with a high willingness to pay for energy conservation would have invested in efficient appliances without a subsidy. Area A therefore represents a transfer to inframarginal households, those consumers who intended to purchase an energy-efficient appliance irrespective of the subsidy. 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, because the program did not change outcomes (and, hence, area *A* represents an economic transfer). While reduced electricity consumption from households in area *A* is not a benefit of the program, it is equally important to emphasize that transfers are not economic costs; they are purely distributional, reflecting a shifting of economic surpluses between groups.

Area *B* in Figure 1 does represent the private benefits attributable to HAT Smart. This triangle captures the additional conservation investment that is directly induced by the subsidy. Notably, for any given subsidy, the economic benefits of the program vary with the slope of the private benefit curve. A flatter demand curve for energy efficiency implies large benefits from utility-provided subsidies. A steeper demand curve for efficiency, in turn, implies that subsidy programs may struggle to induce additional conservation as households are unresponsive along the efficiency margin.

Figure 1 also illustrates the social benefit attributable to the conservation program. This is shown by area *C*. The green social benefit line represents the sum of the private benefits from energy conservation, plus any additional social benefits coming via spillover effects. Reduced emissions and the associated environmental and health improvements imply that private investments in energy efficiency have positive spillover effects, and that subsidies can increase these social benefits in conjunction with the private gains. Spillovers include CO_2e 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. Figure 1 shows a scenario where the subsidy just so happens to equal the marginal damage from CO_2e emissions; however, in practice, these are often not equal. In this analysis, social benefits equal a constant multiplied by the tonnes of emissions abated. As previously stated, 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 CO_2e abated. The constant therefore reflects the constant marginal damage of a tonne of CO_2e emissions, typically referred to as the social cost of carbon (SCC).

Figure 1 outlines the economics of the benefit side of energy efficiency programs. In order to calculate Hat Smart's implications, the underlying parameters (e.g., slope of the demand curve) of this figure are econometrically estimated using data from Medicine Hat.

DEADWEIGHT LOSS FROM SURCHARGE

While Figure 1 outlines the prospective benefits from HAT Smart using a supply and demand graph, Figure 2 illustrates the economic costs of the program. The downward sloping blue curve is the demand curve. In this graph, this represents a household's demand for electricity or gas, rather than energy efficiency. The purple curve refers to the within-month supply functions faced by this household. A household's supply function depends on their total monthly consumption and the threshold at which the ECC kicks in. For example, if a household consumes less than 950 kWh/month, the standard constant rate supply curve applies to all consumption. For those households that exceed 950 kWh 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 energy. 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.

Like with the subsidy-induced benefits, the size of this triangle critically hinges on the responsiveness of demand with respect to the ECC. This responsiveness is encapsulated in a parameter known as the price elasticity of demand – the slope of the green curve. A smaller elasticity of demand (in absolute value) suggests that 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).

Figure 2 | Deadweight Loss Due to ECC

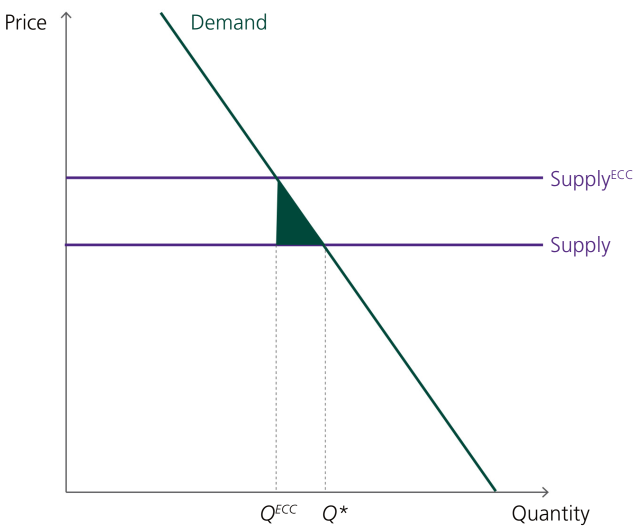

Source: Author created

The deadweight loss, the green triangle in Figure 2, changes 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 energy consumption. Reduced energy consumption attributable to the EC also produces social benefits. Fewer kWh or GJ consumed implies fewer tonnes of CO_2e 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 abated – i.e., the reduction in electricity demand attributable to the ECC – by the constant SCC.

ECONOMETRIC RESULTS

Table 2 summarizes the main conclusions of the analysis of HAT Smart. Panel A presents the econometric estimates of the effect of energy efficiency rebates on average household energy consumption. As described, these estimates should be interpreted relative to a counterfactual scenario where the rebate was not offered. That is, they do not imply that purchasing, say, a new clothes washer will not reduce energy consumption; rather, they show the marginal change in energy consumption that is a direct consequence of the HAT Smart incentive program. Panel B uses these econometric estimates, assumes a social cost of carbon of \$40.70/tCO_2e, and applies the inferred short-run price elasticities of demand from Medicine Hat to calculate deadweight losses.[5] This information is aggregated into a basic benefit-cost analysis of the Hat Smart program.

As Table 2 shows, three rebate classes yielded energy savings. Funds allocated towards air conditioner incentives led to the largest reduction in electricity consumption, about 226 MWh, and created benefits of \$33,335. This reflects slightly less than 1% of total electricity consumed by the average rebate recipient over this period. At prevailing electricity prices, this translates into between \$0.50 and \$1.00 of savings per billing cycle for households that upgraded their air conditioners. Further, assuming each MWh of natural gas-generated electricity yields 0.4 tCO_2e, this implies that 90 fewer tonnes of greenhouse gases were emitted into the atmosphere. At a social cost of carbon of \$40/tCO_2e, the monetized benefits from reduced greenhouse gas emissions equal a paltry \$3,600.

Natural gas dryers conserved 181 MWh of electricity but led to an increase of 1440 GJ of natural gas consumption. The net benefits from natural gas dryer incentives totaled \$5,585. Finally, water heater rebates saved 4800 GJ of natural gas and produced \$16,650 in combined private and social benefits. Payments for clothes washing machines, dishwashers, refrigerators, insulation, furnaces, natural gas stoves, and home energy audits yielded no statistically significant change in energy consumption.

In sum, the energy savings from HAT Smart incentives are small. These savings equal substantially less than 0.01% of the city's total energy consumption. Indeed it turns out that the ECC, the fee used to fund these rebates, yielded the largest reduction in electricity consumption over this period. Between January 2008 and December 2014, the \$0.0074 surcharge on high-consuming households induced an 8400 MWh reduction in electricity consumption, compared with a counterfactual scenario without the fee. This means that the ECC was more than 20 times as effective as the HAT Smart subsidies at reducing electricity consumption. No demand response was found for natural gas, so the ECC did not yield any statistically meaningful energy savings for gas.

Panel B of Table 2 shows that HAT Smart does pass a benefit-cost test, as the benefits are greater than the costs. Still, it is worth unpacking this statement. First, caution is warranted. Few meaningful results are obtained, and measured costs and benefits are small. Ultimately, between 2008 and 2014, HAT Smart yielded net benefits equal to approximately \$50,000 while conserving roughly 8800 MWh of electricity and 3400 GJ of natural gas. Over this period, this is equivalent to a tiny fraction of the city's total energy consumption. Second, the budgeted inflows and outflows of revenues and rebates from HAT Smart dwarf these cost-benefits numbers. It is important to understand why. From an economic perspective, programs that do not change household decision-making are not classified as either costs or benefits; they are instead referred to as transfers, which are neutral in program evaluation and effectively do not count in cost-benefit analysis. This is a strange point to understand for non-economists, but is important in this context as HAT Smart's main outcome has been to transfer money between households in the city. In essence, money moved between accounts, but most of this movement neither improved nor harmed outcomes.[6]

Table 2 | Summary of Econometric Analysis of the HAT Smart Program

Panel A: Effect of HAT Smart Rebates on Energy Conservation			
Incentives with Measurable Energy Conservation Benefits			
Incentive	Benefit ($)	Energy Savings	
		MWh	*GJ*
Air conditioner	$33,335	226	
Natural gas dryer	$5,585	181	-1440
Water heater	$16,650		4800
Incentives with No Measurable Energy Conservation Benefits			
	Clothes washer	Refrigerator	
	Dishwasher	Insulation	
	Furnace	Home energy audits	
	Natural gas stove		
Panel B: Cost-Benefit Analysis of HAT Smart			
Benefits from HAT Smart rebates	$55,570	407	3360
Benefits from financing HAT Smart with ECC		8400	
Economic costs from financing HAT Smart with ECC	$6,825		
Net benefits from HAT Smart program	**$48,745**	**8807**	**3360**

RECOMMENDATIONS FOR FUTURE VERSIONS OF HAT SMART

Given the empirical results from HAT Smart's analysis, it worth considering which program design options may yield better outcomes. Three strategies are discussed.

Targeting and verification. A common recommendation for energy efficiency programs is more precise targeting and verification (e.g., Allcott and Greenstone, 2017). Targeting may mean funds are directed to low income households or towards "energy hogs" – i.e., houses with unusually high consumption for their profile. The hope is 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. Heterogeneity means that rebates can induce some households to invest in energy efficiency and reduce their electricity and natural gas consumption. However, it also means that there are some recipients for whom the rebate is a pure transfer, those for whom a HAT Smart cheque has no effect on energy consumption or behaviour.

Similar to targeting, verification may also improve program performance, especially when beer fridge problems are a potential concern. Verification means that program administrators require evidence that old appliances are removed prior to issuing rebates. Likewise, eligibility for rebates could mimic the Mexican 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

[6] An implicit assumption underlying this statement is that the transfers were "distributionally neutral" with respect to the income distribution. Under the unlikely scenario that low-income households were more likely to pay the ECC and high-income households received rebates, the program could be viewed as distributionally negative, and vice versa. It is not possible to draw distributional conclusions with the data available.

mitigate energy consuming upgrading behaviour.

Despite the appeal of targeting and verification, some 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. 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 $1,000 per household. Finally, administrators must solve a fundamental information asymmetry problem, by predicting which households offer the biggest return. Overcoming the challenge of imperfect information, in practice, may be unfeasible.

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 or rebates 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 BC Hydro 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 would also include behavioural change. For instance, a household that actively reduces its energy consumption by turning down its thermostat is not currently eligible for a HAT Smart payment. Under a redesigned scheme, this household may be able to make a large contribution to the city's conservation goals and should become eligible for rebates.

Many energy conservation programs could easily evolve into a Team Power Smart-type of program but should include two important tweaks. First, programs should offer a sliding scale of rebates, where increasingly larger rewards are granted for larger energy reductions. Second, these programs should auto-enroll all households, rather than having them opt-in. Currently, many programs including HAT Smart require ratepayers to actively participate to obtain benefits. Automatically enrolling households ensures that all utility customers have equal access to program funding.

Higher prices. Finally, if jurisdictions are primarily concerned with improving energy conservation and reducing emissions, they could increase energy prices to match social costs. The city of Medicine Hat appears to have significant scope to increase the price of electricity and natural gas before substantial consumer behavioural changes are undertaken. Higher prices mean that substantial additional revenue would be collected by the utility. These funds could be used to offset other taxes or to fund community projects that may themselves offer conservation benefits. Effectively, jurisdictions would engage in a tax swap, whereby property taxes or user fees are reduced and energy taxes are increased. This type of tax swap is welfare-improving as long the distortion from higher energy prices is less than the distortion from property taxes or user fees.

REFERENCES

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.

Boomhower, J. and L.W. Davis, 2014. *A credible approach for measuring inframarginal participation in energy efficiency programs.* Journal of Public Economics, 113: 67—79.

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.efficiencyalberta.ca/residential/home-efficiency-tool/.

Fowlie, M., M. Greenstone and C. Wolfram, 2018. *Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program.* Quarterly Journal of Economics 133(3): 1597-1644.

Fraser, A., 2017. *Financial Incentives or Persistent Habits? How Households Respond to Financial Rewards for Electricity Conservation.* Working paper, http://www.alastairfraser.ca/.

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

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.

McCarthy, S., 2017. *Ottawa Eyes Energy Retrofits in Buildings.* The Globe and Mail (September 20,2017).

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.

ABOUT THE IVEY ENERGY POLICY AND MANGEMENT CENTRE

The Ivey Energy Policy and Management Centre is the centre of expertise at the Ivey Business School focused on national energy business issues and public policies. It conducts and disseminates first class research on energy policy; and promotes informed debate on public policy in the sector through supporting conferences and workshops that bring together industry, government, academia and other stakeholders in a neutral forum. The Centre draws on leading edge research by Ivey faculty as well as by faculty within Western University.

More information is available at
www.ivey.ca/energy

Energy Policy and Management Centre

AUTHORS

Brandon Schaufele, Assistant Professor
Business, Economics and Public Policy
Ivey Business School

The findings and opinions contained in this report reflect solely those of the author. The Ivey Energy Policy and Management Centre submits reports for external review by academic and policy experts and energy sector stakeholders. The Centre gratefully acknowledges support from organizations and individuals listed on the Centre's website:
https://www.ivey.uwo.ca/energycentre/about-us/supporters