Journal of Public Economics 113 (2014) 67–79

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

A credible approach for measuring inframarginal participation in energy efficiency programs[☆]

Judson Boomhower [a,b,*], Lucas W. Davis [c,b,d]

[a] Agricultural and Resource Economics, University of California, Berkeley, United States
[b] Energy Institute at Haas, University of California, Berkeley, United States
[c] Haas School of Business, University of California, Berkeley, United States
[d] National Bureau of Economic Research, United States

ARTICLE INFO

Article history:
Received 19 December 2013
Received in revised form 18 March 2014
Accepted 20 March 2014
Available online 2 April 2014

JEL classification:
D12
H23
Q40
Q54

Keywords:
Energy efficiency
Regression discontinuity
Additionality

ABSTRACT

Economists have long argued that many recipients of energy-efficiency subsidies may be "non-additional," getting paid to do what they would have done anyway. Demonstrating this empirically has been difficult, however, because of endogeneity concerns and other challenges. In this paper we use a regression discontinuity analysis to examine participation in a large-scale residential energy-efficiency program. Comparing behavior just on either side of several eligibility thresholds, we find that program participation increases with larger subsidy amounts, but that most households would have participated even with much lower subsidy amounts. The large fraction of inframarginal participants means that the larger subsidy amounts are almost certainly not cost-effective. Moreover, the results imply that about half of all participants would have adopted the energy-efficient technology even with no subsidy whatsoever.

1. Introduction

Global energy consumption is forecast to increase 56% by 2040. While the energy mix is becoming somewhat less carbon-intensive, carbon dioxide emissions are still forecast to increase by 45% over the same period.[1] There is a wide agreement among economists that the best policy to reduce carbon dioxide emissions and other negative externalities from energy use would be a Pigouvian tax. Although there has been some recent progress, the vast majority of carbon dioxide emissions worldwide remain untaxed and there are many countries, including the United States, where it seems unlikely that there will be large-scale carbon policy in the near term.

Instead what is receiving much attention is energy efficiency. Electric utilities in the United States, for example, spent $34 billion on energy-efficiency programs between 1994 and 2012.[2] Energy-efficiency measures like appliance replacement, industrial process changes, and weatherization have the potential to greatly reduce energy consumption (National Academy of Sciences et al., 2010). Proponents of energy-efficiency policies argue that these savings are available at very low cost (McKinsey and Company, 2009). Thus, energy-efficiency policies are promoted as "win–win" policies that reduce both private energy expenditures and the externalities associated with energy use.

Despite all of the resources aimed at energy-efficiency programs, there is a surprisingly small amount of direct evidence evaluating their effectiveness. A recent review paper emphasizes this lack of

☆ We are thankful to seminar participants and discussants at UC Berkeley, the U.S. Environmental Protection Agency, the University of Maryland, Michigan State, the National Bureau of Economic Research, and Stanford University, as well as two anonymous reviewers, for their helpful comments. This research was supported in part under a general research contract from the California Energy Commission (500-08-006) to the Energy Institute at Haas and by a National Science Foundation (DGE 1106400) Graduate Research Fellowship to Boomhower. The authors have not received any financial compensation for this project nor do they have any financial relationships that relate to this research.

* Corresponding author. Tel.: +1 510 642 9590; fax: +1 510 643 5180.
E-mail addresses: boomhower@berkeley.edu (J. Boomhower), ldavis@haas.berkeley.edu (L.W. Davis).

[1] These statistics come from the U.S. DOE, EIA, "International Energy Outlook", released July 2013, Figs. 1 and 10. Global energy consumption increased from 350 quadrillion Btu in 1990 to 520 in 2010, and is forecast to increase to 820 by 2040. Energy-related carbon dioxide emissions increased from 20 billion metric tons in 1990 to 30 billion in 2010, and are forecast to increase to 45 billion by 2040.

[2] U.S. Department of Energy, *Electric Power Annual*, 1995–2013. All dollar amounts in the paper are reported in year 2010 dollars. Spending increased every year from 2004 to 2012, with $4.2 billion in 2012.

http://dx.doi.org/10.1016/j.jpubeco.2014.03.009

evidence and goes on to argue that there is, "great potential for a new body of credible empirical work in this area, both because the questions are so important and because there are significant unexploited opportunities for randomized control trials and quasi-experimental designs that have advanced knowledge in other domains" (Allcott and Greenstone, 2012).

We are particularly interested in the question of additionality. Many energy-efficiency programs work by subsidizing households and firms to adopt energy-efficient technologies. A fundamental question in evaluating the cost-effectiveness of these programs is how many of the participants would have adopted these technologies with a lower subsidy, or even with no subsidy at all. Economists have long argued that many participants in energy-efficiency programs may be non-additional or "free riders" (Joskow and Marron, 1992), but demonstrating this empirically has been difficult.[3]

Determining the causal relationship between subsidies and technology adoption is challenging because one must construct a credible counterfactual for adoption in the absence of the policy. Cross-sectional comparisons are misleading because places with generous subsidies are different from places with less generous subsidies. For example, "green" communities like Berkeley, California have more generous subsidy programs but also more eager adopters. Similarly, although programs change over time, it is difficult to separate the causal effect of these changes from other time-varying factors. Changes over time in energy-efficiency subsidies are correlated with changes in technology, pricing, and consumer preferences.

In this paper we address these challenges using a regression discontinuity (RD) analysis. Many energy-efficiency programs have eligibility cutoffs and our paper illustrates how these thresholds can be used to measure inframarginal participation. We apply this approach to a national appliance replacement program in Mexico. We first examine the eligibility thresholds carefully, demonstrating clear discontinuous changes in subsidy amounts and testing for manipulation of the running variable. We then turn to the main analysis, finding that program participation increases noticeably with larger subsidy amounts. For example, when a refrigerator subsidy increases from \$30 to \$110 (both in U.S. 2010 dollars), the number of participants increases by 34%. Thus, the participation elasticity is substantial. However, it is also evident that there are a large number of inframarginal participants. At this threshold, for example, our estimates indicate that about 75% of households would have participated in the program even with the lower subsidy amount. For the four main thresholds in our analysis we find that 65%+ of households are inframarginal. This large fraction of inframarginal households means that the larger subsidy amounts are almost certainly not cost-effective because each actual increased participant costs a large amount in additional program funds.

We next use the observed changes in demand at these four thresholds to infer what fraction of participants would have participated with no subsidy whatsoever. Under reasonable assumptions, the estimates imply that about half of all participants would have replaced their appliances with no subsidy. We then discuss the implications of non-additionality for cost-effectiveness and welfare. These non-additional participants add cost to the program without yielding any actual reductions in energy consumption. When the marginal cost of public funds is larger than one or when there are indirect program costs then it does not make sense to think of these payments as pure transfers. Our results also demonstrate the potential for cost savings if program designers can target subsidies towards groups where the number of likely non-additional participants is low.

Our paper is the first that we are aware of to use RD to study participation in an energy-efficiency program. We see broad potential for applying this approach in evaluating similar programs. Although eligibility requirements vary widely across programs, the desire to simplify program design often results in the kind of discrete thresholds that we exploit here.[4] In addition, energy consumption is typically carefully measured for large numbers of participants and non-participants. Both of these features make RD a natural approach for causal inference in this context. Relative to the alternative of randomized control trials (RCTs), RD is limited by its focus on specific thresholds. However, RD is easier and less expensive. In addition, RD analyses with administrative datasets have more power and thus can measure smaller effects than typical RCTs.

Most previous studies of additionality in similar programs have been of a much smaller scale (see, e.g., Hartman, 1988), or based on stated-choice experiments (Revelt and Train, 1998; Grosche and Vance, 2009; Bennear et al., 2013). Several related papers look at the impact of subsidies on adoption of energy-efficient vehicles (Chandra et al., 2010; Gallagher and Muehlegger, 2011; Sallee, 2011; Mian and Sufi, 2012). There is also a small literature which addresses additionality indirectly by comparing realized aggregate savings at the utility level to engineering estimates (Loughran and Kulick, 2004; Auffhammer et al., 2008; Arimura et al., 2012). Our paper differs from all of these previous studies because of the RD research design. Probably the closest existing study is Ito (2013), which uses an RD analysis to examine a California policy that paid households to reduce their electricity consumption in Summer 2005.

The paper is also related to a broader literature that examines government programs that subsidize socially-beneficial behavior. A key issue with these programs is the need to distinguish between additional and non-additional participants. Examples include tax subsidies for charitable giving (Feldstein and Clotfelter, 1976), subsidies for building low-income housing (Sinai and Waldfogel, 2005), conditional cash transfer programs (De Janvry and Sadoulet, 2006), pollution offset programs (Schneider, 2007), and environmental conservation programs (Sánchez-Azofeifa et al., 2007).[5]

2. Conceptual framework

2.1. Technology adoption with externalities

In this section we propose a simple framework for thinking about the costs and benefits of energy-efficiency subsidies. We illustrate the welfare loss introduced by transfers to inframarginal participants and show how the optimal subsidy amount depends on the relative shares of marginal and inframarginal participants. We focus on the adoption of an energy-efficient technology, but the same basic framework applies to many other types of government programs that subsidize socially-beneficial behavior.

We begin with a simple graphical partial equilibrium analysis. Fig. 1 describes the market for an energy-efficient technology. Along the x-axis is the number of adopters. Demand is given by the

[3] The term "free rider" has long been used in the context of energy-efficiency programs to describe participants who receive a subsidy for doing something they would have done anyway. This is distinct from the use of the term in economics. The well-known "free rider problem" in economics is that individuals underinvest in public goods because they do not internalize the benefits to others. To avoid confusion we use the term "non-additional" throughout the paper.

[4] For instance, the two largest utilities in California offer rebates for energy-efficient heating and cooling equipment that vary across 16 climate zones. These zones were established by California law in 1978 as a function of climate characteristics. Cities can straddle multiple climate zones, and there are large discontinuous changes in rebates at climate zone boundaries. For example, during 2013 Southern California Edison offered three different subsidy amounts (\$550, \$850, and \$1100) for central air conditioners. Other eligibility thresholds that would be amenable to RD analyses include requirements about the vintage of the home, size or characteristics of the households' current equipment, and, for need-based programs, household income.

[5] In this broader literature there are a few studies that use RD. Baum-Snow and Marion (2009) examine the effect of tax credits for building low-income housing, exploiting a discontinuous increase in the credit amount in census tracts where more than 50% of households qualify for means-tested government housing assistance. Filmer and Schady (2011) study a conditional cash transfer program in Cambodia where program eligibility is limited to households scoring below a specified level on a government poverty index.

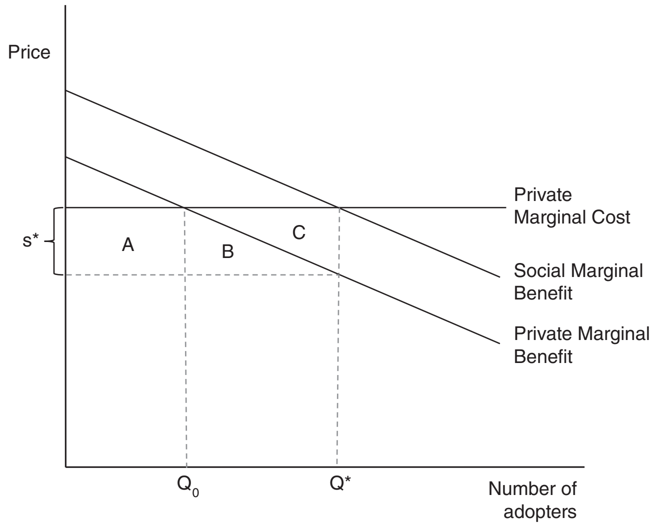

Fig. 1. The market for an energy-efficient technology.

downward-sloping private marginal benefit curve. The benefits of adoption vary across potential adopters due to differences in expected utilization and other factors. Supply is described by the private marginal cost curve.

The privately optimal level of adoption is labeled in the figure as Q_0. These consumers adopt the technology purely on the basis of private benefits, even with no subsidy or other form of government intervention. If there are no externalities, then Q_0 is socially optimal. Once externalities are introduced, however, this is no longer the case. The figure illustrates the case in which there is a positive marginal external benefit from adoption, so the social marginal benefit exceeds private marginal benefit. The socially optimal level of adoption is labeled in the figure as Q^*. This optimum is defined as the intersection of the social marginal benefit and private marginal cost curves. The optimal subsidy is s^*. With this subsidy, adopters between Q_0 and Q^* are additional. They adopt under the subsidy and do not adopt without it.

The total amount paid in subsidies is indicated by the rectangle A/B/C. Rectangle A is a transfer to non-additional participants, i.e. consumers who would have adopted the energy-efficient technology even with no subsidy whatsoever. Triangle B is excess payment to consumers who are induced to adopt because of the subsidy. Most adopters receive a subsidy that is more than the minimum amount necessary to induce them to adopt. And triangle C is the payment required to make adopters between Q_0 and Q^* indifferent between adopting and not adopting.

Before proceeding it is worth highlighting a couple of important assumptions. First, we have assumed that the external benefits from adoption are the same for all potential adopters. When external benefits differ there can be gains from targeting energy conservation policies towards high value participants (Allcott and Mullainathan, 2014; Allcott et al., 2014). We have also assumed constant marginal costs. With increasing marginal costs the analysis is similar but the incidence of the subsidy is partly on sellers. As a result there are "non-additional recipients" on both sides of the market. Subsidies increase the equilibrium price of the good, leading to higher revenues for sellers even for transactions which would have occurred anyway.

2.2. Incorporating pre-existing taxes and other distortions

This partial equilibrium analysis ignores interactions with taxes and other pre-existing distortions. Consider the following welfare function,

$$W = U(Q(s)) - C(Q(s)) + \tau Q(s) + Q(s)s - \eta Q(s)s. \tag{1}$$

Here $Q(s)$ is the quantity of technology adoption, which is a weakly increasing function of the subsidy s. $U(\cdot)$ and $C(\cdot)$ are private benefits and costs from the energy-efficient technology. In the graphical analysis, these correspond to the areas under the private marginal benefit and private marginal cost curves to the left of Q. τ is the constant external benefit of technology adoption derived from, for example, reduced carbon dioxide emissions.

The final two terms reflect general equilibrium effects. The subsidy payments are a transfer from taxpayers to adopters in the amount $Q(s)s$. The efficiency cost of interactions with pre-existing distortions is denoted η. If η is one, then there is no efficiency loss associated with the transfers, and the gains by adopters exactly offset the costs to taxpayers.

The welfare change from a marginal increase in the subsidy is given by,

$$\frac{dQ}{ds}\left[U'(Q(s)) - C'(Q(s)) + \tau - (\eta - 1)s\right] - (\eta - 1)Q(s). \tag{2}$$

The additional adoption induced by the subsidy increase is $\frac{dQ}{ds}$. The left-hand term gives the welfare effect of bringing these marginal participants into the program: private marginal benefits minus private marginal costs, plus external benefits, minus the efficiency cost of financing the subsidy payments to new participants. The right-hand term gives the welfare cost of increased payments to inframarginal participants: The $Q(s)$ participants already adopting the technology each receive an infinitesimal increase in subsidy payment, financed at a cost of η. The welfare effects of increasing the subsidy depend on the relative numbers of marginal and inframarginal participants. If $\frac{dQ}{ds}$ is large relative to $Q(s)$, then the left-hand term matters more than the right-hand term. As the subsidy level increases, $Q(s)$ becomes larger and payments to inframarginal participants become more and more important.

If η is equal to one then Eq. (2) simplifies considerably and it is optimal to set the subsidy equal to marginal external benefits (τ). This is exactly what we described in Fig. 1 with s^* and Q^*. However, if η is greater than one then the optimal subsidy level is below marginal external benefits. The optimal subsidy amount balances the benefits of increased adoption with the full welfare costs, including the general equilibrium efficiency costs of larger transfers.

The value of η is informed by a large literature on the general equilibrium effects of environmental taxes and subsidies, which we quickly summarize here. See, e.g., Bovenberg and Goulder (2002) and references therein. These policies create "tax interaction" and "revenue recycling" effects. Environmental taxes exacerbate pre-existing distortions in the economy, for example, by further decreasing the real wage in the presence of a labor tax (the tax interaction effect). At the same time, environmental taxes also generate revenues, allowing labor and other distortionary taxes to be lower than they would be otherwise (the revenue recycling effect). A series of analytical and numerical studies have concluded that, for taxes, tax interaction is more important than revenue recycling, so that optimal tax rates on externalities are generally below marginal damages (Bovenberg and Goulder, 1996, 2002). One reason for this is that environmental taxes discourage consumption of the taxed good, which erodes the tax base and undermines revenue recycling.

A symmetric set of results holds for environmental subsidies. Subsidies for "clean" goods increase real wages, thereby *decreasing* the distortionary effects of labor taxes (the tax interaction effect). However, subsidies also require labor and other taxes to be *higher* than they would be otherwise, exacerbating distortions (the "revenue-financing" effect). In our framework, η represents the net of these two effects. The literature suggests that, for subsidies, the revenue-financing effect exceeds the tax-interaction effect (Parry, 1998). Thus, η is greater than one and the optimal level of subsidy is positive, but below marginal external benefit, just like the optimal level of tax is positive, but below marginal external damages. Throughout the paper, we refer to η as the efficiency cost of transfers.

Table 1
Subsidy amounts.

Program tier	Direct cash payments	Maximum loan amount	Consumption levels, air conditioners	Consumption levels, refrigerators
1	2200 pesos ($170)	3400 pesos ($270)	251–500	76–175
2	1400 pesos ($110)	4200 pesos ($330)	501–750	176–200
3	400 pesos ($30)	5200 pesos ($410)	751–1000	201–250
4	No cash payment	8700 pesos ($690)	1000+	250+

Notes: This table describes the direct cash payments and maximum loan amounts available to households with different levels of average historical electricity consumption (in kilowatt-hours per month). For further details, see Online Appendix. Dollar amounts are reported in U.S. 2010 dollars using the average exchange rate for 2010 (12.645 pesos per dollar). For expositional clarity, we rounded all dollar amounts to the nearest $10. Households with consumption below the first tier were ineligible for subsidies.

The preceding analysis assumes that the government must pay all adopters the same subsidy. Targeting subsidies to different groups based on expected benefits and costs could increase welfare substantially. In the extreme, perfect price discrimination would pay each individual adopter the minimum amount they require to adopt. There would be no payments to inframarginal participants and the right-hand term in Eq. (2) would disappear. However, in practice, equity concerns, imperfect information, and other factors prevent perfect price discrimination and limit group-level targeting.

The main takeaway from this section is that, in general, the optimal subsidy amount is lower than the marginal external benefits. The welfare effects of a subsidy depend critically on the effect of the subsidy on program participation. If $\frac{dQ}{ds}$ is small relative to the number of existing participants, then the benefits from increased adoption will be small relative to the efficiency costs of the payments to inframarginal participants. Accordingly, this is where we focus our attention in the empirical analyses which follow.

3. Program background and construction of dataset

3.1. Background

Our empirical analysis focuses on a large-scale energy-efficiency program in Mexico. The program was launched in March 2009 and ended in December 2012. During this period, the program subsidized the replacement of 1.9 million refrigerators and air conditioners with energy-efficient models. Davis et al. (forthcoming) compare electricity consumption by program participants before and after appliance replacement, finding that realized savings were considerably smaller than what was predicted by ex ante analyses. There was no attempt in this previous work, however, to distinguish between additional and non-additional participants, nor was there any examination of the programs eligibility thresholds or RD analysis.

To participate in the program a household had to have a refrigerator or air conditioner that was at least 10 years old and agree to purchase a new appliance meeting Mexican energy-efficiency standards. The old appliances were transported to recycling facilities and disassembled. The refrigerator subsidies were available nationwide. For the air conditioner program, a household needed to live in one of four officially-designated climate zones with a mean summer temperature of at least 30 °C (86 °F); this included about one-quarter of all households.

Table 1 describes the subsidies available under the program. The direct cash payments came in three different amounts, approximately corresponding to $30, $110, and $170 (all in U.S. 2010 dollars). Eligibility for the different subsidy amounts depended on a household's average historical electricity consumption, calculated over the previous year. There was a minimum consumption level below which households were ineligible for subsidies. Above this minimum, the cash payment amount decreased with a household's consumption level. This structure was designed to target the larger subsidies to lower-income households.

The program also offered on-bill financing at an annual interest rate of 13.8%, repaid over four years. At the first two thresholds, the increase in maximum loan amount exactly equals the decrease in cash. If households would otherwise have financed the purchases using credit cards, the increase in maximum loan amount offsets about 18% of the decrease in cash subsidies at these thresholds.[6] At the highest consumption threshold, the increase in maximum loan amount greatly exceeds the decrease in direct cash payments. For households with a typical cost of borrowing who take the full loan amount, the economic value of the program actually increases at this threshold.

There are several features of this program that make it particularly conducive to an empirical analysis. First, retailers did not have any discretion in assigning subsidy amounts. Participating retailers determined which subsidy a household was eligible for by entering the household's account number into a website designed for this purpose, and this online record became part of the paperwork necessary for the retailer to be reimbursed. This lack of scope for retailer discretion is important because even a small amount of selection at these thresholds would have been a threat to our identification strategy.

Another nice feature of the program is that participants received these subsidies immediately. In order to participate, a household was required to show a recent electricity bill and an identification card, but there was no paperwork required and no delay in receiving the subsidy. This differs from appliance subsidy programs in the United States which typically require participants to fill out and mail application forms and proofs of purchase, and then wait for a rebate check to arrive in the mail. In programs for which there are "hassle" costs like these, not all eligible households will participate. And the amount of selection depends on the size of the subsidy, making it difficult to interpret differences in participation across subsidy levels.

3.2. Construction of the dataset

A key feature of our analysis is the use of high-quality, household-level microdata, both about the program participants and about the entire pool of potential participants. The fact that we observe eligibility for non-participants is important because, as usual, the objective in the empirical analysis is to construct a credible counterfactual, and this is hard to do without information about the broader pool.

The first component of this database is a two-year panel dataset of household-level electric billing records describing bimonthly electricity consumption for the universe of Mexican residential customers from May 2009 through April 2011. The complete set of billing records

[6] According to Banco de Mexico, "Indicadores Básicos de Tarjeta de Crédito" October 2012, , credit cards in Mexico charged an average interest rate of 25.3% in 2011. This is not a perfect measure. On the one hand, not all households have access to credit cards, and the interest rates on other forms of borrowing will vary. On the other hand, collateralized loans for durable goods purchases typically can be made at lower rates. Most participants took out at least some loans, suggesting that the market cost of borrowing exceeds 13.8% in most cases.

includes data from 25,786,609 households. This represents the entire pool of potential participants in the program.

The second component of this database is a record of all households who participated in the program between March 2009 and June 2011. In the complete dataset there are a total of 1,162,775 participants. We merged this list with the electric billing records using customer account numbers. We used our database to calculate average historical electricity consumption for each household according to the program rules. For details see the Online Appendix.

We focus on the 237,552 participants in 2011 because calculating average historical electricity consumption for earlier participants would require data from before May 2009, the first month in our billing records. For each participant, we know the exact dates of purchase and replacement, whether the appliance was a refrigerator or an air conditioner, and the amount of direct cash payment and loan received.

4. Empirical strategy

4.1. Estimating equation

Our empirical strategy exploits the discrete eligibility thresholds that determined whether a household was eligible for zero subsidy, \$30, \$110, or \$170. There are six total thresholds; three for air conditioners and three for refrigerators. At each of these thresholds, we use a standard RD estimating equation (Lee and Lemieux, 2010):

$$1[Participate]_i = \alpha + f(X_i) + \rho 1[Below\ Threshold]_i + \eta_i \tag{3}$$

where $1[Participate]_i$ is an indicator variable equal to one if a household participated in the program and zero otherwise. We include in the regression $f(X_i)$, a polynomial in average historical electricity consumption, and $1[Below\ Threshold]_i$ an indicator variable equal to one if the household's average historical electricity consumption was below the given threshold. The coefficient of interest is ρ, which measures the discontinuous change in program participation at the threshold. Moreover, we normalize X_i to be equal to zero at the threshold so the coefficient α corresponds to the predicted probability of participating just below the threshold, and $\alpha + \rho$ corresponds to the predicted probability just above the threshold. In terms of the conceptual framework described in Section 2, ρ is the empirical analog of $\frac{dQ}{ds}$, and α is the empirical analog of $Q(s)$.

The error term η_i captures unobserved determinants of the participation decision. An important advantage of RD is that it requires a considerably weaker identifying assumption than other approaches. Hahn et al. (2001) show that identification with RD requires that the conditional mean function $E[\eta_i|X_i]$ is continuous at the discontinuity. In the limit, one is comparing outcomes within an arbitrarily small neighborhood around each threshold and the identifying assumption requires only that there not be a discontinuous change in these other factors that occurs exactly at the eligibility thresholds. Of course, in practice there are few observations within an arbitrarily small neighborhood around these thresholds, and so there is a trade-off between bias and efficiency. Flexibly parameterizing the polynomial $f(X_i)$, allows us to expand the sample to include households farther away from the threshold.

We report results using several different bandwidths. In our preferred specification, we include all households within 100 kWh of the thresholds for air conditioners, and within 50 kWh of the thresholds for refrigerators. The wider bandwidth for air conditioners reflects that these thresholds were much higher (500, 750, and 1000 kWh compared to 175, 200, and 250) and the density of households in that part of the distribution is lower. With refrigerators, the thresholds are close enough together that, in some cases, the bandwidth includes more than one threshold. In the results which follow we use one estimating equation per threshold, but we include intercept terms for any additional thresholds.

4.2. Validity of research design

4.2.1. The discontinuity in subsidy amounts

Fig. 2 plots the fraction of participants who received the larger subsidy as a function of average historical electricity consumption. The dots represent mean values for three kilowatt-hour bins. In all six cases there is a clear discontinuity at the threshold. Almost all households with average historical consumption below the threshold receive the higher subsidy amount and almost all households with average historical consumption above the threshold receive the lower subsidy amount. Fig. 2a is typical of all three air conditioner thresholds. The share of participants receiving the larger subsidy falls from near one to near zero. Even within very narrow bandwidths around these thresholds, we are able to correctly predict subsidy levels for 99%+ of all participants (see the Online Appendix for details).

The discontinuities are less sharp for refrigerators. Fig. 2d is typical of the three refrigerator thresholds. Near the threshold, a small number of participants receive a different subsidy than we would have predicted. This is due to measurement error in our reconstruction of average historical electricity consumption. As we explain in more detail in the Online Appendix, the program rules for refrigerators were especially complicated, introducing a small amount of measurement error for some observations. Battistin et al. (2009) show that this type of measurement error biases sharp RD estimates downward in proportion to the fraction of observations measured with error, but that the fuzzy RD estimator is unbiased as long as the measurement error is uncorrelated with the subsidy amount. In practice, because a small share of observations is measured with error, sharp RD and fuzzy RD produce very similar estimates.

4.2.2. Checking for manipulation of the running variable

A standard concern with RD analyses is manipulation of the running variable. If participants could completely or partially manipulate their treatment status, this would represent a substantial threat to the identifying assumption. Understanding any strategic behavior in response to eligibility thresholds is also of significant independent interest because it may introduce inefficiencies, as agents alter their behavior to qualify for more generous subsidies (Sallee and Slemrod, 2012).

Fig. 3 plots the frequency distribution of average historical electricity consumption for all households. We use three kilowatt-hour bins and include separate plots for air conditioners and refrigerators because the measure of average historical electricity consumption used to determine eligibility was different for the two appliance types. Examining the smoothness of the running variable is a valuable first test for manipulation (McCrary, 2008). If households were changing their behavior to qualify for the more generous subsidy, we would expect to see bunching to the left of the thresholds. For both appliance types, the frequency distributions appear smooth across all eligibility thresholds. This lack of evidence of manipulation is perhaps not surprising given that it is difficult for a household to control its average historical electricity consumption. Perhaps most importantly, this is *historical* consumption, so at the time of participating in the program, there is no scope for the household to go back and change its electricity consumption patterns in the past.

Another standard RD specification test is to look for changes at the threshold in covariates unrelated to the treatment variable. If manipulation of the running variable leads to systematic sorting of households around the threshold, we would expect to see discontinuous differences in household characteristics at the threshold. In our dataset, we do not have any household-level covariates. Instead, we merged our dataset with municipality-average household income from the 2010 Census. Fig. 4 shows that municipality-average income is smooth across all eligibility thresholds, suggesting that there is no discontinuous change at the threshold in the affluence of the places where participants live.

Finally, we consider a more subtle form of strategic behavior. If a household somehow learned that they just missed qualifying for a larger subsidy, they could in theory wait one or more billing cycles, perhaps

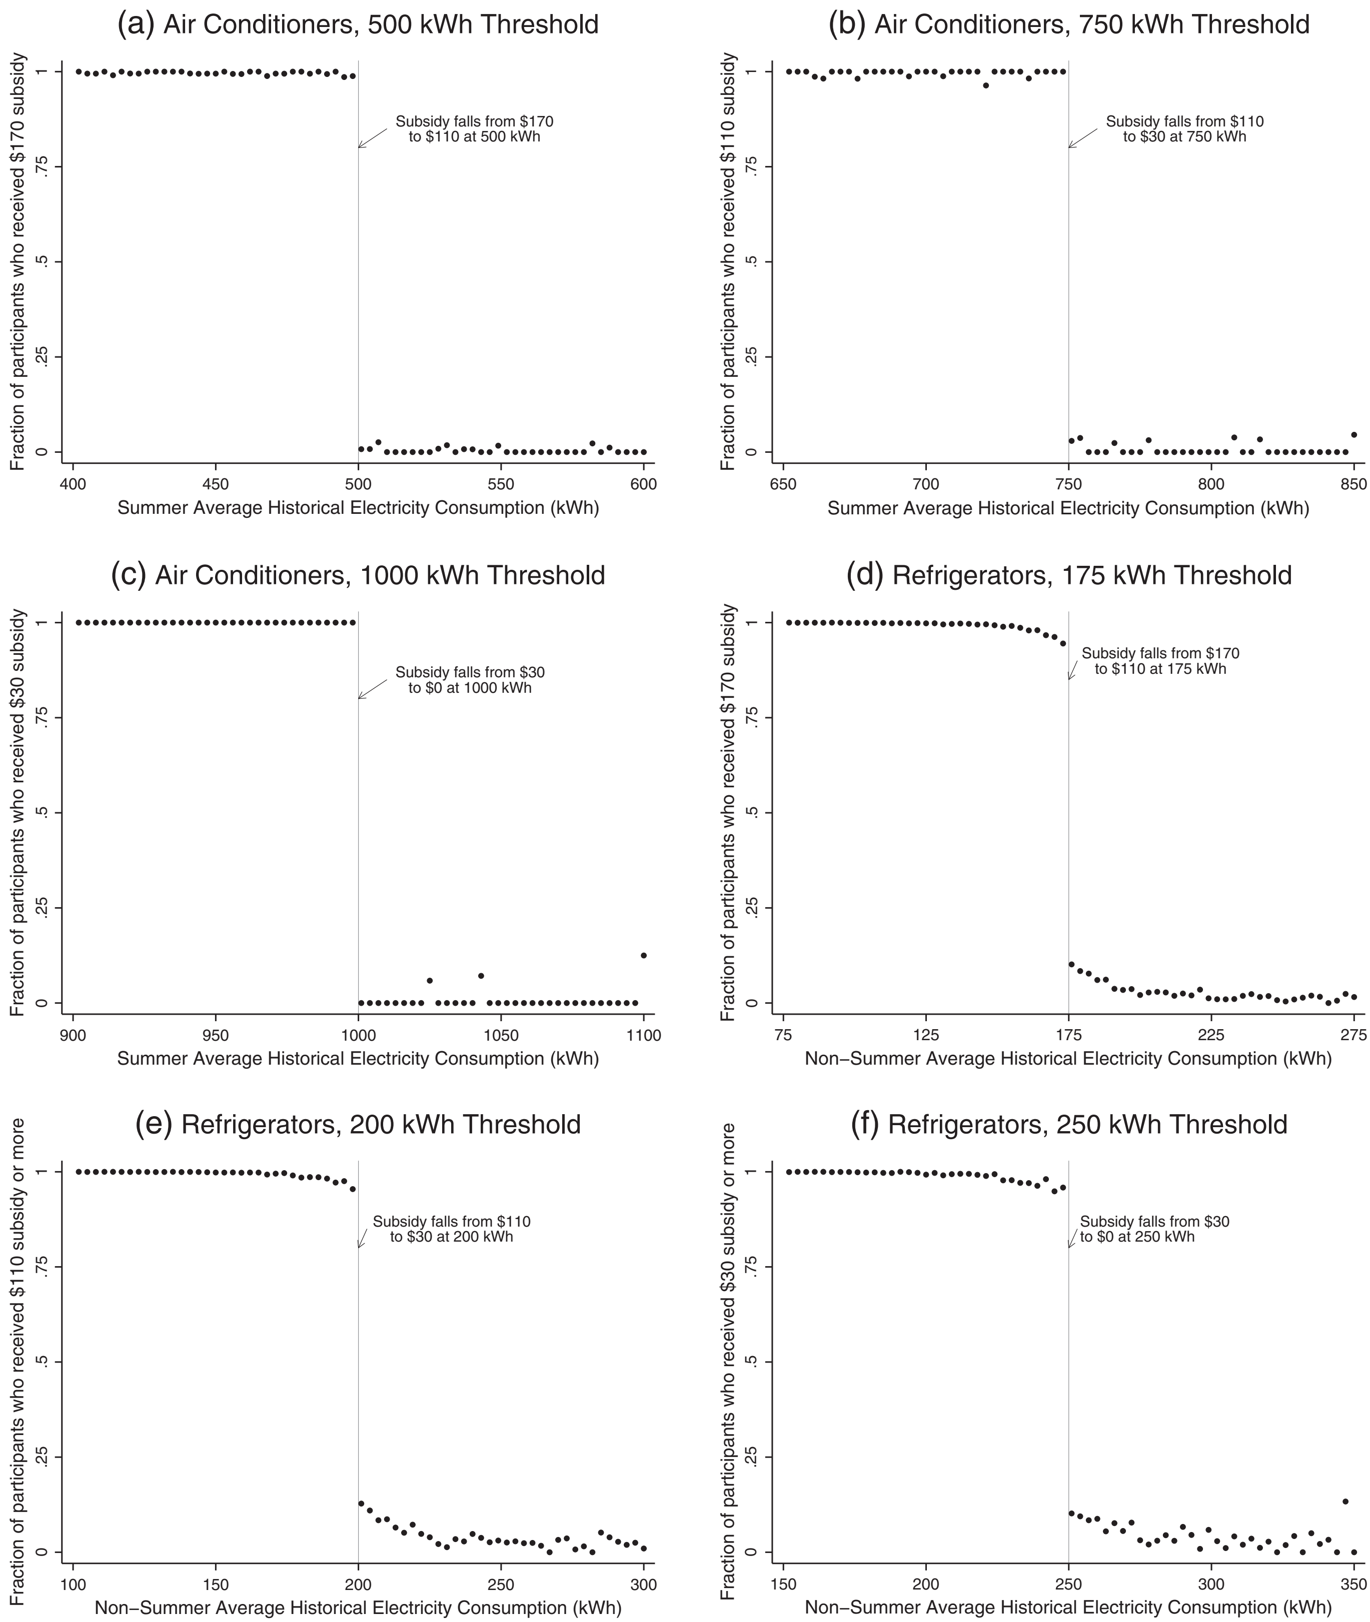

Fig. 2. The discontinuities.

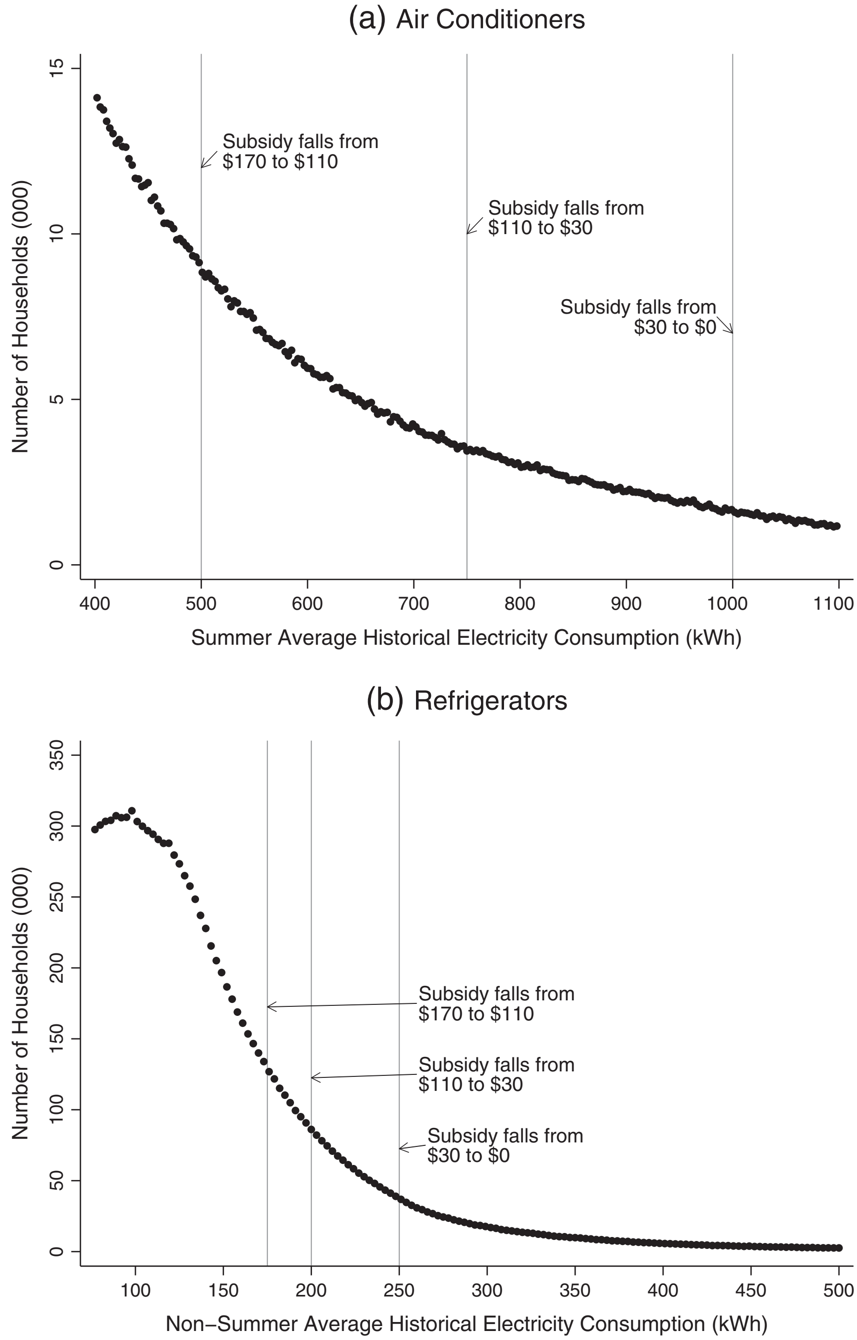

Fig. 3. Smoothness of running variable across subsidy thresholds.

while intentionally reducing electricity consumption, and then reapply. In practice the program structure made this unlikely.[7] Moreover, in the Online Appendix we test for this explicitly by comparing participants' historic average electricity consumption in the months before participation to nonparticipants' historic average electricity consumption over the same months. We find no evidence that participants were more likely than nonparticipants to become eligible for larger subsidies immediately before participating. Thus, there is no evidence of strategic delay.

[7] Participating retailers determined whether a household was eligible by entering the household's account number into a website designed for this purpose. Households could not access this site without a retailer's login and password. The website reported the subsidy level for which a household is qualified, but did not describe the intermediate calculations which determined eligibility or let a household know when it was close to a more generous subsidy level.

5. Results

5.1. Graphical evidence

We now turn to our main results, first presenting graphical evidence and then reporting regression estimates in Section 5.2 and alternative specifications in Section 5.3. Fig. 5a and b plots program participation against average historical electricity consumption for air conditioners and refrigerators, respectively. We again use three kilowatt-hour usage bins. The y-axis in these figures is the percentage of households in each bin that participated in the program during our sample period. For refrigerators the denominator is all Mexican households. For air-conditioners the denominator is all Mexican households living in climate zones that were eligible for the air conditioner program.

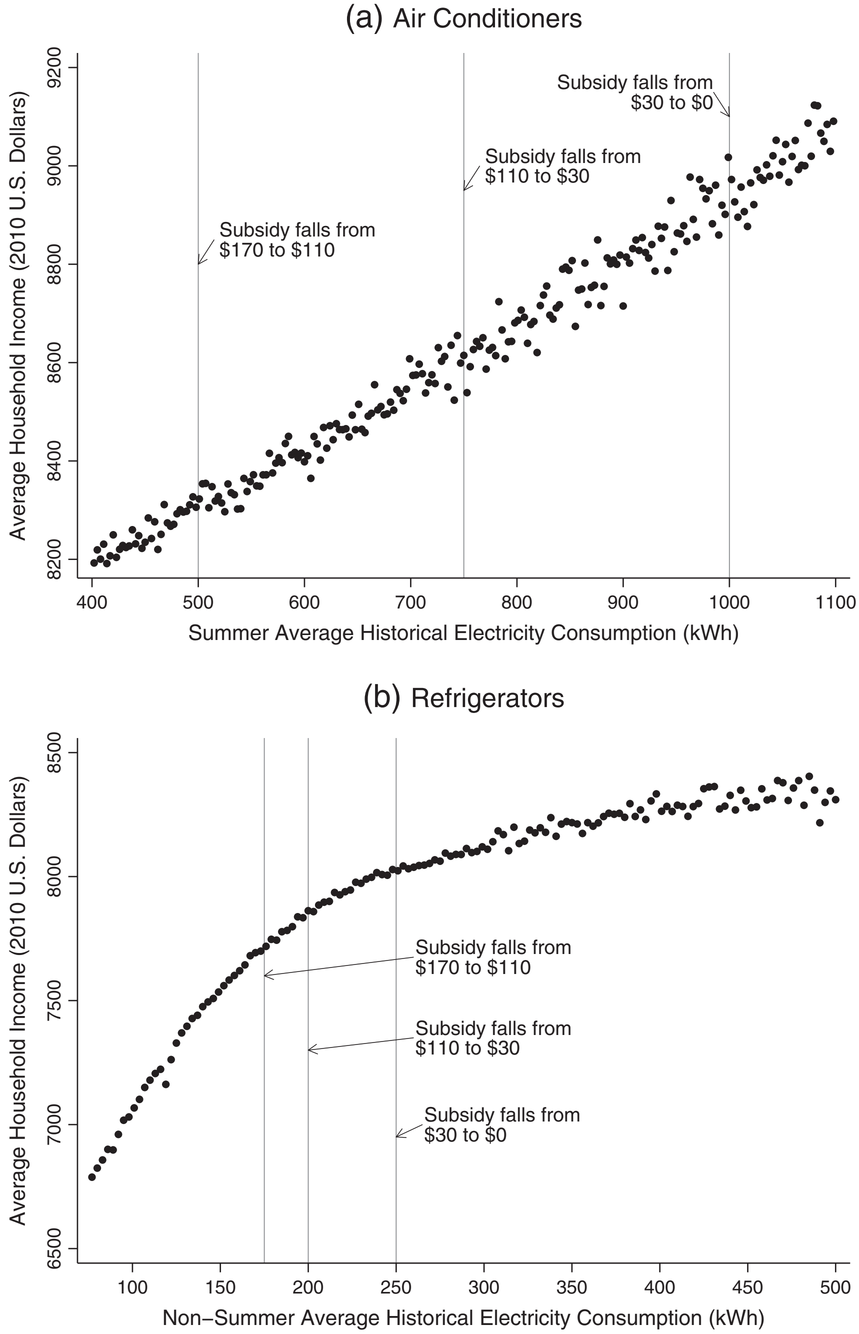

Fig. 4. Smoothness of household income across subsidy thresholds.

It is first worth noting that there is essentially no participation by households who used less than the minimum levels of electricity required for participation. This is reassuring, though not surprising given the way the program was administered. The small number of participating households to the left of the minimum eligibility thresholds for refrigerators reflects a small amount of measurement error in average historical electricity consumption.

For air conditioners, participation increases steadily between 250 and 500 kWh, levels off between 500 and 750, and then declines slowly after 750. Our main interest is in behavior at the 500, 750, and 1000 kilowatt-hour thresholds. In the first two cases, there appears to be a discontinuous decrease in participation at the threshold. The second decrease is particularly visible and appears to occur exactly at the threshold in which the subsidy amount decreases from $110 to $30. It is difficult to make strong statements based on this graphical evidence because the participation rate moves around across bins, but at this threshold the participation rate appears to drop from about 1.5% to about 1%. At the final threshold, where the cash subsidy amount falls from $30 to zero, there does not appear to be any discontinuous change in participation.

For refrigerators, participation follows a similar inverted "U" pattern, peaking at about 1.8% near 150 kWh and then decreasing steadily between 150 and 300. At both the 175 and 200 kilowatt-hour thresholds there are visible discontinuous decreases in participation. At the 250 kilowatt-hour threshold there is no apparent decrease. This general pattern is similar to what is observed for air conditioners, with decreases at the first two thresholds and no visible decrease at the third threshold.

For both appliance types, there is no observed change in participation when the subsidy falls from $30 to $0. As we discussed in

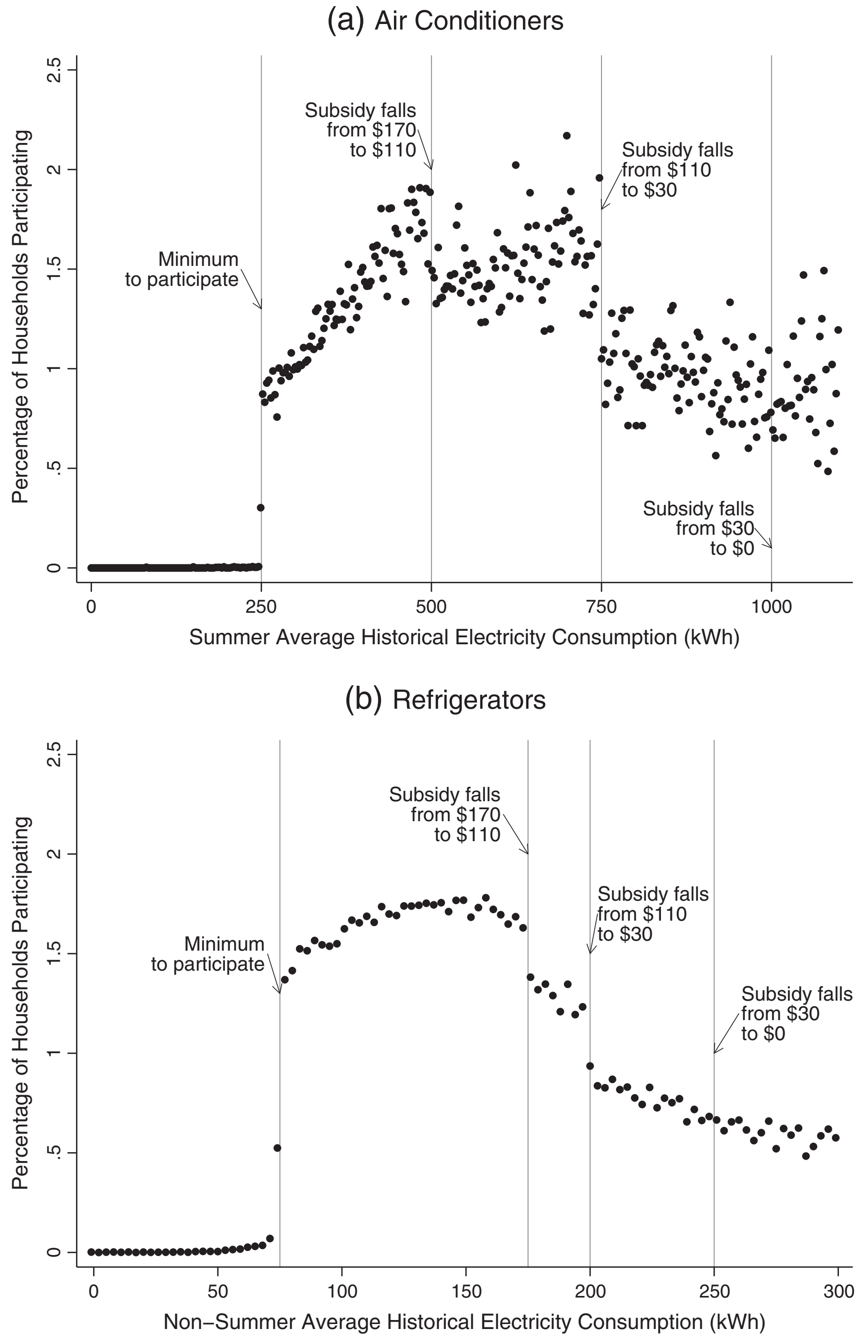

Fig. 5. Program participation.

Section 3.1, this threshold was different from the others in that there was a large offsetting increase in the maximum loan amount. We were expecting to see a much smaller change in participation at this threshold, and the data appear to bear this out. The near zero change in participation implies that, on average, the increase in maximum loan amount had about the same value to households as the $30 decrease in cash. We find this very interesting, but in the regression analysis which follows we focus on the four other thresholds where there is a clear and unambiguous change in the value of the program.

5.2. Regression estimates

Table 2 reports RD estimates and standard errors from four separate regressions. For each threshold we report the percentage of households participating at each side of the threshold as well as the percent change in participation. Because we have normalized the running variable to be equal to zero at the threshold, these statistics come right out of our estimating equation. From each regression, column (2) reports the estimated intercept, column (3) reports our estimate of the intercept plus our estimate of the discontinuous change at the threshold, and column (4) reports the percent change between the two. Columns (5) and (6) report the implied linear slope of demand and price elasticity at each threshold.

Consistent with the graphical evidence, participation increases at all four thresholds. All four changes are statistically significant (three at the 1% level, one at the 5% level). With air conditioners, the increases are 21% and 45%. For refrigerators the estimated changes in participation are similar, 19% and 34%. As with air conditioners, the larger increase corresponds to the subsidy increase from $30 to $110. The estimates

Table 2
RD estimates of the effect of subsidies on program participation.

Subsidy increase	Percent of households participating		Percent change in participation at the threshold	Implied slope	Implied price elasticity
	At lower subsidy amount	At higher subsidy amount			
(1)	(2)	(3)	(4)	(5)	(6)
Panel A. air conditioners					
$110 to $170	1.45	1.75	20.6	0.0054	0.88
	(0.23)	(0.29)	(8.7)	(0.0023)	(0.38)
$30 to $110	1.07	1.55	44.6	0.0069	1.76
	(0.23)	(0.31)	(11.6)	(0.0019)	(0.49)
Panel B. refrigerators					
$110 to $170	1.37	1.63	19.1	0.0047	0.90
	(0.11)	(0.13)	(2.7)	(0.0007)	(0.13)
$30 to $110	0.89	1.20	34.1	0.0044	1.51
	(0.08)	(0.10)	(4.7)	(0.0006)	(0.21)

Notes: This table reports sharp RD estimates of the effect of increased subsidies on program participation from four separate regressions. In each regression, the sample includes all households within our preferred bandwidth. We use a 100 kWh bandwidth for air conditioners, and a 50 kWh bandwidth for refrigerators. All regressions include a cubic polynomial in average historical electricity consumption, normalized to zero at the threshold. Column 2 reports the estimated intercept. Column 3 reports the estimated intercept plus the estimated coefficient on an indicator variable equal to one for households below the eligibility threshold. Column 4 reports the percent change between the previous two columns. Column 5 reports the change in the percent of households participating (as a fraction of all households at the threshold) per dollar of subsidy change. Column 6 reports the implied price elasticities evaluated using the net change in the cost of replacement at each threshold. Standard errors are clustered at the municipality level.

for refrigerators are more precisely estimated because of the large number of households with average historical electricity consumption near these thresholds.

The increases are clear, but the estimates also imply that a large number of participants are inframarginal. Most households who just barely qualified for the $170 subsidy would have participated even if they had only received $110 and most households who just barely qualified for the $110 subsidy would have participated even if they had only received $30. The percent inframarginal can be calculated by dividing column (2) by column (3). For example, when the air conditioner subsidy increases from $110 to $170, our estimates imply that 83% ($\frac{1.45}{1.75} = 0.83$) of households are inframarginal. Across thresholds the percentage inframarginal ranges from 69% to 84%. The estimates are similar for air conditioners and refrigerators, suggesting that the results are not driven by idiosyncratic features of a particular appliance market.

In column (5) we report the implied slope of demand at the threshold. These are calculated for each threshold by dividing the percent change in participation by the subsidy change in dollars. For each $1 of subsidy change, the share of households replacing air conditioners increases 0.0054 to 0.0069 percentage points and the share of households replacing refrigerators increases 0.0044 to 0.0047.[8] These slopes appear quite small but it is important to keep in mind that the base participation rates are very low.

In column (6) we report the implied price elasticities. These are calculated for each threshold by dividing the percent change in participation by the percent change in the price of appliance replacement net of the subsidy. In calculating this net price we use the average appliance price paid by program participants at the threshold.[9] For air conditioners, the elasticity is 0.88 at the first threshold and 1.76 at the second. For refrigerators, the elasticity is 0.90 at the first threshold and 1.51 at the second.[10]

It is important to interpret these elasticities carefully. They describe how demand would change in response to a market-wide price change,

Table 3
Alternative bandwidths and specifications.

Subsidy increase	Panel A: air conditioners Sharp RD				Fuzzy RD
	Cubic Polynomial 125 kWh	Cubic Polynomial 100 kWh	Cubic Polynomial 75 kWh	Local Linear 50 kWh	Cubic Polynomial 100 kWh
$110 to $170	20.7	20.6	29.4	30.7	20.9
	(8.4)	(8.7)	(9.8)	(8.9)	(9.1)
$30 to $110	56.1	44.6	38.2	49.4	45.7
	(13.4)	(11.6)	(11.9)	(13.4)	(12.9)
Subsidy increase	**Panel B: refrigerators Sharp RD**				**Fuzzy RD**
	Cubic Polynomial 75 kWh	Cubic Polynomial 50 kWh	Cubic Polynomial 25 kWh	Local Linear 15 kWh	Cubic Polynomial 50 kWh
$110 to $170	22.5	19.1	15.7	15.0	22.6
	(2.3)	(2.7)	(3.2)	(3.2)	(3.1)
$30 to $110	35.2	34.1	37.5	40.6	41.0
	(4.5)	(4.7)	(6.2)	(6.1)	(5.7)

Notes: This table reports the estimated percent increase in program participation from 20 separate regressions, corresponding to the four main eligibility thresholds. The specification and bandwidth used are indicated at the top of each column. The cubic polynomial estimates using a 100 kWh bandwidth for air conditioners, and a 50 kWh bandwidth for refrigerators are identical to our estimates in Table 2. The columns on either side report estimates from two alternative bandwidths. The fourth column reports estimates using local linear regression with a uniform kernel. The final column reports estimates from a fuzzy RD specification which scales the estimated change in participation by the size of the discontinuity at the eligibility threshold. Standard errors in the first four columns are clustered at the municipality level and in the last column are block bootstrap by municipality with 5000 repetitions. See text for details.

[8] For the calculations in columns (5) and (6) we calculate the change in the value of the subsidy incorporating both direct cash payments and the implied cash value of the on-bill financing (assuming a 25.3% annual interest rate on private borrowing; see Section 3.1).

[9] Specifically, we use the average price paid by participants at the low-subsidy side of each threshold. For air conditioners, these prices were $402 at the 175 kWh threshold and $402 at the 200 kWh threshold. For refrigerators the prices were $425 at the 500 kWh threshold and $427 at the 750 kWh threshold. We calculate all elasticities as arc elasticities, and thus use for the denominator in these calculations the midpoint between the high-subsidy and low-subsidy prices.

[10] We are reporting uncompensated elasticities, but compensated elasticities are likely to be very similar. These subsidies represent a tiny share of the total household budget for these households so income effects are likely negligible. One approach for assessing the potential magnitude of income effects is to test for changes at the thresholds in the price of the appliance that is purchased. We observe no significant change in the price of the appliance purchased at three of the four thresholds. At the fourth, the average price increases by about 2%, and the increase is only weakly statistically significant.

Table 4
Inferring the fraction non-additional.

	Projection based on linear demand (1)	Projection based on elasticities (2)
Fraction	53.6%	43.3%
Non-additional	(4.8)	(6.0)
Average payment	$328	$269
Per induced replacement	(36.7)	(30.2)

Notes: In this table we use the RD estimates from the thresholds to infer what fraction of participants would have replaced their appliances with zero subsidy. For the average payment per induced replacement we divide total subsidy payments by the implied total number of additional participants. See text for details.

not the elasticity of demand for a particular appliance model or for all appliances made by a particular manufacturer. It is also worth emphasizing that this is the elasticity of demand for appliance *replacement*, which is different from demand for initial purchase. Still, the estimates appear quite large. They imply that appliance replacement is price-responsive, and that the program caused a large number of appliance replacements that otherwise would not have happened.

These estimates are valuable not only in assessing energy-efficiency subsidies, but also for predicting appliance replacement more broadly. Wolfram et al. (2012) argue that the demand for residential appliances will have an enormous influence on future energy consumption growth in low- and middle-income countries. Appliance prices have been falling for decades and our estimates imply that continued decreases will accelerate the rate at which appliances are replaced. If households are more quickly replacing appliances this means that improvements in energy-efficiency will more quickly be reflected in the appliance stock.

5.3. Alternative specifications

Table 3 reports regression estimates from five alternative specifications. For each specification we report the estimated percent change in participation at each threshold. In the first three columns, we vary the size of the bandwidth used with the cubic polynomial. The second column reports our baseline estimates, identical to the estimates reported in Table 2. The first and third columns assess the sensitivity of our estimates to larger and smaller bandwidths. In the fourth column, we use local linear regression with a uniform kernel and a small bandwidth. Overall, the results are similar across all four columns. Moreover, there is no consistent pattern. As we move across bandwidths and specifications, some point estimates increase while others decrease.

The last column reports estimates from a fuzzy RD specification. In this specification, we scale the estimates by the size of the discontinuity at the threshold following Hahn et al. (2001) and Battistin et al. (2009). Specifically, we run a first stage regression of an indicator for the larger subsidy (1[*Larger Subsidy*]) on (1[*Below Threshold*]) and a cubic polynomial of average historical consumption, $g(X)$,

$$1[\textit{Larger Subsidy}]_i = \phi + g(X_i) + \gamma 1[\textit{Below Threshold}]_i + \epsilon_i. \quad (4)$$

We then divide our baseline estimates by γ to remove any bias caused by measurement error (see Section 4.2.1 and the Online Appendix). The estimates are very similar with the fuzzy RD specification. The air conditioner estimates are essentially identical to the sharp RD estimates, consistent with the near perfect discontinuity observed in Fig. 2a–c. For refrigerators, the scaling increases the point estimates modestly, consistent with the graphical evidence in Fig. 2d–f, which exhibits a somewhat less perfect discontinuity.

6. Discussion

6.1. Inferring the fraction non-additional

These estimates are directly relevant for program design because they show how adjustments in program generosity would have changed participation levels. We are also interested in what program participation would have been with no subsidy whatsoever. Table 4 reports estimates of the fraction of participants that are non-additional under two different assumptions about the shape of the demand curve.

In Column (1) we calculate the fraction of participating households who are non-additional by using the slope estimates from Table 2 to predict appliance replacement at the unsubsidized price. For participants who received the $170 cash payment, we use the slope corresponding to the threshold between $110 and $170, and for participants who received $30 or $110, we use the slope corresponding to the threshold between $30 and $110. The implied slopes are quite similar across thresholds, however, so the results are not particularly sensitive to which slope we use. Under these assumptions our estimates imply that 54% of participants were non-additional, in that they would have replaced their appliances even with no subsidy whatsoever. With this level of non-additionality, the average payment amount per induced replacement is $328, a little more than twice the average subsidy amount ($152).

In Column (2), we infer the fraction of participants that are non-additional by using the estimated *elasticities*, rather than the estimated slopes. With this approach our estimates imply that 43% of participants are non-additional, so that the average payment per induced replacement is $269. Of the two alternatives we prefer to use the estimated slopes because this assumption about the demand curve better fits the observed behavior at the thresholds. Whereas the estimated slopes are similar across thresholds, the estimated elasticities are not. For both appliance types the estimated elasticities are considerably larger at the $30 to $110 threshold than at the $110 to $170 threshold, which is what one would expect with linear demand.

It is worth emphasizing that both of these approaches rely on strong assumptions about demand. By using our estimates of these slopes and elasticities to predict behavior away from the thresholds, we are assuming that behavior at the thresholds is representative of all households. This is a mild assumption for participants who are close to thresholds but is a considerably stronger assumption for participants far away from a threshold like those at the beginning of the $170 tier. We find it somewhat reassuring that the slope estimates are similar across thresholds. Nonetheless, these projections should be viewed with more caution than the RD estimates from which they are derived.

6.2. Implications for cost-effectiveness and welfare

Depending on whether one uses the slopes or the elasticities, we find that 43–54% of participants are non-additional. So it costs on average $269 to $328 in subsidies per induced replacement, instead of $152 in a naive analysis that treats all participants as additional. Thus, accounting for non-additional participants approximately doubles the program cost per unit of reduced energy consumption. Related measures of cost-effectiveness such as the program cost per ton of carbon dioxide abated would also approximately double.

Non-additionality also affects the full welfare calculation. In this section we provide a brief sketch of such a calculation, following the framework outlined in Section 2. To calculate the benefits of appliance replacement, we value the reductions in greenhouse gases and local air pollutants from each appliance replacement. As we explain in detail in the Online Appendix, we use the pre-program engineering estimates of 2900 kWh in lifetime electricity savings per replacement. This is equivalent to 1.6 tons of avoided carbon dioxide emissions so applying a $34 social cost of carbon as in U.S. IAWG (2013), the climate benefits are $53 per replacement. We also include $54 per replacement of benefits from reduced local air pollution.

The program costs can be divided into categories A, B, and C, as indicated in Fig. 1. Rectangle A represents payments to non-additional participants. Private costs are zero for these participants, since they are doing something they would have done anyway. But there is still the efficiency cost associated with financing the subsidies, which we called η in Section 2. This parameter η represents the net welfare cost of the revenue financing and tax interaction effects. For this simple back of the envelope calculation, we assume that η equals 1.3, following Goulder et al. (1997) and other studies in the literature (see the Online Appendix for details). We find that about half of the participants are non-additional and the average subsidy amount was $152. So, transferring funds to non-additional participants imposed an efficiency cost of about $46 per induced replacement.

Rectangle BC represents payments to induced participants. Financing these payments costs an additional $46 per replacement, again using 1.3 for the efficiency cost. In addition, there are the private costs of adoption. These costs are shown as Triangle C. If demand is linear, then this area is half the total amount of subsidies paid to induced participants. The average subsidy amount was $152, so the private cost of replacement averaged $76. Summing up these back-of-the-envelope benefits and costs, we find that each induced replacement yielded benefits of $107 at a cost of $168.

These calculations highlight the importance of distinguishing between additional and non-additional participants. This can be seen most starkly by comparing these numbers to what one would have calculated with a naive analysis that assumes all participants are additional. In the naive analysis, the efficiency cost of financing the program is much smaller: $46 per replacement rather than $92. Private costs are the same ($76), so the total cost is $122 per replacement. This is much closer to the benefits of $107. Thus, in the naive analysis, the program appears much closer to welfare-improving.

These values should be interpreted carefully because they are based on many strong assumptions. These calculations also ignore some components of benefits and costs, such as the benefits of properly disposing of refrigerants and the administrative costs of the program. They also rely on engineering estimates of electricity savings, which recent work suggest may have been overly generous (Davis et al., forthcoming). The goal of this simple back-of-the-envelope calculation is to tie our empirical results to the economic model and provide an example for how to think about welfare analysis in this setting.

Two of the most important uncertain parameters are the social cost of carbon (SCC) and the efficiency cost of financing the subsidies (η). The program becomes more attractive with a high SCC and low η. Still, it would have taken values near the extremes of the range of available estimates in the literature in order to make the program welfare-improving. Holding constant our other assumptions, the program benefits would exceed the costs only if the SCC were greater than $73 per ton or if η were less than 1.1.[11]

These results raise questions about whether energy-efficiency programs could be designed differently to target payments based on expected additionality. For example, if immutable, verifiable household and firm characteristics could be determined to predict adoption with and without the subsidy, payments could be made conditional on these characteristics.[12] The scope for this type of targeting will differ widely across contexts and there are important constraints that may limit targeting in practice. Income-based targeting, for example, can be difficult and expensive to enforce, and geographic targeting may be unacceptable politically. Still, even in programs where explicit targeting is limited, softer versions may be possible. As a simple example, perhaps program advertisements can be tailored towards demographic segments where adoption in the absence of the program would be low.

6.3. Program-wide effects

RD is well-suited for highly-localized predictions about how participation would have changed under alternative subsidies. But we have also stressed that RD is not a panacea and cannot answer all of the questions that could be answered, for example, with a large-scale RCT. A particularly important weakness is the inability of RD to measure broader program-wide effects. In providing subsidies the government is providing information and an explicit endorsement of particular energy-efficient technologies. This focuses attention on these products, potentially influencing replacement decisions above and beyond the direct impact of the subsidies themselves.

Through program-wide effects even non-participants may have their behavior influenced by a program. For example, potential participants may investigate a program only to learn that they are ineligible for a subsidy. However, in learning about the program they focus their attention on energy-efficiency, potentially becoming more likely to adopt the subsidized technology even if they do not end up receiving any monetary incentive whatsoever.

These broader program impacts are difficult to measure empirically. Ideally, one would measure program-wide effects using a large-scale RCT in which randomization was done not over households, but over geographic areas with subsidy generosity varied across areas. This has been done with cash transfer programs (Baird et al., 2012), and given sufficient resources and public cooperation could be implemented with energy-efficiency programs. Experiments could also be designed to directly measure spillovers through social networks, as in Miguel and Kremer (2004) and similar studies.

7. Conclusion

It is hard to provide incentives for socially-beneficial behavior without substantial transfers to those who would have done these behaviors anyway. Subsidies for energy efficiency are a key example, both because the potential external benefits are large and because the first-best policies seem, for the moment, to be impossible politically. Empirical estimates of additionality are critical, however, because if a large enough fraction of participants is non-additional then a program will not be welfare improving.

Our RD analysis avoids many of the measurement and endogeneity problems in previous studies by focusing on behavior within narrow windows around eligibility thresholds. Although these thresholds make RD a natural approach to causal inference, we are not aware of any previous RD analyses of additionality in this context. We see a broad potential for applying our conceptual framework, estimating equations, and tests of strategic behavior in evaluating similar programs. Although the exact eligibility requirements vary across programs, it is typical to see discontinuous thresholds of the type observed here.

The results are striking. We find that most households would have participated even for much lower subsidy amounts. Across thresholds, more than two-thirds of participants are inframarginal and the estimates imply that about half of all participants would have replaced their appliances even with no subsidy whatsoever. These non-additional participants add substantial cost to the program without yielding any real reduction in energy use.

These findings are relevant to current energy policy around the world, which is focusing increasingly on energy efficiency. Billions of dollars are spent each year on programs like this one that provide

[11] An SCC of $73 per ton is above the central range of values presented in U.S. IAWG (2013), but less than the 95th percentile estimate of $129. For η to be less than 1.1, the MCPF would have to be at the bottom of the range of estimates of 1.11 to 1.56 for the United States (Bovenberg and Goulder, 2002), or the tax interaction effect would have to be large.

[12] De Janvry and Sadoulet (2006) propose such targeting for a conditional cash transfer program in Mexico. Using experimental data, they conclude that making school attendance subsidies a function of child's gender, birth order, and distance traveled to school could decrease the program cost per additional child attending school by about 23%.

subsidies for households and firms who adopt energy-efficient technologies. Reliable empirical estimates of the benefits and costs of these policies are essential and using RD to measure changes in behavior at eligibility thresholds can be an important part of these analyses.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.jpubeco.2014.03.009.

References

Allcott, H., Greenstone, M., 2012. Is there an energy efficiency gap? J. Econ. Perspect. 26 (1), 3–28 (September 2012).
Allcott, H., Mullainathan, S., 2014. Site selection bias in program evaluation. Working Paper.
Allcott, H., Mullainathan, S., Taubinsky, D., 2014. Energy policy with externalities and internalities. J. Public Econ. 112, 72–88.
Arimura, T., Li, S., Newell, R., Palmer, K., 2012. Cost-effectiveness of electricity energy efficiency programs. Energy J. 2, 63–99.
Auffhammer, M., Blumstein, C., Fowlie, M., 2008. Demand-side management and energy efficiency revisited. Energy J. 29 (3), 91–104.
Baird, S., Bohren, A., McIntosh, C., Özler, B., 2012. Designing experiments to measure spillover and threshold effects. UC San Diego Working Paper.
Battistin, E., Brugiavini, A., Rettore, E., Weber, G., 2009. The retirement consumption puzzle: evidence from a regression discontinuity approach. Am. Econ. Rev. 99 (5), 2209–2226 (September 2009).
Baum-Snow, N., Marion, J., 2009. The effects of low income housing tax credit developments on neighborhoods. J. Public Econ. 93 (5–6), 654–666.
Bennear, L., Lee, J., Taylor, L., 2013. Municipal rebate programs for environmental retrofits: an evaluation of additionality and cost-effectiveness. J. Policy Anal. Manag. 32 (2), 350–372.
Bovenberg, A., Goulder, L., 1996. Optimal environmental taxation in the presence of other taxes: general-equilibrium analyses. Am. Econ. Rev. 86 (4), 985–1000.
Bovenberg, A., Goulder, L., 2002. Environmental taxation and regulation. Handbook of Public Economics. , 3, pp. 1471–1545.
Chandra, A., Gulati, S., Kandlikar, M., 2010. Green drivers or free riders? An analysis of tax rebates for hybrid vehicles. J. Environ. Econ. Manag. 60 (2), 78–93.
Davis, L., Fuchs, A., Gertler, P., 2014. Cash for coolers: evaluating a large-scale appliance replacement program in Mexico. Am. Econ. J. Econ. Policy (forthcoming).
De Janvry, A., Sadoulet, E., 2006. Making conditional cash transfer programs more efficient: designing for maximum effect of the conditionality. World Bank Econ. Rev. 20 (1), 1–29.
Feldstein, M., Clotfelter, C., 1976. Tax incentives and charitable contributions in the United States: a microeconometric analysis. J. Public Econ. 5 (1–2), 1–26.
Filmer, D., Schady, N., 2011. Does more cash in conditional cash transfer programs always lead to larger impacts on school attendance? J. Dev. Econ. 96 (1), 150–157.
Gallagher, K., Muehlegger, E., 2011. Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology. J. Environ. Econ. Manag. 61 (1), 1–15.
Goulder, Lawrence H., Parry, Ian W.H., Burtraw, Dallas, 1997. Revenue-raising versus other approaches to environmental protection: the critical significance of preexisting tax distortions. RAND J. Econ. 28 (4), 708–731.
Grosche, P., Vance, C., 2009. Willingness to pay for energy conservation and free ridership on subsidization: evidence from Germany. Energy J. 30 (2), 135–154.
Hahn, J., Todd, P., Van der Klaauw, W., 2001. Identification and estimation of treatment effects with a regression-discontinuity design. Econometrica 69 (1), 201–209.
Hartman, R., 1988. Self-selection bias in the evolution of voluntary energy conservation programs. Rev. Econ. Stat. 448–458.
Ito, Koichiro, 2013. Asymmetric incentives in subsidies: evidence from a large-scale electricity rebate program. Working Paper 19485. National Bureau of Economic Research (September 2013).
Joskow, P., Marron, D., 1992. What does a negawatt really cost? Evidence from utility conservation programs. Energy J. 41–74.
Lee, D., Lemieux, T., 2010. Regression discontinuity designs in economics. J. Econ. Lit. 48, 281–355.
Loughran, D., Kulick, J., 2004. Demand-side management and energy efficiency in the United States. Energy J. 19–43.
McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: a density test. J. Econ. 142 (2), 698–714.
McKinsey, Company, et al., 2009. Unlocking Energy Efficiency in the U.S. Economy.
Mian, A., Sufi, A., 2012. The effects of fiscal stimulus: evidence from the 2009 cash for clunkers program. Q. J. Econ. 127 (3), 1107–1142.
Miguel, Edward, Kremer, Michael, 2004. Worms: identifying impacts on education and health in the presence of treatment externalities. Econometrica 72 (1), 159–217.
National Academy of Sciences, National Academy of Engineering, National Research Council, 2010. Real Prospects for Energy Efficiency in the United States. The National Academies Press.
Parry, I., 1998. A second-best analysis of environmental subsidies. Int. Tax Public Financ. 5 (2), 153–170.
Revelt, D., Train, K., 1998. Mixed logit with repeated choices: households' choices of appliance efficiency level. Rev. Econ. Stat. 80 (4), 647–657.
Sallee, J., 2011. The surprising incidence of tax credits for the Toyota Prius. Am. Econ. J. Econ. Policy 3 (2), 189–219.
Sallee, J., Slemrod, J., 2012. Car notches: strategic automaker responses to fuel economy policy. J. Public Econ. 96, 981–999.
Sánchez-Azofeifa, G.A., Pfaff, A., Robalino, J.A., Boomhower, J., 2007. Costa Rica's payment for environmental services program: intention, implementation, and impact. Conserv. Biol. 21 (5), 1165–1173.
Schneider, L., 2007. Is the CDM Fulfilling its Environmental and Sustainable Development Objectives? An Evaluation of the CDM and Options for Improvement.
Sinai, T., Waldfogel, J., 2005. Do low-income housing subsidies increase the occupied housing stock? J. Public Econ. 89 (11–12), 2137–2164.
U.S. IAWG, 2013. Technical support document: technical update on the social cost of carbon for regulatory impact analysis. Technical Report, U.S. Interagency Working Group on the Social Cost of Carbon for Regulatory Impact Analysis May 2013.
Wolfram, C., Shelef, O., Gertler, P., 2012. How will energy demand develop in the developing world? J. Econ. Perspect. 26, 119–138.