informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 67, No. 11, November 2021, pp. 7091–7112
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Social Learning and Solar Photovoltaic Adoption

Kenneth T. Gillingham,[a,b] **Bryan Bollinger**[c]

[a] Yale University, New Haven, Connecticut 06511; [b] National Bureau of Economic Research, Cambridge, Massachusetts 02138; [c] Stern School of Business, New York University, New York, New York 10012

Contact: kenneth.gillingham@yale.edu, https://orcid.org/0000-0002-7329-2660 (KTG); bryan.bollinger@stern.nyu.edu, https://orcid.org/0000-0001-8596-6418 (BB)

Received: December 9, 2019
Revised: June 20, 2020
Accepted: July 26, 2020
Published Online in Articles in Advance: February 17, 2021

https://doi.org/10.1287/mnsc.2020.3840

Abstract. Growing literature points to the effectiveness of leveraging social interactions and nudges to spur adoption of prosocial behaviors. This study investigates a large-scale behavioral intervention designed to actively leverage social learning and peer interactions to encourage adoption of residential solar photovoltaic systems. Municipalities choose a solar installer offering group pricing and undertake an informational campaign driven by volunteer ambassadors. We find a causal treatment effect of 37 installations per municipality from the campaigns and no evidence of harvesting or persistence. The intervention also lowers installation prices. Randomized controlled trials based on the intervention show that selection into the program is important, whereas group pricing is not. Our results suggest that the program provided economies of scale and lowered consumer acquisition costs, leading to low-cost emission reductions.

History: Accepted by Matthew Shum, marketing.
Funding: This material is based on work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) [Solar Energy Technologies Office Awards DE-EE0006128 and DE-EE0007657].
Supplemental Material: The data files and online appendix are available at https://doi.org/10.1287/mnsc.2020.3840.

Keywords: nonprice interventions • social learning • renewable energy • solar photovoltaic panels • technology adoption • natural experiment

1. Introduction

Climate scientists have strongly argued for substantial action to mitigate emissions to reduce the negative consequences of climate change.[1] One path forward is to transition from fossil fuels to renewable energy, a transition that will likely be difficult to achieve without investments in further reducing the costs of renewables. Economists, marketers, and policymakers have also increasingly turned to behavioral *nudges* to encourage socially beneficial actions, such as energy conservation, charitable giving, and healthy eating, which often provide information that includes social comparisons or prosocial appeals (Thaler and Sunstein 2009, Dubé et al. 2016). There is mounting evidence that social interactions themselves, through social learning and peer effects, influence the adoption of new technologies by overcoming *information failures* inherent in their diffusion (Griliches 1957, Foster and Rosenzweig 1995, Kraft-Todd et al. 2018).

This study examines whether an intensive behavioral intervention in the United States designed to leverage the power of social interactions can address this information gap to increase to adoption of a fast-growing renewable energy technology: solar photovoltaic (PV) installations. The "Solarize" program is a community-level campaign with several key pillars. Treated municipalities who receive the intervention choose a single solar PV installer. In order to become this chosen installer, installers submit bids with a discount group price that is offered to all consumers in that municipality during the program. The intervention begins with a kickoff event and involves roughly 20 weeks of community outreach. Notably, the primary outreach is performed by volunteer resident *solar ambassadors* who encourage their neighbors and other community members to adopt solar PV, effectively providing a major nudge toward adoption. This social interaction-based approach parallels previous efforts to use ambassadors as *injection points* into the social network to promote adoption of agricultural technology (Vasilaky and Leonard 2011, BenYishay and Mobarak 2017) and behavior conducive to improving public health (Kremer et al. 2011, Ashraf et al. 2014) in developing countries. A distinguishing feature of our study is that rather than providing a small amount of information from a single source (as in BenYishay and Mobarak 2017), the intervention explicitly aimed to provide information through a variety of channels that complement each other to bolster the effect on the uptake of a nascent technology.

In this paper, we ask several questions that shed light on consumer and market behavior under the influence of a large-scale behavioral intervention. Is such a program effective at increasing adoption of solar PV and lowering installation prices? Do these effects persist after the intervention? Are there spillovers or positive treatment externalities to nearby communities? How cost-effective is the program for meeting policy goals and is it welfare-improving? Finally, what is the role of social learning? These research questions have important relevance to policymakers and practitioners, because similar solar interventions are currently being implemented in many states, and numerous communities have expressed interest in the program to help meet environmental goals.[2] In fact, there is even a program guidebook for stakeholders interested in implementing a Solarize program (Hausman and Condee 2014).

We first establish the effectiveness of the Solarize program by examining the effects on adoptions and prices for municipalities that apply to join the program. The municipalities that apply to join the program are the marginal municipalities that would first join if the program is expanded elsewhere and thus are highly relevant to practitioners who would consider a further expansion. We use a difference-in-difference strategy with rolling control groups, as in Sianesi (2004) and Harding and Hsiaw (2014). Specifically, the treated municipalities are compared with a control group of towns that applied to join the program later. In our context, whether a municipality chooses to apply to conduct the Solarize campaign is not random, but the exact timing of *when* a municipality chooses to apply is plausibly random. For the treated municipalities, we find that the treatment leads to 37 additional installations over the course of a campaign on average, a greater than 1,000% increase above the control. We also explore whether there are any posttreatment effects from the campaign. We find no evidence of either a harvesting effect reducing posttreatment installations (e.g., as occurred with the well-known Cash for Clunkers program; Mian and Sufi 2012) or an increase in posttreatment installations caused by continued social learning or peer effects.

The program lowers the equilibrium price during the Solarize campaigns by roughly \$0.46 per watt (W) out of a mean price of roughly \$4.63/W in the control municipalities.[3] Because the Solarize intervention essentially replaces traditional customer acquisition by installers, this can be compared with typical estimates of installers' customer acquisition costs of \$0.48/W (Friedman et al. 2013); the similarity in the estimates imply that installers pass on much of the acquisition costs to consumers. The effect on prices lasts only during the campaign and is thus likely to be driven by the elimination of customer acquisition costs for the selected solar installer. We also find that the large increases in solar adoption during the campaigns cannot be purely explained by the price reduction, providing strong evidence of an informational component of the campaigns.

Why did the intervention work so well? To shed light on this question, we ran two randomly controlled trials (RCTs) that manipulated the campaign in different ways. The first RCT involves randomly selected municipalities across Connecticut rather than municipalities that applied to participate. Nearly all the municipalities we approached agreed to join the program. The estimated treatment effect is roughly two-thirds the treatment effect in our primary results, both in terms of installations and prices. This finding provides guidance for policymakers who would consider scaling up beyond the municipalities that self-select by applying. Our second RCT removes the group pricing element from the campaigns to assess the degree to which group pricing helps drive word-of-mouth (WOM) and adoptions. In the campaigns randomly assigned to not use group pricing, installers bid using a single price during the precampaign installer selection process. The campaigns without group pricing perform just as well as those that had group pricing, indicating that group pricing is not essential to the success of the campaigns.

We also leveraged the implementation of a very similar program called the "CT Solar Challenge" (CTSC). CTSC included all the central tenets of the Solarize program except the competitive bidding process and the involvement of the state government. The CTSC program allows us to test whether these components are necessary for the effectiveness of the overall campaign. We estimate a small and not statistically significant effect of CTCS on adoptions and a small effect on prices. Comparing Solarize to CTSC highlights the importance of the installer recruitment process for the greater success of Solarize.

Our empirical approach is similar to Bloom et al. (2013), who study management practices in 28 Indian manufacturing plants after randomly assigning some of them to receive management consulting. They establish the importance of information barriers in management using a survey of manufacturers. Similarly, we survey participants in the Solarize program in order to assess how potential adopters heard about the program and the importance of different factors in their decisions. We find that measures related to social influence are rated as extremely important factors in the decision to install solar. This, combined with the finding that including price does not change the estimated treatment effect on installations, provides suggestive evidence that Solarize works primarily by *leveraging social interactions*, which is exactly the intention of the program.

Our results have clear policy implications. Behavioral interventions based on information, word-of-mouth, persuasion, and other nonprice approaches have become increasingly popular for encouraging prosocial behavior, and community-based interventions are perhaps the latest vanguard of this movement among practitioners (McKenzie-Mohr 2013). With billions of dollars spent each year on energy conservation (Gillingham and Palmer 2014) and billions more by federal and state governments on promoting adoption of solar energy (Bollinger and Gillingham 2019), evaluating the effectiveness, persistence, and cost-effectiveness of these rapidly expanding community-based programs is important for policy development.

We find that the acquisition cost of an additional installation because of Solarize is approximately $860, plus the cost of the price discount incurred by the installer, which is approximately $1,700. In comparison, the $0.48/W acquisition cost found in Friedman et al. (2013) amounts to between $1,500 and $3,000 per installation. Assuming the 2012 carbon intensity of the electric grid, this implies a direct program cost-effectiveness estimate of $21 per ton of CO_2. This estimate is below most estimates of the social cost of carbon, suggesting that the program likely improved social welfare based on the carbon benefits alone.

2. Empirical Setting

Our empirical setting is the state of Connecticut over the period 2012–2015. Connecticut has a small, but fast-growing market for solar PV, which expanded from only 3 installations in 2004 to nearly 5,000 new installations in 2014, and more than 7,000 new installations in 2018.[4] Despite this, the cumulative number of installations still remains a very small fraction of the potential; nowhere in Connecticut is the market penetration more than 5% of the potential market, and in most municipalities, it is less than 1%.[5] The pre-incentive price of a solar PV system has also dropped substantially in the past decade, from an average of $8.39/W in 2005 to an average of $4.44/W in 2014 and less than $3/W in 2019. Figure 1 illustrates the dramatic growth of the solar market.

Figure 1. Solar Adoptions and Prices in Connecticut

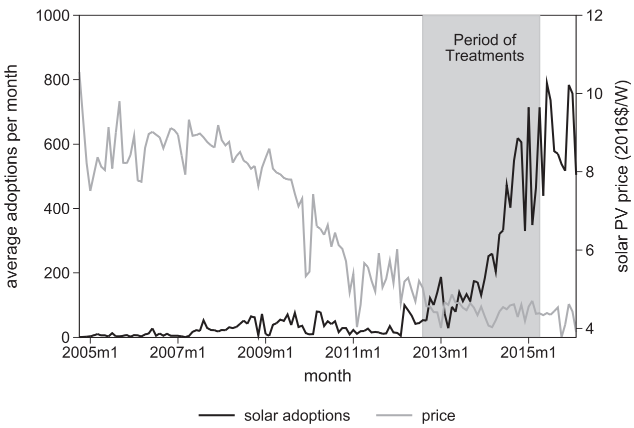

Source: Connecticut Green Bank.

Despite being in the northeastern United States, the economics of solar PV in Connecticut are surprisingly good. Although Connecticut does not receive as much sunlight as other regions, it has some of the highest electricity prices in the United States. Moreover, systems installed in Connecticut are eligible for state rebates, federal tax credits, and net metering.[6] For a typical 4.23-kW system in 2014, we calculate that a system purchased with cash in southern Connecticut would cost just under $10,000 after accounting for state and federal subsidies and would have an internal rate of return of roughly 7% for a system that lasts the expected lifetime of 25 years (see Online Appendix OA.A for more details on this calculation and some sensitivity analysis). Thus, for many consumers, solar PV systems are an ex ante profitable investment from the private perspective alone.

From 2012 to 2015, the Connecticut solar market had 89 installers, ranging in size from small local companies to large national installers. The state rebates, disbursed by the Connecticut Green Bank, began in 2006 at $5.90/W and declined to $1.75/W by the end of 2014. The incentives were held constant during the time periods covered by the treatments in this study. The Connecticut solar market was slow to adopt third-party ownership (e.g., solar leases or power purchase agreements), and most systems during the time period of our analysis were purchased outright.[7] Regardless of ownership, the state rebates are nearly always applied for by the installer and passed on to consumers.[8]

3. The Solarize Intervention

3.1. Why Solarize?

The Solarize program in Connecticut is a behavioral intervention with several components, each motivated by findings in the literature. At its core, the program focuses on facilitating social learning and peer influence. Social learning can include campaigns matching local entrepreneurs to remote coaches with business experience, as in Anderson et al. (2020), to observational learning, as is common in the peer effects literature. Peer effects have been demonstrated to speed the adoption of many new technologies and behaviors, including agricultural technologies (Foster and Rosenzweig 1995, Conley and Udry 2010), information technology (Tucker 2008), criminal behavior (Glaeser et al. 1996, Bayer et al. 2009), health and retirement plan choice (Sorensen 2006, Duflo and Saez 2003), home foreclosure (Towe and Lawley 2013), water conservation (Bollinger et al. 2020), and hybrid vehicle purchases (Narayanan and Nair 2013). Bollinger and Gillingham (2012) and Graziano and Gillingham (2015) find evidence of neighbor or peer influence on the adoption of solar PV technology in California and Connecticut, respectively, and other studies find similar results in Germany and Switzerland (Rode and Weber 2016, Carattini et al. 2018). We study the Solarize program because of considerable policymaker interest in enhancing the uptake of solar energy and the potential for the program to draw on promising approaches from existing work.

In facilitating the Solarize program in Connecticut, we worked with a state agency, the Connecticut Green Bank (CGB), and a nonprofit marketing firm, SmartPower.[9] The intervention is performed at the municipality level. After a municipality is selected to participate, the first component of the campaign is a *competitive bidding process*, in which installers submit bids to be the single vetted installer for a given campaign. The installer works with SmartPower and volunteers at events by providing information and will assist with its own marketing materials. The winning bid is chosen in a beauty contest auction, as in Yoganarasimhan (2015). Municipality leaders rank their choices based on a variety of attributes, including a group price. There is a matching process in which municipalities may not always receive their first choice. This is because a single installer is not permitted to participate in too many campaigns simultaneously for fairness and logistical reasons (usually no more than two).

The second critical component of the Solarize program is the *use of volunteer promoters or ambassadors* to provide information to their community about solar PV. There is growing evidence on the effectiveness of promoters or ambassadors in driving social learning and influencing behavior (Kremer et al. 2011, Vasilaky and Leonard 2011, Ashraf et al. 2014, BenYishay and Mobarak 2017). Why might volunteer community members be effective in Solarize? At the community level, social connectedness is likely to be high, which enhances trust (Glaeser et al. 2000, List and Price 2009). Because the ambassadors are volunteers, they may also be more likely to be seen as trustworthy by other community members. Indeed, Kraft-Todd et al. (2018) find that an ambassador's adoption during a Solarize campaign significantly increases total adoptions, through the effect on second order beliefs—the beliefs of others about the beliefs of the ambassadors regarding the value of adopting solar.[10]

The third major component of the program is the *focus on community-based recruitment*. In Solarize, this consists of mailings signed by the ambassadors, open houses to provide information about panels, tabling at events, banners over key roads, op-eds in the local newspaper, and even individual phone calls to neighbors who have expressed interest by the ambassadors. Jacobsen et al. (2013) use nonexperimental data to show that a community-based recruitment campaign can increase the uptake of green electricity using some (but not all) of these approaches. Kessler (2014) shows that public announcements of support can increase public good provision, which might improve the effectiveness of the ambassadors in Solarize.

The fourth major component is the *group pricing discount* offered to the entire community, in which the final price is a function of the total number of contracts signed as part of the campaign.[11] This provides an incentive for early adopters to convince others to adopt solar as well and reduce the price for everyone. By basing prices on total adoptions, group pricing can help spur word-of-mouth and establish norms around adoption. There is strong evidence from consumer decisions about charitable contributions that indicates consumers are more willing to contribute when others contribute (Frey and Meier 2004, Karlan and List 2007, DellaVigna et al. 2012). Moreover, there is building evidence demonstrating the effectiveness of social norm-based informational interventions to encourage electricity or water conservation (Ferraro et al. 2011, Ferraro and Price 2013, LaRiviere et al. 2014, Brandon et al. 2018). The choice to install solar PV is a much higher-stakes decision than to contribute to a charity or conserve a bit on electricity or water, so it is not obvious that effects seen in lower-stakes decisions apply. However, Coffman et al. (2014) show that provision of social information can have an important impact even on high-stakes decisions such as undertaking teacher training and accepting a teaching job.

The final major component is the *limited time frame* for the campaign. Such a limited time frame may provide a motivational reward effect (Duflo and Saez 2003) because the price discount would be expected to be unavailable after the campaign. Recent reviews (Gneezy et al. 2011, Bowles and Polania-Reyes 2012) suggest that monetary incentives can be substitutes for prosocial behavior, but by providing a prosocial reward that helps all, it is quite possible that the two are complements in this situation.

Thus, the program is designed as a package of all these components. In this sense, our program has a clear parallel to Bloom et al. (2013) in applying many of the best approaches known to make a difference at the problem at once; only in our case, it is about encouraging solar adoption rather than improving firm management.

A standard timeline for the program in Connecticut is as follows:

1. CGB and SmartPower inform municipalities about the program and encourage town leaders to submit an application to take part in the program.
2. CGB and SmartPower select municipalities from those that apply by the deadline.
3. Municipalities issue a request for group discount bids from solar PV installers for each municipality through a publicized request-for-proposals.

4. Municipalities choose a single installer, with guidance from CGB and SmartPower.

5. CGB and SmartPower recruit volunteer solar ambassadors to promote the campaign.

6. A kickoff event begins an approximately 20-week campaign featuring workshops, open houses, local events, and so on, coordinated by SmartPower, CGB, the installer, and ambassadors.

7. Consumers can request site visits, and if the rooftop is viable, the consumer will receive a quote and can choose to install solar PV.

8. After the campaign is over, the installations occur.

This timeline describes our baseline type of campaign, which we call the Solarize "Classic" campaign. Figure 2 shows the Solarize CT website, and Figure 3 provides a few examples of elements of the real-world grassroots Solarize campaigns, including a kickoff meeting, a solar open house, volunteers participating in a parade, and a roadside sign displaying the pricing tier for the campaign.

3.2. Research Design

3.2.1. Research Design for the Effects of Solarize. Similar to the management interventions in Bloom et al. (2013), running the Solarize programs is expensive, which necessitated some tradeoffs. For example, we could only run a relatively limited number of interventions, as each campaign costs approximately $30,000 to run. From September 2012 through November 2014, we ran classic Solarize campaigns in 34 municipalities. Another tradeoff is that manpower and logistical constraints prevent us from running all of the campaigns at once. Thus, we ran the campaigns in a staggered rollout over time in five distinct rounds, similar to the multiple rounds of interventions in Bloom et al. (2013).

Figure 2. (Color online) SolarizeCT.org Website

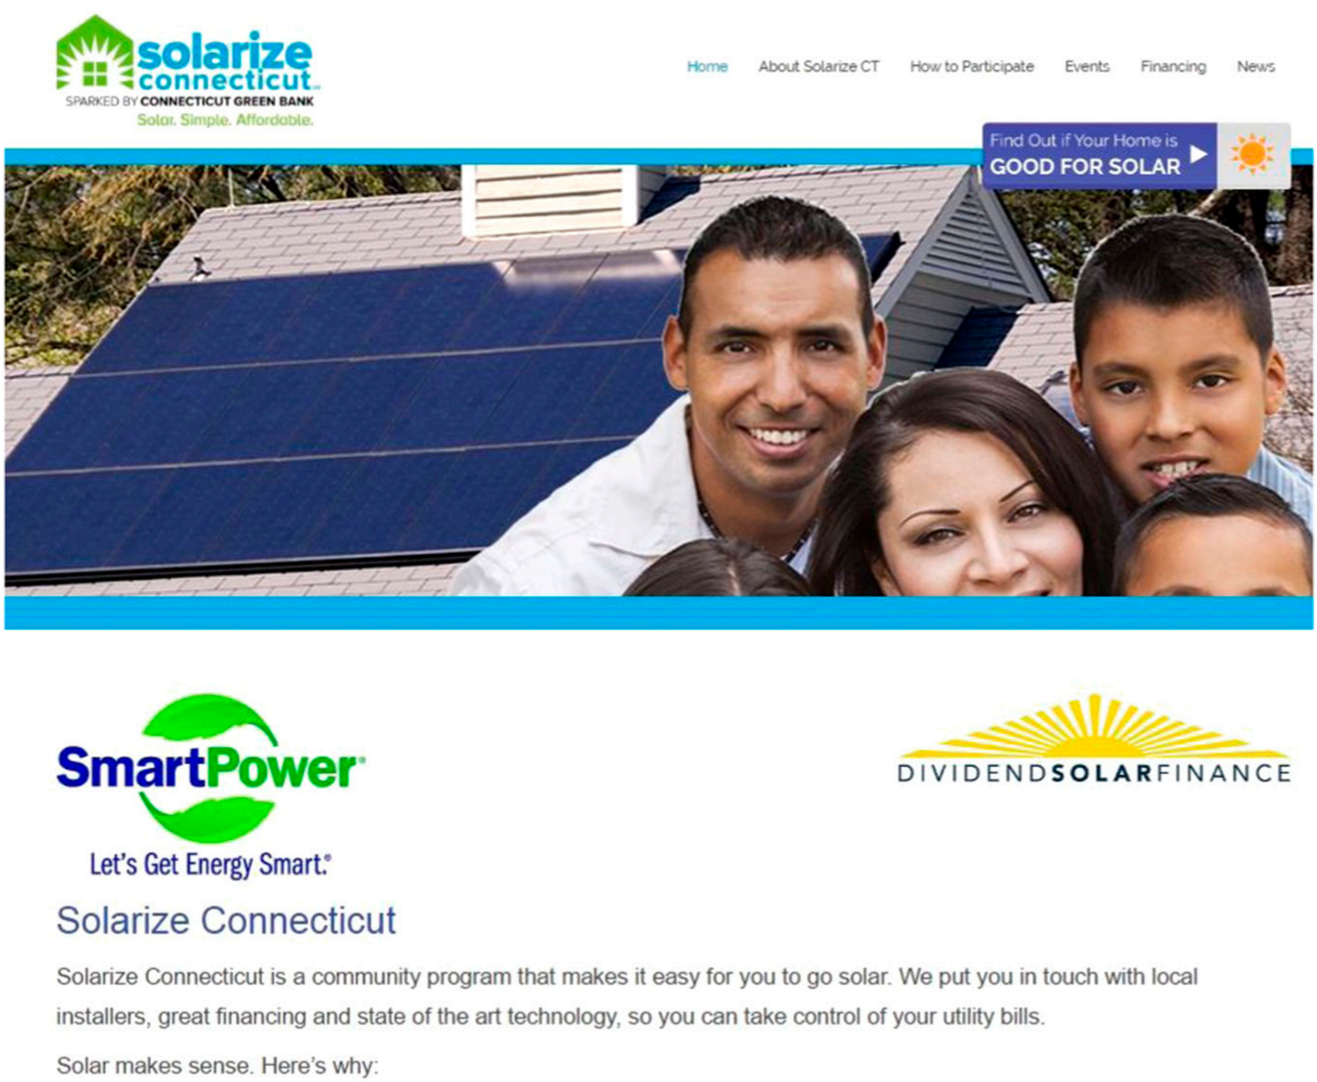

Figure 3. (Color online) Example Photographs from Solarize Campaigns

To recruit municipalities for each of the five rounds, SmartPower reached out to representatives of municipalities in state-wide events and through personal contacts. For example, there are occasional gatherings of members of municipality clean energy task forces, which are groups of residents charged with finding ways to encourage clean energy and energy efficiency. Representatives of a municipality hear about Solarize at these events through conversations with SmartPower staff or other municipality representatives. Thus, there is a random, idiosyncratic element to exactly when any given municipality representative received the information about Solarize. This element depends on who the SmartPower staff end up sitting next to and talking to and who they happened to be introduced to. Once representatives of a municipality are informed about the Solarize campaign, they then must find the time to present it to their municipality board (often in Connecticut, this is a set of town selectmen or selectwomen). In many cases, representatives from municipalities expressed great interest in hosting a campaign, but because of individual-specific reasons (e.g., a major town construction project underway, a health concern from a key representative, a key player in the task force being especially busy at that moment, or a representative stepping down), they preferred to wait some number of months before applying for a campaign. Thus, the exact timing of when any specific municipality chooses to apply for a Solarize campaign differs for highly idiosyncratic factors that influence when municipality representatives are informed about Solarize and when they manage to successfully get approval from the municipality board.

Our primary identification strategy exploits the plausibly random timing of when municipalities apply to receive the treatment. In our study design, untreated municipalities that eventually become treated are first used as controls and later are treated. Therefore, for the first round of the program, the control group consists of all municipalities that apply for a campaign in all of the subsequent rounds. In the final round, we can only include municipalities that apply for Solarize after the end of our study period.[12] This research design is analogous to the approach in Sianesi (2004). We also make sure to exclude municipalities from the control if they are adjacent to one of the treated municipalities because of the possibility of spillovers (a topic we will investigate in Section 4.5). For robustness, we also examine two other possible control groups. The first uses a nearest-neighbor propensity score matching approach with a 0.05 caliper to match each of the treated municipalities to three nearest neighbors based on observed demographics. The second takes advantage of a pre-existing set of municipalities that have chosen to participate in the "Connecticut Clean Energy Communities" program by committing to having a municipality clean energy task force. For this second control group, we use all nontreated municipalities in the Connecticut Clean Energy Communities program. The results from these other two control groups serve as useful robustness checks.[13]

Table 1 shows the balance of observables between the treated and control municipalities. To calculate these statistics, we bring in data from three sources. First, we have data from the Connecticut Green Bank on all solar installations in Connecticut through 2017. In order to receive the rebate, firms report each installation to the Green Bank, including the address, the date the contract was approved, the date the installation was completed, the size of the installation, the preincentive price, the incentive paid, whether the installation is third party-owned (e.g., solar lease or power-purchase agreement), and additional system characteristics. Because the rebate has been substantial over the past decade, nearly all solar PV installations in Connecticut are included in the database.[14] Second, we use municipality-level demographic data from the U.S. Census Bureau's 2009–2013 American Community Survey. Third, we include voter registration data from the Connecticut Secretary of State (SOTS).[15]

In Table 1, we see that that none of the means of the observables are significantly different between the treatment and future controls or propensity score matched controls. The sample sizes are reasonably sized (e.g., 72 municipalities when comparing the treatment to the future controls), so we do not believe that these results are solely because of small sample sizes. A few variables have means that are statistically significantly different (or close to it) between the treatment and the Connecticut Clean Energy Communities control, such as the percentage of the population with a college degree. Thus, we are more cautious about using the Connecticut Clean Energy Communities as the control, although these differences could simply be from random variation.

3.2.2. Research Design for the Analysis of Mechanisms.

To better understand the mechanisms at work in the Solarize campaign, we run two RCTs and exploit related campaigns run by a nonprofit focused around a for-profit solar firm. In our RCT, we compare the classic Solarize campaigns, where municipalities apply themselves to be able to run a campaign to a campaign where we preselected the municipalities and approached them ourselves with an offer to run a campaign. For simplicity, we will refer to the classic Solarize campaigns hereafter as Solarize Classic campaigns.

One might hypothesize that having an eager and engaged set of volunteers who put in the work for the municipality to apply to the program is crucial for the success of the program, and our first RCT tests the importance of this aspect of campaigns. In effect, it is testing the importance of *selection* into the

Table 1. Table of Balance

	All of CT		Treated	Caliper controls		CEC controls		Future controls	
Variable	Mean	Standard deviation	Mean	Mean	Diff. p value	Mean	Diff. p value	Mean	Diff. p value
Number of towns	169	n/a	34	23	n/a	47	n/a	38	n/a
Cumulative preadoptions	13.9	9.4	16.4	13.5	0.29	15.0	0.51	17.6	0.63
Population density	885	1,218	900	933	0.93	1344	0.16	1,041	0.67
Median household income	83,899	26,846	94,095	86,826	0.38	82,148	0.06	90,512	0.61
Percent over 65 years of age	15.7	3.9	14.7	15.3	0.48	15.9	0.11	15.3	0.49
Percent white	88.9	11.9	86.9	89.3	0.48	87.3	0.88	87.3	0.89
Percent black	4.0	7.8	5.9	3.3	0.32	4.8	0.61	4.6	0.57
Percent households with children	71.1	7.5	72.0	71.4	0.80	70.4	0.33	71.6	0.82
Percent commute > 60 miles	8.6	5.8	9.2	8.4	0.66	8.0	0.36	9.1	0.95
Percent below poverty	6.3	5.1	6.3	6.2	0.94	6.7	0.74	6.4	0.94
Percent college degree	47.8	5.3	49.3	48.4	0.53	46.4	0.02	48.4	0.51
Percent unemployed	8.3	2.7	7.7	8.3	0.45	8.8	0.09	7.8	0.93
Percent detached housing	77.0	17.1	78.4	78.3	0.98	72.7	0.15	76.6	0.64
Percent republican voters	23.5	7.4	24.5	25.9	0.57	22.1	0.17	24.1	0.86
Percent democrat voters	30.9	8.9	32.4	29.0	0.17	32.3	0.97	32.0	0.85

Notes. Means are at the municipality level. p values are from two-sided t tests of differences in treatment versus control means. Diff. refers to "difference."

Solarize campaign. To examine this, we randomly draw five municipalities from the pool of all non-Solarize municipalities in Connecticut. We ran these five randomly drawn campaigns alongside Solarize Classic campaigns in seven municipalities during round 4 of the program. These municipalities may not be the marginal municipalities most likely to be the most engaged in clean energy, and thus the success of the campaign in these municipalities provides guidance on how Solarize might work on average if campaigns were run in all municipalities in Connecticut.

In our second RCT, we compare the Solarize Classic campaign described previously to an intervention where we remove the group pricing aspect of the campaigns in order to test whether group pricing is critical to the success of the campaigns. We randomized municipalities that applied to receive a Solarize campaign during round 5 of the program into either the Solarize Classic or no group pricing versions of the campaign.[16] Five municipalities in this round received the Solarize Classic campaign, whereas another four received Solarize without group pricing. In each of the two RCTs, we make sure not to run campaigns in adjacent municipalities at the same time because of the possibility of spillovers across municipality borders.

During the time frame of our study, a for-profit solar installer, Aegis Solar, independently conducted similar campaigns. Specifically, Aegis Solar created and funded the nonprofit CTSC to contact municipalities and encourage them to participate in a Solarize-type campaign. These CTSC campaigns are explicitly modeled after Solarize, because Aegis Solar took part in the first round of Solarize and thus was very familiar with the program. The only substantive differences from the Solarize campaign are that (1) there was no competitive bidding process for the installer (Aegis Solar was the only participating installer), (2) the Connecticut Green Bank and SmartPower were not involved, and (3) the length of the campaigns tended to be slightly longer. Otherwise, the CTSC campaigns are the same as Solarize campaigns. Thus, analyzing these campaigns provides an opportunity to explore the importance of the competitive bidding and trust in the program organizers in the success of the campaigns. We analyze 10 CTSC campaigns conducted during the time frame of our study.

As each of these additional analyses uses a small sample, it is important to keep these results in context. For example, in the RCT where we removed group pricing, we have a sample size of nine and are comparing five in one treatment to four in the other. In such a small-sample analysis, a couple of issues are raised. One is whether we can find statistically significant results. The interventions themselves are intensive enough that we expect to find clear differences from the controls, but we may not find differences across the two types of treatments. A second issue is how to perform statistical inference. As we will discuss, we use small-sample statistical inference approaches on all our analyses.

In total, we analyze campaigns in 53 municipalities in this study. Figure 4 provides a map of the 169 municipalities in Connecticut, illustrating the 53 treated municipalities we examine in this study.[17] Online Appendix Table OA.1 provides a complete list of all of the campaigns we analyze in this paper and the dates that they were run.

4. Impact of the Treatment on Solar Adoption and Prices

4.1. Descriptive Evidence

We begin with some descriptive evidence of the effect of the Solarize treatment. For the big picture, we examine the mean number of adoptions in a municipality over the full length of the campaigns in the raw data. Figure 5 presents the mean number of adoptions in a municipality during a campaign by round for the Solarize treatment group and each of the three potential control groups, where the *future control group* (municipalities that in the future will apply to join Solarize) is our preferred control group.

In Figure 5, we observe a substantial increase in adoptions of solar systems in the Solarize Classic campaigns over any of the control municipalities during the time period of the campaigns. In general, the control municipalities show low rates of adoption during the campaign period, and these rates are generally not statistically different from each other across the control groups. In many cases, they are not statistically different from zero either, although they do tend to increase over time, along with a general upward trend of solar adoptions in the Connecticut market. In contrast, the increase from Solarize Classic is just as large in some of the later campaigns as it is in the earlier ones. This is important, because it provides evidence in support of our identifying assumption in using the future control group. We might have been worried if the municipalities that were treated earlier had greater adoptions because this would have raised questions about whether the timing of when municipalities apply for the program is plausibly random. Fortunately, we observe no discernable pattern in the numbers of adoptions during the campaign across campaign rounds, consistent with plausibly random timing.

By pooling across rounds in the raw data, we observe that the average number of adoptions in a municipality in the Solarize Classic campaign is 48.1. In comparison, the average for the propensity score matching caliper control group is 4.0, whereas for the Connecticut Clean Energy Communities control group,

Figure 4. (Color online) Map of Solarize Campaigns

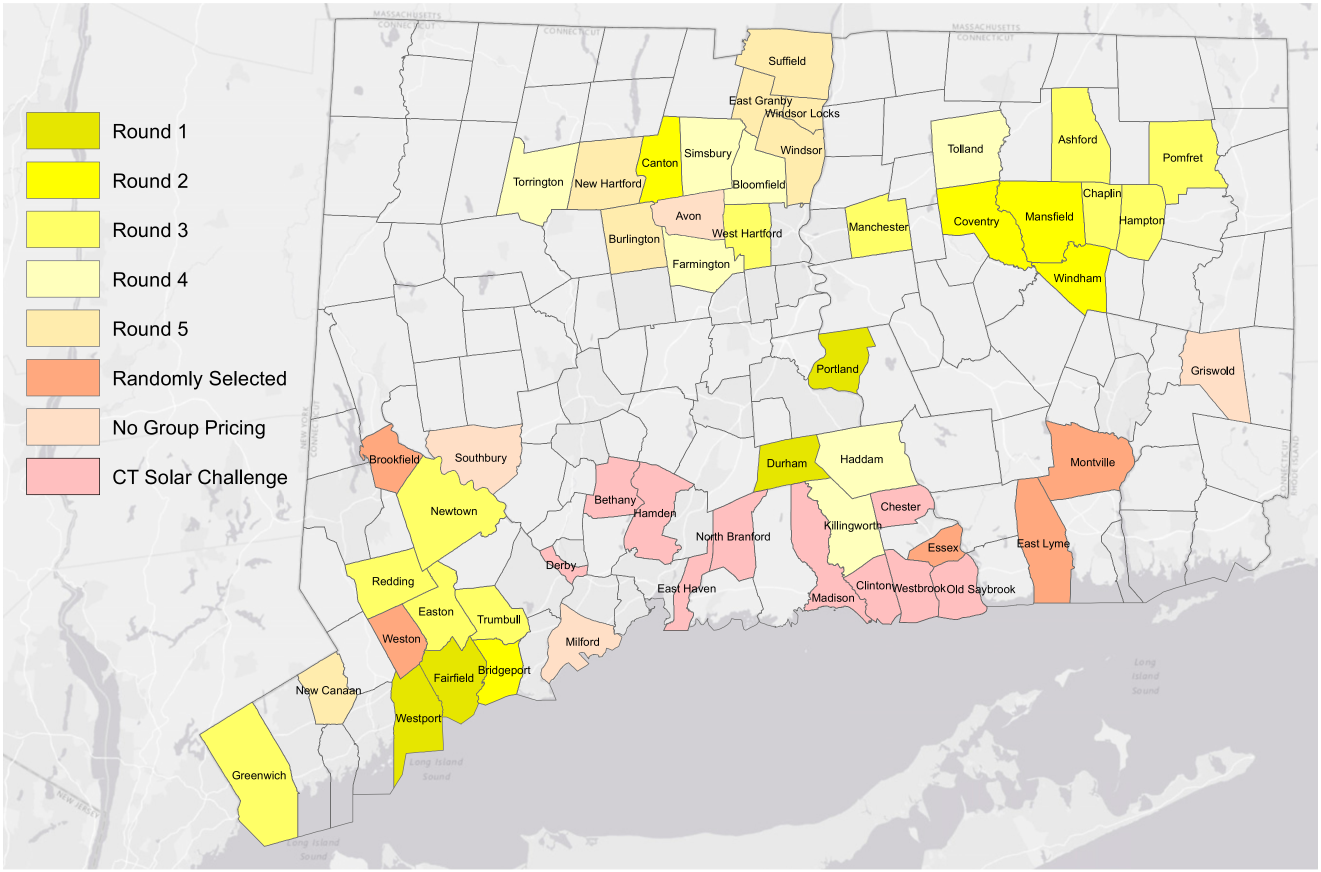

it is 10.2, and for the future control group, it is 9.2. Using a standard t test of differences in means reveals that the Solarize Classic mean is significantly different from any of the control groups. Just taking a simple difference in the raw data suggests that the added lift from Solarize Classic is in the range of 38 to 44 additional solar installations from the campaign. This provides a benchmark to compare our empirical results to.

Figure 6 pools across the rounds and presents the average number of monthly solar adoptions by a municipality (a) and the average prices (b) for the Solarize Classic treatment and the three control groups. The x axis plots the number of months since the start of the Solarize campaigns, and the shaded area refers to length of the Solarize campaigns. Figure 6(a) clearly shows a dramatic spike in adoptions during the Solarize Classic campaigns, whereas each of the control groups show a modest rate of adoption that slowly and steadily increases over time. Moreover, the preperiod adoptions are very similar between the treatment and control groups. Along with Figure 5, this is strong descriptive evidence of a major increase because of the intervention. Another observation from Figure 6(a) is that the posttreatment rate of adoption of the treated municipalities goes back to being similar to the control municipalities. This suggests that there is neither harvesting reducing adoption nor enhanced peer effects increasing adoption in the posttreatment period.

Figure 6(b) shows that the average solar prices are roughly the same in the pretreatment period between the treatment and control municipalities. However, during the campaign, the prices dropped substantially in the treated municipalities. Posttreatment, we see the prices returning to be about the same between

Figure 5. Mean Adoptions by Round

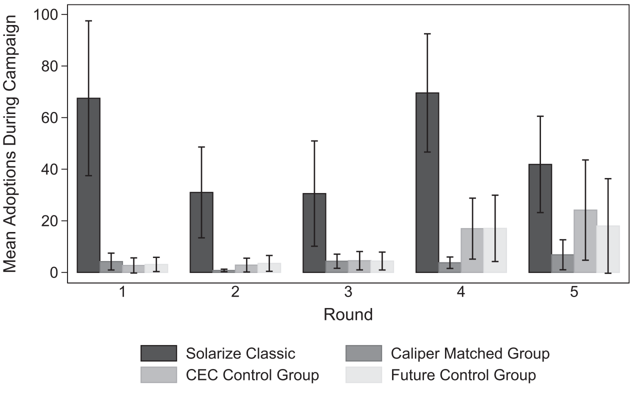

The lines on each bar indicate +/- one standard deviation from the mean

Figure 6. (Color online) Campaign Adoptions and Prices

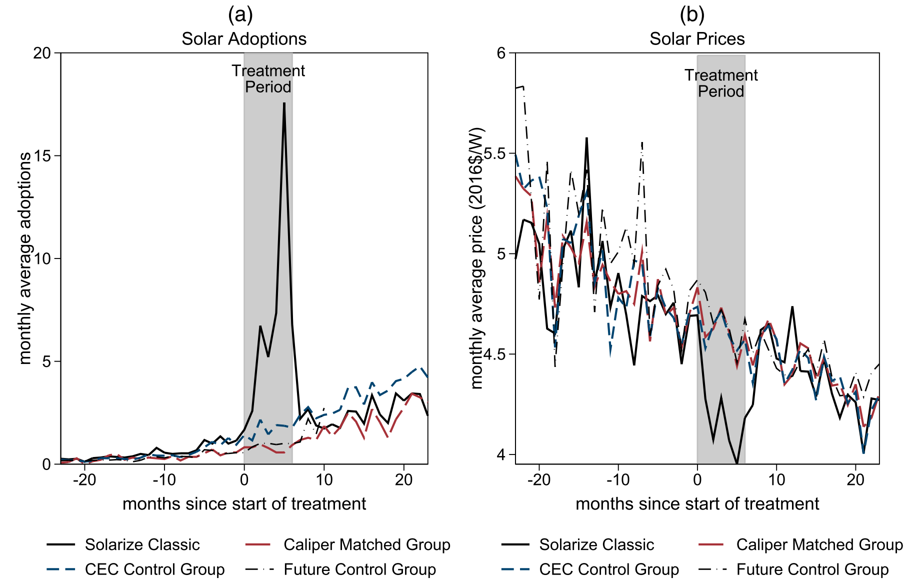

the treatment and controls. This descriptive evidence supports the contention that the intervention did succeed in lowering prices during the campaign through the group pricing deal.

4.2. Causal Effect of the Intervention

The descriptive evidence from the raw data already presents a compelling picture of a strong treatment effect of the Solarize program in increasing adoptions. We now turn to an empirical model to estimate the causal effect of the Solarize intervention after controlling for potential confounders.

We are interested in the average treatment effect on the treated (ATET) for adoptions and prices. Because the treatment is at the municipality level, we estimate the effects at this level and convert our data to a municipality × month panel. Our preferred specification for the causal effect on adoptions in municipality i in month t is given as follows:

$$Adoptions_{it} = \beta T_{it} + \eta_i + \mu_t + \epsilon_{it}. \tag{1}$$

Here $Adoptions_{it}$ refers to the number of adoptions in municipality i and month t[18]; T_{it} is a dummy for the Solarize treatment (i.e., a treated municipality during the treatment period); η_i are municipality × round fixed effects; μ_t are month-of-the-sample dummy variables; and ϵ_{it} is the error term.

To estimate the treatment effect on equilibrium prices, we use a similar specification:

$$Price_{it} = \delta T_{it} + \lambda_i + \pi_t + \varepsilon_{it}. \tag{2}$$

In this specification, $Price_{it}$ is the average price in municipality i and month t. We use the T_{it} again as the dummy for the Solarize treatment, λ_i are municipality × round dummy variables, π_t are month-of-the-sample fixed effects, and ε_{it} is the error term.

Because we have five rounds, we create a stacked data set. For each round, we include all the treatment and control municipalities and two years before the start of the campaigns. We do not include any post-periods in these estimations.[19] Then we stack these data sets together into a single pooled data set. This means that a municipality can be a control in the first round and a treated municipality in a later round. We also run each of the rounds separately, but we face small sample issues in doing so, and thus we prefer the analysis of the pooled data.

These estimations are difference-in-difference estimations, so identification is based on the assumptions of parallel trends and the stable unit treatment value assumption (SUTVA). The parallel trends assumption requires that the control group would have had an identical trend to the treatment group had the treatment not been implemented. If this assumption holds, then any time-varying unobservables will be captured through the trends in the control group. For the parallel trends assumption to hold, we must be confident in the validity of the control group. We argued previously that the timing of municipality applications for Solarize is plausibly random because of the process by which it occurred, which involved chance contacts between individuals and idiosyncratic factors at the municipality level. This contention is strongly supported by the table of balance (Table 1) and by the nearly identical pretrends for adoptions and prices in Figure 6.

SUTVA first requires stable treatments, which means that the treatments are applied the same to all treated municipalities. Our research design assures this. SUTVA also requires noninterference, which implies that there are no spillovers between the treatment and control. To assure that this holds, we drop all municipalities adjacent to the treated municipalities from the control groups. Figure 6 provides descriptive evidence that treatment spillovers are unlikely to have a dominant effect, because there is no discernible change in the trends in the control groups. We further explore the possibility of spillovers shortly after first presenting our primary results and robustness checks.

4.3. Primary Results

We present our primary results of the effect of the Solarize intervention on adoptions and prices in Table 2. The first three columns present the results from estimating the adoption Equation (1), whereas the second three columns present the results from estimating the pricing Equation (2). Columns 1 and 4 present the results using the propensity score caliper matched control group, columns 2 and 5 present the results using the non–Solarize Connecticut Clean Energy Communities control group, and columns 3 and 6 present the results using our preferred future Solarize control group.

One possible concern about inference is that the number of clusters is relatively small. Even with clustering, our baseline treatment effect inference relies on asymptotic arguments, and therefore the municipality-level clustered standard errors may understate the uncertainty in the estimates. Bertrand et al. (2004) perform simulations that show that the cluster-correlated Huber-White estimator can lead to an over-rejection of the null hypothesis when the number of clusters is small, with 50 being a common benchmark. We are

Table 2. Impact of the Solarize Intervention on Solar Installations and Prices

	Dependent variable					
	Installations			*Prices*		
Control group	(1) Caliper	(2) CEC	(3) Future	(4) Caliper	(5) CEC	(6) Future
Intervention_{it}	6.82	6.20	6.63	–0.44	–0.39	–0.46
	(0.77)***	(0.78)***	(0.80)***	(0.08)***	(0.08)***	(0.08)***
Small sample robustness						
Wild cluster bootstrap-t (*p* value)	0.00	0.00	0.00	0.00	0.00	0.00
Randomization inference (*p* value)	0.00	0.00	0.00	0.00	0.00	0.00
No. municipality fixed effects	62	199	153	62	199	153
No. month-of-sample dummies	43	43	43	43	43	43
Observations	1,059	3,705	2,769	1,059	3,705	2,769
R^2	0.51	0.42	0.47	0.29	0.22	0.20
Effect over entire campaign						
ATET per municipality	38.2	34.7	37.1			
Effect in raw data	44.1	37.9	38.9			
Treated town average price ($/W)				4.16	4.16	4.16
Control town average price ($/W)				4.56	4.61	4.63

Notes. An observation is a municipality-round-month. Standard errors in parentheses are block bootstrapped at the municipality level. *Intervention* is a dummy equal to 1 for a Solarize campaign occurring. The sample is stacked over the rounds, so a municipality may be a control for an earlier round and treated in a later round. Once treated, municipalities are removed from the sample. *Caliper* refers to propensity score matching of the treated municipalities to the three nearest neighbors with a 0.05 caliper. *CEC* refers a control group of all Connecticut Clean Energy Community municipalities except those treated or adjacent to treated. *Future* refers to a control group of municipalities that in future rounds opted-in to a Solarize campaign. *No. municipality fixed effects* reports the number of municipality × round fixed effects. *No. month-of-sample dummies* reports the number of month-of-sample dummies (one or more are dropped because of collinearity). *Wild cluster bootstrap-t* reports the *p* value for testing the null hypothesis that the treatment has no effect using the wild cluster bootstrap-t procedure from Cameron et al. (2008). *Randomization inference* reports the *p* value for testing the same null hypothesis using randomization inference. The *Average Treatment Effect on the Treated (ATET) per municipality* calculates the average effect over an entire campaign of 5.6 months on average. The *Effect in raw data* calculates the difference between the total installations during a campaign for treatment and control groups, averaged over all municipalities. The final two rows calculate the average price during campaigns (in 2014$/W) in treatment and control groups.

***Significant at 1%.

above this threshold in all columns (see the number of municipality dummies). However, for robustness, we consider other methods of inference. Recently, the Cameron et al. (2008) wild bootstrap method has found wide application in recent empirical studies (Giné and Yang 2009, Ben-David et al. 2013, Bloom et al. 2013, Elberg et al. 2019). As another alternative, Cohen and Dupas (2010) and Bloom et al. (2013) use randomization inference, which does not require asymptotic arguments or distributional assumptions (Fisher 1935, Rosenbaum et al. 2002). The first standard errors in parentheses are simply block bootstrapped at the municipality level to allow for any within-municipality correlation in errors, but we implement both small-sample methods and report the p values for our hypotheses. These small-sample methods become all the more important when we use our smaller-sample randomized controlled trials in later sections.

The primary results in Table 2 demonstrate a strong and significant causal effect of the Solarize intervention on adoptions. In column 3, we see that the treatment leads to an increase in adoptions of 6.63 solar systems per municipality in a month. This implies an average treatment effect of 37.1 over the entire campaign per municipality, which is more than a 1,000% increase from the control group adoption rate. This is just a bit below the effect in the raw data when we simply subtract the mean adoptions in the control municipalities (0.6) from the mean adoptions in the treated municipalities (7.9). For further context, if we estimate our specification in column 3 using monthly adoptions per owner-occupied household, we find a treatment effect of 0.2 percentage points, which is a massive increase over the monthly adoptions per household in the control (0.02%). We also find a similar strong and statistically significant effect in Table 2 for prices. Our result in column 6 indicates that the Solarize intervention decreased prices by −\$0.46 per watt during the campaign, a 10% decline,[20] which is in line with expectations, because Friedman et al. (2013) find installer's customer acquisition costs to be \$0.48/W, in the absence of a Solarize campaign.

We can also estimate the treatment effects over time by modifying Equation (1) slightly by interacting the treatment dummy with a dummy variable for each of the months since the beginning of the treatment. In this estimation, we also extend the sample out to two years before treatment and two years after treatment. Figure 7 shows the treatment effect on adoptions over time using the future Solarize control group. As expected, there is no statistically significant effect in the before period, but there is a dramatic spike during the treatment, mirroring what we observed in the raw data in Figure 6. After the treatment, we see no

Figure 7. Treatment Effects on Adoption Over Time

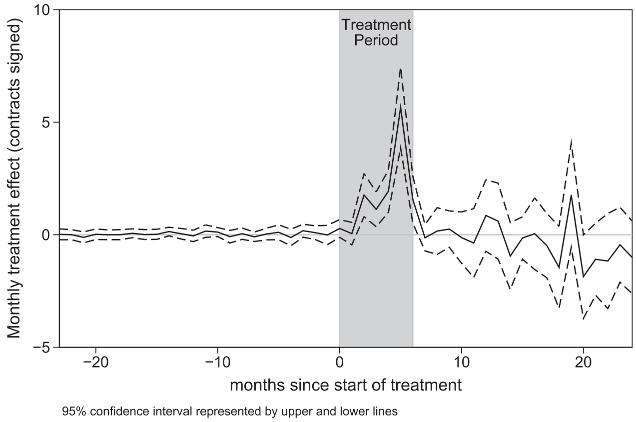

statistically significant treatment effect, indicating that the treated municipalities largely returned to the rate of adoption of the control municipalities.

These results provide strong evidence of the success of the Solarize campaign in increasing adoptions and lowering prices. This success raises a several questions. Did the campaigns spill over to adjacent municipalities? What is the effect on randomly selected municipalities rather than the marginal municipalities that opted-in to the campaign? What are the mechanisms underpinning the success of the campaigns? Before addressing these questions, we first present several robustness checks.

4.4. Robustness Checks

In Table 2, we show the robustness of our primary results to two other potential control groups. However, we also explore the robustness of our results to several alternative assumptions and perform a placebo test to further support our identification arguments.

In our first robustness check, we run the model described by Equation (1), only we vary the length of the pretreatment period that we include. In our primary specifications, we include two years of pretreatment period. If we include a longer pretreatment period, we begin to include a time frame that is less relevant for pinning down the trends in the treatment period. If include a shorter pretreatment period, we have less data available to pin down these pretreatment trends. In Online Appendix Table OA.2, we compare our primary results using the future control group (column 1) to those with only one-year pretreatment period (column 2) and only two months of a pretreatment period (column 3). We find very similar results, suggesting that the somewhat arbitrary assumption we made about the length of the pretreatment period does not affect our results.

In our second robustness check, we estimate Equation (1) using a negative binomial model rather than using ordinary least squares. The number of adoptions in a municipality may be thought of as a count variable, and thus a negative binomial model may be appropriate. Column 4 in Online Appendix Table OA.2 presents the results using the negative binomial model, showing that the main findings do not substantially change. Column 5 in Online Appendix Table OA.2 reports a zero-inflated negative binomial to flexibly account for the many municipality-months that have zero adoptions. Similarly, one might be concerned that municipalities have different potential market sizes, so the number of adopters is not fully reflecting the change in propensity to adopt across municipalities. This should not be an issue in our research design because the municipalities in the treatment and control have similar potential market sizes, but to assure that it is not an issue, we estimate a log-odds model based on the share of the potential market in each municipality (see Online Appendix OA.D). The results from the log-odds model are again very similar to our preferred specification and are found in column 6 of Online Appendix Table OA.2.

We also perform a set of robustness checks to examine different assumptions that might affect our price specifications. In our preferred specification, we impute missing solar installation prices for municipalities that do not have any installations in that month. This is done by using the average price in the same county as the municipality, and for any remaining missing prices where there are no installations in the county during that month, we use the state-wide price (only a very small percentage of municipality-month observations). This interpolation could lead to some measurement error, so to examine the robustness of our results, we estimate the model in Equation (2) on the sample for which price is not missing. Online Appendix Table OA.3 presents these results in columns 1–3, whereas the primary results from Table 2 are presented in columns 4–6 for comparison. We observe slightly larger treatment effects on prices when we remove the missing observations, which makes sense because the sample is slightly different, with low-adoption municipalities weighted more heavily in our primary results. The difference could also be caused by attenuation bias from measurement error, but it is impossible to separate out the two hypotheses. We find it comforting that the coefficients did not change substantially and that the coefficients indicate that our primary results are conservative.

Finally, we perform a series of placebo or falsification tests where we shift the dates of the intervention in our analysis to before they actually happened. Online Appendix Table OA.4 presents the results from the placebo tests changing the intervention period to the six months before the start of the campaigns. In these estimations, we only include the sample before the beginning of the treatments. As suggested by our descriptive statistics, there is no discernable difference between the treatment and control group adoption rates or prices during the pretreatment period, so it is not surprising that we find that all of the coefficients are close to zero and/or not statistically significant.

4.5. Spillovers to Adjacent Municipalities?

If the Solarize treatment leads to additional installations through social interactions and word-of-mouth, we might expect nearby municipalities to also experience some treatment effect, because social networks extend across municipal borders. Such spillovers or *treatment externalities* have been exhibited in other field experimental settings (Miguel and Kremer 2004) and can contribute positively to the cost-effectiveness of the program. They also could pose a challenge to the experimental design if they lead to violations of SUTVA. As mentioned previously, we exclude adjacent municipalities from the control groups in our primary analysis to address this potential concern.

We estimate the models in Equations (1) and (2), only using a treatment dummy that refers to a municipality adjacent to a treated Solarize municipality. In these regressions we include only the adjacent municipalities and the controls (the treated Solarize or CTSC municipalities are excluded). Again, we include the two years before the beginning of the interventions. Our results show no statistically significant evidence of spillovers at the municipality level in either adoptions or prices. Online Appendix Table OA.5 provides the results using the future Solarize control group. The coefficients suggest a small and positive effect on adoptions and small and negative effect on prices, but the coefficients are not precisely estimated and cannot be distinguished from zero.

This analysis thus far was based at the municipality level, but it is possible that spillovers happen only in the area immediately adjacent to a treated Solarize municipality. Thus, to look more deeply into spillovers, we also performed a spatial analysis, similar to the study of Anderson et al. (2018) of small business owner training in South Africa. Using geographic information system software (ArcMap), we created a buffer zone around each Solarize municipality and then a further zone with the same width in the interior of the adjacent municipality, as is illustrated in Online Appendix Figure OA.1. Then we calculated the rate of adoption of solar in the immediately adjacent buffer zone and compared it with the rate of adoption in the equivalently sized zone in the interior of the municipality. We examined several different buffer zone sizes, including 1 mile and 0.5 miles. Similarly, we examined several different values for the gap between the adjacent buffer zone and the interior zone, such as 2 and 5 miles. Our results from this detailed analysis are similarly inconclusive; we perform a set of t tests comparing the mean adoption rate during the campaign in the two zones and find no statistically significant results. The lack of significant spillover effects suggests that the social learning is predominantly confined within the area of the campaign.

4.6. How Important Is Selection into the Program?

We thus far demonstrated a substantial causal effect from the Solarize treatment on the municipalities that chose to apply for the treatment. One could view these municipalities as the *cream-of-the-crop* for such behavioral programs because enthusiastic individuals in these communities went out of their way to select into the program. This raises the natural question of what the effects of the Solarize intervention would be if the program is scaled up to *all* municipalities in Connecticut rather than only the most enthusiastic. We thus test this with the randomly assigned set of municipalities that SmartPower approached with the opportunity to participate. Fortunately, the approached municipalities agreed to participate, but there was not always complete buy-in from the town council in all municipalities, and it was more difficult finding campaign volunteers.

For our analysis here, we use two control groups for comparison. The first control group consists of *all* municipalities in Connecticut that did not yet (by the end of the fifth round of Solarize campaigns) receive a Solarize or CTSC intervention. Recall that this is the pool of municipalities that we randomly drew from. This is a very appropriate control group for this analysis. However, because of the small sample size of randomly drawn campaigns that we were able to run, it is likely that some of the observable municipality-level characteristics may differ on average between this control group and the randomly drawn treatment group. Thus, to be conservative, we also run an analysis using a propensity score matched caliper control group, where we again match each of the randomly drawn municipalities to three nearest neighbors (again using a 0.05 caliper). In Online Appendix Table OA.6, we show that observable demographics and pretreatment cumulative adoptions are generally similar between the randomly drawn treatment and control groups, although there are a few statistically significant differences, which may not be surprising given the relatively small number of treated municipalities.

Table 3 presents the results showing the causal effect of the Solarize intervention on randomly selected municipalities. Columns 1 and 3 present the results using the broad control group of all municipalities that did not yet receive a Solarize or CTSC campaign. Columns 2 and 4 show the results using the propensity score matched caliper control group. The first two columns show the results for adoptions, whereas the second two show the results for prices. Again, we include two years before the beginning of the interventions as a pre period. The results suggest an increase in adoptions of 3.66 to 5.22 adoptions per municipality per month, which translate into 20.5 and 29.2 over the entire campaigns. Prices decline by \$0.26/W to \$0.29/W from an average of about \$4.60. These coefficients are all statistically significant despite the smaller sample size.

However, when we compare these results to those for Solarize Classic in Table 2, it is clear the effects are muted in the randomly selected group. There are fewer adoptions, and the price decline is substantially lower. This also holds if we compare the results in Table 3 only to the Solarize Classic campaign in round 4, with an even greater difference in the number of adoptions and a similar difference for the price decline. Specifically, the Solarize Classic treatment effect in round 4 is 8.97 adoptions per municipality per month, with a cost decline of \$0.42/W (both are statistically significant). This comparison shows that, although the Solarize program can still be effective in randomly selected municipalities, selection into the program matters.

This sheds further light on the mechanisms underlying the effectiveness of the program. Selecting into the Solarize campaign by applying for it is usually the result of one or two key ambassadors or municipality leaders who are particularly interested in promoting solar to their community. Having these key promoters at the center of a campaign is a primary difference between the campaigns in the randomly drawn municipalities and the municipalities that selected into the program.

5. Mechanisms

5.1. Is Solarize Just a Discount Pricing Scheme?

5.1.1. Prices or Information? We can examine whether the installation treatment effect is simply the result of the price decline by exploring whether prices or information are the primary determinants of the increased adoption from the campaign. We re-estimate Equation (1), only we include the average installation price as a covariate. In order to identify the causal effect of prices, we need to address price endogeneity, especially given the fact that the Solarize installers' bids will reflect the expected demand lift from the campaign. To do so, we leverage cross-sectional and time series variation in the Bureau of Labor Statistics county-level roofing wage rate, which is driven in large part by the availability of roofing labor. Roofers and solar installer employees have a very similar skill set, and so the roofing wage rate is a good proxy for installers' labor costs (this instrument has also been used in previous papers, including Gillingham and Tsvetanov 2016).

Table 3. Impact of Solarize in Randomly Selected Towns

	Dependent variable			
	Installations		*Prices*	
Control group	(1) Non-Solarize	(2) Caliper	(3) Non-Solarize	(4) Caliper
$Intervention_{it}$	3.66 (1.04)***	5.22 (1.02)***	–0.26 (0.09)***	–0.29 (0.11)***
Small sample robustness				
Wild cluster bootstrap-t (*p* value)	0.07	0.00	0.01	0.03
Randomization inference (*p* value)	0.00	0.01	0.01	0.01
No. municipality fixed effects	126	9	126	9
No. month-of-sample dummies	17	17	17	17
Observations	2,268	162	2,268	162
R^2	0.26	0.46	0.14	0.19
Effect over entire campaign				
ATET per municipality	20.5	29.2		
Effect in raw data	22.2	32.3		
Treated town average price ($/W)			4.38	4.38
Control town average price ($/W)			4.64	4.62

Notes. Regressions are run on the subsample of randomly selected Solarize and control towns. An observation is a municipality-round-month. Standard errors in parentheses are block bootstrapped at the municipality level. *Intervention* is a dummy equal to 1 for a Solarize campaign occurring. *Non-Solarize* refers to the control group of all nontreated municipalities. *Caliper* refers to propensity score matching of the treated municipalities to the three nearest neighbors with a 0.05 caliper. *No. municipality fixed effects* reports the number of municipality × round fixed effects. *No. month-of-sample dummies* reports the number of month-of-sample dummies (one or more are dropped due to collinearity). *Wild cluster bootstrap-t* reports the *p* value for testing the null hypothesis that the treatment has no effect using the wild cluster bootstrap-t procedure from Cameron et al. (2008). *Randomization inference* reports the *p* value for testing the same null hypothesis using randomization inference. The *ATET per municipality* calculates the average effect over an entire campaign of 5.6 months on average. The *Effect in raw data* calculates the difference between the total installations during a campaign for treatment and control groups, averaged over all municipalities. The final two rows calculate the average price during campaigns (in 2014$/W) in treatment and control groups.

***Significant at 1%.

In columns 1–3 of Table 4, we present the ordinary least squares (OLS) results, using each of the three control groups, and in columns 4 through 6, we instrument for price using the roofing wage rate. We see no effect of price in the OLS regressions, but this is as expected because of the price endogeneity. In the instrumental variables (IV) regressions, we can only imprecisely estimate the effect of price when using the caliper control group, given the smaller sample size. However, when using the other two control groups, we find that price has a significant, negative effect on adoptions. For instance, using the CEC control group, we find that a $1/W price increase leads to a decline of 7.5 installations per month, and with the future control group, we find a reduction of 5.8 installations per month.

With a price decline of $0.46/W because of the campaigns, our estimates suggest that the decrease in price can only explain a lift from Solarize of approximately 2.5 installations per month. This is less than half of the total treatment effect of more than six installations per month, which suggests that most of the treatment effect from the Solarize campaign cannot be explained by the price reduction alone. Other elements of the campaigns, such as the solar ambassadors and the community-based recruitment, must be playing a more important role in the increased adoptions from the campaigns than the price discount.

5.1.2. How Important Is Group Pricing? We can further explore the importance of pricing by examining the role of group pricing. Theory suggests that with group pricing, installers may benefit from the dissemination of information through peer effects from additional WOM. Given the many psychological drivers of WOM (Zhang et al. 2013, Berger 2014), we might expect the effectiveness of WOM to be altered with the inclusion of group pricing, given that group pricing includes extrinsic motivation for WOM. The presence of a group buy may also change the dynamics within a campaign. Kauffman and Wang (2001) and Kauffman et al. (2010) show that one should expect inertia in group buys, with the greatest uptake at the end of the deal. This intuitively makes sense because consumers have an incentive to wait for more information regarding what the final price would be, but it could also suppress WOM in the first part of the campaigns.

Recall that we randomized municipalities that applied to receive a Solarize campaign during round 5 of

Table 4. Effect of Prices vs. Information

	OLS			IV		
Control group	(1) Caliper	(2) CEC	(3) Future	(4) Caliper	(5) CEC	(6) Future
$Intervention_{it}$	6.76 (0.76)***	6.17 (0.78)***	6.61 (0.79)***	−2.48 (7.61)	3.31 (1.27)***	3.95 (1.48)***
Price per watt $(\$/W)_{it}$	−0.12 (0.12)	−0.09 (0.05)	−0.04 (0.03)	−20.9 (17.64)	−7.48 (3.05)**	−5.76 (3.10)*
No. municipality fixed effects	62	199	153	62	199	153
No. month-of-sample dummies	43	43	43	43	43	43
Observations	1,059	3,705	2,769	1,059	3,705	2,769
First stage *F* statistic				75.76	52.45	19.87
p value on instrument				0.28	0.002	0.06
Effect over entire campaign						
ATET per municipality	37.9	34.5	37.02	−13.9	18.5	22.1
Effect in raw data	44.1	37.9	38.9			

Notes. An observation is a municipality-round-month. The first three column present ordinary least squares (OLS) results, while the second two present instrumental variables (IV) results where we instrument for price with the roofer wage rate. Standard errors in parentheses are block bootstrapped at the municipality level for the OLS results. Because of limited variation across municipalities in the instrument, we show standard errors clustered on municipality for the IV results. *Intervention* is a dummy equal to 1 for a Solarize campaign occurring. The sample is stacked over the rounds, so a municipality may be a control for an earlier round and treated in a later round. Once treated, municipalities are removed from the sample. *Caliper* refers to propensity score matching of the treated municipalities to the three nearest neighbors with a 0.05 caliper. *CEC* refers a control group of all Connecticut Clean Energy Community municipalities except those treated or adjacent to treated. *Future* refers to a control group of municipalities that in future rounds opted-in to a Solarize campaign. *No. municipality fixed effects* reports the number of municipality × round fixed effects. *No. month-of-sample dummies* reports the number of month-of-sample dummies (one or more are dropped due to collinearity). The *ATET per municipality* calculates the average effect over an entire campaign of 5.6 months on average. The *Effect in raw data* calculates the difference between the total installations during a campaign for treatment and control groups, averaged over all municipalities.

***, **, and *Significant at 1%, 5%, and 10%, respectively.

the program into either the Solarize Classic or no group pricing versions of the campaign in order to test the importance of group buys. Online Appendix Figure OA.2 shows the number of installations during the campaign for the treatment municipalities and the control groups, suggesting a similar increase in adoptions regardless of whether group pricing is included.

To estimate causal treatment effects, we again use our primary specification given in Equation (1). Results are shown in columns 1–3 in Table 5. We find that the treatment effect for the number of installations per month is between 3.22 and 5.34, depending on which control group is used. To provide an even more direct comparison, we estimate the treatment effect for just the round 5 Classic municipalities, shown at the bottom of Table 5, which range between 2.95 and 5.63. In both the CEC and future control groups, the point estimate is actually higher when group pricing is removed. We also directly compare the two types of campaigns using only the experimental variation by restricting our sample to using *only* the set of round 5 Solarize Classic and no group pricing campaigns. We find that the coefficient on the Solarize Classic dummy is not significant (−0.73 with a standard error of 2.24), suggesting that the number of adoptions is not appreciably influenced by the group pricing.

We can also examine how the presence of group pricing influences the effect on equilibrium prices. Columns 4–6 in Table 5 show the treatment effect on prices. We observe that the decline in prices because of the campaign without group pricing is smaller than the price declines from the Solarize Classic campaigns in Table 2, but the treatment effect is very noisily estimated and is generally not statistically significant. When we estimate the model using data that include only round 5 Classic and no group pricing campaigns, we find that the coefficient on the Solarize Classic dummy is 0.23 with a block bootstrapped standard error of 0.11, suggesting that the prices are lower *without* group pricing. In hindsight, this makes sense. Group pricing can be advantageous for firms because of its effect on information sharing (Jing and Xie 2011, Chen and Zhang 2014), but the Solarize campaigns are effective in disseminating information even without group pricing. Our findings suggest that in our context, group pricing does not appear to appreciably

Table 5. Impact of the Solarize Intervention Without Group Pricing

	Dependent variable					
	Installations			Prices		
Control group	(1) Caliper	(2) CEC	(3) Future	(4) Caliper	(5) CEC	(6) Future
Intervention_{it}	3.22	3.68	5.34	–0.13	–0.13	–0.18
	(2.42)	(2.14)*	(2.18)**	(0.32)	(0.06)**	(0.19)
Small sample robustness						
Wild cluster bootstrap-t (*p* value)	0.28	0.21	0.07	0.58	0.07	0.39
Randomization inference (*p* value)	0.18	0.04	0.02	0.36	0.19	0.23
No. municipality fixed effects	9	36	9	9	36	9
No. month-of-sample dummies	15	15	15	15	15	15
Observations	135	540	135	135	540	135
R^2	0.51	0.41	0.53	0.10	0.21	0.24
Effect over entire campaign						
ATET per municipality	18.1	20.6	29.9			
Effect in raw data	20.9	24.1	34.7			
Treated town average price ($/W)				4.34	4.34	4.34
Control town average price ($/W)				4.45	4.48	4.45
Round 5 Classic results for comparison						
Intervention_{it}	5.63	2.95	4.62	–0.13	0.10	0.05
	(1.16)***	(1.28)**	(1.50)***	(0.14)	(0.12)	(0.21)

Notes. An observation is a municipality-round-month. Standard errors in parentheses are block bootstrapped at the municipality level. *Intervention* is a dummy equal to 1 for a Solarize campaign occurring. The sample is stacked over the rounds, so a municipality may be a control for an earlier round and treated in a later round. Once treated, municipalities are removed from the sample. *Caliper* refers to propensity score matching of the treated municipalities to the three nearest neighbors with a 0.05 caliper. *CEC* refers a control group of all Connecticut Clean Energy Community municipalities except those treated or adjacent to treated. *Future* refers to a control group of municipalities that in future rounds opted-in to a Solarize campaign. *No. municipality fixed effects* reports the number of municipality × round fixed effects. *No. month-of-sample dummies* reports the number of month-of-sample dummies (one or more are dropped due to collinearity). *Wild cluster bootstrap-t* reports the *p* value for testing the null hypothesis that the treatment has no effect using the wild cluster bootstrap-t procedure from Cameron et al. (2008). *Randomization inference* reports the *p* value for testing the same null hypothesis using randomization inference. The *ATET per municipality* calculates the average effect over an entire campaign of 5.6 months on average. The *Effect in raw data* calculates the difference between the total installations during a campaign for treatment and control groups, averaged over all municipalities. The final two rows calculate the average price during campaigns (in 2014$/W) in treatment and control groups.

***, **, and *Significant at 1%, 5%, and 10%, respectively.

lower prices or increase the number of adoptions, further underscoring that other elements of the campaign, rather than pricing, are the more dominant mechanisms leading to the outcomes we observe.

5.2. How Important Is the Municipality Selection Process?

The CTSC program included all the central tenets of the Solarize program except the competitive bidding process for the installer and the involvement of SmartPower and the Green Bank. In this respect, CTSC provides a useful example of how Solarize could work if it is run without government involvement.[21] In our analysis, we again use the future Solarize controls because the CTSC municipalities selected into the program in the same way as the classic Solarize program. The balance of covariates is shown in Online Appendix Table OA.7. We observe no significant differences across the treatment and control groups in observables. Because of the limited effectiveness of the campaigns, CTSC extended their campaigns by an additional two months over the standard five-month length of a Solarize campaign.

The estimated treatment effects of the CTSC campaigns are displayed in Table 6. Although not statistically significant, our point estimate suggests that the CTSC led to 0.58 additional installations per month (column 1), which is a substantially smaller lift than the more than six additional installations from the classic Solarize campaigns. One explanation for this smaller effect is that the CTSC did not provide as substantial of a price discount. We see this in column 2 of Table 6, where there is a much smaller price decline from the CTSC relative to the Solarize treatment. This finding shows that when we remove the competition at the bidding stage, regardless of whether we have group pricing, there is less of a price decline.

Table 6. CTSC Treatment Effects

	Dependent variable	
	(1) Adoptions	(2) Prices
Treatment$_{it}$	0.58	−0.28
	(0.45)	(0.11)**
Small sample robustness		
Wild cluster bootstrap-t (*p* value)	0.19	0.02
Randomization inference (*p* value)	0.08	0.01
No. municipality fixed effects	85	85
No. month-of-sample dummies	35	35
Observations	1,611	1,611
R^2	0.32	0.13
Effect over entire campaign		
ATET per municipality	3.25	
Effect in raw data	12.05	
Treated town average price ($/W)		4.35
Control town average price ($/W)		4.59

Notes. An observation is a municipality-round-month. Standard errors in parentheses are block bootstrapped at the municipality level. These regressions use the same future controls as in our primary regressions and all small sample robustness rows are the same. There are 10 CTSC-treated towns. The *ATET per municipality* calculates the average effect over an entire campaign of seven months on average (although some CTSC campaigns were longer). The *Effect in raw data* calculates the difference between the total installations during a campaign for treated and control towns, averaged over all towns. The final two rows calculate the average price during campaigns (in 2014$/W) in treated and control towns, respectively.

***, **, and *Significant at 1%, 5%, and 10%, respectively.

The smaller price decline during the CTSC is unlikely to explain the entire difference relative to the Solarize Classic campaigns in the number of installations. For the Classic campaigns, we found an additional treatment effect of over three installations per month *after* controlling for the price reduction. Aegis Solar was also extremely effective in Solarize round 1, so it is unlikely that the difference is because of a lower-quality installer who did not know how to run the Solarize intervention. This leads to a final possibility: that trust in the program is a critical element. The inclusion of third parties, the Green Bank and SmartPower, along with the competitive bidding process, provided potential customers with more trust in the installer and thus more trust in the process. Grayson et al. (2008) explicitly show that customers' trust in the firm(s) is a necessary mediator for trust in the market context. Indeed, in the Solarize program, the installers were selected by the municipalities and referred to as vetted installers.

5.3. Social Learning and WOM

To more deeply understand the mechanisms driving the treatment effects, we surveyed solar PV adopters after each Solarize round. This survey was performed through the Qualtrics survey software and was sent to respondents via email, with an iPad raffled off as a reward for responding. The email addresses came from Solarize event signup sheets and installer contract lists. Approximately 6% of the signed contracts did not have an email address. All others were contacted one month after the end of the round, with a follow-up to nonrespondents one month later. The overall response rate across the five rounds of Classic was 42%; this is a very high response rate for an online survey, which is a testament to the enthusiasm of the adopters in solar and the Solarize program.

We are especially interested in how solar adopters found out about the program. One question in our survey provides 14 possible factors that influence the decision to install solar. The question asked respondents to "rate the importance of each factor in your decision to install solar PV," with the following possible answers: extremely important, very important, somewhat important, not at all important. Figure 8 shows the number of survey respondents that rated the information sources as extremely important and very important, for each of the 14 information sources. Several of these sources rely on social learning. Indeed, the top five sources of information listed as either extremely or very important (not including "other") involve social learning: the "town information event," a "friend or neighbor's recommendation," a "recommendation of someone you interact with in your town," the "solar ambassador," and "seeing solar on another home or business."[22] The only social information not rated highly is a "recommendation of someone you work with," which is not surprising because the campaigns leverage social interactions within the communities where people live rather than workplaces. These survey results provide evidence that the Solarize intervention may be working exactly as intended: by fostering social learning.

6. Cost-Effectiveness and Welfare

In this section, we assess cost-effectiveness. SmartPower provided us with their cost breakdown of the five round of Solarize Classic, which is mostly covering staff time. For the 34 Classic campaigns, the total cost to SmartPower was $800,000. In addition, we surveyed the Solarize installer firms, who reported costs of approximately $5,000 per municipality (e.g., mailers, additional staff time). These lead to a combined cost figure of $28,500 per town. With an estimated treatment effect of 6.63 installations per month (using the future Solarize town control group), the direct program cost for each new installation from the program is $860.

From the solar installer firm's perspective, there was also the discount of $0.46/W. For an average system size of 4.23 kW, this amounts to $1,690 per installation. Adding in the direct marketing expenditures

Figure 8. (Color online) Importance of Information Channels

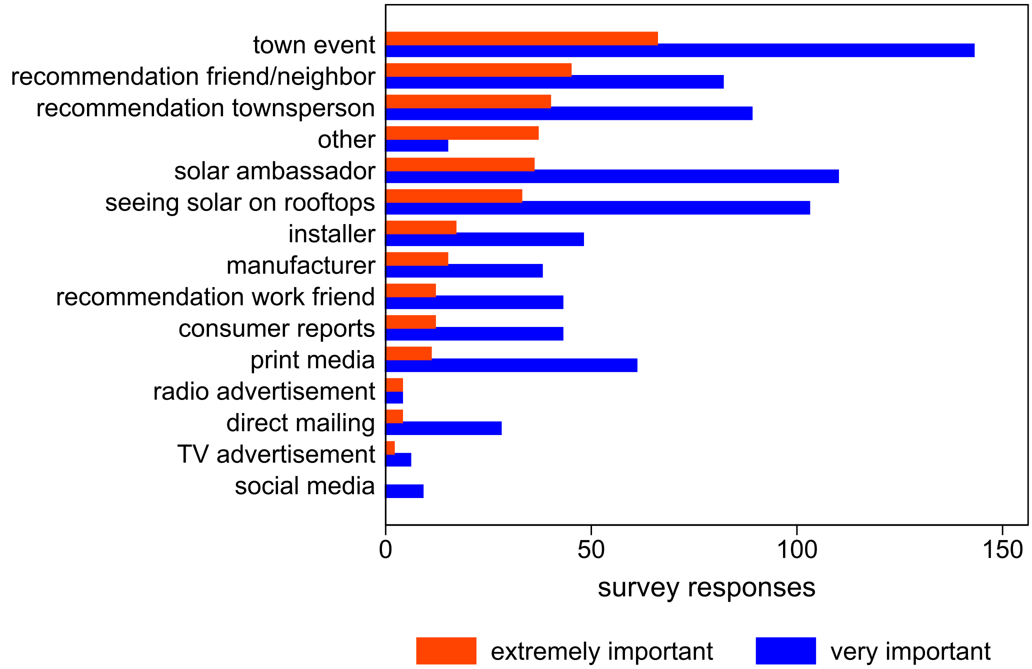

by the installers, this means that the solar firm spent roughly $1,840 per installation installed through Solarize. This compares favorably to installers' reported consumer acquisition costs of $1,500–$3,000 per installation when the installers are finding customers on their own. The main advantage to installers of participating in the Solarize campaigns is that they provide many more leads and adoptions, allowing for more business for the installers, which also provides greater economies of scale. Of course, the major reason for policymakers to support Solarize campaigns is because of the environmental benefits. In 2012, electricity on the Connecticut grid had a carbon intensity of 547 lb CO_2/MWh. Using estimates of expected solar electricity generation at the municipality level from the Green Bank, we find that, in total, the 1,127 additional installations from the Solarize Classic campaigns led to a reduction of 3,683 lb carbon annually. Assuming the same carbon intensity over time, the total reduction over the 25-year lifespan of the solar panels is 92,100 lb.

Applying our estimates of direct program cost and carbon reductions yields a cost per ton of avoided CO_2 from the Solarize program of $20.61 (the cost per ton would be higher if the New England electricity grid decarbonizes). In comparison, the 2019 central estimate of the social cost of carbon is about $50/ton (in 2019 dollars) (Environmental Protection Agency 2016), although some recent work has put the price at more than $100/ton (Daniel et al. 2019). Thus, the Solarize program very likely increases social welfare based on the environmental benefits alone.[23]

7. Conclusions

This paper contributes to the literature on prosocial behavioral interventions. The Solarize program, which draws on several theoretical and empirical findings in previous work, is expanding rapidly and could be applied to other energy-saving or renewable energy technologies. We find very strong treatment effects from the program: an increase in installations by 37 per municipality, which is more than a 1,000% increase from the control group rate, and a decrease in preincentive equilibrium prices of $0.46/W. We use a field experiment to demonstrate that these programs can also increase installations in randomly selected municipalities (although with a slightly smaller lift), providing guidance on the external validity of our results.

Our research also delves into the mechanisms underlying this result. We show that the discount pricing is a secondary factor influencing the success of the campaigns, and in a second side field experiment with the same campaigns but without the group pricing, we show that group pricing is not essential to the lift in adoptions. We also examine a similar campaign with all the central tenets of Solarize, only without competition for the chosen installer and run without the participation of SmartPower and the Green Bank. This campaign led to far fewer installations, suggesting the value of both the installer selection process and trust in the campaign organizers.

Our survey results highlight the importance of social learning and information provision within the campaign. Our calculations reveal that the program is surprisingly cost-effective, with a direct program cost of \$21/ton of CO_2 reduced, which is less than half of common estimates of the social cost. Although residential rooftop solar alone will not solve the world's dependency on fossil fuels, the efficacy of the Solarize campaigns underscores the promise of leveraging social learning in other energy-related behaviors for reducing climate risks.

Acknowledgments

The authors contributed equally to this manuscript. The authors thank J. R. DeShazo, Pedro Gardete, Matthew Harding, Stefan Lamp, Josh Graff-Zivin, Arndt Reichert, Charles Towe, and the participants at many seminars for useful comments. They also thank Hilary Staver, Brian Keane, Toni Bouchard, Lyn Rosoff, Kate Donnelly, Bernie Pelletier, Bob Wall, Robert Schmitt, Stuart DeCew, Jen Oldham Rogan, and the Yale student team. All data and code for this paper can be accessed at https://doi.org/10.3886/E120503V1. This report was prepared as an account of work sponsored by an agency of the U.S. Government. Neither the U.S. Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.

Endnotes

[1] This is according to both the U.S. National Climate Assessment and the UN Intergovernmental Panel on Climate Change (IPCC).

[2] States that have implemented Solarize include Oregon, Washington, California, Colorado, South Carolina, North Carolina, Ohio, Pennsylvania, New York, Rhode Island, Massachusetts, and Vermont.

[3] All dollars in this paper are 2014\$.

[4] See https://www.ctgreenbank.com/wp-content/uploads/2019/01/RSIP-Legislative-Report-2019.pdf.

[5] Estimates based on authors' calculations from solar installation data and potential market data based on satellite imaging from Geostellar (2013). The potential market data are focused on the shading of households, but accounts for the possibility of some ground-mounted systems. Ground-mounted systems are more expensive, and they make up only a small percent of the systems.

[6] Net metering allows excess solar PV production to be sold back to the electric grid at retail rates, with a calculation of the net electricity use occurring at the end of each month. Any excess credits remaining on March 31 of each year receive a lower rate.

[7] As of 2014, roughly 37% of all systems installed were third party-owned, and these third party-owned systems were distributed across Connecticut and not concentrated in any particular municipalities. From 2015 to 2017, third-party ownership increased greatly but since then has been decreasing.

[8] Gillingham and Tsvetanov (2016) estimate the pass-through of state rebates in the Connecticut solar market during a similar time period and find that only roughly 16% of the rebates are captured by firms.

[9] The programs were funded by the Connecticut Green Bank, The John Merck Fund, The Putnam Foundation, and a grant from the U.S. Department of Energy.

[10] Kraft-Todd et al. (2018) does not estimate the treatment effect of the Solarize campaigns.

[11] If a rooftop or system requires additional effort, installers are permitted to include Green Bank–approved adders to the price.

[12] There were six Solarize campaigns run after the end of our study period, and we use these as controls.

[13] In our first round of Solarize, we had more municipalities submit applications than we could accept and thus we could randomize. However, the sample was too small to perform an adequately powered analysis.

[14] The exception is any installation performed in the three small municipal utility regions of Wallingford, Norwich, and Bozrah. Given their ineligibility for the state rebates, we expect few installations in these areas.

[15] These data include total voter registration and the number of active and inactive registered voters in each political party (CT SOTS 2015).

[16] During this round, we also randomized some municipalities into a campaign based around the online platform of EnergySage, but we do not cover this arm of the experiment in this paper.

[17] Some contiguous municipalities are run as joint campaigns, such as Mansfield and Windham in order to reduce costs. However, both municipalities still receive the full treatment.

[18] This is an OLS fixed-effects specification, but we find similar results using a negative binomial specification.

[19] We also explore other assumptions, such as including a longer preperiod, no preperiod, or a one-year postperiod. We find extremely similar results.

[20] For reference, the average price in the control municipalities is \$4.63/W, so this is a substantial price decline.

[21] Although Aegis Solar created and funded CTSC, it is technically a nonprofit organization.

[22] "Other" is the third ranked source of information when ranking only by the number of respondents rating the source as "extremely important."

[23] Our calculations are not a full social welfare analysis, which would also account for the other subsidies for solar and all of the other externalities. The subsidies are a transfer from the government to consumers, but there may be a social cost from raising the tax revenue. If we assume a 10% marginal social cost of public funds to pay for the subsidies, the cost per ton of avoided CO_2 rises to \$44.47.

References

Anderson SJ, Chandy R, Zia B (2018) Pathways to profits: The impact of marketing vs. finance skills on business performance. *Management Sci.* 64(12):5559–5583.

Anderson SJ, Chintagunta P, Vilcassim N (2020) Stimulating marketing strategy innovation with entrepreneurs in Uganda: Examining the impact of skype-aided business coaching on firm sales. Working paper, University of Chicago, Chicago, Illinois.

Ashraf N, Bandiera O, Jack BK (2014) No margin, no mission? A field experiment on incentives for public service delivery. *J. Public Econom.* 120:1–17.

Bayer P, Pintoff R, Pozen D (2009) Building criminal capital behind bars: Peer effect in juvenile corrections. *Quart. J. Econom.* 124(1):105–147.

Ben-David I, Graham JR, Harvey CR (2013) Managerial miscalibration. *Quart. J. Econom.* 128(4):1547–1584.

BenYishay A, Mobarak AM (2017) Social learning and incentives for experimentation and communication. *Rev. Econom. Stud.* 86(3): 976–1009

Berger J (2014) Word of mouth and interpersonal communication: A review and directions for future research. *J. Consumer Psychol.* 24(4):586–607.

Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust difference-in-differences estimates. *Quart. J. Econom.* 119(1): 249–275.

Bloom N, Eifert B, Mahajan A, McKenzie D, Roberts J (2013) Does management matter? Evidence from India. *Quart. J. Econom.* 128(1):1–51.

Bollinger B, Burkhardt J, Gillingham K (2020) Peer effects in water conservation: Evidence from consumer migration. *Amer. Econom. J. Econom. Policy* 12(3):107–133.

Bollinger B, Gillingham K (2012) Peer effects in the diffusion of solar photovoltaic panels. *Marketing Sci.* 31(6):900–912.

Bollinger B, Gillingham K (2019) Learning-by-doing in solar photovoltaic installations. Working paper, Yale University, New Haven, CT.

Bowles S, Polania-Reyes S (2012) Economic incentives and social preferences: Substitutes or complements? *J. Econom. Literature* 50(2):368–425.

Brandon A, List J, Metcalfe R, Price M, Rundhammer F (2018) Testing for crowd out in social nudges: Evidence from a natural field experiment in the market for electricity. *Proc. National Acad. Sci. USA* 116(12):5293–5298.

Cameron C, Gelbach J, Miller D (2008) Bootstrap-based improvements for inference with clustered errors. *Rev. Econom. Statist.* 90(3):414–427.

Carattini S, Péclat M, Baranzini A (2018) Social interactions and the adoption of solar pv: Evidence from cultural borders. Working paper, London School of Economics, London.

Chen Y, Zhang T (2014) Interpersonal bundling. *Management Sci.* 61(6):1456–1471.

Coffman L, Featherstone C, Kessler J (2014) Can social information affect what job you choose and keep? Working paper, Ohio State University, Columbus, OH.

Cohen J, Dupas P (2010) Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Quart. J. Econom.* 125(1):1–45.

Conley T, Udry C (2010) Learning about a new technology: Pineapple in Ghana. *Amer. Econom. Rev.* 100(1):35–69.

CT SOTS (2015) Registration and enrollment statistics data. Accessed June 1, 2015, http://www.sots.ct.gov/sots.

Daniel KD, Litterman RB, Wagner G (2019) Declining co2 price paths. *Proc. National Acad. Sci. USA* 116(42):20886–20891.

DellaVigna S, List J, Malmendier U (2012) Testing for altruism and social pressure in charitable giving. *Quart. J. Econom.* 127(1):1–56.

Dubé JP, Luo X, Fang Z (2016) Self-signaling and prosocial behavior: a cause marketing experiment. *Marketing Sci.* 36(2):161–186.

Duflo E, Saez E (2003) The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quart. J. Econom.* 118(3):815–842.

Elberg A, Gardete PM, Macera R, Noton C (2019) Dynamic effects of price promotions: A large-scale field experiment. *Quantitative Marketing Econom.* 17(1):1–58.

Environmental Protection Agency (2016) Technical update of the social cost of carbon for regulatory impact analysis under executive order 12866. Technical report, Interagency Working Group on Social Cost of Greenhouse Gases, Washington, DC.

Ferraro P, Price M (2013) Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment. *Rev. Econom. Statist.* 95(1):64–73.

Ferraro P, Miranda JJ, Price M (2011) The persistence of treatment effects with norm-based policy instruments: Evidence from a randomized environmental policy experiment. *Amer. Econom. Rev.* 101(3):318–322.

Fisher AR (1935) *The Design of Experiments* (Oliver and Boyd, London).

Foster A, Rosenzweig M (1995) Learning by doing and learning from others: Human capital and technical change in agriculture. *J. Political Econom.* 103(6):1176–1209.

Frey B, Meier S (2004) Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment. *Amer. Econom. Rev.* 94(5):1717–1722.

Friedman B, Ardani K, Feldman D, Citron R, Margolis R, Zuboy J (2013) Benchmarking non-hardware balance-of-system (soft) costs for U.S. photovoltaic systems using a bottom-up approach and installer survey-second edition. National Renewable Energy Laboratory Technical Report, National Renewable Energy Laboratory, Golden, CO.

Geostellar (2013) The addressable solar market in connecticut. Report for CEFIA, Geostellar and Martinsberg, West Virginia.

Gillingham K, Palmer K (2014) Bridging the energy efficiency gap: Policy insights from economic theory and empirical analysis. *Rev. Environ. Econom. Policy* 8(1):18–38.

Gillingham K, Tsvetanov T (2016) Hurdles and steps: Estimating demand for solar photovoltaics. Working paper, Yale University, New Haven, CT.

Giné X, Yang D (2009) Insurance, credit, and technology adoption: Field experimental evidence from Malawi. *J. Developed Econom.* 89(1):1–11.

Glaeser E, Sacerdote B, Scheinkman J (1996) Crime and social interaction. *Quart. J. Econom.* 111(2):507–548.

Glaeser E, Laibson D, Scheinkman J, Soutter C (2000) Measuring trust. *Quart. J. Econom.* 115(3):811–846.

Gneezy U, Meier S, Rey-Biel P (2011) When and why incentives (don't) work to modify behavior. *J. Econom. Perspective* 25(4): 191–210.

Grayson K, Johnson D, Chen DFR (2008) Is firm trust essential in a trusted environment? How trust in the business context influences customers. *J. Marketing Res.* 45(2):241–256.

Graziano M, Gillingham K (2015) Spatial patterns of solar photovoltaic system adoption: The influence of neighbors and the built environment. *J. Econom. Geography* 15(4):815–839.

Griliches Z (1957) Hybrid corn: An exploration in the economics of technological change. *Econometrica* 25(4):501–522.

Harding M, Hsiaw A (2014) Goal setting and energy conservation. Working paper, Duke University, Durham, NC.

Hausman N, Condee N (2014) Planning and Implementing a Solarize Initiative: A Guide for State Program Managers. *Clean Energy States Alliance Guidebook, Clean Energy States Alliance, Montpelier, VT*.

Jacobsen G, Kotchen M, Clendenning G (2013) Community-based incentives for environmental protection: The case of green electricity. *J. Regulatory Econom.* 44:30–52.

Jing X, Xie J (2011) Group buying: A new mechanism for selling through social interactions. *Management Sci.* 57(8):1354–1372.

Karlan D, List J (2007) Does price matter in charitable giving? Evidence from a large-scale natural field experiment. *Amer. Econom. Rev.* 97(5):1774–1793.

Kauffman RJ, Wang B (2001) New buyers' arrival under dynamic pricing market microstructure: The case of group-buying discounts on the Internet. *J. Management Inform. Systems* 18(2): 157–188.

Kauffman RJ, Lai H, Ho CT (2010) Incentive mechanisms, fairness and participation in online group-buying auctions. *Electronic Commerce Res. Appl.* 9:249–262.

Kessler J (2014) Announcements of support and public good provision. Working paper, University of Pennsylvania, Philadelphia.

Kraft-Todd GT, Bollinger B, Gillingham K, Lamp S, Rand DG (2018) Credibility-enhancing displays promote the provision of a non-normative public good. *Nature* 563(7730):245.

Kremer M, Miguel E, Mullainathan S, Null C, Zwane AP (2011) Social engineering: Evidence from a suite of take-up experiments in Kenya. Working paper, Harvard University, Cambridge, MA.

LaRiviere J, Price M, Holladay S, Novgorodsky D (2014) Prices vs. nudges: A large field experiment on energy efficiency fixed cost investments. Working paper, University of Tennessee, Knoxville, TN.

List J, Price M (2009) The role of social connections in charitable fundraising: Evidence from a natural field experiment. *J. Econom. Behav. Organ.* 69:160–169.

McKenzie-Mohr D (2013) *Fostering Sustainable Behavior: An Introduction to Community-Based Social Marketing* (New Society Publishers, Gabriola Island, BC).

Mian A, Sufi A (2012) The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. *Quart. J. Econom.* 127(3):1107–1142.

Miguel E, Kremer M (2004) Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72(1):159–217.

Narayanan S, Nair H (2013) Estimating causal installed-base effects: A bias-correction approach. *J. Marketing Res.* 50(1):70–94.

Rode J, Weber A (2016) Does localized imitation drive technology adoption? A case study on rooftop photovoltaic systems in Germany. *J. Environ. Econom. Management* 78:38–48.

Rosenbaum P, Duflo E, Mullainathan S (2002) Covariance adjustment in randomized experiments and observational studies. *Statist. Sci.* 17(3):286–327.

Sianesi B (2004) An evaluation of the swedish system of active labor market programs in the 1990s. *Rev. Econom. Statist.* 86(1): 133–155.

Sorensen A (2006) Social learning and health plan choice. *RAND J. Econom.* 37(4):929–945.

Thaler R, Sunstein C (2009) *Nudge: Improving Decisions About Health, Wealth, and Happiness* (Penguin, London).

Towe C, Lawley C (2013) The contagion effect of neighboring foreclosures. *Amer. Econom. J. Econom. Policy* 5(2):313–335.

Tucker C (2008) Identifying formal and informal influence in technology adoption with network externalities. *Management Sci.* 55(12):2024–2039.

Vasilaky K, Leonard K (2011) As good as the networks they keep? improving farmers' social networks via randomized information exchange in rural Uganda. Working paper, Columbia University, New York.

Yoganarasimhan H (2015) Estimation of beauty contest auctions. *Marketing Sci.* 35(1):27–54.

Zhang Y, Feick L, Mittal V (2013) How males and females differ in their likelihood of transmitting negative word of mouth. *J. Consumer Res.* 40(6):1097–1108.