Low-Income Residential Solar: A SASH Evaluation

Jeff Knaide

Under the Supervision of:

Prof. Christopher Timmins

Prof. Kent Kimbrough

Honors thesis submitted in partial fulfillment of the requirements for Graduation with Distinction

Department of Economics, Duke University

Durham, North Carolina

2017

Acknowledgements: I would like to thank Professor Christopher Timmins for guiding me through the research process, offering both sound advice and strong ideas at every stage of this process. I would also like to thank Professor Kent Kimbrough for guiding my thoughts and reading countless drafts as this paper developed throughout the year. This paper would not have been possible without the support and experience that you two have offered me throughout this long and educational process.

Abstract: In this paper, I examine the impacts of California's Single-Family Affordable Solar Housing (SASH) subsidy on the rate of adoption of residential solar power. The SASH program looks to provide low-income families with a sizeable subsidy to install residential solar panels. Eligibility for the program depends on income, among a few other factors. This work represents part of a small body of energy justice literature, and the only existing evaluation of SASH. Zip-codes with higher eligibility (based on income levels) showed a significantly higher number of adoptions when controlling for important characteristics, specifically median income. While this policy did generate low-income adoptions, it does not offer a strong carbon abatement strategy – low-income households require greater financial support than might higher-income households.

JEL: Q29

Keywords: solar, energy justice, environmental justice, Single-Family Affordable Solar Housing

I. Introduction

Residential solar represents a growing trend in today's energy market. Solar photovoltaic (PV) panels are among its most commercially accessible forms. In recent years, the number of firms offering panels and contractors installing panels in California, the leading state in solar energy production, has become quite large. However, several hurdles to solar PV adoption remain, the primary being initial financial investment. Financial barriers may have differential impacts on income groups, with high-income households able to shoulder significantly greater upfront financial costs than otherwise similar low-income households. Energy and Environmental Economics, Inc. (2013) found that, in 2013, the median income of solar adopters connected to the grid was $91,210, well above the overall median California household income of $54,283. Moreover, Kelly (2014), found that low-income households account for only 2.5% of all solar installations. These two figures paint a jarring picture of differential adoption of residential solar PV. Low-income households have made the decision to participate directly in clean energy generation far less frequently than have high-income households. The California Public Utilities Commission has launched a number of campaigns to assist low-income homes in overcoming financial barriers, one of which will be the focus of my empirical analysis. The Single-Family Affordable Solar Housing intends to create a greater space for low-income participation in the solar market by offering a subsidy to qualifying low-income households. The main concern of my research will be how its late 2007 launch bolstered adoption rates in low-income zip codes as compared with rates in high-income zip codes.

A significant issue that emerges with differential solar take-up across income groups couches itself in the justice issues surrounding electricity pricing. California, the center of my analysis, charges consumers entirely volumetrically (i.e. by the number of kilowatt hours consumed). The costs of energy production, however, come largely from capital and infrastructure costs that do not hinge on quantity of electricity provided. In her paper *Clean Energy Justice*, Shelley Welton (2016) explains that Net Energy Metering (NEM) policies allow those who with solar PV to sell energy back to the grid in order to “run the meter backwards,” reducing their energy bill, sometimes all the way to zero. Welton (2016) also discusses the interaction between differential take-up and energy cost structures: “lower-income consumers who cannot afford the panels are left shouldering a mounting proportion of grid maintenance costs.” Environmental justice issues arise at this juncture: whether intentional or not, imposing costs on the poor in order to lessen costs for the wealthy raises an ethical dilemma. High-income solar owners are essentially cross-subsidized by the low-income ratepayers who are unable to afford the upfront costs of adoption. Alternative structures that charge a base rate for access to the grid in addition to volumetric charges have been employed in some contexts, such as Arizona. However, as Welton points out, these structures can stymie the growth of clean energy programs by reducing the financial benefits of adoption. When considering the broader goal of increasing renewable energy’s share of our energy portfolio (to decrease greenhouse gas emissions), minimizing energy justice issues by simply minimizing overall adoption is not wholly satisfactory.

In discussing the cross-subsidization through grid infrastructure, it is valuable to understand how energy provider cost and pricing structures compare. Blackburn, Magee, and Rai (2014) investigate modern utility response to residential solar PV, finding that nine of the thirteen utilities with above-median solar PV penetration recover 50 percent or more of their fixed costs through volumetric charges. Because firms regain the bulk of their fixed costs through volumetric charges, as solar PV consumers purchase less energy, remaining rate-payers each shoulder a larger share of fixed costs. Net Energy Metering (NEM) heightens this issue by allowing solar PV users to further reduce energy costs by selling all surplus energy back to the grid. Solar PV holders use the grid both when they are producing too much and too little energy but do not pay for the maintenance costs vital to such an exchange. With NEM acting as a strong performance subsidy for solar, firms must shift their expenses to non-solar customers. Blakburn, Magee, and Rai (2014) note this, citing California utilities' plans to shift sums varying from $200 million to $700 million in annual costs onto non-solar customers. Given that solar adopters are disproportionately high-income, subsidizing their solar usage with increased costs to ratepayers draw environmental justice concerns.

If it is a goal of policy is to minimize the difference between high-income and low-income adoption, researchers and policymakers must first understand the barriers that prevent participation. Balcombe et al (2014) performed best-worst surveys to determine the degree to which different factors influenced respondent adoption decisions in the United Kingdom. The authors found that the two largest blockages are financial and informational. According to Balcombe et al (2014), the key financial barriers to adoption, which will function as the focus of my empirical work, are high upfront capital costs as well as potential financial loss upon moving. These financial barriers call for a solution that turns the risk away from consumers, something that rings especially true for low-income households with limited ability to support increased total monthly expenses.

In this paper, I will build and execute a model focused on the effects of reducing the financial barrier to adoption in low-income communities. The mechanism for reducing or eliminating that barrier is the SASH subsidy. The specific target of my examination will be how increasing access to this subsidy within a zip-code, conditional on constant median income, affects total number of adoptions. Utilizing the effect sizes output by my model, I will extrapolate implications in terms of energy justice and carbon abatement. In doing so, I will show that the subsidy provides significant energy justice support, but is incredibly weak from a cost of carbon abatement perspective.

The remainder of the paper is organized as follows: section 2 provides background information on the SASH subsidy; section 3 offers a review of existing literature in solar adoption and energy justice; section 4 outlines the data used in this analysis; section 5 explains my strategy for identifying those affected by the subsidy and empirical model; section 6 analyzes my regression results; section 7 extrapolates policy implications; and section 8 contains a concluding discussion of my results and their implications.

II. Single-Family Affordable Solar Housing (SASH)

The major focus of my research will be California's Single-Family Affordable Solar Housing (SASH) program. Evident in their espoused goals, regulators concerned themselves primarily with alleviating environmental justice and equity concerns related to a growing Solar PV base. Approved by the California Public Utilities Commission (CPUC) in 2007, SASH targeted solar PV adoption among low-income households by providing incentives based on the expected performance of applicant systems. To ensure the benefits reached their targeted group, administrators imposed a number of eligibility requirements for participation in the program.

SASH intended to lessen the solar ownership demographic gap that had emerged as a result of large capital barriers associated with Solar PV adoption. The program handbook proposed a number of smaller goals directed at achieving that objective. The first three goals of the project were directed at similar space: subsidization as a tool to ease financial burdens through adoption. These goals, presented on the SASH program web page, were to "reduce energy bills without increasing monthly expenses;" to "provide full and partial incentives for solar systems for low-income participants;" and to "decrease the expense of solar ownership with a higher incentive." These goals are interrelated in their designation of eliminating financial blocks as a key to solving equity issues in residential solar PV. As such, they signal that SASH intended to generate greater access to solar, making worthwhile an examination of its impacts on low-income adoption rates.

In order to effectively push solar ownership toward increasingly equitable income distribution, SASH had to set eligibility requirements. These requirements focused largely on ensuring that applicants fell into the income group of interest. The foremost precondition for SASH participation was household income no greater than 80 percent of the area median income (AMI). The Department of Housing and Community Development updates the AMI yearly with median income by household size. The area level of focus is counties, of which there are fifty-eight in California. The AMI has a large degree of variation, with the Government of Housing and Community Development (2015) reporting 2015 four-person household medians ranging from $57,900 to $106,300, opening the door for higher income households to qualify under this engagement. In order to ensure that benefits reached the intended group, the CPUC included an additional requirement, noted in the SASH Handbook (2015), fulfilled by residence in a home with an equity-sharing agreements. Both the financial and equity-sharing requirements must be met in order to qualify for the subsidy. The strength of these requisite applicant characteristics inspires confidence that these subsidies enabled low-income households to overcome the financial barriers to adoption, with little money spent on less-than-needy households. I look to draw conclusions about low-income households' elasticity of adoption with respect to the income cutoff by studying the effects of SASH on adoption. The final two requirements designate an eligible applicant as one who is owner and occupier of a home connected to one of the three major Investor-Owned Utilities (IOUs): South California Edison (SCE), San Diego Gas & Electricity (SDG&E), or Pacific Gas & Electric (PG&E).

Should the applicant fit into the eligibility requirements, regulators administer SASH funding once the solar system is purchased, installed, interconnected, and inspected. A "funding reservation" offers the purchaser assurance that reserved funds will be available when the payment claim is made, likely easing some of the liquidity issues that might otherwise emerge. The SASH handbook (2015) indicates that money is given in an "Expected Performance Based Buydown (EPBB) structure to help reduce the cost of installation." Because the program is intended to mitigate the upfront capital costs barring low-income participation, this payment structure makes sense. The size of SASH payouts currently sits at $3 per Watt of the system. This incentive is significant; the median system size in my sample was 4.7 kilowatts. The median price per watt for residential solar panels in my sample was $7.26, so the subsidy is approximately 41% of the total cost of a system. The price of solar panels has been falling throughout that time period, though, with a median price of $5.52 in the final year of my sample (2013).

My empirical analysis will attempt to draw conclusions about the causal effects of these programs by examining the rates of adoption in different zip-codes starting in 2008. The primary reason for this temporal cutoff is fit: observations in 2008 are the earliest defined by the 2010 census boundaries used in later years. The focus of my evaluation is the period following 2008, so this is the sensible time period to use. Unfortunately, this prevents any comparison between adoption before and after the subsidy opened.

III. Literature Review

There is little, if any, empirical work that focuses on the intersection of solar PV adoption and environmental justice. Many facets of solar adoption have seen immense research, while this one remains relatively understudied. Shelley Welton's paper *Clean Energy Justice* has acted as a major motivating force for my research. Meanwhile Paul Balcombe, Dan Rigby, and Adisa Azapagic's paper *Investigating the importance…uptake in the UK* (2014) enabled a better understanding of the barriers to adoption further investigated in my paper. Meanwhile has provided information on the effects of other solar PV subsidy programs.

Welton (2016) provided significant insight into the effects of differential solar take-up, while also highlighting some of the issues inequity creates in solar policy. Her most relevant exploration can be found in her discussion of Net Energy Metering (NEM). Welton (2016) explains that, though there is mixed evidence about whether net metering accurately compensates solar, "it is not clear whether remaining grid customers, as opposed to society as a whole, should shoulder the costs of attaining solar's non-monetized environmental benefits." This question of cost-bearing is an important one in determining the equitable distribution of residential solar PV. Welton (2016) discusses the influence of a particularly effective anti-NEM campaign on Arizona regulators: this campaign focused on the inequity in forcing ratepayers (specifically those of the lower-income and working classes) to "cover the costs for people who choose to add solar panels." Inequity slowed the growth of solar panels, highlighting the need to develop policy solutions to ensure groups are equally represented in the community of solar owners. In my research, I aimed to determine whether SASH significantly moved California's ownership toward a more even distribution.

In order to gauge value of these programs, one must gain an understanding of the drivers of solar PV uptake. Balcombe et al (2014) researched the major motivations for and barriers to microgeneration uptake, though their study focused primarily on the UK. Balcombe et al (2014) refer to financial factors as the most important barriers, with participating non-adopters finding the prospect of losing money when moving homes to be of the most importance. Balcombe et al (2014) find a second dominant barrier to adoption to be access to trustworthy information.

Another valuable piece of insight developed by Molly Podolefsky addressed the impacts of subsidization on take-up rates. Though this paper did not direct its efforts at the divide between low- and high-income rates, it does provide some information on how subsidies effect solar adoption. When Podolefsky (2013) examined a 7 percent increase in mean rebate rate, the number of installations per day rose 11 to 15 percent. This finding follows economic intuition: one would expect adoption rates to increase as prices fall. It will be interesting to examine how this impact rises or falls when considering only low-income groups. I hypothesize that the increase will be greater, as lower-income groups likely have limited access to credit, among other tools to help overcome financial barriers. That being said, it is possible that, because low-income households need greater assistance, they are less responsive to small subsidies and more responsive to very high subsidies.

IV. Data

In putting together this dataset, I've collected items from a number of different sources. In order to isolate the effects of the SASH subsidy, I've controlled for variation in the demographic, economic, and social characteristics of residents of different zip-codes. I've gathered demographic data from the American Community Survey's (ACS) 5-year moving averages, taking from the 2006-2010, 2007-2011, 2008-2012, 2009-2013, 2010-2014, and 2011-

2015. This data is at the census tract level, using definitions from the 2010 decennial census. I collected the adoption data from the California Solar Initiative at the zip-code level. I converted characteristics to zip-code level, requiring a few key assumptions, specifically that zip-codes feel spillover that does not alter demographics and have homogenous populations, in order to remedy incongruity with adoption data. I'll discuss these assumptions in detail later in this section.

The two centerpieces of my research have been provided by the California Solar Initiative (CSI) and the American Community Survey (ACS). The CSI dataset provides information on each adoption, including the date and size of installation and zip code of the home. I have used these features to generate panel data in which each observation has the number of installations for a given zip-code for one year between 2008 and 2013. I have pulled ACS data from 5-year moving average surveys, with each survey represented by its middle year (e.g. the 2006-2010 survey is used for 2008). Because there was a change in the districting of California with the 2010 census, any moving averages that ended before that year don't fit well with the rest of my data. As a result, I chose to focus my study primarily on the years 2008 through 2013. Little relevant is lost with this decision, since the policy went into effect in late 2007. However, it would have been helpful to get a sense of what solar adoption looked like prior to the implementation of this program.

In order to perform my analysis, I needed panel data for 2008-2013 for each zip-code in California. To build this, I aggregated individual adoption data into counts by zip code and year, such that each zip-code has six observations (one for each year). Building a fixed effects regression will require this panel data in order to sort out the unobservable effects specific to each zip-code over time. I also put together a variable on the cumulative past solar installations in each zip code each year (*cuminstalls*). This covariate is valuable in controlling for small zip-codes that see adoptions curtail with increasing residential saturation. Alternatively, it will be valuable in addressing potential increases in available information: additional installs in the zip-code provide another point of reference for non-adopters to gain trustworthy information about the benefits of solar.

As noted above, the ACS data that I have put together is assigned to the years 2008 through 2013, though it takes survey data from as early as 2006 and as late as 2015. This data provides the characteristics at the tract level across a number of social and economic traits. The reason that I chose to use different surveys for each year (as opposed to using the 5-year moving averages to represent each of the 5 years they contain), was to generate some degree of variation within the demographic variables in my dataset. Had I just used a single survey, these characteristics would have been filtered out by the fixed effects, such that I would not have been able to control for their influence on adoption. Based on an initial analysis of my variables and their effects on adoption, I chose to include median household income (*medhhincome*), average household size (*avghhsize*), median number of rooms (*medhhsize*), and a suite of education covariates as control variables. The means, standard deviations, and ranges of these control variables are presented in Table 1. Other controls that I considered, such as race variables, were insignificant statistically or economically (even when considered jointly) when including for the variables mentioned above.

Area median incomes (AMI) are an important county-level piece of data that I have gathered. The AMI for each county sets the threshold for participation in the SASH program and as such is vital to the identification strategy I will describe in section V. The California Housing and Urban Development agency publishes the AMI figures for each county annually. I have pulled the AMIs for each year between 2008 and 2013 and calculated the 80 percent mark used in determining SASH eligibility. Based on these figures and the income buckets provided in the ACS data, I was able to construct an estimation of the percentage of people in each zip-code that were financially eligible for the subsidy. I used this percentage as the primary covariate (*eligible*) of interest in my regressions. I've provided a more detailed discussion about the building of this variable and its distribution in the following section (Identification Strategy).

I moved the demographic data from census tract to zip-code level so that it would align properly with my data on solar installations. Zip-codes are generally geographically larger than census tracts, so it is most feasible to move from tracts to zips in connecting my data. This conversion required an assumption of equal spillover of populations and tracts across zip-codes, uncorrelated with household characteristics. Ideally, each census tracts would exist within a single zip code, without spilling over into another. However, in many cases, boundaries do not align, and multiple tracts correspond to a single zip code and one tract also corresponds to multiple zips. The two solutions to this issue are to either double count tracts that fall under into two zips or to leverage areal containment strategies using ArcGIS. For my research, I intend to simply double count tracts that occur in two different zip codes. The major weakness of double counting comes with weighting each tract by population size when aggregating into the zip-code level. Of the approximately 8,000 census tracts in California, around half exist within only one zip-code, while another eighth exist within only two. The greatest number of zip-codes within one census tract is 14. The averages will be incorrect if a large portion of one census tract's population lives in one zip code, while a small portion lives in a second zip code. Here, the population is counted equally in determining the average demographics of both zip-codes, while it is representative of very little of the second zip-code. This issue is compounded if specific factors in the region represented by the second zip-code attract specific demographics. In order to circumvent this issue, I will assume that the populations are distributed evenly across each census tract, such that their characteristics are not significantly different in the different zip-codes they inhabit.

The actual process of moving from census tract to zip-code level data required the use of a crosswalk file containing all overlapping geographies. To complete this shift, I had to merge the crosswalk file with my demographic dataset along the census tract. From this point, there were a number of repeated zip-codes, so I collapsed to the mean of variables by zip code, weighting by the total population of each census tract. These processes gave a dataset providing information on the demographics at a zip-code level easy to merge with the installation data.

In order to use this data, I have also assumed that census tracts and zip-codes are made up of relatively homogenous populations surrounding the geographic average. This assumption is important for drawing conclusions about the effects of SASH on individual household decisions. Without this assumption, zip-code level data cannot provide information on the households that were receiving the subsidy. Moreover, any sort of discontinuity analysis surrounding SASH eligibility falls apart without an assumption of homogeneity within zip-codes. While this assumption is a necessity of the available data, it is likely true because of the relatively small size of zip-codes.

I have assembled data from several ACS moving averages, the CSI, and some county level data. The ACS and CSI data merged together across a total of 2,422 zip-codes. From the CSI dataset, I have created a single variable, *installs,* focused on the number of adoptions in each zip-code by year. Table 1 presents summary statistics for the variables central to my model and upcoming analysis. It is important to note that *installs*, the primary outcome variable in my upcoming analysis, has a mean of 7.758 installations per zip-code per year and a standard deviation of 19.65. This wide standard deviation is the result of a both a large number of 0 observations, as well as a number of extremely high adopting zip-codes. When excluding observations with no installations, the mean number of installs jumps from 7.758 to 17.488 per zip-code per year. Dealing with the large number of observations with no adoptions was an important part of my decision to use fixed effects regressions – there may be some unobservable factor driving the large number of zip-codes that have no adoptions throughout the six-year time frame of this study.

Table 1: Summary Statistics

VARIABLES	N	mean	sd	min	max
Installs	14,573	7.758	19.65	0	351
% Eligible	14,533	50.00	14.66	0	100
Median HH Income	14,519	60,197	25,801	2,373	196,939
Year	14,573	-	-	-	-
ZIP	14,585	-	-	-	-
Avg HH Size	14,524	2.777	0.592	1.230	5.400
% with less than HS degree	14,550	17.55	13.79	0	77.34
% with at least a HS degree	14,550	82.45	13.79	22.67	100
% with at least some college	14,550	60.90	17.79	7.270	100
% with at least a Bachelor’s	14,550	29.76	18.24	0	90.87
% with at least a Master’s	14,550	11.11	9.203	0	59.06
% with at least a Professional degree	14,550	3.829	3.966	0	51.72
% with a PhD	14,550	1.539	2.079	0	51.72
Median Gross Rent	14,499	1,184	371.1	269	3,430
Median number of Rooms	14,522	5.007	0.868	1.300	9
Total Cost of solar panels	6,465	96,315	313,030	5,398	8.458e+06
Total Population	10,549	21,185	21,329	0	105,549
Cumulative Installs	14,585	14.87	40.35	0	616
ln(Median HH Income)	14,519	10.91	0.436	7.772	12.19
Positive Installs (1 or 0)	14,585	0.444	0.497	0	1
Number of ZIP	2,422	-	-	-	-

V. Identification Strategy and Estimation

An integral part of understanding the effectiveness of the SASH subsidy will be identifying those that are eligible for the subsidy. Because my intention in this study is to consider the way that access to this subsidy influenced decision-making among low-income households, it is important to identify variation in access to the subsidy (while holding income levels relatively constant). In order to accomplish this end, I have developed an Interpolated Distribution Function using the income buckets provided in the ACS data. I then compared that function to the Area Median Incomes to create a variable representing the number of households eligible for the subsidy in each census tract.

The Interpolated Distribution Function depicts a piecewise cumulative distribution of incomes with points at each of several buckets. The ten buckets provided by the ACS are: less than $10,000; $10,000-$14,999; $15,000-$24,999; $25,000-$34,999; $35,000-$49,999; $50,000-$74,999; $75,000-$99,999; $100,000-$149,999; $150,000-$199,999; and $200,000 or more. From these buckets, I was able to determine the number of households falling under the upper bound of each (e.g., the sum of households in each of the first three buckets provides us with the number of households with incomes up to $25,000). I assigned every household in the $200,000 or more income group to the $200,000 group, as there is no upper bound imposed. An example Interpolated distribution function built from the distribution in one census tract can be seen in Graph 1. From this graph, I will be able to project the number of households at any income level less than $200,000 in the given zip-code. While California does have a large quantity of high income households, my analysis will be unaffected by an inability to distinguish between incomes above $200,000 because of its focus on the lower-income population.

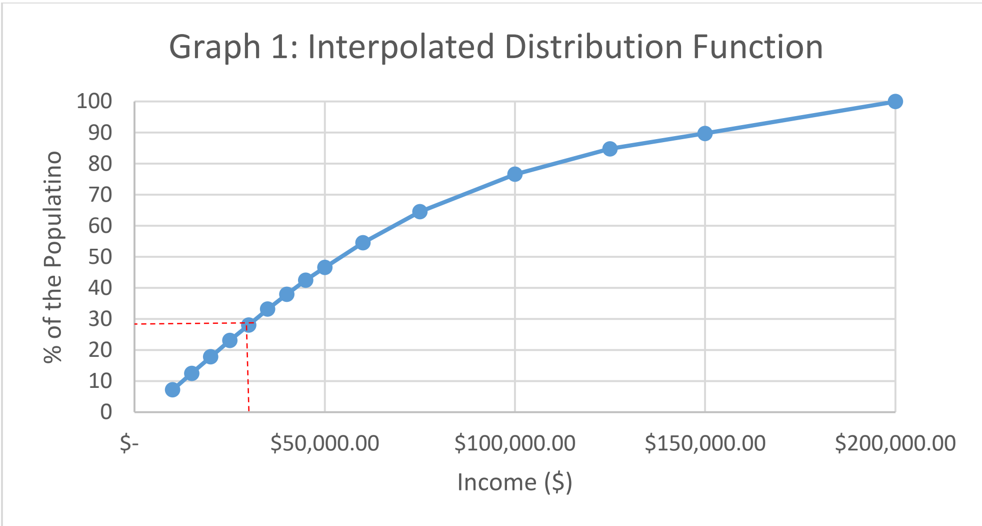

I have assembled interpolated distribution functions for each available year that SASH was in effect. I built these functions from the buckets included in each ACS 5-year moving average, starting with the 2006-2010 survey and concluding with the 2011-2015 survey. Because these surveys each cover 5 years, I used each to represent the middle year of the range (eg 2005-2009 represents 2007). This strategy allowed me to cover the years 2008 (the year SASH was introduced) through 2013, while continuing to work with the same census tract boundaries. While this makes things convenient, it is also the most logical way to connect years to surveys – the midyear middle of each survey is likely the most similar to the greatest number of other years contained within the survey.

From the interpolated distribution functions, I calculated the number of households in each census tract falling within 80% of the Area Median Income (AMI). As noted earlier in this paper, the AMI, provided by the California HUD, is one requirement for SASH eligibility: household income must not exceed 80% of AMI in order to receive funds. In order to identify the number of households in each census tract that fall within this threshold, I calculate 80% of the AMI for each county, then tracked the percentage of households projected to make as much or less. In Graph 1 above, the red lines show that if 80% of the AMI in the county is $30,000, the interpolated distribution function of this census tract would have projected around 28% of households to be eligible under the income requirement. The percent of households in a zip-code eligible for the subsidy based on their income (some may fail to meet other eligibility requirements) ranges all the way from 0 percent to 100 percent. That large range is accompanied by a mean of 50% eligibility and a standard deviation of 14.63 (approximately 30% of the mean figure). While there is a good deal of variation in the percentage of the eligible population, *eligible* does have a fairly normal distribution, as shown in Histogram 1.

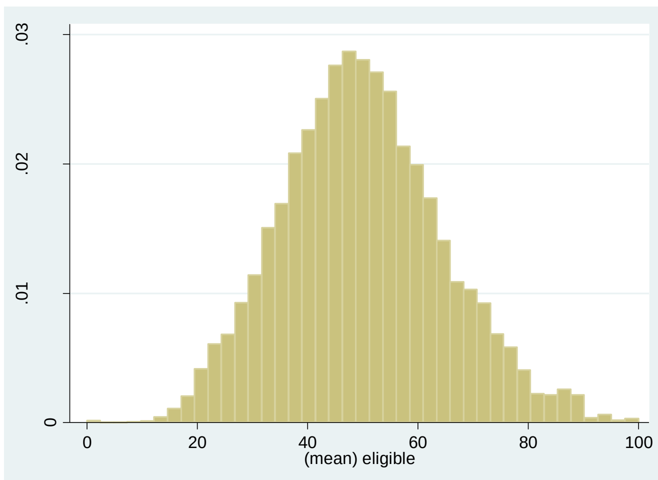

Histogram 1: Eligibility Distribution

The most interesting element of *eligible*, though, is the way that it varies alongside the Area Median Income when holding a zip-code's median income constant. I am trying to simulate an experiment that examines the effects when one of two otherwise identical zip-codes is chosen to receive a subsidy. In order to effectively target this situation, I designed my regression to hold constant several factors, most notably median income, while varying the size of eligibility. This structure ensures that a factor not directly related to the zip-code, specifically Area Median

Income, governs the variation in the subsidy. As a result, variation in the number of adoptions in a zip-code can be attributed to changing percentage of eligible population.

Worth noting is that, while my data for my demographic variables (e.g., *medhhincome*, *avghhsize*, etc.) is primarily taken from a few 5-year moving averages, I built the data for percent of population that is eligible (based on income) using observations attributed to each year. This decision was motivated by a desire for greater variation in the eligible population when considering the effects on installations.

Now that I've identified the determinant of interest using the above outlined strategy, I will be able to build a regression model to estimate the effects of SASH subsidy access on solar adoption rates. My study is centered around the following specification:

$$\mathrm{Installs}_{it} = \beta_0 + \beta_1[\mathrm{eligible}]_{it} + \beta_2[\mathrm{medhhinc}]_{it} + \sum_{t=1}^{n}\sum_{i=1}^{n} a_{it}\, Z_{it} + \varepsilon_{it}$$

The left hand side in this equation, $Installs_i$, is installations per household where *i* indicates zip-code and *t* represents year. The major determinant of interest in this first empirical specification is $eligible_{it}$, the variable for the percentage of households in zip-code *i* and year *t* that are SASH eligible based on the income cutoff. When including $medhhinc_{it}$ as median household income at the zip-code level, the size of β_1 estimates the effect of an additional percentage point of eligibility on the number of installations per household, holding zip-code median income constant. When controlling for median household income at the zip-code level (along with a vector of other controls), X_{1it} picks up the effects of variation in subsidy availability without variation in income. In this specification, Z_{it} represents a vector of control variables. I chose control variables based on an analysis of the effects of demographic variables on installation rates. I will discuss several of the more noteworthy control variable options in the following section.

One of the key elements to my model has been the inclusion of zip-code fixed effects to control for unobservable characteristics that may be biasing my estimates. Because I have a series of observations for each zip-code, I have been able to create time series data to eliminate any of the unobserved drivers of adoption that remained constant over time. When failing to include ZIP code fixed effects, the regressions generally provide counterintuitive results, suggesting that unobserved biases are eliminated by these fixed effects.

VI. Analysis

The most obvious, and perhaps glaringly flawed, way to examine the effects of SASH eligibility under the income requirement on adoption would be to run a standard OLS regression on installations, without including an income variable. Regardless of the additional controls included in the regression, any attempt to discern the effects of *eligible* without controlling for income essentially tracks the rate of adoptions as zip-code median income falls. If this mistake is made, negative coefficients emerge for *eligible*, as seen in columns 1-3 of table 2. Even when including the median household income, though, the effects of increasing eligibility for the subsidy remain negative. This result likely stems from zip-code specific effects that remain unaccounted for in these straightforward OLS regressions. It is conceivable that the highest eligibility occurs in zip-codes with low home ownership or high residency in apartment buildings that are surrounded by higher income zip-codes (e.g. more urban areas). The surrounding zip-codes increase the AMI, such that this zip-code has a higher eligibility than others with the same median household income. While such zip-codes may have more eligibility for the subsidy, residents would have little ability to install residential solar panels. This issue is largely averted when including zip-code fixed effects, as they account for characteristics that remain unchanged over the time period (e.g. lack of roof space or home ownership).

In order to confirm whether the fixed effects were important to controlling for the counterintuitive effect signs, I created a simple regression. The regressor, here, being a dummy variable for non-zero adoption in a given zip-code and year (*posinstalls* = 1 if *installs* > 0). Because fixed effects corrected the counterintuitive sign, I concluded that their inclusion was important to isolating the effects of *eligible*. Table 3 shows that fixed effects flip the sign of *eligible*'s effect to the expected direction. It is evident that controlling for the zip-code fixed effects makes a significant difference in achieving results in line with the economic theory – a subsidy would generally not create an additional barrier to solar panel adoption.

The next step in my analysis was to determine the appropriate controls to include in my fixed effects regressions. I considered the impacts of different demographic covariates on zip-code level adoptions, with all years lumped into one count and demographic data from the 2005-2009 ACS. My decision to include or exclude hinged on economic and statistical significance, as well as availability of the covariates in later ACS data. The primary covariates of interest, per this initial study were education (several different education variables), house size (*mednumrooms*), family size (*avghhsize*), and income (*medhhincome*). I considered but chose to exclude a suite of race variables, as they were individually and jointly insignificant when controlling for median household income; length of residence in one's home (dummy variable equal to 1 if it has been more than one year), as they were unavailable in later ACS data; and median home value, as the relative home value across years is largely factored into the zip-code fixed effects.

Table 2. Standard OLS Regressions

VARIABLES	(1) *installs*	(2) *installs*	(3) *installs*	(4) *installs*
eligible	-0.330***	-0.0416***	-0.0451***	-0.0658***
	(0.0122)	(0.00646)	(0.00910)	(0.0128)
medhhincome				-3.40e-05***
				(1.06e-05)
cuminstalls		0.380***	0.381***	0.382***
		(0.0126)	(0.0129)	(0.0130)
mednumrooms		1.626***	1.304***	1.562***
		(0.136)	(0.164)	(0.153)
avghhsize		1.398***	1.924***	2.119***
		(0.128)	(0.194)	(0.207)
o.pctlesshs			-	-
pcthspl			0.0602***	0.0586***
			(0.0136)	(0.0137)
pctsomcolpl			0.000118	-0.00219
			(0.0151)	(0.0150)
pctbspl			0.0357**	0.0648***
			(0.0168)	(0.0167)
pctmspl			-0.103***	-0.0919**
			(0.0368)	(0.0357)
pctprofpl			-0.136**	-0.131**
			(0.0564)	(0.0569)
pctphd			0.00587	-0.0449
			(0.0726)	(0.0735)
Constant	24.26***	-7.829***	-11.90***	-11.32***
	(0.732)	(0.751)	(1.541)	(1.590)
Observations	14,533	14,517	14,517	14,511
R-squared	0.060	0.663	0.664	0.665

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3: Positive installations with and without ZIP Fixed Effects

VARIABLES	(1) posinstalls	(2) posinstalls
eligible	-0.00428***	0.00246***
	(0.000420)	(0.000885)
medhhincome	1.07e-06***	4.33e-07
	(2.67e-07)	(6.78e-07)
cuminstalls	0.00442***	3.47e-05**
	(0.000180)	(1.68e-05)
Constant	0.528***	0.295***
	(0.0354)	(0.0786)
Observations	14,519	14,519
R-squared	0.191	0.001
Number of ZIP		2,423
ZIP Fixed Effects	NO	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In putting together my final set of regressions, I cycled through some of these controls to determine a primary set to include. Median household income (*medhhincome*) is essential to my analysis, and thus included in each regression in table 3. As discussed in my identification section, this control is critical to isolating the effects of *eligible*, as *eligible* tracks the number of lower income families, relative to the county mean. Without controlling for the median household income, this might simply act as an inverse measure of wealth. By introducing a number of controls into the regression, I am able to simulate an experiment comparing nearly identical zip-codes when provided differential access to the SASH subsidy. Median household income is the most important of these controls, because this characteristic actively influences the size of the subsidy. Given that position, it is important to consider the effects of varying eligibility conditional on a constant median income – this way, eligibility moves in response to fluctuations in external factors (specifically the Area Median Income).

Table 3, below, provides a series of regressions focused on *eligible*, with different combinations of control variables. My expectation in these regressions is to see a positive and significant effect of *eligible* when controlling for median household income. In table 4, *eligible* is positive and significant in each of the regressions, consistently sitting around 0.08. This coefficient can be interpreted as saying that a 1 percentage point increase in the percentage of households eligible for the SASH subsidy in a zip-code and year results in an increase in the number of adoptions of 0.08 in that same zip-code, holding all other controls constant. While the size of this effect appears to be small on an absolute scale, it comes to around 1% of the mean number of adoptions per zip code per year. In other words, a 1 percentage point increase in the percentage of eligible households in a zip-code in a given year leads to a 1% increase in the number of solar panel adoptions in that zip-code and year (relative to the mean). This is a sizeable effect, and marks a promising elasticity of adoption with respect to income eligibility. I will discuss some consequences of this result in greater detail in the Policy Implications section.

Table 4. Fixed Effects Regressions				
	(1)	(2)	(3)	(4)
VARIABLES	*installs*	*installs*	*installs*	*installs*
eligible	0.0826***	0.0794***	0.0823***	0.0803***
	(0.0295)	(0.0292)	(0.0295)	(0.0298)
cuminstalls	0.276***	0.276***	0.276***	0.276***
	(0.0196)	(0.0195)	(0.0195)	(0.0197)
medhhincome	-0.000124***	-0.000121***	-0.000124***	-0.000126***
	(3.18e-05)	(3.19e-05)	(3.17e-05)	(3.11e-05)
avghhsize	1.345*		1.367*	1.516*
	(0.781)		(0.767)	(0.805)
mednumrooms	0.141	0.299		0.113
	(0.538)	(0.525)		(0.536)
pctlesshs				-0.0596**
				(0.0258)
o.pcthspl				-
pctsomcolpl				0.0460
				(0.0311)
pctbspl				-0.0693
				(0.0447)
pctmspl				0.0227
				(0.0551)
pctprofpl				-0.142
				(0.107)
pctphd				0.167
				(0.129)
Constant	2.582	5.467	3.226	2.783
	(4.342)	(3.899)	(3.862)	(4.844)
Observations	14,511	14,516	14,513	14,511
R-squared	0.398	0.398	0.398	0.398
Number of ZIP	2,422	2,423	2,422	2,422
ZIP Fixed Effects	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In terms of the controls I included in the regressions shown in table 4, the most notable are *cuminstalls* (cumulative installations) and *medhhincome* (median household income). I chose to include cumulative installations in order to control for a few potential effects. Primarily, as discussed earlier in this paper, I had some concerns about potential saturation within smaller zip-codes – if the population is small, all of those with the desire to adopt may have already done so. This would impact the apparent effects of the subsidy. These results, though, provide a *positive* effect of cumulative installations. This might be a result of some sort of network effects – as there are more adoptions in a given zip-code or neighborhood, other residents of the area are more exposed to residential solar and have increasing access to reliable information about it. Median household income behaves a bit differently than one might expect, here, as well, with consistently negative effects. The negative coefficient on *medhhincome* may be a function of the conditional statement implicit in the regression. In order for *eligibility* to stay constant while *medhhincome* varies means that either the distribution of incomes within the zip-code shifted in an unexpected way or the Area Median Income shifted proportionately (holding the percentage of eligible population constant). A shift in the income distributions, such as a slight rise in incomes among those sitting right around the zip-code median income and a dip in income for those at the bottom of the distribution, could increase median incomes without impacting the eligible population. Moreover, the lowest income group could have especially little access to home ownership, so zip-codes with this sort of income distribution do not see higher adoption with increasing median income.

In addition to those primary control variables, I also included several less important controls: average household size in terms of family members (*avghhsize*), median number of rooms to proxy for home size (*mednumrooms*), and a suite of education attainment variables. Larger family size results in significantly increased adoptions when holding income and eligibility constant. This follows expectations – larger families are likely to consume more energy, increasing the benefit of adopting residential solar. Median number of rooms, too, behaves as anticipated. As with large family sizes, large homes are likely to consume more energy, increasing the value of residential solar. The insignificance of this coefficient may be a function of the fact that it is an imperfect proxy for home size. It is possible that large and small homes differ primarily on room size, not number of rooms. Unfortunately, a more direct measure of household size is not available in the ACS datasets. I included the education variables in order to control for potential correlation between higher education and adoption. Given the general insignificance of the education variables, there does not appear to have been much, if any, effect of education when controlling for the other items included in the regression. However, they were jointly significant, based on an F-test. Likewise, an F-test including all of the control variables showed significance.

It is worth noting that I have not been especially concerned with the R-squared values in my regressions – it is highly possible that they do not fully capture the fluctuations in adoptions. Because I am chiefly concerned with the specific effects of *eligible*, as long as I have controlled for the things that might bias this estimate, then the regression suits my needs well. Closely predicting the overall number of adoptions that a zip-code might see in a given year is *not* one of the aims of my analysis. Rather, my concern is the number of additional adoptions produced by the SASH subsidy. In considering the policy implications of this program, I made specific efforts to isolate the changes in adoption that occur as a result of SASH.

VII. Policy Implications

Understanding the climate change and environmental/energy justice impacts exists as an important final step in my analysis of the SASH program. As a subsidy, SASH functioned as an effort to provide low-income households with more affordable energy, while also generating some carbon abatement. The two major ways to understand the effectiveness of this policy are the rate at which low-income consumers adopted the subsidy when given the opportunity, and the cost at which the subsidy attained various amounts of carbon abatement.

The first of these lenses, rate of adoption, is best illustrated by elasticity of adoption with respect to the size of the eligible population. Determining this elasticity requires a few "back of the envelop" calculations. Taking the figures from California's sample, the mean level of eligibility is 50% and the mean number of adoptions is 7.758 (as noted in Table 1). A 1 percentage point increase in eligible households is equal to a 2% increase in eligibility and generates 0.08 additional adoptions, equivalent to just above a 1% rise. These give the following equation:

$$elasticity = \frac{\frac{Adoptions_1 - Adoptions_0}{Adoptions_0}}{\frac{Eligible_1 - Eligible_0}{Eligible_0}} = \frac{\frac{7.838 - 7.758}{7.758}}{\frac{51 - 50}{50}} = \frac{.0103}{.02} = .515$$

From this equation, one can see that the elasticity of adoption with respect to SASH eligibility (based on income) is around .515. This means that for every 1 percent increase eligibility provided, there is a corresponding increase in adoptions of .515 percent.

To provide further context to the effectiveness of this program as a tool to push low-income households to adopt solar, it is relevant to consider the size of the increase in adoptions. With a percent of households eligible for SASH subsidies in each zip-code in each year (based on income, this does not account for the equity sharing requirement), I was able to slot these numbers into my regression equations from the Results section. I chose to use the most complete regression that excluded education levels (regression 1 from table 4):

$$Y_{it,e} = 2.582423 + .0826118 * eligible_{ite} + .1409207 * mednumrooms_{it} - .0001242 * medhhincome_{it} + 1.344536 * avghhsize_{it} + .2756647 * cuminstalls_{it} + \varepsilon_{it}$$

Here, I estimated the number of installations per zip-code per year at a given level of subsidy ($Y_{it,\,e}$) using the medians of all variables other than *cuminstalls* (cumulative installations up to the year *t*) and *eligible* (percent of eligible households in zip-code *i* and year *t*). For *cuminstalls*, I input the mean, as the median fell at zero. For *eligible*, I input the figures that I calculated for each of the varying cutoff levels, where *e* is a given cutoff level. The goal of this calculation was to find the additional adoptions generated by increasing the cutoff. Assuming that these adoptions come from those newly able to access the subsidy, graph 2 shows the distribution of households with incomes greater than 60% of AMI that will adopt given access. Moreover, these houses all lie at a point less than the county median incomes, with at most 95% of that figure. As one would expect, increases in the cutoff result in increases in adoption, suggesting that this program effectively encourages low-income residential solar adoption. When moving the cutoff from 60% to 61% of AMI, 737 additional adoptions are generated. Alternatively, raising the cutoff to 95% of AMI generates 22,159 additional adoptions, relative to a 60% cutoff. At an eligibility cutoff of 0% of AMI (essentially eliminating the subsidy entirely), this model estimates approximates roughly 82,000 residential solar panel adoptions over the course of the 6-year window I've examined. Shifting that cutoff up to the 80% level increases that estimate up to roughly 120,000, a shift of nearly 40,000 residential solar panels provided directly to low-income households. Because the estimate of adoptions assuming an eligibility cutoff of 0% of AMI is likely incorrect – the regression that I've used to make this estimate has an R-squared around 0.4 (table 4 column 1) – the relative size of this effect is likely somewhat exaggerated. Moreover, the SASH budget is relatively limited, so this huge effect size likely hasn't been totally realized across California. This high level of additional low-income adoptions suggests that SASH has effectively met its goals of providing low-income households with residential solar energy.

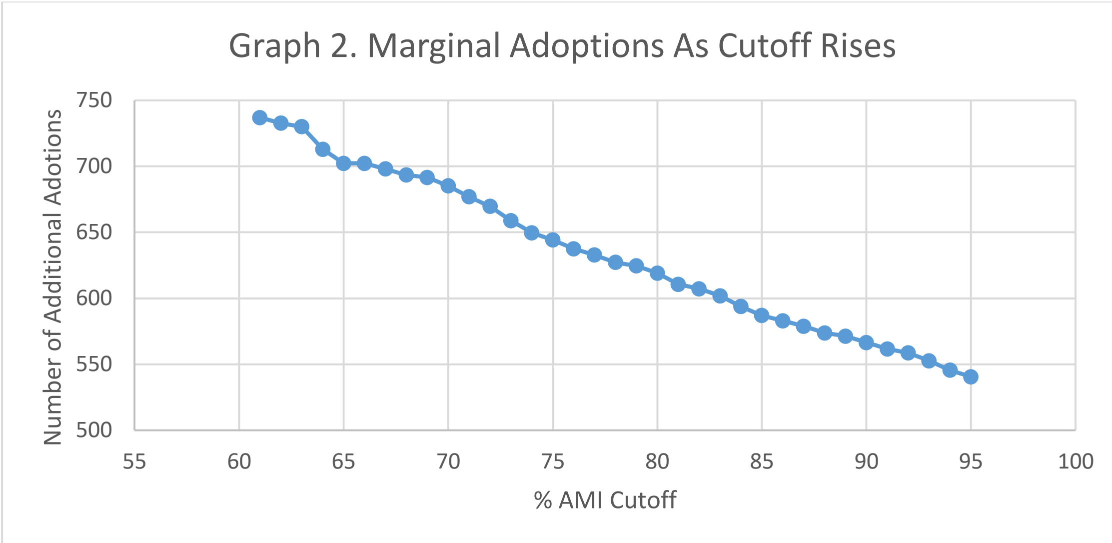

It is also important to understand how effective this program was at reducing carbon emissions, by way of providing homeowners with a source of renewable energy. Marginal cost of carbon abatement provides the most informative number when considering the policy from this perspective. To determine these figures, I began by finding the percent of households in each zip-code eligible for the subsidy under the income requirement when varying the cutoff from 60% to 95% of the Area Median Income (the actual program cutoff is 80%). To perform the calculations required to develop these curves, I assumed the value of a few parameters necessary to building a model for carbon abatement. I will note these assumptions as they arise in the following explanation. In order to create an optimistic estimation, I chose to assume the most beneficial in each case. Doing so does not change my conclusion, but rather provides a lower bound on marginal cost that still sits far above the social cost of carbon provided by the EPA.

The next step in building marginal cost of abatement and marginal abatement per additional dollar spent on the subsidy curves was to determine the cost of this program at varying sizes of the eligible population. I employed the following formula to do so:

$$Cost_e = \sum_{i} \sum_{i,t=1}^{i,t=n} Y_{it,e} * [system\ size\ (kW)] * 1000 * (3)$$

The lefthand side variable of this equation provides an estimate of the cost of SASH subsidy for the panels purchased at a given cutoff (relative to AMI, ranging from 60-95%), e. The figure $Y_{it,e}$ comes from the same calculations that I used to put together graph 2. Again, system size is assumed to sit at the median. Because system size is provided in kilowatts and the subsidy is granted per watt, I multiplied by 1000. Lastly, I multiplied by $3, the size of the subsidy per watt. This figure assigns the $3 subsidy to each adoption, so only the marginal cost of moving from one eligibility cutoff to another is relevant.

From here, I considered the amount of energy produced by each solar panel in production, using the equation:

$$Energy_e = \sum_{i} \sum_{t=1}^{n} Y_{it} * [system\ size] * \left[\frac{\frac{kWh}{kW}}{year}\ factor\right] * [years\ in\ use]$$
$$* [average\ production\ considering\ degradation]$$

The lefthand side variable of this equation provides an estimate of the energy produced over the lifetime of the panels purchased at a given cutoff, e. Working through this equation required a number of assumptions – for the system size, I chose to use the median from my data; for the kWh/kW-year factor, I estimated the median based on a map provided by U.S. Department of Energy providing factors based on the geography; and, lastly, I assumed 25 years of use at an average of 95% efficiency. The kWh/kW-year factor provides a conversion from the amount of sunlight predicted to reach an area per year to the amount of energy generated in that year. With these assumptions in tow, I was able to estimate the total energy production at each cutoff level.

This figure will clearly overestimate the energy created by the SASH subsidy, so the main item of interest is the *marginal* energy production gained when moving from one cutoff to another. This isolates the effects of the variation in access to SASH subsidies.

After determining clean energy production created by SASH, I had to determine the level amount of carbon generated in the production of energy in the state of California. This figure illustrates an estimate of the emissions abated by shifting to solar energy. The U.S. Energy Information Agency publishes the average number of kilograms of carbon produced per million Btu of energy output for each year. After converting these figures into tons of carbon per kilowatt hour, I compared them to the amount of marginal power generation created by varying SASH program eligibility cutoff. The results of this comparison provide the marginal abatement. When comparing these, in turn, with the marginal cost of varying SASH program eligibility, I was able to generate Graph 3.

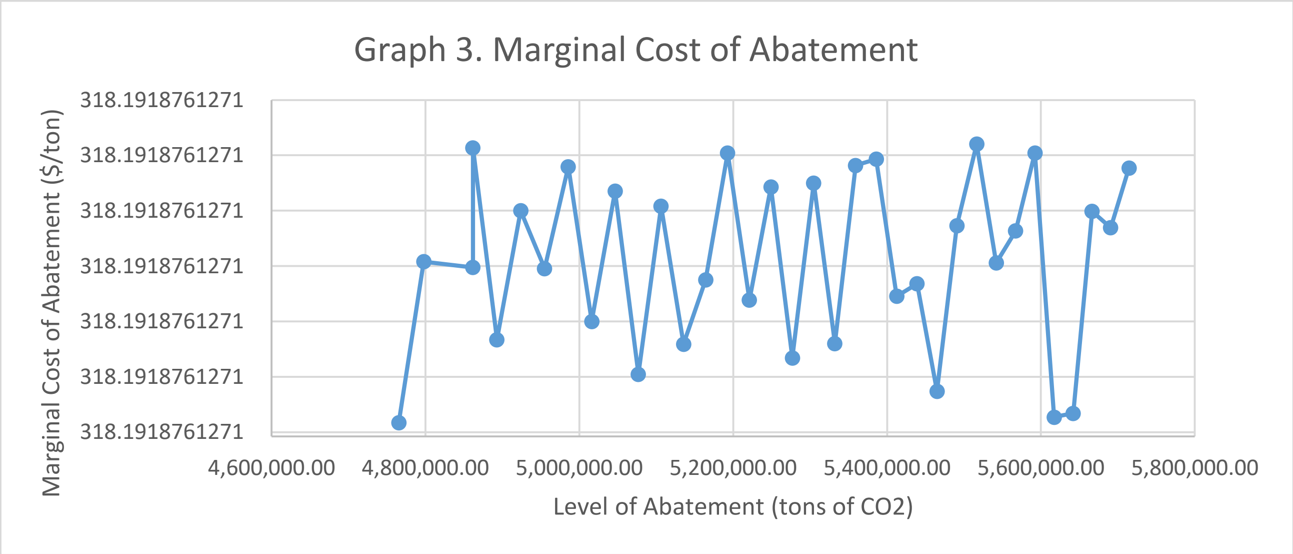

The vertical axis of Graph 3 shows the estimated marginal cost of an additional ton of carbon, while the horizontal axis shows the level of abatement at which that marginal cost occurs. While this graph does reveal some variance, its great magnification exaggerates the level of that variance - the y-axis varies only at an extremely small scale. Overall, the calculations I described above generated almost no variance in this figure. The high marginal abatement per marginal dollar for this program suggests that it is *not* a good carbon abatement strategy. Moreover, I have intentionally assumed optimistic numbers for relevant figures, such as kWh/kW-year factor, years in use, and rate of degradation. With these optimistic assumptions, my resulting estimate of the marginal cost of abatement is likely *less* than the actual figure. The EPA publishes the social cost of carbon, a measure of the present value of long-term damages done by a ton of carbon dioxide emissions emitted in a given year. For 2015 and at the lowest discount rate included (providing the highest estimate of the social cost of carbon), that number comes to only $56. Given this context, the estimated marginal abatement per marginal dollar of approximately $319 provided by this program is exceptionally high. Even though the assumptions made in calculating the marginal cost of abatement are as favorable as possible, the program is still far from a viable abatement tool.

VIII. Conclusions

California's Single-Family Affordable Solar Housing (SASH) program provides subsidization for low-income households to purchase residential solar panels. Qualification requires income below a given cutoff, alongside participation in a home equity sharing agreement. Throughout this paper, I have examined the effects of varying access to the subsidy through changing incomes, though I have been unable to collect data regarding home equity sharing. The primary stated goal of SASH is to open access to affordable electricity through the subsidization of residential solar PV. The move to support low-income homeowners through renewable energy subsidization follows with much of California's renewable-slanted policy.

The SASH program has proven itself to be a success at its primary goal of encouraging low-income solar adoption. In my analysis, I have found that increasing the percentage of financially eligible households in a zip-code by 1 percentage point leads to 0.08 additional solar PV adoptions per year. This provides an elasticity of .515 at median eligibility and adoption. Moreover, when considering the number of low-income adoptions generated by this subsidy, relative to the number of existing higher-income adoptions, it is clear that this program makes significant strides in closing the existing gap between the two. Promoting a capital intensive investment such as residential solar among low-income households will never be a low cost endeavor, however. Limited funds, bureaucratic slowdowns in reading and accepting subsidy applications, or lack of awareness may be responsible for this elasticity and low-income adoption numbers sitting lower than they might otherwise reach.

While it may not have been a primary stated goal of SASH, this program does not provide an affordable method of abating carbon emissions. The effective price of each ton of carbon emissions abated through this policy is approximately $319, which sits well above other programs designed to reduce carbon emissions in energy production. This should not come as a great surprise: low-income households need great assistance overcoming the upfront financial barriers to participation in residential solar markets. A much more efficient allocation of government funds (from a carbon abatement perspective) could be achieved by simply providing higher income families with significantly smaller incentives to install solar. As a result of this extremely expensive reduction in carbon emissions, it is difficult to positively evaluate this program in terms of its effects on carbon emissions.

Overall, the effectiveness of the SASH program depends on one’s perspective. From an environmentalist perspective concerned solely with carbon abatement, this program falls flat. It reduces carbon emissions at a social cost far higher than might be seen through other initiatives, such as improvements in energy efficiency. It would undoubtedly be cheaper to offer small subsidies to wealthy households than to offer large subsidies to poorer households, while both would achieve the same abatement. However, this program succeeds from an environmental or energy justice perspective. SASH opens a door for low-income households that would not otherwise have had access to solar to purchase residential panels, and a relatively high percentage of those financially eligible are taking advantage of that opportunity. It is important for the evaluation of this program that the primary motivator be the egalitarian desire to assist low-income households in managing rising energy prices. If that energy justice view is the driving force behind one's assessment of the program, it is a success.

Works Cited

Balcombe, P., Rigby, D., Azapagic, A. (2014, June 20). *Investigating the importance of motivations and barriers related to microgeneration uptake in the UK. Applied Energy, 130,* 403-418.

Blackburn, G., Magee, C. & Rai, V. (2014, Jan/Feb). *Solar Valuation and the Modern Utility's Expansion into Distributed Generation*. *The Electricity Journal*, *27,* 18-32.

California Public Utilities Commission. (2016). CSI Single-Family Affordable Solar Homes Program. http://www.cpuc.ca.gov/General.aspx?id=3043

California State Government. California Public Utilities Commission. (2015, July). *MASH Handbook.* California.

California State Government. California Public Utilities Commission. (2015, May). *SASH Handbook.* California.

California State Government. Department of Housing and Community Development. (2015, April 15). *State Income Limits for 2015*. Sacramento, California.

Energy and Environmental Economics, Inc. (2013, Oct. 28*) Introduction to the California Net Energy Metering Ratepayer Impacts Evaluation. Prepared for the CPUC Energy Division*.

Enkvist, P., Naucler, T. & Rosander, J. (2007, February). *A cost curve for greenhouse gas reduction.* McKinsey Quarterly. http://www.mckinsey.com/business-functions/sustainability-and-resource-productivity/our-insights/a-cost-curve-for-greenhouse-gas-reduction

Environmental Protection Agency. (n.d.). *The Social Cost of Carbon: Estimating the Benefits of Reducing Greenhouse Gas Emissions*. https://www.epa.gov/climatechange/social-cost-carbon

Kelly, K. (2014, February). *Low-Income Renewable Energy Programs: Case Studies of State Policy in California and Massachussetts* (Master's thesis). Retrieved from Google Scholar. (UMI No. 1552973)

Podolefsky, M.C. (2013). *Three Essays in Energy, Environmental and Public Economics*. (Doctoral Dissertation). Retrieved from ProQuest. (UMI No. 3607346)

Rai, V., Sigrin, B. (2013, February 19). *Diffusion of environmentally-friendly energy technologies: buy versus lease differences in residential PV markets. Environmental Research Letters, 8,* 1-8.

U.S. Department of Energy, Energy Efficiency and Renewable Energy. (2003). *A Consumer's Guide: Get Your Power from the Sun*. http://www.nrel.gov/docs/fy04osti/35297.pdf

U.S. Energy Information Administration. (2017, January 17). *Energy-Related Carbon Dioxide Emissions at the State Level 2000-2014*. https://www.eia.gov/environment/emissions/state/analysis/

Welton, S. (2016, March). *Clean Energy Justice*. *University of Colorado Law Review (Forthcoming).*

Data Sources

American Community Survey (2009) *Selected Social Characteristics in the United States: 2005-2009*. Retrieved from www.census.gov/programs-surveys/acs/data.html

American Community Survey (2009) *Selected Economic Characteristics: 2005-2009*. Retrieved from www.census.gov/programs-surveys/acs/data.html

American Community Survey (2009) *Selected Housing Characteristics: 2005-2009*. Retrieved from www.census.gov/programs-surveys/acs/data.html

American Community Survey (2009). *ACS Demographic and Housing Estimates: 2005-2009*. Retrieved from www.census.gov/programs-surveys/acs/data.html

California Solar Initiative (2016). *Working Data Set*. Retrieved from www.Californiasolarstatistics.ca.gov/data_downloads

California Solar Initiative (2016). *NEM Currently Interconnected Data Set*. Retrieved from www.Californiasolarstatistics.ca.gov/data_downloads