Hurdles and Steps: Estimating Demand for Solar Photovoltaics

Kenneth Gillingham*
Yale University and NBER

Tsvetan Tsvetanov†
University of Kansas

July 8, 2017

Abstract

This paper estimates demand for residential solar photovoltaic (PV) systems using a new approach to address three empirical challenges that often arise with count data: excess zeros, unobserved heterogeneity, and endogeneity of price. Our results imply a price elasticity of demand for solar PV systems of -0.65. Counterfactual policy simulations indicate that reducing state financial incentives in half would have led to 9 percent fewer new installations in Connecticut in 2014. Calculations suggest a subsidy program cost of \$364/tCO_2 assuming solar displaces natural gas. Our Poisson hurdle approach holds promise for modeling the demand for many new technologies.

Keywords: count data; hurdle model; fixed effects; instrumental variables; Poisson; energy policy.
JEL classification codes: Q42, Q48, C33, C36

*Corresponding author: School of Forestry & Environmental Studies, Department of Economics, School of Management, Yale University, 195 Prospect Street, New Haven, CT 06511, phone: 203-436-5465, e-mail: kenneth.gillingham@yale.edu, and the National Bureau of Economic Research.

†Department of Economics, University of Kansas, 1460 Jayhawk Boulevard, Lawrence, KS 66045, phone: 785-864-1881, e-mail: tsvetanov@ku.edu.

1 Introduction

The market for rooftop solar photovoltaic (PV) systems has been growing rapidly around the world in the past decade. In the United States, there has been an increase in new installed capacity from under 500 MW in 2008 to over 4,500 MW in 2013, along with a decrease in average (pre-incentive) PV system prices from over $8/W in 2008 to just above $4/W in 2013 (in 2014 dollars) (Barbose et al. 2014). These major changes in the market have come over a period of considerable government support at both the state and federal levels, ranging from state-level rebates to a 30 percent federal tax credit. For example, in Connecticut (CT), the 2008-2014 average combined state and federal incentives equaled 50 percent of the system cost for most PV system purchasers.[1] Yet such substantial government support is being slowly reduced in many markets throughout Europe and the United States.

This paper develops a new approach to addressing key empirical challenges in estimating demand with count data and applies this approach to model the demand for residential PV systems in CT. We use rich installation-level data over the period 2008 to 2014 to shed new light on a small, but fast-growing market that is similar to many others around the world. We estimate a price elasticity of demand for a PV system of -0.65, a finding useful to both policymakers and firms. Policymakers are often interested in how changes in PV system prices – whether due to policy or other factors – influence the sales of PV systems. Such knowledge is essential for assessing the impacts of solar PV policies. Given the evidence that PV system installation markets are often imperfectly competitive (Bollinger and Gillingham 2016; Gillingham et al. 2016; Pless and van Benthem 2017), firms may be able to use this elasticity to inform pricing decisions and market forecasts.

The empirical challenges that motivate our approach are threefold: possible unobserved heterogeneity at a fine geographic level, excess zeros, and endogeneity of price (due to simultaneity). Existing studies have developed count data methods that can separately address

[1]This calculation is made based on an average system price in 2008-2014 in CT of $36,607, average state rebate amount of $10,288, and tax credit (taken post-incentive) of $7,896 for consumers with at least this much tax burden (all in 2014 dollars).

the presence of excess zeros (e.g., Pohlmeier and Ulrich 1995; Santos Silva and Windmeijer 2001; Winkelmann 2004) or endogeneity of covariates without unobserved heterogeneity (e.g., Mullahy 1997; Windmeijer and Santos Silva 1997; Terza 1998). Windmeijer (2008) goes further in addressing endogeneity in the presence of unobserved heterogeneity, but does not address excess zeros. This paper is the first to address all three challenges in a count data setting with clear policy significance. We use panel data on the count of annual solar PV systems installed in a Census block group. Such panel data allow us to address likely unobserved heterogeneity in environmental preferences or other block group-specific factors with fixed effects. Aggregating at a higher level would leave this unobserved heterogeneity unaddressed. However, 74 percent of the observations have zero values for the number of installations in a block group in a year. Thus, classic models for count data are problematic, such as the Poisson and negative binomial, which likely misspecify the underlying data generating process. Moreover, the negative binomial model cannot readily accommodate block group fixed effects and endogeneity.

To address this problem, we employ a hurdle model, which is recognized as an effective tool for dealing with the presence of excess zeros in count data settings (see, e.g., Cameron and Trivedi 2013). We estimate a hurdle model based on two data generating processes: a standard logit for whether a block group has at least one adoption and a zero-truncated Poisson that models the rate of adoptions conditional on a block group having an adoption. This hurdle model has a clear behavioral interpretation in our setting: the first installation in an area is a rare event, but once there are multiple installations, installers can focus marketing on the area, so we have a count process. In order to tackle unobserved heterogeneity and endogeneity of the price variable, we extend the hurdle model to accommodate fixed effects and instrumental variables.

At the basis of our approach is the conditional maximum likelihood (CMLE) estimator for fixed effects logit models introduced in a sequence of works by Rasch (1960, 1961), Andersen (1972), and Chamberlain (1980). Majo and van Soest (2011) show that this conditional maximum likelihood approach can also be applied to the zero-truncated Poisson framework and present an application of this in a two-period setting. We generalize the model of Majo and van Soest (2011) to a setting with an arbitrary number of periods and possible endogeneity. In contrast to Majo and van Soest (2011), our approach uses a generalized method of moments (GMM) estimator that is based on the first-order conditions of the truncated Poisson CMLE. By using a GMM approach (as in Windmeijer (2008) and several other papers), we can draw upon the established procedures for addressing endogeneity using instrumental variables. To the best of our knowledge, this instrumental variables Poisson hurdle model with fixed effects is new to the literature, and we both prove the consistency of our estimator and verify it with a Monte Carlo simulation. This approach is particularly useful for solar PV markets, but we expect that it is more broadly applicable to many other settings with similar empirical challenges, such as the demand for many early-stage technologies.

Identification in our setting is based on deviations from block group and year means in both the number of installations and PV system prices, after controlling for a variety of potential confounders and instrumenting for price. As instruments, we use a set of supply shifters: local roofing contractor wage rates and state incentives for PV systems. After controlling for income at a localized level, county-level roofer wages act as a valid contractor marginal cost shifter. Solar PV incentives in CT are given directly to the installing firm rather than the consumer, so the consumer sees the post-incentive price at the bottom of any contract to install a PV system. Thus, the incentives also act as a valid marginal cost shifter.

Our preferred estimate of the price elasticity of PV system demand of -0.65 is the first estimate we are aware of for CT. This result is comparable to the existing literature, which uses data from California and very different empirical strategies. For example, Rogers and Sexton (2014) find a rebate elasticity of approximately -0.4, while Hughes and Podolefsky (2015) find an estimate of approximately -1.2. In contrast to the previous papers, which both use reduced form approaches, Burr (2014) estimates a dynamic discrete choice model of demand, but does not specifically estimate a price elasticity of demand. The dynamic discrete choice approach may seem well-suited for our context, since subsidies may be changing over time and a household usually only installs a solar system once (although a household can always add capacity later if there is available roof space). However, we provide survey and descriptive evidence in Appendix A suggesting that solar PV demand in CT is more similar to the many other contexts where consumers do not appear to treat adoption as a dynamic "buy-or-wait" decision.

Using our results, we perform policy counterfactual simulations to examine the impact of state financial incentives and permitting policies on PV system adoption. We first perform a simple analysis of the pass-through of the incentives to consumers, following Sallee (2011), and find the pass-through rate of 84 percent. Our simulation results suggest, under a 50 percent reduction of all financial incentives for purchasing PV systems, the number of new installations in CT in 2014 would have been 9 percent percent less than observed. This would result in up to 1.3 MW less added PV capacity in 2014. Simple calculations suggest a cost-effectiveness of the program of \$364 per avoided ton of CO_2 (\$594 if the federal tax credit is included), assuming that solar power displaces natural gas-fired generation.

The remainder of this paper is structured as follows. Section 2 provides a background on policies in CT targeted at stimulating solar demand. Section 3 describes the data used in our analysis. Sections 4 and 5 outline the estimation methodology. Section 6 presents our empirical results. Section 7 presents the pass-through analysis, counterfactual policy simulations, and a discussion of cost-effectiveness and welfare. Finally, Section 8 concludes.

2 Background on Solar PV Policies in Connecticut

Despite receiving fewer hours of sun than more southerly regions,[2] CT has a robust and growing market for solar PV systems, due to high electricity prices, many owner-occupied

[2] See solar insolation maps at http://www.nrel.gov/gis/solar.html.

homes, and considerable state support for PV systems. A $5 per watt (W) rebate for residential solar PV systems (up to 5 kW) was available in CT as early as 2004, as shown in Table 1. The structure of incentives changed on July 1, 2011, with the passage of Connecticut Public Act 11-802, directing the newly established CT Energy Finance and Investment Authority, which has since been renamed the Connecticut Green Bank (CGB), to develop a residential solar investment program that would result in at least 30 MW of new residential PV installations by the end of 2022 (Shaw et al. 2014).[3]

Starting on March 2, 2012, two types of financial incentives were offered under the "Residential Solar Investment Program" (RSIP). For households that purchase a solar PV system, CT offers an upfront rebate, under an "expected performance based buy-down" (EPBB) program (replaced by the similar "homeowner performance based incentive" (HOPBI) rebate program after July 11, 2014). The rebate per W decreases based on the size of the system, with lower incentives per W for systems larger than 5 kW and even lower for systems larger than 10 kW. Most systems in CT fall under this incentive program until near the end of our time period, when a higher percentage of systems are not purchased outright. For such third party-owned systems (e.g., installed under a solar lease or power purchase agreement), CT offers quarterly incentive payments based on production (in kWh) over six years, under a "performance-based incentive" (PBI) program (Shaw et al. 2014). All incentives are being reduced over time in a series of steps as the solar PV market grows. Unlike other states, where learning-by-doing was an explicit policy motivation (Bollinger and Gillingham 2016), in CT this declining step schedule is primarily motivated by budget constraints and a broad desire for a self-sustaining market. Table 2 provides details of the programs and the step schedule. Appendix A gives more detail on the path of subsidies and consumer adoption of solar PV, providing evidence that the subsidy declines were unanticipated by consumers.

Besides direct financial incentives, CT has several additional programs to promote solar PV systems. One program is "net metering," which allows owners of solar PV systems to

[3]This target had already been reached by the end of 2014. See http://www.ctcleanenergy.com.

return excess generated electricity to the grid, offsetting electricity that is used during times of non-generation, so the final electricity bill includes a charge for only the net electricity usage. PV system owners can also carry over credits for excess production for up to one year. On March 31 of each year, the utility pays the PV system owners any net excess generation remaining at the avoided cost of wholesale electricity (a lower rate than the retail rate).

Another program in CT is an effort starting in 2011 to streamline municipality and utility permitting, inspection, and interconnection processes. Several municipalities in CT have reduced their permitting costs and reduced the permitting time in response to this effort. Similar efforts have occurred in a number of other states, such as Arizona, California, and Colorado (CEFIA 2013).

Perhaps the most important non-rebate program for solar PV system adoption in CT is the CGB-sponsored "Solarize CT" grassroots marketing program in select municipalities.[4] Solarize CT involves local campaigns with a municipality-chosen installer, roughly 20 week time frame, and pre-negotiated group discount pricing for all PV system installations in the town. According to **?**, this program led to a substantial increase in demand in these municipalities: on average roughly 30 additional installations per municipality over the length of the program. The program is also associated with a roughly $0.40 to $0.50 per W decrease in installation prices.

3 Data

Our primary dataset contains nearly all residential solar PV system installations in CT from the period 2008-2014. Each installation that receives a rebate in the two major investor-owned electric utility regions in CT, United Illuminating and Eversource Energy (formerly Connecticut Light & Power) is entered into a database by CGB, with information on price, rebates granted, system size (in kW), technical characteristics, financing arrangements, ad-

[4]Solarize CT is funded by foundations, ratepayers, and other grants.

dress, date of application processing, and date of installation.[5] Until late 2014, there were few third party-owned (TPO) systems in CT (i.e., solar leases or power-purchase agreements) and the price data for these TPO systems are well-known to be less reliable than owned-system price data. Thus, we exclude these TPO systems from our primary analysis.[6] Our raw dataset contains 5,070 residential PV installations approved for the rebate by the CGB between January 1, 2008 and December 31, 2014.

We geocode the installations and match each to the U.S. Census block group they reside in. We then collect U.S. Census data at the block group level from the 2006-2010, 2007-2011, 2008-2012, and 2009-2013 waves of the American Community Survey (ACS). We include data on total population, median household income, median age, and education level. We obtain a measure of block group population density by dividing population by land area. To generate panel data for each variable, we use the average value for each variable across all waves available for that year. So, for example, the value for population in 2010 would be the average of the 2006-2010, 2007-2011, 2008-2012, and 2009-2013 values for this variable. We use the 2009-2013 values for 2014.[7] In addition to Census data, we also draw town-level voting registration data from the Office of CT's Secretary of the State (http://www.ct.gov/sots) for each year in our study period. Finally, we bring in annual county-level roofing contractor wage data from the U.S. Bureau of Labor Statistics (http://www.bls.gov).

3.1 Preparation of the Panel Dataset

Due to the possibility of unobserved heterogeneity at the localized level, we convert our data to a panel dataset, where an observation is at the Census block group-year level. This leads

[5] The only other utilities in CT are small municipal utilities in Bozrah, Norwich, and Wallingford (Graziano and Gillingham 2015).

[6] As a robustness check, we also analyze the full dataset of purchased and third-party-owned (TPO) systems, controlling for third party ownership in the regressions. The results, presented in Appendix H, should be interpreted with great caution due to the known issues with over-reporting of TPO system prices. They are suggestive of a different consumer decision-process for TPO systems.

[7] We also employed an alternative approach, using only the mid-year of each ACS wave and interpolating for all missing years, but this made no practical difference to our results. More generally, as shown in Section 6.2, our results are also robust to the exclusion of all demographic variables.

to a balanced panel of 10,738 observations in 1,534 block groups. For each block group-year, we take the average of the system price, system size, and incentive level granted. We also create an indicator variable for a Solarize campaign occurring in the given block group-year.

Using this panel dataset for our empirical analysis necessitates one further step. Many observations refer to block group-year combinations where there are no installations. In fact, our hurdle model approach is motivated by these excess zeros in our dataset. However, we still have to determine the relevant installation price, system size, and incentive for observations with no recorded contracts. This is a common issue in the empirical literature and we take a conservative stance by examining two different approaches.

In our primary approach, we fill in the missing price and other variables with the average annual value for the same municipality, and if this is not possible, we use the average annual value within the county of the installation. Given that our adoption data in these block groups is censored at zero, this approach may underestimate the price. Thus, as shown in Section 6.2, we perform several robustness checks and find that our results are robust to this choice. Our post-incentive PV system price per W variable is based on the ratio of the average block group-year PV system prices and the average system size.

3.2 Trends and Summary Statistics

Figure 1 displays the overall trend in installations and PV system prices (in 2014 dollars) during our study period. Between 2008 and 2014, the real average (post-incentive) price falls by more than 40 percent, with a brief spike in 2009. This spike in the post-incentive price is commonly attributed to a roughly 50 percent drop in the incentive between 2008 and 2009. Consistent with larger trends in the global PV market (Barbose et al. 2013), real installation costs have been steadily decreasing during the entire study period. Figure 1 also shows a substantial increase in PV system installations after 2011. Some of this increase after 2011 consists of installations under Solarize programs, which involve both lower prices and additional solar marketing.

Table 3 presents summary statistics for our panel dataset of 10,738 block group-year observations. The number of PV system installations is a count variable, with a mean of 0.48 and a variance of 1.19. Hence, these data exhibit clear overdispersion relative to standard count models, such as the Poisson model, in which an underlying assumption is the equality of mean and variance of the count variable (see Section 4.2). Figure 2 reveals why this might be the case: the distribution of installations in a block group in a given year is very strongly skewed, with most of the mass (over 74 percent) centered at zero. This visually illustrates the issue of excess zeros in our dataset.

Table 4 presents summary statistics for the truncated dataset with positive installations. Comparing Tables 3 and 4 reveals some differences. Most notably, there is a much stronger presence of Solarize campaigns in the truncated sample (12 percent versus 6 percent). The truncated sample also has a lower average population density, consistent with Graziano and Gillingham (2015), who show that most installations in CT occur in suburban or rural areas. The truncated sample also tends to have slightly larger systems, likely because it more heavily samples block groups that are slightly wealthier and have a greater unobserved preference for PV systems.

4 Empirical Specification: Preliminaries

Let Y_{it} denote the number of PV installations in block group i purchased in year t. We use $\boldsymbol{W}_{it}$ to denote a vector containing the installation price p_{it} and demand shifters. Demand for solar is therefore represented by the following general function:

$$Y_{it} = D_{it}(\boldsymbol{W}_{it}, \boldsymbol{\Theta}), \tag{1}$$

where $\boldsymbol{\Theta}$ is a vector of parameters.

Before we describe the Poisson hurdle model, which is our preferred empirical specification, we first briefly review two alternative specifications: linear and Poisson. This exposition lays the foundation for our Poisson hurdle model in Section 5.

4.1 Linear

In this specification, $D_{it}(\cdot)$ is modeled as a linear function and (1) becomes:

$$Y_{it} = \boldsymbol{W}_{it}^{'}\boldsymbol{\Theta} + \epsilon_{it}, \tag{2}$$

where ϵ_{it} denotes an idiosyncratic error term. We can expand (2) further as follows:

$$Y_{it} = \alpha_i + \boldsymbol{w}_{it}^{'}\boldsymbol{\theta} + \mu_t + \epsilon_{it},$$

where α_i captures block group-specific effects, μ_t is a year fixed effect that captures time-varying demand shocks common across block groups, and $\boldsymbol{w}_{it}$ contains all remaining observed regressors.

The advantage of using a linear model is that it can easily accommodate fixed effects to control for region-specific unobserved heterogeneity. Furthermore, the linear model also allows for the straightforward implementation of an instrumental variables procedure to address the endogeneity of prices. However, it is well-recognized in the literature (e.g., King 1988; Wooldridge 2002) that a linear model is poorly suited for a count data setting. A linear model not only misspecifies the count data generating process, but it also often predicts negative and non-integer outcome values.

4.2 Poisson

In this specification, we model the data generating process as a Poisson process, which is perhaps a better model for our count data. More specifically, if Y_{it} is modeled as a Poisson random variable with parameter $\lambda_{it} > 0$, its probability mass function, denoted by Po, is given by:

$$Po(y|\lambda_{it}) = Pr[Y_{it} = y|\lambda_{it}] = \frac{\lambda_{it}^y e^{-\lambda_{it}}}{y!} \quad \text{, where } y = 0, 1, 2, ..., \infty.$$

Furthermore, it follows that

$$\mathrm{E}[Y_{it}|\lambda_{it}] = \mathrm{Var}[Y_{it}|\lambda_{it}] = \lambda_{it}, \tag{3}$$

i.e., the conditional mean and variance of the Poisson random variable are equal. The parameter λ_{it} is modeled as an exponential function of the covariates:

$$\lambda_{it} = e^{\boldsymbol{W}'_{it}\boldsymbol{\Theta}}. \tag{4}$$

Therefore, in this specification, the conditional expectation of (1) is:

$$\mathrm{E}[Y|\boldsymbol{W}_{it}, \boldsymbol{\Theta}] = e^{\boldsymbol{W}'_{it}\boldsymbol{\Theta}} \equiv e^{\alpha_i + \boldsymbol{w}'_{it}\boldsymbol{\theta} + \mu_t}. \tag{5}$$

To facilitate discussion of this common model, we can then further rewrite (5) using $\boldsymbol{X}$ to denote the vector of observed regressors and time dummies, and $\boldsymbol{\delta}$ to denote the vector of all slope parameters (i.e., *excluding the time-invariant fixed effects parameters*):

$$\mathrm{E}[Y|\boldsymbol{X}_{it}, \boldsymbol{\delta}, \alpha_i] = e^{\alpha_i + \boldsymbol{X}'_{it}\boldsymbol{\delta}}. \tag{6}$$

The presence of individual-specific fixed effects usually poses a challenge in the estimation of nonlinear panel models, such as the one presented in equation (6), since these parameters cannot be eliminated from the model through any of the standard transformations used in linear models. If heterogeneity is left completely unrestricted and the model is estimated directly with the fixed parameters included, the estimates would suffer from the "incidental parameters problem" first noted by Neyman and Scott (1948). The problem arises since the estimator of the fixed effects parameter α_i is based only on the number of observations for each $i \in \{1, ..., N\}$, and that number remains fixed even as $N \to \infty$. Hence, the estimator $\hat{\alpha}_i$ is inconsistent, while the estimator of $\hat{\boldsymbol{\delta}}$ is only consistent if it can be expressed separately from $\hat{\alpha}_i$. Fortunately, as shown by Blundell et al. (2002), the latter is the case in the Poisson panel model. In particular, when this model is estimated by maximizing the log-likelihood function, the first-order conditions can be rearranged so that $\hat{\boldsymbol{\delta}}$ does not depend on the estimates of the fixed effects coefficients. Therefore, the maximum likelihood estimator (MLE) for $\boldsymbol{\delta}$ in our model with block group-specific constants does not suffer from the incidental parameters problem.[8]

Yet even with time-invariant unobserved heterogeneity accounted for, $\hat{\boldsymbol{\delta}}_{MLE}$ may still be inconsistent due to the endogeneity of the price variable in $\boldsymbol{W}_{it}$. One potential solution to the presence of an endogenous regressor in a Poisson model is the control function (CF) approach. Imbens and Wooldridge (2007) lay out a CF approach for addressing endogeneity in a cross-sectional Poisson model, which can be easily generalized to a panel data setting.

First, re-express the vector of covariates as $\boldsymbol{W} = (p, \boldsymbol{Z}_1)'$ and the corresponding vector of parameters as $\boldsymbol{\Theta} = (\gamma, \boldsymbol{\Gamma})'$. Furthermore, let $\boldsymbol{Z}$ be a vector of exogenous variables, such that $\boldsymbol{Z}_1 \subset \boldsymbol{Z}$. Combining (3) and (4) and dropping all subscripts for simplicity, we can flexibly re-write the conditional mean as:

$$\mathrm{E}[Y|\boldsymbol{Z}, p, \zeta] = \exp(\gamma p + \boldsymbol{Z}_1' \boldsymbol{\Gamma} + \zeta), \tag{7}$$

where ζ is an unobserved component of the conditional expectation that is correlated with

[8]Note that if the Poisson model contains unobserved market-level demand shocks that vary both cross-sectionally and by time, the corresponding parameters α_{it} can no longer be substituted out in the MLE's first-order conditions in a way that would allow solving separately for $\boldsymbol{\delta}$. Consistent estimation of $\boldsymbol{\delta}$ would then require integrating out the demand shocks. Interestingly, such a random Poisson model, in which the shocks exhibit substantial variance, can produce a large number of zero outcomes and overdispersion relative to the standard Poisson distribution, which is consistent with our data. Nonetheless, we demonstrate in Appendix E.1 by simulating a random Poisson model that even in that setting the hurdle model provides very close estimates to the true elasticities. We thank an anonymous referee for pointing us to this finding.

p. Let the endogenous variable p be written as:

$$p = \boldsymbol{Z}'\Pi + r. \tag{8}$$

Applying the law of iterated expectations to (7) and (21),

$$\mathrm{E}[Y|\boldsymbol{Z}, p] = \mathrm{E}[Y|\boldsymbol{Z}, r] = \mathrm{E}[\mathrm{E}[Y|\boldsymbol{Z}, \zeta, r]|\boldsymbol{Z}, r] = \exp(\gamma p + \boldsymbol{Z}_1'\boldsymbol{\Gamma})\mathrm{E}[\exp(\zeta)|\boldsymbol{Z}, r]. \tag{9}$$

Note that both ζ and r are independent of $\boldsymbol{Z}$. The standard assumption here is joint normality of ζ and r, and under this assumption it can be shown that $\mathrm{E}[\exp(\zeta)|r] = \exp(\rho r)$ for some scalar ρ. Hence, (9) becomes

$$\mathrm{E}[Y|\boldsymbol{Z}, p] = \exp(\gamma p + \boldsymbol{Z}_1'\boldsymbol{\Gamma} + \rho r) = \exp(\boldsymbol{W}'\boldsymbol{\Theta} + \rho r),$$

which suggests a straightforward two-stage estimation procedure. In the first stage, we estimate a linear regression of price on all excluded and included instruments in order to obtain the residual $\hat{r}$. In the second stage, we estimate a Poisson model with $\boldsymbol{W}$ and $\hat{r}$ as covariates using MLE, as described earlier. Of course, since this is a two-stage procedure, bootstrapping is the easiest way to obtain valid standard errors.

The Poisson model can thus be used to address unobserved heterogeneity with fixed effects and endogeneity of price. Moreover, it is likely a much better model for count data than the linear specification. However, by virtue of its underlying distribution, the Poisson model predicts that only a small proportion of all counts in the data equal zero and is likely misspecified when there are excess zeros in the outcome variable, so the finite-sample bias from running quasi-maximum likelihood would be expected to be large.

5 Empirical Specification: Hurdle Model

5.1 General Form

In situations with very large numbers of zeros, the data generating process is perhaps better thought of as a two-stage process. A hurdle model, originally formulated in Mullahy (1986), is designed to model such a two-stage process.[9] The first stage determines whether the count variable has a zero or positive realization. A positive realization means that the “hurdle” is crossed, in which case the exact outcome value is modeled by a truncated count distribution. The two stages are functionally independent (Cameron and Trivedi 2013). For a given non-negative count variable Y, let the first-stage process be driven by a distribution function f_1, while the second stage governed by f_2. Then, the complete distribution of Y is given by:

$$Pr(Y = y) = \begin{cases} f_1(0) & \text{if } y = 0, \\ (1 - f_1(0))\, f_2(y) & \text{if } y = 1, 2, ..., \infty. \end{cases}$$

It is common to model the first stage using a logistic distribution, i.e., $f_1 = 1 - Pr[Y > 0]$, where $Pr[Y > 0]$ is estimated through a logit regression model, and the second stage as a zero-truncated Poisson, i.e., $f_2 = \frac{Po(y)}{(1-Po(0))}$ (e.g., Min and Agresti 2005; Gallop et al. 2013; Neelon et al. 2013). Depending on the setting, the explanatory variables in the first- and second-stage regressions may differ.

Adapting this framework to our setting,

$$Pr[Y_{it} > 0|\boldsymbol{W}_{it}, \boldsymbol{\Theta}_1] = \frac{e^{\boldsymbol{W}'_{it}\boldsymbol{\Theta}_1}}{1 + e^{\boldsymbol{W}'_{it}\boldsymbol{\Theta}_1}}\ ,\ f_2(y|\lambda_{it}) = \frac{\lambda_{it}^{y}}{y!\,(e^{\lambda_{it}} - 1)}\ , \text{ and } \lambda_{it} = e^{\boldsymbol{W}'_{it}\boldsymbol{\Theta}_2}.$$

Note that while we expect the same determinants to play a role in both stages of the

[9] A zero-inflated Poisson (ZIP) model is a similar model, but not as readily extendable to a panel data context with fixed effects and instrumental variable estimation. Gilles and Kim (2017) develop a quasi-conditional likelihood method for estimating ZIP with fixed effects only in the count part that is consistent under strict exogeneity of all regressors. Kitazawa (2014) derives moment conditions that allow for the estimation of a fixed effects ZIP model with exogenous regressors in the logit component and predetermined explanatory variables in the Poisson component.

demand model, there is no behavioral reason to presume *a priori* that the effects of these determinants would be identical. Therefore, we allow $\mathbf{\Theta}_1$ to be different than $\mathbf{\Theta}_2$.

Let $\iota_{it} = I(Y_{it} > 0)$, where $I(\cdot)$ denotes the indicator function. The log-likelihood function of this model can be expressed as follows:

$$\begin{aligned} L(\mathbf{\Theta}_1, \mathbf{\Theta}_2) &= \\ & \sum_{i=1}^{N}\sum_{t=1}^{T}\left\{(1-\iota_{it})\log\left[\frac{1}{1+\exp(\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_1)}\right] + \iota_{it}\log\left[\frac{\exp(\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_1)}{1+\exp(\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_1)}\right]\right\} \\ & + \sum_{i=1}^{N}\sum_{t=1}^{T}\left\{\iota_{it}\left[Y_{it}\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_2 - \log(Y_{it}!) - \log\Big(\exp(\exp(\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_2)) - 1\Big)\right]\right\} \\ & \equiv L_l(\mathbf{\Theta}_1) + L_t(\mathbf{\Theta}_2), \end{aligned} \tag{10}$$

where $L_l(\mathbf{\Theta}_1)$ is the log-likelihood function for the logit model and $L_t(\mathbf{\Theta}_2)$ is the log-likelihood for the truncated Poisson model. Therefore, maximizing $L(\mathbf{\Theta}_1, \mathbf{\Theta}_2)$ is equivalent to maximizing the two terms separately, which implies that estimating the hurdle model effectively reduces to separately estimating a binary logit model and a zero-truncated Poisson model.

This model is convenient both for estimation of the parameters and for estimation of the price elasticity of demand. To see this, first note that the conditional mean of the outcome variable in this model is given by:

$$\begin{aligned} \mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \mathbf{\Theta}] &= Pr[Y_{it} > 0|\boldsymbol{W}_{it}, \mathbf{\Theta}_1]\mathrm{E}_t[Y_{it}|\boldsymbol{W}_{it}, \mathbf{\Theta}_2] \\ &= \frac{\exp(\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_1)}{1+\exp(\boldsymbol{W}_{it}^{'}\mathbf{\Theta}_1)}\frac{\lambda_{it}}{1-\exp(-\lambda_{it})}, \end{aligned} \tag{11}$$

where $\mathbf{\Theta} = (\mathbf{\Theta}_1, \mathbf{\Theta}_2)^{'}$ and $\mathrm{E}_t[\cdot]$ is the mean in the truncated Poisson model. Then, the price elasticity η at the mean values of p and Y can be derived through the following expression:

$$\eta = \frac{\partial\left[\mathrm{E}[Y_{it}]\right]}{\partial p_{it}}\frac{\mathrm{E}[p_{it}]}{\mathrm{E}[Y_{it}]} = \frac{\partial\left[Pr[Y_{it} > 0]\right]}{\partial p_{it}}\frac{\mathrm{E}[p_{it}]}{Pr[Y_{it} > 0]} + \frac{\partial\left[\mathrm{E}_t[Y_{it}]\right]}{\partial p_{it}}\frac{\mathrm{E}_t[p_{it}]}{\mathrm{E}_t[Y_{it}]} \equiv \eta_l + \eta_t.$$

In other words, η is a sum of the price elasticity from the logit and truncated Poisson components of the hurdle model.

An important characteristic of this model is that, unlike the Poisson specification, it can accommodate both overdispersion and underdispersion in the full dataset. In fact, it can be shown that

$$\mathrm{Var}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}] = \mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}] + \mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}] \left(\lambda_{it} - \mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}]\right).$$

Therefore, depending on the relative magnitudes of λ_{it} and $\mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}]$, the model could result in $\mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}] \leq \mathrm{Var}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}]$ or vice versa. In particular, $\mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}]$ in (11) is a linear function of $Pr[Y_{it} > 0]$, so an excessive number of zero outcomes implies that $Pr[Y_{it} > 0]$ (and, hence, $\mathrm{E}[Y_{it}|\boldsymbol{W}_{it}, \boldsymbol{\Theta}]$) is relatively small, which leads to overdispersion relative to the Poisson model.

Thus, the hurdle model is a promising approach for our setting. But the challenge of handling unobserved heterogeneity and endogeneity remains. Other approaches, such as the negative binomial model, can only handle an extensive set of fixed effects and endogeneity with great difficulty and restrictive assumptions. Furthermore, while studies have extended the traditional Poisson model to accommodate individual-specific fixed effects and instrumental variables (e.g., Windmeijer 2008), there is little in the current literature on such extensions of the hurdle model.

5.2 Fixed Effects

As suggested by equation (10), estimating the hurdle model is equivalent to separately estimating a logit model and truncated Poisson model. This allows us to address the challenge of accommodating block group fixed effects separately in each of the two components of the hurdle model.

Logit

With fixed effects, the logit model becomes

$$Pr[Y_{it}>0|\boldsymbol{W}_{it},\boldsymbol{\Theta}_1] = Pr[\iota_{it}=1|\boldsymbol{W}_{it},\boldsymbol{\Theta}_1] = \frac{\exp(\alpha_i^l+\boldsymbol{w}_{it}'\boldsymbol{\theta}_1+\mu_t^l)}{1+\exp(\alpha_i^l+\boldsymbol{w}_{it}'\boldsymbol{\theta}_1+\mu_t^l)} \equiv \frac{b_i\exp(\boldsymbol{X}_{it}'\boldsymbol{\delta}_1)}{1+b_i\exp(\boldsymbol{X}_{it}'\boldsymbol{\delta}_1)}, \tag{12}$$

where $b_i\equiv\exp(\alpha_i^l)$ is a block group-specific coefficient, μ_t^l is the year fixed effect, while $\boldsymbol{X}$ and $\boldsymbol{\delta}_1$ again denote the vectors of time variables and observed regressors and their respective coefficients. As discussed in Section 4.2, the presence of N block group-specific coefficients can be problematic in a nonlinear setting, unless $\hat{\boldsymbol{\delta}}_{\boldsymbol{1}}$ can be derived independently of $\hat{b}_i$. For the case of logit with fixed effects, early work by Rasch (1960, 1961), Andersen (1972), and Chamberlain (1980) established that model parameters can be estimated using a conditional maximum likelihood estimator (CMLE). In particular, this estimator uses the sum of outcomes within each group i as a minimal sufficient statistic for i's group-specific parameter. As a result, groups with all equal outcomes (all positive or all zeros) are not used in the estimation. Using ℓ_i^l to denote the conditional log-likelihood for all observations in block group i, it can be shown that:

$$\ell_i^l(b_i,\boldsymbol{\delta}_1|\sum_{t=1}^{T}\iota_{it}) = \ell_i^l(\boldsymbol{\delta}_1|\sum_{t=1}^{T}\iota_{it}). \tag{13}$$

In other words, the conditional log-likelihood does not depend on the fixed effects parameter b_i. The remaining parameters are then estimated by maximizing $\sum_{i=1}^{N}\ell_i^l$. The conditional maximum likelihood estimator for this model, $\hat{\boldsymbol{\delta}}_{1,CMLE}$, is consistent under mild restrictions on the rate at which the sequence of b_i's grows to infinity (Andersen 1970; Chamberlain 1980).[10]

[10]However, as noted by Greene and Hensher (2010), the conditional likelihood of the logit is based on a restricted dataset that excludes all groups with equal outcomes over time, as well as all group-specific parameters, and is therefore not comparable to the unrestricted likelihood of alternative model specifications. This prevents the implementation of model selection tests based on likelihood ratio statistics, including selection tests among non-nested models, such as the Vuong test (Vuong 1989) and its extensions (e.g., Chen et al. 2007).

Truncated Poisson

The truncated Poisson model is estimated using only observations for which $Y_{it} > 0$. Allowing for unobserved block group heterogeneity, the truncated Poisson parameter λ_{it} for each one of these observations can be expressed as follows:

$$\lambda_{it} = \exp(\alpha_i^t + \boldsymbol{w}_{it}^{'}\boldsymbol{\theta}_2 + \mu_t^t) \equiv c_i \exp(\boldsymbol{X}_{it}^{'}\boldsymbol{\delta}_2) \equiv c_i\beta_{it}, \tag{14}$$

with year fixed effects represented by μ_t^t, $c_i \equiv \exp(\alpha_i^t)$, $\boldsymbol{\delta}_2$, as before, denoting a vector of parameters for the time variables and all remaining observed regressors, and $\beta_{it} \equiv \exp(\boldsymbol{X}_{it}^{'}\boldsymbol{\delta}_2)$. Unlike the fixed effects Poisson model, estimating a zero-truncated Poisson with individual fixed effects through maximum likelihood does not allow for $\boldsymbol{\delta}_2$ to be estimated independently of the fixed effects coefficients. Therefore, MLE is inconsistent in this specification.

As shown by Majo and van Soest (2011), conditional maximum likelihood can be employed in this setting, in a similar fashion to the fixed effects logit model, which eliminates the fixed effects parameters from the estimation. Majo and van Soest (2011) demonstrate this method for a two-period panel dataset. In what follows, we generalize their procedure to any panel with an arbitrary number of longitudinal observations.

Let $\boldsymbol{Y}_i = (Y_{i1}, ..., Y_{iT})^{'} \in \mathcal{Y}$ denote the T-dimensional vector of outcomes in block group i. Also, let $\boldsymbol{X}_i = (\boldsymbol{X}_{i1}, ..., \boldsymbol{X}_{iT})^{'}$ denote the matrix of regressors. Conditional on $\boldsymbol{X}_i$ and the model parameters c_i and $\boldsymbol{\delta}_2$, the outcomes in i have a joint truncated Poisson distribution with probabilities

$$Pr(\boldsymbol{Y}_i = \boldsymbol{y}_i | \boldsymbol{X}_i = \boldsymbol{x}_i, c_i, \boldsymbol{\delta}_2) = \prod_{t=1}^{T} \frac{\lambda_{it}^{y_{it}}}{y_{it}!\,(e^{\lambda_{it}} - 1)}.$$

We now derive an expression for the joint distribution of $\boldsymbol{Y}_i$ conditional not only on $\boldsymbol{X}_i$ and the parameters, but also on $\sum_{t=1}^{T} Y_{it}$. Through an application of the Bayes' rule, it can be shown that

$$g(\boldsymbol{y}_i|\boldsymbol{x}_i, n_i, c_i, \boldsymbol{\delta}_2) = Pr(\boldsymbol{Y}_i = \boldsymbol{y}_i|\boldsymbol{X}_i = \boldsymbol{x}_i, \sum_{t=1}^{T} Y_{it} = n_i, c_i, \boldsymbol{\delta}_2) = \frac{1}{h_i}\prod_{t=1}^{T}\frac{n_i!}{y_{it}!}\lambda_{it}^{y_{it}}, \tag{15}$$

where $h_i = h(T, n_i, \lambda_{i1}, ..., \lambda_{iT})$ is an expression that depends on the total number of observations with positive Y_{it} for a given i. In order to represent the general form of h, we simplify notation by dropping the i subscript. Then, for any arbitrary $T \in (1, \bar{T}]$, where $\bar{T} > 5$, the general expression for h takes the following form:

$$\begin{aligned} h(T, n, \lambda_1, ..., \lambda_T) \; = \; & \left(\sum_{t=1}^{T}\lambda_t\right)^n - \sum_{l=1}^{T}\left(\sum_{t=1}^{T}\lambda_t - \lambda_l\right)^n \\ & + I(T>2)\sum_{l=1}^{T}\sum_{\substack{m=1\\ m\neq l}}^{T}\left(\sum_{t=1}^{T}\lambda_t - \lambda_l - \lambda_m\right)^n \\ & - I(T>3)\sum_{l=1}^{T}\sum_{\substack{m=1\\ m\neq l}}^{T}\sum_{\substack{p=1\\ p\neq l\\ p\neq m}}^{T}\left(\sum_{t=1}^{T}\lambda_t - \lambda_l - \lambda_m - \lambda_p\right)^n \\ & + I(T>4)\sum_{l=1}^{T}\sum_{\substack{m=1\\ m\neq l}}^{T}\sum_{\substack{p=1\\ p\neq l\\ p\neq m}}^{T}\sum_{\substack{q=1\\ q\neq l\\ q\neq m\\ q\neq p}}^{T}\left(\sum_{t=1}^{T}\lambda_t - \lambda_l - \lambda_m - \lambda_p - \lambda_q\right)^n \\ & - \; \ldots \; + I(T=\bar{T})(-1)^{\bar{T}+1}\sum_{t=1}^{\bar{T}}\lambda_t^n \; . \end{aligned} \tag{16}$$

Note that, following from (14), and again dropping the i subscript for simplicity, we have $\lambda_t = c\beta_t$. So, λ_t is a linear function of c. Then, it is easy to show that $h(T, n, \lambda_1, ..., \lambda_T) = h(T, n, c\beta_1, ..., c\beta_T) = c^n h(T, n, \beta_1, ..., \beta_T)$, i.e., the function $h(\cdot)$ is homogeneous of degree n in $(\beta_1, ..., \beta_T)'$. So, from (15),

$$\begin{aligned} \frac{1}{h(T, n_i, \lambda_{i1}, ..., \lambda_{iT})}\prod_{t=1}^{T}\frac{n_i!}{y_{it}!}\lambda_{it}^{y_{it}} \; &= \; \frac{1}{c_i^{n_i} h(T, n_i, \beta_{i1}, ..., \beta_{iT})}\prod_{t=1}^{T}\frac{n_i!}{y_{it}!}c^{y_{it}}\beta_{it}^{y_{it}} \\ &= \; \frac{1}{h(T, n_i, \beta_{i1}, ..., \beta_{iT})}\prod_{t=1}^{T}\frac{n_i!}{y_{it}!}\beta_{it}^{y_{it}}, \end{aligned}$$

i.e., $g(\boldsymbol{y}_i|\boldsymbol{x}_i, n_i, c_i, \boldsymbol{\delta}_2) = g(\boldsymbol{y}_i|\boldsymbol{x}_i, n_i, \boldsymbol{\delta}_2)$, and, hence, the conditional distribution does not depend on the nuisance parameter c_i. The conditional log-likelihood then takes the following form:

$$L(\boldsymbol{\delta}_2) = \sum_{i=1}^{N} \ell_i^t(\boldsymbol{\delta}_2 | \sum_{t=1}^{T} Y_{it}) =$$
$$\sum_{i=1}^{N} \left\{ \log[(\sum_{t=1}^{T} Y_{it})!] - \sum_{t=1}^{T} \log[Y_{it}!] + \sum_{t=1}^{T} Y_{it} \log(\beta_{it}) - \log(h_i) \right\}, \tag{17}$$

where ℓ_i^t denotes the conditional log-likelihood for all observations from block group i. Note that, because we condition on the sum of outcomes in a block group, only groups with more than one positive outcome during the study period contribute to the log-likelihood.

This procedure is rather general and can be applied to an arbitrarily large number of longitudinal panel observations and any number of model parameters. The parameter estimator of this fixed effects truncated Poisson model is obtained as $\hat{\boldsymbol{\delta}}_{2,CMLE} = \arg\max_{\boldsymbol{\delta}_2} L(\boldsymbol{\delta}_2)$. Under strict exogeneity of the regressors $\boldsymbol{X}$, $\hat{\boldsymbol{\delta}}_{2,CMLE}$ can be shown to be a consistent estimator of $\boldsymbol{\delta}_2$ (See Appendix B).

5.3 Endogeneity

Both the logit and truncated Poisson conditional likelihood estimators, discussed in Section 5.2, are consistent, provided that the respective models are specified appropriately. However, this would no longer be the case with endogeneity of price. We thus extend the hurdle model further in order to accommodate the implementation of a suitable instrumental variable procedure. Once again, we address the problem separately in the logit and truncated Poisson portions of the model.

Logit

In the discrete choice literature, a number of methods have been developed to tackle endogeneity in demand settings. These include the product-market control (and instrumental variable) approach of Berry et al. (1995), a simulated maximum likelihood approach developed by Gupta and Park (2009), and Bayesian methods employed by Yang et al. (2003) and Jiang et al. (2009). More recently, Petrin and Train (2010) propose a control function approach that is arguably easier to estimate and more flexible than the above methods, as it does not require invoking equilibrium or imposing strict distributional assumptions for the identification of demand parameters.

As in Section 4.2, dropping subscripts for simplicity, let $\boldsymbol{W} = (p, \boldsymbol{Z}_1)'$ and $\boldsymbol{Z}_1 \subset \boldsymbol{Z}$, where $\boldsymbol{Z}$ is a vector of exogenous variables. Furthermore, let $\boldsymbol{\Theta}_1 = (\gamma_1, \boldsymbol{\Gamma}_1)'$. Suppose that the purchase of any positive number of PV systems in a given block group generates utility u, given by:

$$u = \gamma_1 p + \boldsymbol{Z}_1' \boldsymbol{\Gamma}_1 + \zeta_1, \tag{18}$$

where ζ_1 denotes an idiosyncratic term that is correlated with price p. As in Section 4.2, let

$$p = \boldsymbol{Z}' \Pi + r. \tag{19}$$

Since ζ_1 is correlated with p but not with Z, there exists some function $CF(r, \rho_1)$, where ρ_1 is a parameter, such that $\zeta_1 = CF(r, \rho_1) + \tilde{\zeta}_1$ and $\tilde{\zeta}_1$ is uncorrelated with p. We can then re-write (18) as

$$u = \gamma_1 p + \boldsymbol{Z}_1' \boldsymbol{\Gamma}_1 + CF(r, \rho_1) + \tilde{\zeta}_1 = \boldsymbol{W}' \boldsymbol{\Theta}_1 + CF(r, \rho_1) + \tilde{\zeta}_1. \tag{20}$$

Therefore, if we estimate our model including the control function $CF(r, \rho_1)$ with valid instruments, we can consistently estimate $\boldsymbol{\Theta}_1$.

Petrin and Train (2010) suggest several simplifying assumptions. First, as an alternative to specifying a joint distribution for both error terms in (18) and (19), one could enter r flexibly in the utility and then choose a distributional assumption for $\tilde{\zeta}_1$. Second, the control function can be approximated as a linear function. Accordingly, we specify $CF(r, \rho_1)$ as $\rho_1 r$ in (20) and assume that $\tilde{\zeta}_1$ is distributed i.i.d. type I extreme value.[11] The probability of $Y > 0$ is now given by:

$$Pr[Y_{it} > 0|\boldsymbol{W}_{it}, \boldsymbol{\Theta}_1] = \frac{b_i \exp(\boldsymbol{X}'_{it}\boldsymbol{\delta}_1 + \rho_1 r)}{1 + b_i \exp(\boldsymbol{X}'_{it}\boldsymbol{\delta}_1 + \rho_1 r)}.$$

As in our earlier discussion from Section 4.2, this CF specification implies a two-stage estimation procedure. In the first stage, we run a linear regression of price on all excluded and included instruments. The residuals from this stage are then included as a covariate in the second-stage estimation, which, following the discussion from Section 5.2, is carried out using conditional maximum likelihood in order to eliminate the nuisance parameters. Finally, we bootstrap the standard errors to ensure that our inference accounts for both stages.

Truncated Poisson

To our knowledge, we are the first to address endogeneity of one or more of the regressors in a fixed effects truncated Poisson model. In what follows, we develop a generalized method of moments (GMM) procedure for the consistent estimation of the vector of truncated Poisson slope parameters $\boldsymbol{\delta}_2$ in cases where $\boldsymbol{X}$ is no longer strictly exogenous.

Recall that the function $h(T, n_i, \beta_{i1}, ..., \beta_{iT})$, described in Section 5.2, is homogeneous of degree n_i in $(\beta_{i1}, ..., \beta_{iT})'$. This implies that $\sum_{t=1}^{T} \frac{\partial h(\cdot)}{\partial \beta_{it}} \beta_{it} = n_i h(\cdot)$, or $\sum_{t=1}^{T} \frac{\partial h(\cdot)}{\partial \beta_{it}} \frac{\beta_{it}}{h(\cdot)} = n_i$

For convenience, we now define a non-linear function ϕ_{it} of all regressors and slope parameters in the truncated Poisson model:

$$\phi_{it}(\boldsymbol{X}_i, \boldsymbol{\delta}_2) \equiv \frac{\frac{\partial h(T, \sum_{t=1}^{T} Y_{it}, \beta_{i1}, ..., \beta_{iT})}{\partial \beta_{it}} \beta_{it}}{h(T, \sum_{t=1}^{T} Y_{it}, \beta_{i1}, ..., \beta_{iT})}.$$

By definition, it follows that $\sum_{t=1}^{T} \phi_{it} = \sum_{t=1}^{T} Y_{it}$ since both are equal to n_i. Thus, $\frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \phi_{it} =$

[11] Alternatively, $CF(r, \rho_1)$ can be specified as a quadratic function (e.g., Olley and Pakes 1996). We also re-estimated our model using a quadratic control function and found the results to be quite robust.

$\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} Y_{it} \xrightarrow{p} \mathrm{E}_t[Y_{it}]$.[12] Note that this asymptotic property follows directly from the functional form of ϕ_{it} and does not require any assumptions about the exogeneity of $\boldsymbol{X}_i$. Hence, in our context ϕ_{it} can be interpreted as the predicted number of installations in block group i at time t, conditional on the occurrence of a positive number of installations.

Based on this interpretation, we proceed to construct the following model of the zero-truncated demand for solar PV:

$$Y_{it} = \phi_{it}(\boldsymbol{X}_i, \boldsymbol{\delta}_2) + \xi_{it}, \tag{21}$$

where ξ_{it} is the econometrician's error, which represents block group and year-specific idiosyncratic shocks that influence the number of adoptions.

Let $\boldsymbol{\xi}_i = (\xi_{i1}, ..., \xi_{iT})'$. As demonstrated in Appendix C, if $\boldsymbol{X}$ is strictly exogenous in the demand model, the resultant orthogonality with the error term $\boldsymbol{\xi}_i$ provides a moment condition $\mathrm{E}[\boldsymbol{X}_i'\boldsymbol{\xi}_i] = \mathbf{0}$ that can equivalently be derived from the expected value of the score for the conditional log-likelihood function in (17). Hence, under strict exogeneity of all regressors, the GMM estimator from the sample analog of $\mathrm{E}[\boldsymbol{X}_i'\boldsymbol{\xi}_i] = \mathbf{0}$ is equivalent to $\hat{\boldsymbol{\delta}}_{2,CMLE}$, which was shown to be a consistent estimator of $\boldsymbol{\delta}_2$ in Appendix B.

However, with one or more endogenous regressors in the model, $\mathrm{E}[\boldsymbol{X}_i'\boldsymbol{\xi}_i] = \mathbf{0}$ no longer holds and the estimator $\hat{\boldsymbol{\delta}}_{2,CMLE}$ is inconsistent. In that case, a vector $\boldsymbol{Z}$, comprised only of variables that are exogenous in the model, would still be orthogonal to the error term in (21), i.e.,

$$\mathrm{E}[\boldsymbol{Z}_i'\boldsymbol{\xi}_i] = \mathbf{0}, \tag{22}$$

where $\boldsymbol{Z}_i = (\boldsymbol{Z}_{i1}, ..., \boldsymbol{Z}_{iT})'$. Then, if $\boldsymbol{\delta}_2$ is a P-dimensional vector, we can use the sample analog of (22) in order to estimate $\boldsymbol{\delta}_2$ through a GMM estimator, as long as there are a total of at least P exogenous variables in $\boldsymbol{Z}$. Suppose $\boldsymbol{Z}_{it} \in \mathcal{Z} \subset \mathbb{R}^Q$, where $Q > P$, and let

[12] Recall that in the truncated Poisson model we restrict our attention only to $Y_{it} > 0$, and, hence, by the Law of Large Numbers, $\bar{Y}_t \xrightarrow{p} \mathrm{E}_t[Y_{it}]$, where $\bar{Y}_t$ is the mean of Y_{it} in the truncated sample.

$\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2) = \boldsymbol{Z}_i^{'}\boldsymbol{\xi}_i$. Then,

$$\hat{\boldsymbol{\delta}}_{2,GMM} = \arg\max_{\boldsymbol{\delta}_2} \left[\frac{1}{N}\sum_{i=1}^{N}\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2)\right]^{'} \hat{\boldsymbol{\Xi}} \left[\frac{1}{N}\sum_{i=1}^{N}\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2)\right], \tag{23}$$

where $\hat{\boldsymbol{\Xi}}$ is an optimal weighting matrix. In Appendix B, we demonstrate that under the standard GMM assumptions $\hat{\boldsymbol{\delta}}_{2,GMM}$ is a consistent estimator of $\boldsymbol{\delta}_2$.

5.4 Deriving the Price Elasticity

The estimators derived in Sections 5.2 and 5.3 do not allow us to estimate the fixed effects parameters in the model. As a result, we are not able to recover the predicted conditional expectations or calculate the exact marginal effects, needed for the derivation of price elasticity.[13] Instead, we develop, in the spirit of Kitazawa (2012), a procedure for obtaining average elasticity estimates that is shown to converge to the true average elasticity values as $N \to \infty$.

Logit

The average price elasticity in the logit model is given by $\eta_l = (1 - Pr[Y_{it} > 0])\,\gamma_1 \mathrm{E}[p_{it}]$, where γ_1 is the coefficient on price.

Note that $Pr[Y_{it} > 0|\boldsymbol{W}_{it}, \boldsymbol{\Theta}_1] = \frac{b_i \exp(\boldsymbol{X}_{it}^{'}\boldsymbol{\delta}_1)}{1+b_i \exp(\boldsymbol{X}_{it}^{'}\boldsymbol{\delta}_1)}$. Since we do not have an estimate of b_i, we cannot use the post-estimation predicted probabilities to calculate η_l. However, by the Law of Large Numbers, $\bar{\iota} = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\iota_{it}$ is a consistent estimator of $Pr[Y_{it} > 0]$. Similarly, $\bar{p} \xrightarrow{p} \mathrm{E}[p_{it}]$, where $\bar{p}$ is the sample average of price. Therefore, we can obtain a consistent estimate of η_l through the following expression:

$$\hat{\eta}_l = (1 - \bar{\iota})\hat{\gamma}_1\bar{p},$$

where $\hat{\gamma}_1$ can be obtained through a CF approach in our setting.

[13] Furthermore, this implies that model selection tests based on conditional expectations (e.g., Silva et al. 2015) cannot be implemented.

Truncated Poisson

Similarly, we derive an estimator for the average elasticity in the zero-truncated Poisson portion of the hurdle model. First, it can be shown that η_t is given by

$$\eta_t = (1 - \exp(-\lambda_{it})\mathrm{E}_t[Y_{it}])\,\gamma_2 \mathrm{E}_t[p_{it}]. \tag{24}$$

By re-arranging the expression for $\mathrm{E}_t[Y_{it}]$, we obtain

$$-\exp(-\lambda_{it}) = \frac{\lambda_{it}}{\mathrm{E}_t[Y_{it}]} - 1.$$

Hence, (24) becomes $\eta_t = (1 + \lambda_{it} - \mathrm{E}_t[Y_{it}])\,\gamma_2 \mathrm{E}_t[p_{it}]$. Just as in our earlier discussion, $\bar{Y}_t$ and $\bar{p}_t$, the means of Y and p from the truncated Poisson sample, are consistent estimators of $\mathrm{E}_t[Y_{it}]$ and $\mathrm{E}_t[p_{it}]$, respectively.

We also need to derive a separate estimator of λ, since we do not estimate c_i and $\lambda_{it} = c_i \exp(\boldsymbol{X}_{it}^{'}\boldsymbol{\delta}_2)$. Let

$$\mathrm{E}_t[Y_{it}] = \frac{\lambda_{it}}{1 - \exp(-\lambda_{it})} \equiv m(\lambda_{it}).$$

Then, $\lambda_{it} = m^{-1}(\mathrm{E}_t[Y_{it}])$. As shown in Appendix D, $m(\cdot)$ is a monotonic function over the relevant range of λ_{it} values in this model, implying that $m^{-1}(\cdot)$ is a one-to-one mapping from $\mathrm{E}_t[Y_{it}]$ to λ_{it}. Furthermore, since $\bar{Y}_t \xrightarrow{p} \mathrm{E}_t[Y_{it}]$, it follows that $m^{-1}(\bar{Y}_t)$ is a consistent estimator of λ_{it}. Hence,

$$\hat{\eta}_t = \left(1 + m^{-1}(\bar{Y}_t) - \bar{Y}_t\right)\hat{\gamma}_2 \bar{p}_t,$$

where $\hat{\gamma}_2$ is obtained through GMM in our setting.

5.5 Monte Carlo Simulations

While we have already proven the consistency of our estimator, we conduct a set of Monte Carlo simulations to evaluate its performance. In short, we simulate data from our data generating process and then apply our estimator and alternative estimators to the simulated data. We use a simple panel data framework, in which the outcome count variable is determined by a set of unobserved time-invariant and time-varying cross-section-specific factors and a single endogenous regressor. The model, data generation process, and parameter values used in these simulations are described in Appendix E.2.

Our results are encouraging. Table 5 shows the results from one representative simulation. The bias on each of the parameters of interest using the logit CF and truncated Poisson GMM is small: 3 percent for the logit CF and 0.2 percent for the truncated Poisson. Converting this to an elasticity, we find that our estimated (combined) elasticity is within 1.2 percent of true elasticity value. Compared to the results obtained using linear two-stage least squares (2SLS) and Poisson CF approach, we find that our estimated elasticity is much closer to the true value. We repeat this procedure with several different sets of parameter values and find similar results.[14]

6 Results

6.1 Primary Results

Our identification strategy relies on the validity of our instruments. As described earlier, we use two marginal cost shifters as instruments for the post-incentive price: average incentive levels in \$/W for the first 5 kW of installed capacity in each block group-year and county-year average roofer wages. Incentives are given directly to installers in CT, so they act as a shifter of the firm marginal cost. Due to the declining incentives and time lag between the submission and approval of contract applications, there is both temporal and cross-sectional variation in incentive levels.[15] County-year average roofer wages are used as a proxy for PV

[14]We also run another set of Monte Carlo simulations (described in Appendix E.3), which shows that even in the case where the true data generation process is Poisson with fixed effects and an endogenous regressor (i.e., it does not have excess zeros), the hurdle model still performs very well in recovering the true values.

[15]The time lag provides cross-sectional variation because the incentives granted are based on the approval date of the contract.

system installer labor costs, and are a valid shifter after controlling for income.[16] Appendix F contains the results of the first-stage regression, demonstrating that we do not have to worry about weak instruments.

Table 6 presents the results from estimating a linear model, Poisson model, and Poisson hurdle model with instruments for price.[17] Our preferred specification is the instrumental variables hurdle model, consisting of a logit regression estimated with a control function approach in column (3) and a trucated Poisson estimated by GMM in column (4). All columns include block group fixed effects and year dummies. Standard errors are clustered at the town level. At the bottom of the table, we present estimates of the price elasticity of demand for each specification.

We are most interested in the statistical and economic significance of the price coefficient. We find negative and statistically significant coefficients on the price variable in all model specifications. Rather than interpreting the coefficient directly, we find it more instructive to consider the price elasticity implied by the coefficient taken at the mean of our sample.[18] Estimating the model using a Poisson specification implies a considerably higher (in absolute value) price elasticity relative to the linear specification. A linear 2SLS regression yields a price elasticity of -0.62. Fitting a Poisson model generates an elasticity coefficient that is almost twice as high. However, both of these empirical specifications are unsuited for our data with excess zeros, as discussed in Section 4.

As shown earlier, the total price elasticity of the hurdle model, which is our preferred specification, is the sum of the elasticities from the logit and truncated Poisson. Summing

[16] While over-identification tests are never definitive, we use a Hansen's J over-identification test to examine the validity of the instruments in the truncated Poisson GMM specification of our hurdle model. Comfortingly, we find that we fail to reject the null of valid instruments, with a Chi-squared test statistic of 0.47 and a p-value of 0.49. We find similar results for other specifications.

[17] For completeness, Appendix G also presents output from the same model specification without the use of instrumental variables. We view these results with great caution as the coefficients are most likely biased due to the endogeneity of price.

[18] The price elasticity in the linear model is estimated as $\eta = \frac{\hat{\gamma}\bar{p}}{\bar{Y}}$, where $\hat{\gamma}$ is the two-stage least squares (2SLS) estimate of the price coefficient and $\bar{p}$ and $\bar{Y}$ are the sample means for price and number of installations. The price elasticity in the Poisson model is estimated as $\eta = \hat{\gamma}\bar{p}$, where $\hat{\gamma}$ is the Poisson CF estimate of the price coefficient.

the price coefficients in columns (3) and (4) shows that the hurdle model implies a price elasticity of solar PV system demand of -0.65. At the average system price and number of installations in our sample, this elasticity estimate suggests that a $1/W decrease in the installation price (well within the variation in our data) would lead to an increase in demand of approximately 0.083 additional PV systems in each block group during the respective year. In our data, there are 1,534 block groups in any given year. Thus, a $1/W decrease in system price translates into 127 additional installations demanded statewide in that year.

The other coefficients are of less interest to us, but we are reassured to see that the signs are generally consistent and make sense. One of the more interesting of these is the coefficient on the indicator for whether a block group has a Solarize campaign, which is positive and highly statistically significant in all specifications. This is consistent with results in **?**, which uses quasi-experimental and experimental approaches to find a large treatment effect of the Solarize program. The other statistically significant coefficients largely make sense in sign. Solar demand is higher in less densely populated areas and in areas with a younger population. Political views also appear to have a positive effect on the decision to adopt in the linear and truncated Poisson specifications. The general lack of statistical significance for the demographic and voting variables is likely attributable to the lack of sufficient time series variation in these variables.

6.2 Robustness Checks

We perform several robustness checks, which are both reassuring and provide insight into the variation behind our results. Table 7 shows the estimated elasticities from each of these robustness checks and includes the elasticity estimates from our preferred specifications in columns (3) and (4) of Table 6 for comparison.

First, we are concerned that our results in the logit regression may be driven by the method we use to fill in missing price data in block group-years where no contracts were signed. We therefore re-estimate the model using the highest, rather than average, recorded prices as proxies, preserving the order of our approach outlined in Section 3.1 (column I). We find that our logit CF estimates are hardly affected by this change in the interpolation approach, a very reassuring result.

As an alternative, we also estimate a hedonic equilibrium price equation for each year using the subsample with positive installation counts at the block group level and including county fixed effects. The regressors are a Solarize dummy and the set of demographic and voting variables from the baseline specification. We then use the estimated coefficients to predict prices in block groups with zero installations. We are careful in interpreting the estimates in this robustness check (shown in column II) due to the likely selection bias that we cannot control for in the hedonic equation. It is nonetheless reassuring to note that the elasticity we obtain of -0.78 is relatively close to our baseline result of -0.65.

Next, we explore the extent to which the changing market size for solar may affect our results. In particular, the pool of potential buyers continuously declines as more households adopt solar, resulting in a sequential truncation in the distribution of willingness to pay for the consumers remaining in the market.[19] To test the effects of this endogenous exit of adopters, we estimate an alternative specification in which we control for the market size. We scale the number of Census occupied housing units in each block group-year by the percentage of solar-viable homes in the area calculated from satellite data by GeoStellar, and then subtract the cumulative number of installations in the block group up to that year. This provides a time-varying measure of market size. While the results from this specification should be viewed with caution, as the market size regressor may not be strictly exogenous but only predetermined, they are quite supportive of the robustness of our main model. As shown in column III of Table 7, our elasticity estimate is almost identical to the baseline result. Furthermore, the market size variable is not statistically significant effect in either the logit or truncated Poission, with p-values of 0.23 and 0.91 respectively.

In addition, we test the sensitivity of our results to an alternative specification without

[19]We thank an anonymous referee for suggesting this robustness check.

demographic and voting variables among the regressors (column IV). Excluding these variables, we obtain almost identical elasticity estimates as in our baseline logit and truncated Poisson runs. This result is not surprising, given the lack of substantial temporal variation in these variables, as noted earlier.

Lastly, we examine the effect of excluding observations from Solarize campaigns from the analysis (column V). This robustness check also provides insight into the origins of our results. We find an elasticity estimate from the truncated Poisson GMM that is of smaller magnitude than the baseline value and is no longer significant, while the logit CF is roughly unchanged. The overall implied elasticity from the hurdle model is quite close to our baseline elasticity.

These results underscore the importance of the Solarize campaigns for our observations with multiple installations in a block group-year. Figure 3 shows a histogram of observations with at least one installation. Non-Solarize observations with positive outcomes are largely centered at low counts, while higher-count outcomes are mostly Solarize observations. Hence, when we drop all Solarize observations, our outcome variable is effectively reduced to a binary variable, enabling us to capture most of the variation in the data through the logit component of the hurdle model. This suggests that the hurdle model, as a mix of components that can exploit both binary and zero-truncated count data variation, offers a flexible approach for estimating solar demand in various settings: from emerging markets with few installations to booming higher-demand markets.

7 Policy Analysis

In this section, we highlight what our results imply for policies in the solar PV market in CT through a set of simple counterfactual simulations. We run three policy counterfactuals: a reduction in the state financial incentives down to Step 6 (recall Table 2), a reduction in state incentives to half of their average level in 2014 (well within the variation in the data), and a waiving of all permitting fees for solar PV systems. These simulations are particularly policy relevant. CT plans is in the process of phasing out financial incentives entirely and at the same time has made efforts to convince municipalities to reduce or entirely waive permitting fees (CEFIA 2013).

7.1 Simple Pass-through Analysis

As a first step, the above analysis necessitates obtaining a measure of the pass-through of installation costs to customers. While there are some estimates of pass-through in the California solar PV system market, the CT market may be quite different than the California market. Moreover, these estimates based on California data vary widely. For example, Henwood (2014) finds a low pass-through rate, while Dong et al. (2014) and Pless and van Benthem (2017) find nearly complete pass-through. These working papers use similar data, but very different empirical strategies.

Thus, we perform our own simple pass-through analysis for CT, following the standard approach for estimating pass-through of a subsidy or tax on consumer prices (see, e.g., Sallee 2011).[20] We observe the rebate received by each of the 5,070 purchased installations in 2008-2014. As noted earlier, the time lag between the submission and approval of contract applications gives rise to cross-sectional variation in rebate levels. We exploit this variation, along with the temporal variation due to declining incentive levels over time, in order to identify the effect of rebates on pre-rebate price, after controlling for unobserved region and time effects.

Using installation-level data, we estimate a linear regression of pre-rebate price (in \$/W) on the rebate level for the first 5 kW of installed capacity (in \$/W), a Solarize campaign indicator variable, municipality fixed effects, and year dummies. Our results are shown in

[20] Note that we do not account for the federal investment tax credit in our pass-through analysis, unlike Pless and van Benthem (2017), because in CT the tax credits are "state rebates" rather than "utility rebates," so they are not viewed as taxable income (see http://solaroutreach.org/wp-content/uploads/2015/03/ResidentialITC_Factsheet_Final.pdf).

Table 8.[21] The estimated coefficient on the rebate variable can be interpreted as the fraction of a \$1/W increase in the rebate level that is captured by the installer. A coefficient equal to zero would suggest complete pass-through: the rebate is entirely captured by consumers. A coefficient equal to one would suggest zero pass-through: the rebate is entirely captured by firms. We find a highly statistically significant coefficient on the rebate level of 0.16, implying a pass-through rate of 84 percent. In other words, a \$1/W decrease in installer costs translates into an 84 cents/W decrease in the PV system price.

7.2 Policy Simulations

For each policy scenario, we quantify the counterfactual number of installations in CT in 2014 and 2015, and compare these to the observed number of installations in 2014 and a projected number of installations in 2015. We limit our simulations to this two-year period in order to ensure that the demand structure and system characteristics are relatively stable– as in any new technology the distribution of willingness-to-pay is likely to change over the longer run. In this rapidly changing market, extrapolating too far out is not likely to be a useful exercise. These simulations have an important caveat: a model of firm pricing is outside the scope of this paper, so we do not model competition in the market. Thus, just as in any time when a demand elasticity is applied, the simulations assume that competition remains constant in this market with the relatively short-run price changes we examine.[22]

Our methodology is straightforward. The CGB database contains installer-reported data on the module, inverter, permitting, and labor costs for each installation.[23] Since the statutory incidence of the policies falls on the installing firms, we model the policy counterfactuals as a change in the installer's marginal cost. We find the percentage change in the sum of the module, inverter, permitting, and labor costs due to the policy, adjust this by a pass-through

[21] We also re-ran the estimation, aggregating all data at the block group-year level and replacing municipality fixed effects with block group fixed effects. We obtained an almost identical rebate coefficient.

[22] Firms in this market generally compete throughout the state, although they may sometimes focus on certain more narrow regions. We assume that the pass-through rate is constant across the state.

[23] We are not entirely confident in these reported cost estimates in the dataset, so we also use the module and inverter price indices instead of the reported module and inverter costs. We find very similar results.

rate to estimate the change in price, and then use our estimated elasticity (-0.65) to calculate the percent change in the number of installations. This is then converted into total number of new installations and additional installed capacity in MW. Further details of our approach used for the counterfactual simulations are in Appendix I.

The results of our simulations are shown in Table 9. In the baseline, state policies are assumed to remain unchanged and system prices are affected only by the falling global costs of PV components. Hence, projected installations and added capacity for 2015 are slightly higher than the observed numbers in 2014. In the "Small Reduction" counterfactual, the rebate offered by the state is reduced down to Step 6 in the rebate schedule in Table 2, lowering the average incentive amount in 2014 from the observed \$0.935/W to \$0.675/W. In the "Moderate Reduction" counterfactual, state incentives are cut in half, resulting in an average rebate of \$0.468/W. Finally, the "Permit Fees" counterfactual simulates a decrease in average PV system costs from waiving all municipal permit fees. This amounts to a reduction in average system costs by approximately \$0.052/W in 2014.

The "Small Reduction" simulation reveals a 4.8 percent decrease in the number of new PV installations in CT during 2014 relative to the number observed, and 4.9 percent in 2015. This would imply 95 fewer installations in 2014 and 97 fewer in 2015, and is equivalent to a reduction in added PV capacity of 0.74 MW in 2014 and 0.76 MW in 2015.[24] The "Moderate Reduction" simulation shows much larger impacts on new installations and capacity. The results suggest a decrease in adoption by 171 installations in 2014 and 175 installations in 2015 (1.33 MW in 2014 and 1.37 MW in 2015). If CT succeeded in eliminating municipality permit fees, the number of installations would have been increased by roughly 1 percent in both 2014 and 2015. This amounts to approximately 51 additional installations in each year (0.15 MW in 2014 and 0.16 MW in 2015). Of course, this may be an underestimate of the effect of eliminating permitting fees, since it does not account for the reduced paperwork

[24] The simulation results assume that the increased adoptions are not simply "pulled-forward" from future adoptions, which is reasonable since there is a large base of potential adopters and the subsidies continued through 2015.

and associated labor costs from expedited permitting, which likely would occur with a reduction or elimination of the fees. Such additional cost reductions may also be passed on to consumers, increasing the effect.

7.3 Cost-Effectiveness of Subsidies

Our above results can be used to provide insight into the cost-effectiveness and welfare implications of the subsidy policy in CT. We focus on direct program costs (dollar value of subsidies), as in Davis et al. (2014). A quick calculation reveals that the cost of the subsidy is \$3.03 of state dollars per additional watt installed (in 2014 dollars). Assuming all consumers take the full \$7,500 federal tax credit, this implies a total state and federal program cost of \$4.95 per watt.[25] To determine the benefits from a watt of installed solar capacity, we calculate the average annual electricity generation from a PV system in CT. We use the National Renewable Energy Laboratory's "PVWatts" solar electricity calculator for four different locations across CT and average the resultant output.[26] We assume a 25-year lifespan of an average installed system, which is in line with both the standard warranty offered by most manufacturers[27] and the assumptions made by the CGB in calculating lifetime electricity generation for an average system. This yields an average lifetime electricity generation of 32.26 kWh per watt of installed capacity, implying a short-run program cost of \$0.094/kWh generated (\$0.153/kWh when including the federal tax credit).

There are two sources of avoided pollution from installing solar PV systems: greenhouse gas emissions and local air pollutants. The amount of avoided pollution depends on assumptions about the type of generation that is displaced by solar PV both now and for the next 25 years. The average carbon dioxide emissions rate in the Northeast is estimated to be 0.000258 tons of CO_2 per kWh (Graff Zivin et al. 2014). Assuming this holds for the next

[25]Note that state policymakers find this to be irrelevant.

[26]See http://pvwatts.nrel.gov. The locations we use are New London, North Canaan, Putnam, and Stamford. The calculations account for reduction in output due to periods of no sunlight, as well as loss of generation during conversion from direct to alternating current power.

[27]See, e.g., http://energyinformative.org/solar-panel-warranty-comparison.

25 years, the lifetime cost-effectiveness would be \$364/tCO_2 (\$594/tCO_2 when including the federal credit). If we assume that the CT electricity grid will continue to become less carbon intensive over time, this estimate will be correspondingly higher. While there are numerous caveats to this simple calculation, including the social cost of public funds, the utility consumers receive, profits for installers, intermittency issues (Gowrisankaran et al. 2016), uncertainty in future electricity generation, and spillovers in learning-by-doing at the localized level (Bollinger and Gillingham 2016), it is notable that this estimate is much above the IAWG (2013) central value of the social cost of carbon of \$42/tCO_2 (in 2014 dollars).[28]

8 Conclusions

This study estimates the demand for solar PV systems using a new empirical approach: a Poisson hurdle model with fixed effects and instrumental variables. This approach allows us to tackle several key challenges that arise in modeling count data in the diffusion of any new technology. Specifically, it addresses unobserved heterogeneity at a fine geographic level, excess zeros in the outcome variable, and the endogeneity of price. In addition to the adoption of new technologies, we also expect this approach to be useful in a variety of other settings, such as the demand for health care in hospital units.

We estimate the price elasticity of demand for solar PV systems in CT over 2008-2014 to be -0.65. This estimate is valuable to both policymakers and firms. As module prices continue to drop, it provides useful guidance for forecasting the number of new installations, absent policy changes. It is also very useful for examining changes in policy, in light of continuing policy discussions about phasing out the financial incentive program and reducing municipal permit fees. After estimating a pass-through rate of 84 percent, we perform counterfactual policy simulations with less generous state incentives. We find that dropping incentives to Step 6 would have reduced the number of installations by 5 percent in 2014, while reducing incentives in half would have led to a 9 percent drop in installations in

[28]Including benefits from reduced criteria air pollutants does little to change this calculation.

2014. Simple calculations suggest that the direct program cost is \$364/tCO_2 (\$594/tCO_2 including the federal tax credit), significantly higher than the central estimate of the social cost of carbon used by the U.S. government. Other market failures, such as innovation market failures (van Benthem et al. 2008), must be significant for the policy to be social welfare-improving.

For firm decision-making, our finding of a price elasticity of -0.65 suggests that consumers are not very price sensitive. Firms can use this knowledge for short-run forecasting of the expected growth of the market. Moreover, in a market with imperfect competition, our estimated pass-through rate of 84 percent can provide important guidance to firms for optimal price-setting to maximize profits.

Our results also provide insight into how other policies influence the solar PV market. For example, we find a highly statistically significant effect of the Solarize grassroots campaigns on installations. Moreover, we find that when we exclude all Solarize installations from our analysis, our elasticity estimate remains the same, but the variation identifying it is almost exclusively from the logit specification. This indicates that without Solarize, the data can be treated as binary rather than count data, a finding due to the fact that the Solarize program accounts for nearly all installations in a block group-year after the first contract. This may not be surprising, due to evidence of neighbor effects in the diffusion of solar PV systems (Bollinger and Gillingham 2012; Graziano and Gillingham 2015).

Stepping back, this study of demand in the CT solar PV market highlights the intense policy effort to promote the technology, which mirrors efforts in many other U.S. states, European countries, and other countries around the world. As our approach is applicable to evaluating many solar PV incentive programs, we view further policy analyses in similar settings as a promising area for future research.

Acknowledgments. The authors would like to thank the Connecticut Green Bank for providing the data used in this analysis. We also thank Bryan Bollinger, CG Dong, Robert Mendelsohn, Corey Lang, Jim Sallee, Steve Sexton, Ted Juhl, Arthur van Benthem, and numerous seminar audiences for helpful discussions. Finally, we acknowledge funding from US Department of Energy award DE-EE0006128.

Figure 1: Trends in Installations and Average Post-Incentive System Price

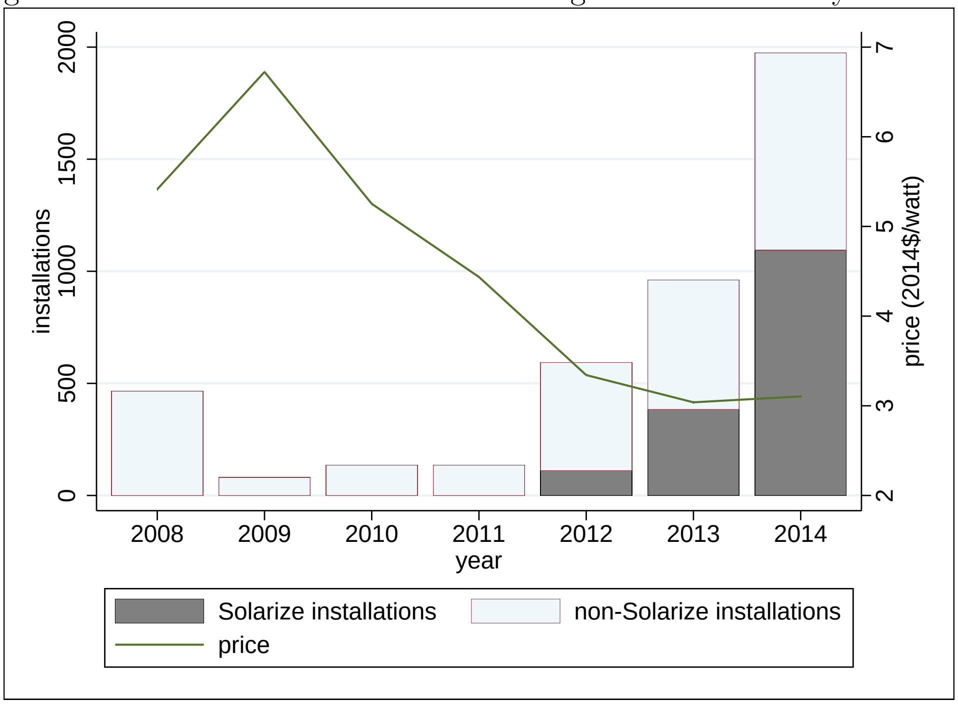

Figure 2: Histogram of the Count of Installations in the Full Sample

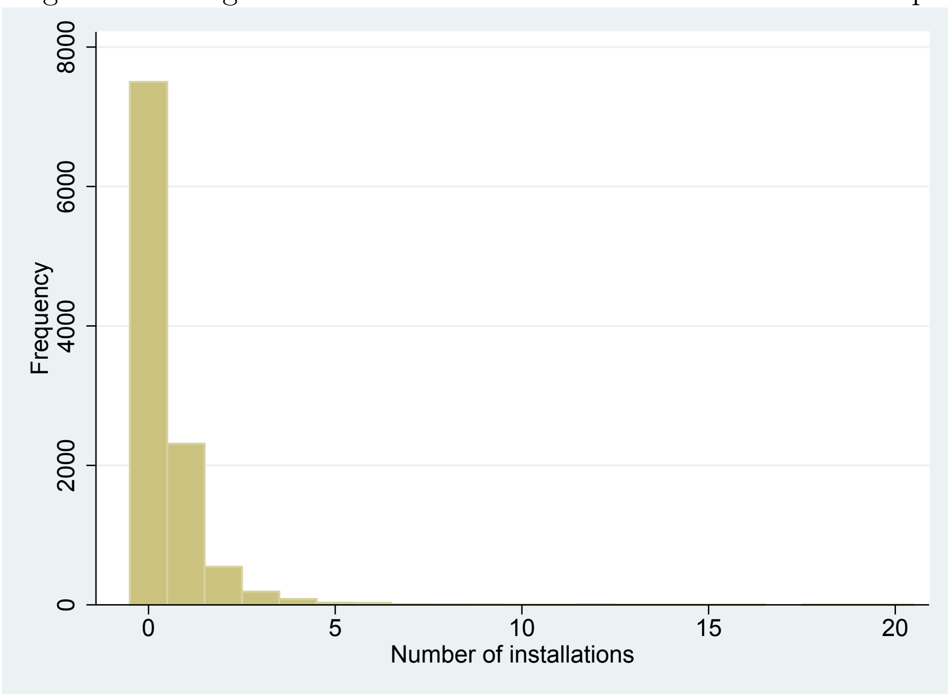

Figure 3: Histogram of the Count of Installations in the Truncated Subsample

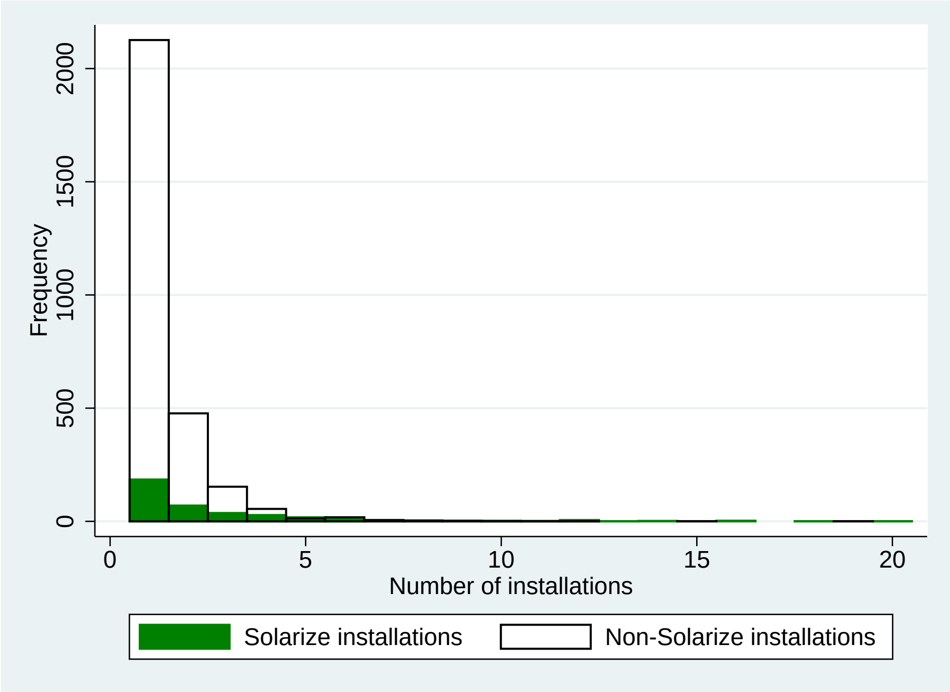

Table 1: Timeline of Pre-RSIP Solar Rebates

Rebate Structure	Program Opens	Change	Change	Closed	Reopened
Date	July 1, 2004	Jan 29, 2007	Oct 27, 2008	Nov 19, 2008	May 18, 2009
$\leq$ 5 kW	\$5.00/W	\$5.00/W	\$4.00/W	-	\$1.75/W
$>$ 5 kW and $\leq$ 10 kW	\$0.00/W	\$4.30/W	\$2.50/W	-	\$1.25/W

Source: The Connecticut Green Bank

Table 2: Residential Solar Investment Program Timeline

Incentive Type	Step 1	Step 2	Step 3	Step 4	Step 5	Step 6
EPBB/HOPBI						
Start Date	March 2, 2012	May 18, 2012	Jan 4, 2013	Jan 6, 2014	Sept 1, 2014	Jan 1, 2015
Incentive Design:						
$\leq$ 5 kW	\$2.450/W	\$2.275/W	\$1.750/W	\$1.250/W	\$0.800/W	\$0.675/W
$>$ 5 kW and $\leq$ 10 kW	\$1.250/W	\$1.075/W	\$0.550/W	\$0.750/W	\$0.800/W	\$0.675/W
$>$ 10 kW and $\leq$ 20 kW	-	-	-	-	\$0.400/W	\$0.400/W
PBI						
Start Date	March 2, 2012	May 18, 2012	April 1, 2013	Jan 6, 2014	Sept 1, 2014	Jan 1, 2015
Incentive Design:						
$\leq$ 10 kW	\$0.300/kWh	\$0.300/kWh	\$0.225/kWh	\$0.125/kWh	\$0.180/kWh	\$0.080/kWh
$>$ 10 kW and $\leq$ 20 kW	-	-	-	-	\$0.600/kWh	\$0.060/kWh

Source: The Connecticut Green Bank

Table 3: Summary Statistics for the Full Sample

Variable	Mean	St. Dev.	Min	Max
Number of PV installations	0.4781	1.0891	0	20
System capacity (kW)	6.972	1.8417	0.7	22.1
Post-incentive system price ($/W)	3.7645	1.4369	0.0171	20.5824
Solarize campaign	0.0595	0.2366	0	1
Population density (per km^2)	853.651	1249.742	0	21784
Median household income (in 1,000$/year)	88.7752	41.7775	2.499	250.001
Median age	42.6938	7.0762	14.8	80.6
% population above 25 with some college or college degree	46.7736	9.9282	0	81.203
% population above 25 with graduate or professional degree	18.3589	12.1701	0	80.7786
% Republican voters	21.9084	7.8536	3.69	51.14
% Democrat voters	34.0913	10.5664	16.92	72.3
Incentive level ($/W)	3.5288	2.0203	0.6824	6.4879
Roofing contractor wage ($/week)	1031.203	171.964	548.894	1312.658

Notes: All variables have 10,738 observations. All dollars in 2014 dollars.

Table 4: Summary Statistics for the Subsample with Positive Installations

Variable	Mean	St. Dev.	Min	Max
Number of PV installations	1.5855	1.4758	1	20
System capacity (kW)	7.2405	2.5195	0.7	22.1
Post-incentive system price ($/W)	3.8945	1.4433	0.0171	19.0565
Solarize campaign	0.1167	0.3212	0	1
Population density (per km^2)	607.418	943.207	0	14530.95
Median household income (in 1,000$/year)	91.8035	40.2957	10.833	250.001
Median age	43.3453	6.795	17.6	79.2
% population above 25 with some college or college degree	47.4923	9.4677	0	77.8978
% population above 25 with graduate or professional degree	19.2437	11.89	0	80.7786
% Republican voters	22.7355	7.1854	3.69	50.19
% Democrat voters	32.5461	9.2511	16.92	71.78
Incentive level ($/W)	3.0427	2.0825	0.6824	6.4879
Roofing contractor wage ($/week)	1010.557	175.821	548.894	1312.658

Notes: All variables have 3,238 observations. All dollars in 2014 dollars.

Table 5: Monte Carlo Simulation Results

Specification	Parameters				Elasticity		
	True	Estimated			True	Implied	
	Value	Mean	Bias	MSE	Value	Value	Bias
Logit CF	δ_1 = -0.1	-0.0961	0.0039	0.0003	η_1 = -0.0767	-0.0780	-0.0013
Tr. Poisson GMM	δ_2 = -0.1	-0.0999	0.0001	0.0001	η_2 = -0.1716	-0.1712	0.0004
Poisson hurdle	n/a	n/a	n/a	n/a	η = -0.2483	-0.2493	-0.0010
Linear 2SLS	n/a	-0.2257	n/a	n/a	η = -0.2483	-0.2654	-0.0171
Poisson CF	n/a	-0.1103	n/a	n/a	η = -0.2483	-0.2659	-0.0176

Notes: See Appendix E.2 for details about construction of the simulated data. Output is based on 5,000 replications. Sample means and parameter values, averaged over the 5,000 replications, are used to compute elasticity under each specification. Elasticity values in the hurdle model are calculated using the formulas in Section 5.4. The elasticity in the linear model is derived from $\hat{\eta} = \frac{\hat{\delta}\bar{w}}{\bar{y}}$. The elasticity in the Poisson model is derived from $\hat{\eta} = \hat{\delta}\bar{w}$.

Table 6: Primary Estimation Results

Variable	Linear	Poisson	Hurdle	
			Logit	Trun. Poisson
	2SLS[i]	CF[ii]	CF[ii]	GMM[i]
	(1)	(2)	(3)	(4)
Price	-0.079***	-0.286***	-0.201**	-0.292**
	(0.0292)	(0.0982)	(0.0999)	(0.132)
Solarize	0.929***	0.913***	0.863***	0.877***
	(0.242)	(0.15)	(0.204)	(0.161)
Pop. density	-5.6×10^{-6}	-0.0002***	-0.0002**	0.0004
	(9.4×10^{-6})	(0.00007)	(0.00009)	(0.0005)
Income	0.0008	0.0004	-0.0014	0.0013
	(0.00077)	(0.00154)	(0.00152)	(0.00267)
Age	-0.003	-0.012**	-0.0096*	-0.008
	(0.0026)	(0.0047)	(0.0053)	(0.012)
% (some) college	0.0003	0.0002	0.002	0.005
	(0.001)	(0.003)	(0.004)	(0.007)
% grad/prof degree	0.001	-0.0003	0.006	-0.005
	-0.002	-0.004	-0.005	(0.008)
% Republican	0.063***	0.077*	0.036	0.153***
	(0.024)	(0.040)	(0.041)	(0.059)
% Democrat	0.027**	0.041*	0.011	0.102***
	(0.0126)	(0.0222)	(0.0248)	(0.0384)
BG FE	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes
Instruments	yes	yes	yes	yes
Price elasticity[iii]	-0.621***	-1.076***	-0.528**	-0.123**
	(0.2301)	(0.3698)	(0.2626)	(0.0554)
Observations	10,738	10,738	10,738	3,238

Notes: Dependent variable is number of residential PV installations. Unit of observation is block group-year. BG FE refers to block group fixed effects. Year FE refers to year fixed effects. The instruments used for all IV specifications are the EPBB/HOPBI state financial incentives given to the installers and the county roofing contractor wage rate. $p < 0.1$ (*), $p < 0.05$ (**), $p < 0.01$ (***).

[i] Clustered standard errors at the town level in parentheses.

[ii] Block bootstrapped standard errors (100 replications), clustered at the town level, in parentheses.

[iii] Standard errors for the price elasticity obtained by the delta method.

Table 7: Elasticity Values under Different Specifications

Model Specification	Baseline	Robustness Checks				
		I	II	III	IV	V
Logit CF[i]	-0.528**	-0.478*	-0.661*	-0.530**	-0.521**	-0.543**
	(0.2626)	(0.2626)	(0.2498)	(0.2635)	(0.2698)	(0.2622)
Tr. Poisson GMM[ii]	-0.123**	-0.123**	-0.123**	-0.122**	-0.118**	-0.067**
	(0.0554)	(0.0554)	(0.0554)	(0.0553)	(0.0522)	(0.0426)
Combined elasticity[iii]	-0.651**	-0.601**	-0.784**	-0.652**	-0.639**	-0.61**
	(0.2684)	(0.2833)	(0.3542)	(0.2692)	(0.2748)	(0.2656)
BG FE	yes	yes	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes	yes	yes
Market size	no	no	no	yes	no	no
Demographics/voting	yes	yes	yes	yes	no	yes
Solarize included	yes	yes	yes	yes	yes	no
Missing price proxy	average price	highest price	hedonic price	average price	average price	average price

Notes: Dependent variable is number of residential PV installations. Unit of observation is block group-year. All other variables are the same as in Table 6. Standard errors are obtained by the delta method. $p < 0.1$ (*), $p < 0.05$ (**), $p < 0.01$ (***).

[i] Block bootstrapped standard errors (100 replications), clustered by town.

[ii] Robust standard errors, clustered by town.

[iii] Standard errors of combined elasticity coefficients obtained assuming independence of the data generating processes.

Table 8: Pass-through of Rebates

Variable	I
Rebate level	0.156***
	(0.045)
Solarize	-0.547***
	(0.057)
Town FE	yes
Year Dummies	yes
Observations	5,070

Notes: Dependent variable is pre-rebate price. Specification is a linear least squares regression. The sample includes only solar contracts with EPBB/HOPBI rebates. Each contract is a separate observation. Town FE refers to town fixed effects. Standard errors are clustered by town. $p < 0.1$ (*), $p < 0.05$ (**), $p < 0.01$ (***).

Table 9: Policy Simulations

	Baseline		Small Reduction		Moderate Reduction		Permit Fees	
Year	2014	2015	2014	2015	2014	2015	2014	2015
New installations	1974	1994	1879	1896	1803	1819	1993	2013
Added capacity (MW)	15.43	15.59	14.69	14.10	14.22	12.08	15.59	15.74
Δ installations			-95	-97	-171	-175	19	19
% change in installations			-4.81	-4.89	-8.66	-8.78	0.98	0.98
Δ capacity (MW)			-0.74	-0.76	-1.33	-1.37	0.15	0.16

Notes: Simulations use a price elasticity estimate of -0.65 from Table 6 and a pass-through rate of 84 percent, as estimated in Table 8. The "Small Reduction" counterfactual decreases state incentives down to the Step 6 level (see Table 2) in January 2014. The average rebate in this counterfactual is \$0.675/W. The "Moderate Reduction" counterfactual reduces the subsidy to half of its average 2014 level of \$0.935/W in January 2014. The "Permit fees" counterfactual entirely removes municipal permit fees for solar installations in January 2014, reducing average system costs by \$0.05/W. Further details on the counterfactual simulation methodology are in Appendix I.

Appendix A Consumer Expectations, Dynamics, and Subsidies

While several prominent recent papers on solar PV adoption have used reduced form approaches, rather than dynamic discrete choice models (see, e.g., Rogers and Sexton 2014; Hughes and Podolefsky 2015), many economists may instinctively consider the purchase of a solar PV system as a "buy-or-wait" decision that is best modeled with a dynamic discrete choice model. For example, consumers may recognize that subsidies are about to change in the near future and time their purchase to ensure the higher subsidy. Indeed, in California, Rogers and Sexton (2014) finds intriguing evidence of such dynamics in consumer behavior, with a considerable "bunching" of adoptions just before a step decline in the subsidy. Papers such as Hendel and Nevo (2006) point out that static demand elasticities are over-estimated in the context of temporary sales or price reductions that lead to large increases in the quantity sold due to consumer recognition of the temporary nature of the price increases. Burr (2014) and Bollinger and Gillingham (2016) implement structural dynamic discrete choice models of solar PV demand to attempt to model such features of consumer decision making California.

However, it is quite likely that the California setting may not transfer to many other settings around the country and around the world. In California the subsidies were phased out in a way that depended on the total amount of installed PV capacity, which allows consumers and firms to reasonably anticipate the timing of subsidy declines. Moreover, the subsidy changes were large and firms occasionally advertised using the upcoming subsidy change as the key message. Such features do not exist in Connecticut (or most other states).

The changes in rebates in Connecticut were abrupt and not easily forecastable. In some cases, they changed with votes in the state legislature. In others, they changed due to the budget for that set of incentives running out. These end dates are very plausibly random. Equally importantly, unlike in California, consumers and firms could not precisely time purchases, since consumers and firms had little information about when the change would occur. In Connecticut, there was no messaging or advertising about future subsidy declines. The only case where timing of any sort occurred is when CGB explicitly timed two rounds of Solarize programs to end just before a change in the incentive (others ended at a different time).

This appendix provides two types of evidence supporting the contention that dynamics are a dominant force in this context. First, we show that there is no obvious bunching in the Connecticut data that would come about due to consumers treating the solar PV adoption decision as an optimal stopping problem. Second, we provide evidence from a survey of consumers that highlights the key factors in the solar adoption decision process and suggests that expectations of future changes in subsidies–or other dynamic factors–are not likely to be key influences in the decision process for most solar PV adopters in Connecticut.

A.1 No obvious bunching in Connecticut

Figure A1 shows the number of installations per month (the most disaggregated our data permits with any accuracy). The red lines in the figure indicate the dates of changes in incentives. After early 2012, the figure shows several months of very large numbers of installations, which correspond to months at the end of Solarize programs. Sometimes they correspond to steps that decline, but more often they do not, with months of high sales before, during, and after changes in steps. In general, there is no clear pattern of bunching before changes in incentives in Connecticut. The high levels of installations in the Solarize months are due to the Solarize programs, rather than the changing incentives (there is no uptick in non-Solarize municipalities). These high levels of installations during those months underscore the value of performing our analysis both with and without the Solarize programs included.

Examining the data in this way is useful for it demonstrates that there is no obvious bunching or "harvesting" effect (Rogers and Sexton 2014) in our data, which could confound identification of the demand elasticity. Furthermore, by aggregating our data to the yearly level, as we do for our estimation, we can be even more confident that harvesting is not an identification concern in our setting.

Figure A1: The number of installations over time does not show bunching at the steps of incentive declines (red lines) except for during times of Solarize campaigns.

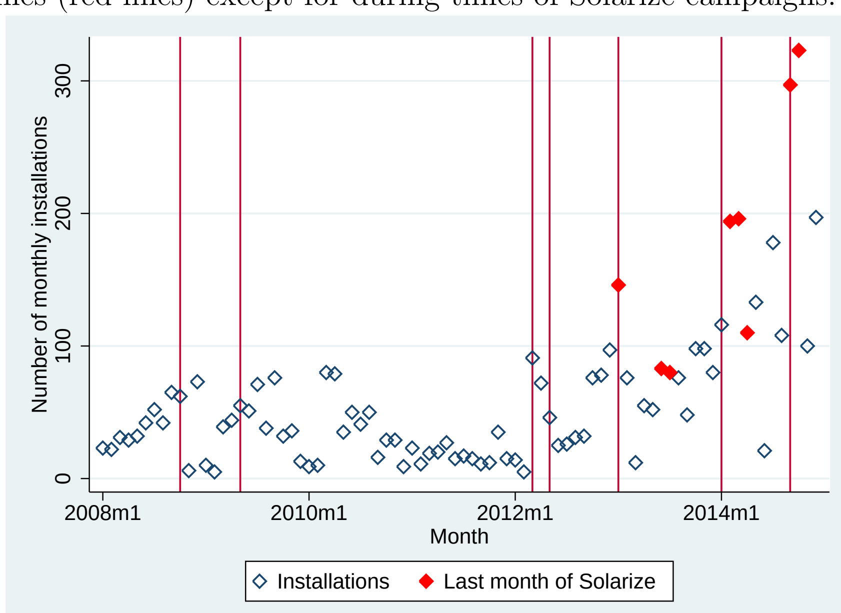

A.2 Consumer expectations

Our evidence of a lack of bunching indicates that consumers in Connecticut are not obviously making decisions considering the expectations of future subsidy declines. We can also examine survey evidence to better understand how consumers are making decisions.

We use survey data that was conducted in Connecticut just shortly after each of the first three Solarize rounds. The surveys were conducted online in March 2013, September 2013, and June 2014 using the Qualtrics survey tool. Respondents were contacted via e-mail. The e-mail addresses were obtained from the Solarize campaigns. The response rate for those who had solar or signed a contract to install was approximately 40 percent. The response rate for those who expressed interest, but did not adopt, was approximately 17 percent. We received 1,392 full responses, 36 percent of whom either had installed or had signed a contract to install solar PV.

The responses provided deep insight into how consumers make decisions about solar PV. A few survey questions are most relevant for the question at hand here: do consumers in Connecticut appear to time their decision to adopt? While this evidence is certainly only suggestive, the short answer is no, it appears that timing is not a dominant factor in the decision process.

Our first evidence of this is in an open-ended qualitative question at the end of the survey where respondents were asked to provide any further thoughts about solar PV. Typical responses from those who did not install a panel are "My house faces in wrong direction and my age (79) would not see return on investment" or "Had my home evaluated - roof wasn't large enough, so the only other option was a large array structure in my backyard" or "I just can't justify cutting beautiful, old trees down to become more environmentally friendly" or "took too long to get back investment in solar." Typical responses from those who did install are "LOVE IT!" or "Great program! Keep up with incentives so that more people can access solar power" or "Communication by the town is critical in getting the word out about the Solarize program." Notably, not a single response mentions anything about timing the installation in any way (e.g., before incentives dropped). Some responses mentioned that they are interested but the numbers do not work out for solar to be financially attractive on their home right now and that they would consider solar PV in the future if the cost comes down. But this is very different than households anticipating the expiration of subsidies and timing their installations. If such timing was a critical factor in this market, it stands to reason that at least one consumer would have mentioned this in their comments.

Further evidence is from a question that asks consumers who did not install to rate the importance of different factors in their decision not to install. Respondents could rate each of the factors on a scale from "not at all important" to "extremely important." The two factors that rose to the top: whether their home is suitable for solar and the current cost of solar. 46 percent of the respondents indicated that whether their home is suitable for solar was extremely important for their choice not to install solar PV. 41 percent of the respondents indicated that the current cost of solar is extremely important for their choice not to install solar. In contrast, only 23 percent of the respondents indicated that future costs of solar were extremely important. 100 percent of those who indicated that future costs of solar PV were extremely important also indicated that the current costs of solar were extremely important in their decision. This evidence suggests that at most only a relatively small group of potential consumers did not install due to expectations of lower future prices, and that these consumers also simply saw current prices as a dominant factor.

While this survey evidence is by no means definitive, when considered along with the evidence of a lack of bunching just before subsidy declines and the fact that there was no easy way for consumers to know that the subsidies were about decline, it builds a reasonable case that our estimated static demand elasticities are not substantially biased from neglecting dynamics.

Appendix B Consistency of the IV Fixed Effects Truncated Poisson Estimator

$\boldsymbol{Y}_i$ is a vector of outcome variables, such that $\boldsymbol{Y}_i \in \mathcal{Y} \subset \mathbb{R}_+^T$ and each element Y_{it} is independently distributed with truncated Poisson distribution. Let $\boldsymbol{X}_{it} \in \mathcal{X} \subset \mathbb{R}^P$ and $\boldsymbol{\delta}_2 \in \Delta \subset \mathbb{R}^P$. Let $g_0(\boldsymbol{y}|\boldsymbol{x}, n)$ denote the true conditional distribution of $\boldsymbol{Y}_i$ given $\boldsymbol{X}_i$ and $\sum_{t=1}^{T} Y_{it}$. If our model is correctly specified, then, for some $\boldsymbol{\delta}_2^0 \in \Delta \subset \mathbb{R}^P$, $g(\cdot\,|\boldsymbol{x}, n, \boldsymbol{\delta}_2^0) = g_0(\cdot\,|\boldsymbol{x}, n)$, for all $\boldsymbol{x} \in \mathcal{X}$ and $n \in \mathcal{N}$. We now show that the conditional maximum likelihood estimator, obtained from the maximization of (17), is a consistent estimator of $\boldsymbol{\delta}_2^0$.

Lemma B.1. *Let* $Q(\boldsymbol{\delta}_2) = \mathrm{E}[\log\ g(\boldsymbol{Y}_i|\boldsymbol{X}_i, n_i, \boldsymbol{\delta}_2)]$ *and* $Q_N(\boldsymbol{\delta}_2) = \frac{1}{N}\sum_{i=1}^{N} \log\ g(\boldsymbol{Y}_i|\boldsymbol{X}_i, n_i, \boldsymbol{\delta}_2)$. *Under the following assumptions:*
A1. Δ *is a compact set*
A2. *For each* $(\boldsymbol{y}, \boldsymbol{x}, n) \in \mathcal{Y} \times \mathcal{X} \times \mathcal{N}$, $\log\ g(\boldsymbol{y}|\boldsymbol{x}, n, \cdot\,)$ *is a continuous function on* Δ
A3. *For all* $\boldsymbol{x} \in \mathcal{X}$ *and* $n \in \mathcal{N}$, $Q(\boldsymbol{\delta}_2) \neq Q(\boldsymbol{\delta}_2^0)$ *if* $\boldsymbol{\delta}_2 \neq \boldsymbol{\delta}_2^0$
A4. *For all* $\boldsymbol{x} \in \mathcal{X}$, $n \in \mathcal{N}$, *and* $\boldsymbol{\delta}_2 \in \Delta$, $\sum_{\boldsymbol{y}\in\mathcal{Y}} g(\boldsymbol{y}|\boldsymbol{x}, n, \boldsymbol{\delta}_2) = 1$
A5. *The Uniform Weak Law of Large Numbers holds*
A6. *The sequence* $\{c_i\}_{i=1}^N$ *is bounded, i.e.,* $\lim_{N\to\infty} c_N < \infty$,

$\hat{\boldsymbol{\delta}}_{2,CMLE} \xrightarrow{p} \boldsymbol{\delta}_2^0$, *where* $\hat{\boldsymbol{\delta}}_{2,CMLE} = \arg\max_{\boldsymbol{\delta}_2\in\Delta} Q_N(\boldsymbol{\delta}_2)$.

Proof. From A1 and A2, it follows that the problem $\max_{\boldsymbol{\delta}_2\in\Delta} Q(\boldsymbol{\delta}_2)$ always has a solution. A3 implies that this solution is unique. Using A4 and Jensen's inequality, it can easily be shown that $Q(\boldsymbol{\delta}_2)$ is maximized at $\boldsymbol{\delta}_2^0$. From A1 and A2, we also know that the problem $\max_{\boldsymbol{\delta}_2\in\Delta} Q_N(\boldsymbol{\delta}_2)$ always has a solution. Let $\hat{\boldsymbol{\delta}}_{2,CMLE} = \arg\max_{\boldsymbol{\delta}_2\in\Delta} Q_N(\boldsymbol{\delta}_2)$. Under A5, $|Q_N(\hat{\boldsymbol{\delta}}_{2,CMLE}) - Q(\boldsymbol{\delta}_2^0)| \xrightarrow{p} 0$, by the Uniform Weak Law of Large Numbers. Therefore, $\hat{\boldsymbol{\delta}}_{2,CMLE} \xrightarrow{p} \boldsymbol{\delta}_2^0$. Note that A6 is needed to ensure that the sequence of unknown parameters that we do not estimate, $c_1, ..., c_N$, does not contain a high number of extremely large values, which could result in inconsistency of $\hat{\boldsymbol{\delta}}_{2,CMLE}$. □

However, if $\boldsymbol{X}$ is not exogenous, CMLE is no longer consistent. We next show that, in such a setting, the GMM estimator derived from (23) is a consistent estimator of $\boldsymbol{\delta}_2^0$.

Lemma B.2. *Suppose* $\boldsymbol{Z}_{it} \in \mathcal{Z} \subset \mathbb{R}^Q$, *where* $Q > P$. *Let* $\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2) = \boldsymbol{Z}_i^{'}\boldsymbol{\xi}_i$ *and* $E[\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2^0)] = \mathbf{0}$. *Furthermore, let* $G(\boldsymbol{\delta}_2) = [\mathrm{E}[\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2)]^{'}\ \boldsymbol{\Xi}\ [\mathrm{E}[\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2)]$ *and* $G_N(\boldsymbol{\delta}_2) = \left[\frac{1}{N}\sum_{i=1}^{N}[\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2)\right]^{'}$

$\hat{\boldsymbol{\Xi}}\left[\frac{1}{N}\sum_{i=1}^{N}[\psi(\boldsymbol{Z}_i, \boldsymbol{\delta}_2)\right]$, *where* $\hat{\boldsymbol{\Xi}}$ *is a symmetric positive semidefinite weight matrix and* $\boldsymbol{\Xi}$ *is a symmetric and positive definite matrix. Under the following assumptions:*

A1. Δ *is a compact set*

A2. *For each* $\boldsymbol{z} \in \mathcal{Z}$, $\psi(\boldsymbol{z}, \cdot)$ *is a continuous function on* Δ

A3. *For all* $\boldsymbol{z} \in \mathcal{Z}$, $\psi(\boldsymbol{z}, \boldsymbol{\delta}_2) \neq \psi(\boldsymbol{z}, \boldsymbol{\delta}_2^0)$ *if* $\boldsymbol{\delta}_2 \neq \boldsymbol{\delta}_2^0$

A4. $\hat{\boldsymbol{\Xi}} \xrightarrow{p} \boldsymbol{\Xi}$

A5. *The Uniform Weak Law of Large Numbers holds,*

$$\hat{\boldsymbol{\delta}}_{2,GMM} \xrightarrow{p} \boldsymbol{\delta}_2^0, \text{ where } \hat{\boldsymbol{\delta}}_{2,GMM} = \arg\min_{\boldsymbol{\delta}_2 \in \Delta} G_N(\boldsymbol{\delta}_2).$$

Proof. A1 and A2 imply that the problem $\min_{\boldsymbol{\delta}_2 \in \Delta} G_N(\boldsymbol{\delta}_2)$ always has a solution. Let $\hat{\boldsymbol{\delta}}_{2,GMM} = \arg\min_{\boldsymbol{\delta}_2 \in \Delta} G_N(\boldsymbol{\delta}_2)$. By definition, we know that $\boldsymbol{\delta}_2^0 \in \Delta$ solves the problem $\min_{\boldsymbol{\delta}_2 \in \Delta} G(\boldsymbol{\delta}_2)$. By A3 and because $\boldsymbol{\Xi}$ is a positive definite matrix, $\boldsymbol{\delta}_2^0$ is a unique solution to $\min_{\boldsymbol{\delta}_2 \in \Delta} G(\boldsymbol{\delta}_2)$. Then, using A4 and A5, $|G_N(\hat{\boldsymbol{\delta}}_{2,GMM}) - G(\boldsymbol{\delta}_2^0)| \xrightarrow{p} 0$. Therefore, $\hat{\boldsymbol{\delta}}_{2,GMM} \xrightarrow{p} \boldsymbol{\delta}_2^0$. □

Appendix C Equivalence of CMLE and GMM Estimators under Strict Exogeneity

In what follows, we review the equivalence of CMLE and GMM estimators that only holds under strict exogeneity of the vector of regressors. Suppose $\boldsymbol{X}_{it} \in \mathcal{X} \subset \mathbb{R}^P$. Then, starting from (17), we can express the P-dimensional score vector of derivatives of the log-likelihood corresponding to block group i as:

$$\boldsymbol{S}_i(\boldsymbol{\delta}_2) = \nabla_{\boldsymbol{\delta}_2} \ell_i^t(\boldsymbol{\delta}_2 | \sum_{t=1}^{T} Y_{it}) = \sum_{t=1}^{T} \left[Y_{it} - \phi_{it}(\boldsymbol{X}_i, \boldsymbol{\delta}_2) \right] \boldsymbol{X}_{it},$$

where $\phi_{it} \equiv \frac{\frac{\partial h_i}{\partial \beta_{it}} \beta_{it}}{h_i}$. Assuming that the model has been correctly specified, $\sum_{\boldsymbol{y} \in \mathcal{Y}} g(\boldsymbol{y}|\boldsymbol{x}, n, \boldsymbol{\delta}_2) = 1$ for all $\boldsymbol{x}$, n, and $\boldsymbol{\delta}_2$, and it can be shown that the score of the log-likelihood function, evaluated at the true parameter vector $\boldsymbol{\delta}_2^0$, has a zero conditional mean. Let $S_i(\boldsymbol{\delta}_2^0)^{(p)}$ denote the p-th element of the score vector, corresponding to $\frac{\partial}{\partial \delta_2^p} \ell_i(\boldsymbol{\delta}_2) \Big|_{\boldsymbol{\delta}_2 = \boldsymbol{\delta}_2^0}$. Then,

$$\mathrm{E}[S_i(\boldsymbol{\delta}_2^0)^{(p)} | \boldsymbol{X}_i, n_i] = \sum_{\boldsymbol{y} \in \mathcal{Y}} \frac{\partial}{\partial \delta_2^p} \ell_i(\boldsymbol{\delta}_2) \Big|_{\boldsymbol{\delta}_2 = \boldsymbol{\delta}_2^0} g(\boldsymbol{y}|\boldsymbol{X}_i, n_i, \boldsymbol{\delta}_2^0) =$$

$$\sum_{\boldsymbol{y} \in \mathcal{Y}} \frac{\partial}{\partial \delta_2^p} g(\boldsymbol{y}|\boldsymbol{X}_i, n_i, \boldsymbol{\delta}_2) \Big|_{\boldsymbol{\delta}_2 = \boldsymbol{\delta}_2^0} = \frac{\partial}{\partial \delta_2^p} \left(\sum_{\boldsymbol{y} \in \mathcal{Y}} g(\boldsymbol{y}|\boldsymbol{X}_i, n_i, \boldsymbol{\delta}_2) \right) \Bigg|_{\boldsymbol{\delta}_2 = \boldsymbol{\delta}_2^0} = 0.$$

Since the above is true for any element of the score, it follows that

$$\mathrm{E}[\boldsymbol{S}_i(\boldsymbol{\delta}_2^0) | \boldsymbol{X}_i, n_i] = \boldsymbol{0}. \tag{C.1}$$

Note that the log-likelihood function is derived after conditioning on the vector of regressors. Therefore, this result would not hold if one or more of the variables in $\boldsymbol{X}$ are endogenous in the model.

Let $\xi_{it} \equiv Y_{it} - \phi_{it}(\boldsymbol{X}_i, \boldsymbol{\delta}_2^0)$, and let $\boldsymbol{\xi}_i = (\xi_{i1}, ..., \xi_{iT})'$. By the Law of Iterated Expectations, (C.1) implies that

$$\mathrm{E}[\boldsymbol{X}_i' \boldsymbol{\xi}_i] = \boldsymbol{0},$$

i.e., the expected value of the score leads to the same moment condition as the orthogonality of the regressors to the error term in (21) under strict exogeneity of $\boldsymbol{X}$. Thus, with no endogenous regressors, the sample analog of $\mathrm{E}[\boldsymbol{X}_i' \boldsymbol{\xi}_i] = \boldsymbol{0}$ is identical to the first-order conditions of the conditional likelihood from (17) and yields a GMM estimator that is equivalent to $\hat{\boldsymbol{\delta}}_{2,CMLE}$.

Appendix D Monotonicity of $m(\lambda)$

In order to ensure that $m^{-1}(\cdot)$ is a one-to-one function, we need to show that $m(\lambda)$ is monotonic over the relevant range of λ values. In what follows, we prove that, as long as λ is positive, the function $m(\lambda)$ is strictly increasing.

Lemma D.1. *For any $\lambda > 0$, $m^{'}(\lambda) > 0$.*

Proof. Dropping all subscripts for simplicity, $m^{'}(\lambda) = \left[\frac{1}{\lambda} - \frac{e^{-\lambda}}{1-e^{-\lambda}}\right] m(\lambda) = \left[\frac{e^{\lambda}-1-\lambda}{\lambda(e^{\lambda}-1)}\right] m(\lambda)$. Note that, by the properties of the truncated Poisson model, $m(\lambda) = E_t(Y) > 0$ for all λ. Furthermore, $\lambda > 0$, which implies that $\lambda(e^{\lambda} - 1) > 0$. Hence, we need to ensure that $e^{\lambda} - 1 - \lambda$ is either positive or negative over the relevant range of λ.

Let $h_1(\lambda) = e^{\lambda}$ and $h_2(\lambda) = \lambda + 1$ and note that $h_1(0) = h_2(0)$. Also note that $h_1^{'}(0) = h_2^{'}(0) = 1$, while, for any $\lambda > 0$, $h_1^{'}(\lambda) > 1 = h_2^{'}(\lambda)$. Hence, $h_1(\lambda) > h_2(\lambda)$ for all $\lambda > 0$, implying that $m^{'}(\lambda) > 0$ for all $\lambda > 0$, which is the relevant range of λ in a truncated Poisson model.

□

Appendix E Monte Carlo Simulations

In what follows, we describe three different data generation processes used in our Monte Carlo simulations. In particular, we generate panel data from a random Poisson model, a hurdle model with individual fixed effects and an endogenous regressor, and a Poisson model with individual fixed effects and an endogenous regressor.

E.1 Random Poisson Model

We use the following model in this set of Monte Carlo simulations:

$$y_{it} \sim Po(\lambda_{it}), \tag{E.1}$$

$$\lambda_{it} = \exp(\alpha_{it} + \delta x_{it}), \tag{E.2}$$

$$x_{it} = 2 + u_{it}, \tag{E.3}$$

$$\alpha_{it} = -5 + e_{it}, \tag{E.4}$$

$$u_{it} \perp\!\!\!\perp e_{it},\ u_{it} \sim Unif(0,2),\ e_{it} \sim Unif(0,10), \tag{E.5}$$

where $i = 1, ..., N$, and $t = 1, ..., T$. In other words, the Poisson parameter λ_{it} is affected by an individual- and time-specific shock α_{it} which is unobserved, but also uncorrelated with the observed regressor x_{it}.

To generate data, we first draw u_{it} and e_{it} from the distributions shown in (E.5). We then construct the variables α_{it} and x_{it} as shown in (E.4) and (E.3), respectively. Next, we calculate λ_{it} using the formula in (E.2). Finally, we draw y_{it} from the Poisson distribution shown in (E.1). We set the parameter values as δ = -1, N = 5000, and T = 4, and replicate the estimation 5,000 times, each time re-drawing different u_{it} and e_{it}.

Such data generation process can give rise to a distribution of the outcome count variable y with most of its mass centered around zero and variance vastly exceeding the sample mean. Figure E1 displays the distribution of the data from one of the random draws. The thick line traces the true distribution of the outcome variable under the random Poisson data generation process. Note that more than 70 percent of the outcomes are zeros. Furthermore, the sample mean for the data is 0.86 and the variance 4.85, indicating overdispersion relative to the standard Poisson model. Hence, in general, a random Poisson model can yield a distribution similar to the data in our analysis. Figure E1 also shows how a hurdle model distribution would fit through the generated data. Aside from a slightly lower mass at low counts, the hurdle model comes quite close to the actual distribution. We found this trend to hold under a number of different parameter values.

Figure E1: Data distribution under a random Poisson specification

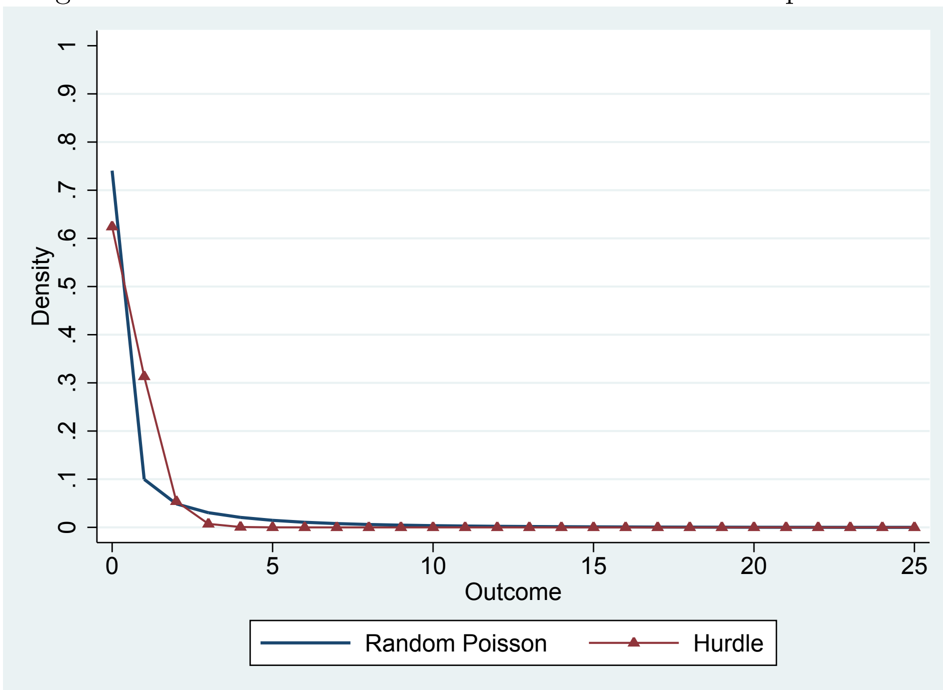

In each replication of the above data generation process, we estimate the logit and truncated Poisson components of a hurdle model with a single exogenous regressor x_{it}. Using the estimated coefficient on the regressor and the formulas in Section 5.4, we then derive the implied elasticity from each component of the model, as well as the combined hurdle model elasticity. As shown in Table E1, the average elasticity estimate from the 5,000 replications is very close to the true value (also averaged over 5,000 replications) of -3. The bias is less than 0.6 percent, which lends support to the use of a Poisson hurdle model even in instances where the true model is a random Poisson.

Table E1: Random Poisson Simulation Results

Specification	Elasticity		
	True Value	Implied Estimated	Implied Bias
Logit	n/a	-1.1727	n/a
Truncated Poisson	n/a	-1.8091	n/a
Poisson hurdle	-3	-2.9818	0.0182

E.2 Hurdle Model with Fixed Effects and Endogeneity

We use the following model:

$$\iota_{it} = \begin{cases} 1 & \text{if } l_{it} > \kappa_{it}, \\ 0 & \text{if } l_{it} \leq \kappa_{it}, \end{cases} \tag{E.6}$$

$$l_{it} = \frac{\exp(\alpha_i + \delta_1 w_{it} + u_{it})}{1 + \exp(\alpha_i + \delta_1 w_{it} + u_{it})}, \tag{E.7}$$

$$\kappa_{it} \sim Unif[0,1], \tag{E.8}$$

$$y_{it} = 0, \text{ if } \iota_{it} = 0, \tag{E.9}$$

$$y_{it} \sim trPo(\lambda_{it}), \text{ if } \iota_{it} = 1, \tag{E.10}$$

$$\lambda_{it} = \exp(\beta_i + \delta_2 w_{it} + v_{it}), \tag{E.11}$$

$$w_{it} = \tau_1 \alpha_i + \tau_2 \beta_i + \pi z_{it} + \rho_1 u_{it} + \rho_2 v_{it} + e_{it}, \tag{E.12}$$

$$u_{it} \perp\!\!\!\perp v_{it},\ (u_{it}, v_{it}, e_{it}) \perp\!\!\!\perp z_{it},$$

$$u_{it} \sim N(0, \sigma_u),\ v_{it} \sim N(0, \sigma_v),\ e_{it} \sim N(0, \sigma_e), \tag{E.13}$$

where $i = 1, ..., N$, and $t = 1, ..., T$. Note that w_{it} is correlated both with the error terms u_{it} and v_{it}, as well as with the individual-specific parameters α_i and β_i, with the magnitude of correlation determined by parameters ρ_k and τ_k, $k \in \{1, 2\}$, respectively. The variable z_{it} is exogenous to the model and is used to instrument for w_{it}. The parameter π measures the strength of this instrument.

The data generation process is as follows. First, we draw each α_i and β_i from a uniform distribution with support $[0, 1]$. Next, we generate a random variable $z_{it} \sim Unif[0, 5]$. After drawing the random error terms u_{it}, v_{it}, and e_{it} from the distributions shown in (E.13), we plug them into (E.12), (E.11), and (E.7) to obtain w_{it}, λ_{it}, and l_{it}, respectively. Finally, we proceed to generate the outcome variable in each of the two stages of the hurdle model. We draw κ_{it} from the distribution shown in (E.8) and then follow the decision rule in (E.6) to generate ι_{it}. We set y_{it} to equal zero for all observations where $\iota_{it} = 0$, as indicated in (E.9). Then, if $\iota_{it} = 1$, y_{it} is drawn from the truncated Poisson distribution in (E.10). To obtain the truncated Poisson draws, we follow Cameron and Trivedi (2013) by drawing from a Poisson distribution with parameter λ_{it} and then replacing each zero draw with another draw from the same distribution until all draws are positive.

The parameters to be chosen are δ_1, δ_2, τ_1, τ_2, π, ρ_1, ρ_2, σ_u, σ_v, σ_e, N, and T. We set $\delta_1 = \delta_2 =$ -0.1, $\tau_1 = \tau_2 = 0.2$, $\pi = 0.8$, $\rho_1 = \rho_2 = 0.5$, $\sigma_u = \sigma_v = \sigma_e = 0.5$, $N = 5000$, and $T = 4$. We replicate the estimation 5,000 times, each time re-drawing different u_{it}, v_{it}, and e_{it}.

E.3 Poisson Model with Fixed Effects and Endogeneity

This Monte Carlo simulation shows that even if the true data generating process was simply Poisson (with endogeneity), our hurdle model performs well in recovering the true values. We use the following model:

$$y_{it} \sim Po(\lambda_{it}), \tag{E.14}$$

$$\lambda_{it} = \exp(\alpha_i + \delta w_{it} + u_{it}), \tag{E.15}$$

$$w_{it} = \tau\alpha_i + \pi z_{it} + \rho u_{it} + v_{it}, \tag{E.16}$$

$$u_{it} \perp\!\!\!\perp v_{it},\ (u_{it}, v_{it}) \perp\!\!\!\perp z_{it},$$

$$u_{it} \sim N(0, \sigma_u),\ v_{it} \sim N(0, \sigma_v), \tag{E.17}$$

where $i = 1, ..., N$, and $t = 1, ..., T$. Note that w_{it} is correlated with the error term u_{it} and with the individual-specific parameter α_i, with the magnitude of correlation determined by parameters ρ and τ, respectively. The variable z_{it} is exogenous to the model and is used to instrument for w_{it}. The parameter π measures the strength of this instrument.

The data generation process is as follows. First, we draw each α_i from a uniform distribution with support $[-1, 1]$. Next, we generate a random variable $z_{it} \sim Unif[0, 5]$. After drawing u_{it} and v_{it} from the distributions shown in (E.17), we plug them into (E.16) and (E.15) to obtain w_{it} and λ_{it}, respectively. Finally, we draw y_{it} from the Poisson distribution shown in (E.14). The parameters to be chosen are δ, τ, π, ρ, σ_u, σ_v, N, and T. We set δ = -0.5, τ = 0.2, π = 0.8, ρ = 0.5, $\sigma_u = \sigma_v = 0.5$, $N = 5000$, and $T = 4$. We replicate the estimation 5,000 times, each time re-drawing different u_{it} and v_{it}.

In each replication, we estimate a fixed effects logit CF and a fixed effects truncated Poisson GMM model with the data, using z_{it} as an instrument for the endogenous variable w_{it}. We then calculate the corresponding elasticity values in each of the two components of the hurdle model and add them up to obtain the total elasticity implied by the model. The true elasticity from the Poisson model, averaged over 5,000 replications, is approximately -1. As shown in Table E2, the average elasticity estimate from the hurdle model lies very close to the true value, with bias of less than 0.7 percent.

Table E2: FE Poisson with Endogeneity Simulation Results

Specification	Elasticity		
	True	Implied	
	Value	Estimated	Bias
Logit CF	n/a	-0.7776	n/a
Tr. Poisson GMM	n/a	-0.2293	n/a
Poisson hurdle	-1.0001	-1.0069	-0.0068

Appendix F First-stage Instrumental Variable Results

To demonstrate the strength of our instruments, Table F1 shows the results of a first-stage regression of the post-incentive PV system price per W on all instruments for both the full and truncated sample. The coefficients on incentives are highly statistically significant. A joint F-test of statistical significance of the excluded instruments provides a test statistic of 43.06 in the full sample and 117.52 in the truncated sample.

Table F1: First-stage Instrumental Variable Regression Results

Variable	Full Sample	Truncated Sample
Incentive level	-0.453***	-0.574***
	(0.0515)	(0.038)
Wage rates	0.00036**	0.00008
	(0.00017)	(0.00045)
Solarize	-0.134	-0.27***
	(0.0838)	(0.0805)
Pop. density	7×10^{-6}	0.00016
	(3×10^{-5})	(0.00014)
Income	-0.0007	-0.00179
	(0.001)	(0.00136)
Age	-0.00217	-0.00666
	(0.00331)	(0.00427)
% (some) college	-0.00049	0.00198
	(0.00193)	(0.00385)
% grad/professional	-0.00043	0.00541
	(0.00309)	(0.00448)
% Republican	0.0154	0.0159
	(0.0285)	(0.0271)
% Democrat	-0.00886	-0.00158
	(0.0214)	(0.0219)
BG FE	yes	yes
Year Dummies	yes	yes
F-statistic	43.06	117.52
Observations	10,738	3,238

Notes: Dependent variable is post-incentive PV system price per W. Unit of observation is block group-year. Specification is a linear least squares regression. BG FE refers to block group fixed effects. Standard errors clustered on town in parentheses. $p < 0.1$ (*), $p < 0.05$ (**), $p < 0.01$ (***).

Appendix G Non-instrumented Regression Output

Table G1: Non-instrumented Regression Results

Variable	Linear	Poisson	Hurdle	
			Logit	Truncated Poisson
	OLS[i]	MLE[ii]	CMLE[ii]	CMLE[i]
	(1)	(2)	(3)	(4)
Price	0.119***	0.223***	0.572***	-0.123**
	(0.0162)	(0.0497)	(0.0787)	(0.0574)
Solarize	0.953***	0.962***	0.948***	0.897***
	(0.24	(0.15)	(0.2)	(0.161)
Pop. density	-0.00001	-0.00019**	-0.00021**	0.00024
	(0.00001)	(0.00008)	(0.00009)	(0.00042)
Income	0.00101	0.00084	-0.00055	0.00273
	(0.000738)	(0.00139)	(0.00142)	(0.00269)
Age	-0.0026	-0.0099**	-0.0084	-0.00926
	(0.0025)	(0.0045)	(0.0052)	(0.0119)
% (some) college	0.00077	0.00193	0.0045	0.00348
	(0.00129)	(0.00299)	(0.00365)	(0.00626)
% grad/prof degree	0.00108	-0.0002	0.00582	-0.00743
	(0.00153)	(0.00405)	(0.00483)	(0.00791)
% Republican	0.0542**	0.0545	-0.0005	0.159***
	(0.0239)	(0.0397)	(0.0443)	(0.0592)
% Democrat	0.0263**	0.0433*	0.0107	0.109***
	(0.0125)	(0.0222)	(0.0252)	(0.0377)
BG FE	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes
Instruments	no	no	no	no
Price elasticity[iii]	0.935***	0.84	1.504***	-0.052**
	(0.128)	(0.186)	(0.207)	(0.024)
Observations	10,738	10,738	10,738	2,636

Notes: Dependent variable is number of residential PV installations. Unit of observation is block group-year. BG FE refers to block group fixed effects. $p < 0.1$ (*), $p < 0.05$ (**), $p < 0.01$ (***).

[i] Clustered standard errors at the town level in parentheses.

[ii] Block bootstrapped standard errors (100 replications), clustered at the town level, in parentheses.

[iii] Standard errors of price elasticity coefficients obtained by the delta method.

Appendix H Including Third-party-owned Systems

Table H1: Regression Results with Purchased and Third-party-owned Systems

Variable	Linear	Poisson	Hurdle	
			Logit	Truncated Poisson
	2SLS[i]	CF[ii]	CF[ii]	GMM[i]
	(1)	(2)	(3)	(4)
Price	0.037	-0.0503	-0.00323	-0.156
	(0.0429)	(0.104)	(0.109)	(0.133)
Solarize	1.294***	0.724***	0.678***	0.794***
	(0.333)	(0.111)	(0.2)	(0.119)
TPO	-0.335***	-0.210**	0.155	-0.429***
	(0.107)	(0.0912)	(0.102)	(0.126)
Pop. density	3.7×10^{-5}***	-9.9×10^{-5}*	-0.00015	0.00079**
	(8.8×10^{-6})	(0.00006)	(0.00009)	(0.00038)
Income	0.00259***	0.00049	-0.0016	0.00313
	(0.000827)	(0.00124)	(0.00154)	(0.00252)
Age	0.00035	-0.0101**	-0.0119**	-0.00149
	(0.00244)	(0.00412)	(0.00553)	(0.0097)
% (some) college	0.00067	0.00057	0.00283	0.00173
	(0.00133)	(0.00293)	(0.00441)	(0.00503)
% grad/prof degree	-0.00169	-0.00803**	0.00321	-0.0201***
	(0.00173)	(0.0032)	(0.00498)	(0.0073)
% Republican	0.0942***	0.0417	0.0203	0.0806*
	(0.0363)	(0.0364)	(0.0418)	(0.0484)
% Democrat	0.0341*	0.0307	0.00735	0.0694**
	(0.0197)	(0.023)	(0.0284)	(0.0313)
BG FE	yes	yes	yes	yes
Year Dummies	yes	yes	yes	yes
Instruments	yes	yes	yes	yes
Price elasticity[iii]	0.204	-0.194	-0.008	-0.097
	(0.2364)	(0.4025)	(0.2787)	(0.083)
Observations	13,510	13,510	13,510	4,513

Notes: Dependent variable is number of residential PV installations. Unit of observation is block group-year. TPO measures the fraction of new installations that are third-party-owned. The number of observations increases from our primary results because TPO systems are included. All other variables are the same as in Table 6. $p < 0.1$ (*), $p < 0.05$ (**), $p < 0.01$ (***).

[i] Clustered standard errors at the town level in parentheses.

[ii] Block bootstrapped standard errors (100 replications), clustered at the town level, in parentheses.

[iii] Standard errors of price elasticity coefficients obtained by the delta method.

Appendix I Inputs for the Policy Simulations

Table I1 lists the values of the main parameters and variables used in our policy simulations in Section 7. As an estimate of system costs in 2015, we assume that the trend of declining module and inverter costs continues. In order to derive the predicted change in pre-incentive average system costs (i.e., the sum of installer reported module, inverter, labor, and permitting costs) during 2015, we extrapolate the trend from 2008 to 2014. More specifically, we calculate the year-to-year percentage change in the average pre-incentive cost per W (calculated as the ratio of the average cost to the average system size) over this time period and fit a simple exponential curve through these percentage changes. We then take the extrapolated 2015 value as our forecasted cost per W in 2015. This value is shown in the top row of Table I1.

The average 2014 values of all installation-related variables are obtained directly from our core dataset. In addition, the average EPBB/HOPBI rate is derived using the mean system capacity value in 2014 at Step 4 and Step 5 incentive levels and taking the average of the two resultant rates. Lastly, our data contains information on the structure of municipal solar PV system permit fees. While some CT towns charge a flat fee per installation, in most towns the fee depends on the capacity of the installation. In those towns, we use our system size data for each installation in 2014 to obtain the fee for that installation. We average across all installations to derive an average town fee per installation and then average across all towns. After dividing the resultant number by average system capacity, we obtain the average town fee per W, reported in the last row of Table I1.

Table I1: Input Values for Policy Simulations

Parameter/Variable	Value
Change in average system cost from Jan 1, 2015 to Dec 31, 2015 (%)	-5.30
Number of block groups	1534
Average number of installations per block group in 2014	1.287
Average system capacity (kW) in 2014	7.819
Average system price ($/W) in 2014	2.946
Average rebate rate ($/W) in 2014	0.935
Average town permit fee ($/W) in 2014	0.0518

References

Andersen EB. 1970. Asymptotic properties of conditional maximum-likelihood estimators. *Journal of the Royal Statistical Society, Series B* **32**: 283–301.

Andersen EB. 1972. The numerical solution of a set of conditional estimation equations. *Journal of the Royal Statistical Society, Series B* **34**: 42–54.

Barbose G, Darghouth N, Weaver S, Wiser R. 2013. *Tracking the Sun VI: An Historical Summary of the Installed Price of Photovoltaics in the United States from 1998 to 2012*. Berkeley, CA: Lawrence Berkeley National Laboratory.

Barbose G, Weaver S, Darghouth N. 2014. *Tracking the Sun VII: An Historical Summary of the Installed Price of Photovoltaics in the United States from 1998 to 2013*. Berkeley, CA: Lawrence Berkeley National Laboratory.

Berry ST, Levinsohn JA, Pakes A. 1995. Automobile prices in market equilibrium. *Econometrica* **63**: 841–889.

Blundell R, Griffith R, Windmeijer F. 2002. Individual effects and dynamics in count data models. *Journal of Econometrics* **108**: 113–131.

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

Bollinger B, Gillingham K. 2016. Learning-by-doing in solar photovoltaic installations. *Yale University Working Paper* .

Burr C. 2014. Subsidies, tariffs and investments in the solar power market. *University of Colorado Working Paper* .

Cameron AC, Trivedi PK. 2013. *Regression Analysis of Count Data*. Cambridge: Cambridge University Press, second edition.

CEFIA. 2013. Sunrise new england: Final project report. *Clean Energy Finance and Investment Authority Report* .

Chamberlain G. 1980. Analysis of covariance with qualitative data. *Review of Economic Studies* **47**: 225–238.

Chen X, Hong H, Shum M. 2007. Nonparametric likelihood ratio model selection tests between parametric likelihood and moment condition models. *Journal of Econometrics* **141**: 109–140.

Davis L, Fuchs A, Gertler P. 2014. Cash for coolers: Evaluating a large-scale appliance replacement program in mexico. *American Economic Journal-Economic Policy* **6**: 207–238.

Dong C, Wiser R, Rai V. 2014. Incentive pass-through for residential solar systems in california. *Lawrence Berkeley National Laboratory Report* .

Gallop RJ, Rieger RH, McClintock S, Atkins DC. 2013. A model for extreme stacking of data at endpoints of a distribution: Illustration with w-shaped data. *Statistical Methodology* **10**: 29–45.

Gilles R, Kim S. 2017. Distribution-free estimation of zero-inflated models with unobserved heterogeneity. *Statistical Methods in Medical Research* **26**: 1532–1542.

Gillingham K, Deng H, Wiser R, Darghouth N, Barbose G, Nemet G, Rai V, Dong C. 2016. Deconstructing solar photovoltaic pricing: The role of market structure, technology, and policy. *The Energy Journal* **37**: 231–250.

Gowrisankaran G, Reynolds S, Samano M. 2016. Intermittency and the value of renewable energy. *Journal of Political Economy* **124**: 1187–1234.

Graff Zivin J, Kotchen MJ, Mansur ET. 2014. Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies. *Journal of Economic Behavior and Organization* **107A**: 248–268.

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

Greene WH, Hensher DA. 2010. *Modeling Ordered Choices: A Primer*. Cambridge University Press.

Gupta S, Park S. 2009. A simulated maximum likelihood estimator for the random coefficient logit model using aggregate data. *Journal of Marketing Research* **46**: 531–542.

Hendel I, Nevo A. 2006. Measuring the implications of sales and consumer inventory behavior. *Econometrica* **74**: 1637–1673.

Henwood K. 2014. Subsidy pass-through in residential solar markets. *University of Chicago Working Paper* .

Hughes JE, Podolefsky M. 2015. Getting green with solar subsidies: Evidence from the california solar initiative. *Journal of the Association of Environmental and Resource Economists* **2**: 235–275.

IAWG. 2013. *Inter-Agency Working Group on Social Cost of Carbon, Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866*. Washington, DC.

Imbens G, Wooldridge JM. 2007. What's new in econometrics? lecture 6: Control function and related methods. NBER Summer Institute.

Jiang R, Manchanda P, Rossi PE. 2009. Bayesian analysis of random coefficient logit models using aggregate data. *Journal of Econometrics* **149**: 136–148.

King G. 1988. Statistical models for political science event counts: Bias in conventional procedures and evidence for the exponential poisson regression model. *American Journal of Political Science* **32**: 838–863.

Kitazawa Y. 2012. Hyperbolic transformation and average elasticity in the framework of the fixed effects logit model. *Theoretical Economics Letters* **2**: 192–199.

Kitazawa Y. 2014. Consistent estimation for the full-fledged fixed effects zero-inflated poisson model. *Kyushu Sangyo University Discussion Paper No. 66* .

Majo MC, van Soest AHO. 2011. The fixed-effects zero-inflated poisson model with an application to health care utilization. *Tilburg University Discussion Paper 2011-083* .

Min Y, Agresti A. 2005. Random effect models for repeated measures of zero-inflated count data. *Statistical Modelling* **5**: 1–19.

Mullahy J. 1986. Specification and testing of some modified count data models. *Journal of Econometrics* **33**: 341–365.

Mullahy J. 1997. Instrumental-variable estimation of count data models: Application to models of cigarette smoking behavior. *Review of Economics and Statistics* **79**: 586–593.

Neelon B, Ghosh P, Loebs PF. 2013. A spatial poisson hurdle model for exploring geographic variation in emergency department visits. *Journal of the Royal Statistical Society: Series A* **176**: 389–413.

Neyman J, Scott EL. 1948. Consistent estimates based on partially consistent observations. *Econometrica* **16**: 1–36.

Olley GS, Pakes A. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* **64**: 1263–1297.

Petrin A, Train K. 2010. A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research* **47**: 3–13.

Pless J, van Benthem A. 2017. The suprising pass-through of solar subsidies. *NBER Working Paper 23260* .

Pohlmeier W, Ulrich V. 1995. An econometric model of the two-part decisionmaking process in the demand for health care. *Journal of Human Resources* **30**: 339–361.

Rasch G. 1960. *Probabilistic Models for Intelligence and Attainment Tests*. Danish Institute for Educational Research.

Rasch G. 1961. On general laws and the meaning of measurement in psychology. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* **4**: 321–333.

Rogers E, Sexton S. 2014. Effectiveness of subsidies for residential rooftop solar adoption: Lessons from california. *North Carolina State University Working Paper* .

Sallee J. 2011. The surprising incidence of tax credits for the toyota prius. *American Economic Journal: Economic Policy* **3**: 189–219.

Santos Silva JMC, Windmeijer F. 2001. Two-part multiple spell models for health care demand. *Journal of Econometrics* **104**: 67–89.

Shaw S, Fahey R, Solomon P. 2014. Progress report: Cefia residential solar investment program. *The Cadmus Group, Inc. Report* .

Silva JMCS, Tenreyro S, Windmeijer F. 2015. Testing competing models for non-negative data with many zeros. *Journal of Econometric Methods* **4**: 29–46.

Terza JV. 1998. Estimating count data models with endogenous switching: Sample selection and endogenous treatment effects. *Journal of Econometrics* **84**: 129–154.

van Benthem A, Gillingham K, Sweeney J. 2008. Learning-by-doing and the optimal solar policy in california. *The Energy Journal* **29**: 131–151.

Vuong QH. 1989. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* **57**: 307–333.

Windmeijer F. 2008. Gmm for panel data count models. In Mátyás L, Sevestre P (eds.) *The Econometrics of Panel Data*. Berlin Heidelberg: Springer-Verlag.

Windmeijer F, Santos Silva JMC. 1997. Endogeneity in count data models: An application to demand for health care. *Journal of Applied Econometrics* **12**: 281–294.

Winkelmann R. 2004. Health care reform and the number of doctor visits - an econometric analysis. *Journal of Applied Econometrics* **19**: 455–472.

Wooldridge J. 2002. *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

Yang S, Chen Y, Allenby GM. 2003. Bayesian analysis of simultaneous demand and supply. *Quantitative Marketing and Economics* **1**: 251–275.