Subsidies, Cost Trends, and Pass-Through: Evidence from the Residential Solar Market

Xiaolin Zhou*

November 1, 2026

Click here for the latest version

Abstract

Clean energy subsidies are one of the main instruments in climate policy. Their values are designed to fall as technology adoption grows, but costs generally change over the same horizon, confounding the estimates of pass-through. I study two Massachusetts programs that pay homeowners per kilowatt-hour based on energy produced. To overcome timing challenges, I employ a novel identification strategy: estimating pass-through from differences in the subsidy across roofs that the same installer equips at the same time, because the subsidy depends on each roof's generation. I find that buyers receive about 85 percent of the subsidy under both designs, whereas a comparison over time gives a biased estimate. Installers raise prices only slightly for roofs that earn larger subsidies, which in this setting bounds the markup from roof-specific price discrimination at 2–7 percent of price.

*University of Massachusetts Amherst, Department of Resource Economics. I am extremely grateful to my advisors, Christine Crago, Yongjoon Park, and Matt Woerman, for their guidance and support. I thank Shanjun Li, Arthur van Benthem, and Jacquelyn Pless for helpful conversations, and participants in numerous seminars for their comments. All errors are my own. Email: xiaolinzhou@umass.edu.

1 Introduction

Subsidies have become one of the main instruments of climate policy. In the United States, the Inflation Reduction Act of 2022 relies on tax credits and other subsidies for clean electricity, electric vehicles, and home energy (Bistline, Mehrotra and Wolfram, 2023). Whether these subsidies lower the cost of adoption depends on pass-through, the share of each subsidy dollar that reaches buyers rather than sellers. Pass-through determines who gains from a program and how much adoption each public dollar buys. Residential solar shows the scale and variety of these subsidies: the federal government alone spent $7.5 billion on solar subsidies in fiscal year 2022 (U.S. Energy Information Administration, 2023), and states add rebates, net metering, and payments for each kilowatt-hour a system generates (Borenstein, 2017).

Measuring pass-through is difficult because it's rare to observe only subsidy changes while everything else stays fixed. Existing estimates for solar identify pass-through from the timing of the subsidy within one state, including step-downs in the California Solar Initiative rebate (Pless and van Benthem, 2019; Dong, Wiser and Rai, 2018), step-downs in Connecticut's rebate (Gillingham and Tsvetanov, 2019), and a change in the federal investment tax credit (Podolefsky, 2013). Like many clean energy subsidies, solar subsidies are scheduled to decline as deployment grows. Installation costs also change over time. They tend to fall as hardware becomes cheaper and installers learn, but they also fluctuate with input prices and general inflation. A regression of prices on the subsidy mixes how the subsidy is shared with how costs and the subsidy move together. Studies of tax incidence in other markets address common cost shocks with a comparison group that the tax does not reach (Kosonen, 2015; Kopczuk et al., 2016; Butters, Sacks and Seo, 2022), but a solar subsidy reaches every installation in a utility territory at a given time.

This paper instead compares buyers who receive different subsidies in the same market at the same time. I study two production-based programs in Massachusetts. Under the Solar Carve-Out II program (2014–2018), owners earned SRECs sold at market prices. Under its successor, the Solar Massachusetts Renewable Target (SMART) program, they receive a fixed rate per kWh that is set when the application is approved and declines as capacity reaches respective targets. Because both programs pay per kWh, a project's subsidy depends on how much electricity its roof will generate. I merge regulatory data from the Massachusetts Department of Energy Resources (DOER) with Lawrence Berkeley National Laboratory's Tracking the Sun database for more than 39,000 host-owned residential systems installed between 2014 and 2023, geocode each project's street address, and model each system's annual generation with the National Renewable Energy Laboratory's PVWatts calculator

from the project's location, size, tilt, azimuth, and equipment. About 91% of the variance in modeled capacity factors is within towns. Roofs that the same installer equips in the same month therefore can earn very different subsidies, and this variation does not depend on the policy calendar. My preferred specification compares roofs within installer-by-year, town-by-year, and month cells. This design is available because the subsidy is paid for production, while an upfront rebate per watt is nearly the same for every roof. Because generation is modeled rather than metered, the subsidy is measured with error. The outcome, the net price, is the gross price minus this subsidy, so the error enters both sides of the regression and pulls estimated pass-through toward one.

The first finding is that the two designs pass through the subsidy at the same rate. Comparing roofs, pass-through is 0.86 under Solar Carve-Out and 0.85 under SMART, and the two are statistically indistinguishable. The programs differ in the form of the payment, in how predictable it is, and in size. The average Solar Carve-Out subsidy is about 3.4 times the average SMART subsidy per kW ($1,923 versus $566 in present value), about 45% and 15% of the price. Buyers nonetheless capture about 85 cents of each dollar under both. Error in modeled generation pulls both estimates toward one, so true pass-through is lower. Because the comparison is across roofs, this rate is the pass-through of a subsidy specific to one roof. With log-concave demand it is an upper bound on the pass-through of a subsidy change common to all buyers, which a timing-based design would only recover if costs remained constant (Section 6.1). To my knowledge, these are the first estimates of the pass-through of a production-based subsidy and the first comparison of a certificate-based and a fixed-rate design within one market.[1]

The second finding is that estimates that rely on the timing of the subsidy are confounded by cost trends. Under SMART, the subsidy kept falling after 2020 while prices rose, which biases a timing-based estimate up. A naive comparison with year fixed effects gives 1.02. A stacked event study around the step-downs in SMART rates, a design meant to avoid trends, gives 1.27. Both lie above the cross-roof estimate. Error in modeled generation cannot explain this gap, because it pulls estimates toward one but not above. Pass-through above one is also inconsistent with the log-concave demand I estimate unless marginal cost falls with volume. Under Solar Carve-Out, the subsidy and installation prices fell together, which biases a timing-based estimate down. The naive estimate, 0.76, lies below the cross-roof estimate of 0.86, as predicted. Error in modeled generation also pulls the cross-roof estimate toward one, and the estimate identified from roof orientation alone is 0.76, so this

[1]Comparisons in other settings find that the form of an incentive can affect pass-through: car buyers capture 70 to 90 percent of a customer rebate but 30 to 40 percent of a dealer discount (Busse, Silva-Risso and Zettelmeyer, 2006), and the point at which a diesel tax is collected changes its pass-through (Kopczuk et al., 2016).

comparison does not separate the timing bias from the pull of error. Timing-based estimates themselves depend on how much of the cost trend their window includes. A regression discontinuity design at the January 2017 cut in the subsidy for new small systems falls from 1.03 to 0.71 as the bandwidth widens from 90 to 365 days that includes more of the decline in costs.

These results also bound the market power from price discrimination across roofs according to the subsidy each roof earns. Under a uniform price, a subsidy specific to one roof would be passed through in full. The shortfall of cross-roof pass-through below one therefore measures how much more installers charge for roofs that earn more subsidy. Installers quote each job individually, and prices differ by about \$450–490 per kW across buyers of the same installer in the same month. Yet a roof that earns \$59–137 more modeled subsidy per kW is charged only \$6–20 more. Combined with log-concave demand under SMART and the sufficient-statistics framework of Weyl and Fabinger (2013), this shortfall bounds the markup from price discrimination across roofs at 2–7% of price. Demand curvature is positive in every SMART specification and significant at the 10% and 5% levels in the two main estimates (Section 6.3). A markup that installers charge across the board leaves cross-roof pass-through at one and lies outside the bound (Section 6.4).

This paper contributes to three literatures. The first studies the incidence of taxes and subsidies and the variation that identifies it. Reduced-form studies identify incidence from changes in taxes and subsidies over time, from rebate step-downs and tax credit changes (Sallee, 2011; Busse, Silva-Risso and Zettelmeyer, 2006) to tax changes compared with untaxed goods or regions (Kosonen, 2015; Kopczuk et al., 2016; Butters, Sacks and Seo, 2022), from national credit prices (Lade and Bushnell, 2019), and from before–after comparisons (Kenkel, 2005). The variation also determines what pass-through measures: in U.S. oil refining, pass-through ranges from near zero for a cost shock specific to one firm to about one for a shock common to the industry (Muehlegger and Sweeney, 2022), and California's electric vehicle rebates, which differ by income, identify the price recipients pay relative to non-recipients (Muehlegger and Rapson, 2018).[2] I introduce a new source of identifying variation for markets where the subsidy schedule tracks costs and no untreated group exists. A subsidy paid for production creates its own comparison group: buyers served by the same installer at the same time who receive different subsidies because their roofs generate different

[2]Related work studies hybrid vehicle subsidies with endogenous quality and costly bargaining (Gulati, McAusland and Sallee, 2017), electric vehicle subsidies across countries (Barwick et al., 2023), the incidence of agricultural subsidies on farmland rents (Kirwan, 2009), and the welfare effect of a consumer subsidy restricted by price ceilings (Fan and Zhang, 2022). Incomplete pass-through can also reflect local costs and markup adjustment (Nakamura and Zerom, 2010), and firms that set prices nationally may respond differently to local cost shocks (Butters, Sacks and Seo, 2022).

amounts. The same design applies wherever a subsidy scales with a measured characteristic of the buyer or the product, such as home energy rebates that scale with predicted energy savings or vehicle tax credits that rise with battery capacity. I also show that timing-based designs, including a regression discontinuity and an event study, are confounded by cost trends in this setting.

The second literature estimates the pass-through of solar subsidies and studies how subsidy design affects buyers. Most of this evidence comes from upfront rebates, and it finds that buyers capture most but not all of the subsidy: 78% for host-owned systems in California (Pless and van Benthem, 2019), 84% in Connecticut (Gillingham and Tsvetanov, 2019), and close to complete pass-through in other work on California (Dong, Wiser and Rai, 2018). Solar adoption responds strongly to rebates (Hughes and Podolefsky, 2015; Crago and Chernyakhovskiy, 2017), buyers discount production-based payments heavily (De Groote and Verboven, 2019), and upfront subsidies induce more adoption while production subsidies are more efficient (Burr, 2016).[3] I provide estimates for production-based subsidies paid through certificates and through fixed rates, and I find that the two designs pass through the same share once the timing of the subsidy is absorbed. Massachusetts in 2018 and New Jersey in 2021 replaced market-priced certificates with administratively set payments. In Massachusetts, buyers received about the same share of the subsidy before and after the switch, so the choice between the two designs can rest on their cost, risk, and administration.

The third literature uses pass-through to learn about market power. Pass-through tends to rise with competition, from 13% in the least competitive Medicare Advantage markets to 74% in the most competitive (Cabral, Geruso and Mahoney, 2018) and, for fuel taxes on Greek islands, from 0.4 in monopoly markets to 1 in markets with four or more competitors (Genakos and Pagliero, 2022). It has been used to measure incidence under imperfect competition (Ganapati, Shapiro and Walker, 2020) and, with convex demand, to detect market power from more-than-complete pass-through (Pless and van Benthem, 2019), while high pass-through can also arise from cost shocks common to all firms (Fabra and Reguant, 2014). I bound the markup from price discrimination across roofs using roof-specific pass-through, which removes the supply side from the pass-through formula, together with estimated log-concave demand. The bound matters for subsidy design. Paying for production gives better roofs larger subsidies, which installers who know each roof's output could capture by charging those roofs more. The bound shows that this form of market power is small, so a production-

[3]Other studies compare the effectiveness of residential solar incentives (Matisoff and Johnson, 2017), estimate how residential solar demand depends on net metering and electricity prices (Ros and Shetty, 2023), and characterize optimal subsidies for residential solar (Colas and Reynier, 2026). Consumer subsidies can also lower production costs through induced innovation (Gerarden, 2023), and the phase-out of a subsidy shapes the dynamics of demand (Hu, Yin and Zhao, 2025).

based subsidy can reward better roofs without much of the reward going to installers.

Section 2 describes the two programs, the data, and the construction of the subsidy. Section 3 sets out the empirical strategy, explains why timing variation is confounded, and compares pass-through across the two programs. Section 4 addresses threats to the cross-roof design, Section 5 shows that the estimates that rely on timing are confounded by cost trends, and Section 6 draws the implications for market power. Section 7 concludes.

2 Institutional Background and Data

2.1 The Two Programs

When a household in Massachusetts installs a solar system, one of two production-based subsidies is available, depending on the policy regime at the time of installation. If a qualifying system entered commercial operation between January 2008 and November 2018, the owner was eligible to receive certificates under the Solar Carve-Out program, which launched in 2010 and covered systems operating since 2008. If it entered commercial operation after November 2018, the owner is entitled to an incentive payment through the electricity bill, on top of net metering, under the SMART program. The Solar Carve-Out is part of Massachusetts's Renewable Portfolio Standard (RPS) and requires that a certain percentage of electricity come from solar resources. The program launched on January 1, 2010, with a capacity target of 400 MW, and the state expanded it in April 2014 with Solar Carve-Out II, which set a total goal of 1,600 MW. Energy suppliers must purchase Solar Renewable Energy Certificates (SRECs) to demonstrate compliance or make an alternative compliance payment (ACP). The NEPOOL Generation Information System issues SRECs quarterly to system owners at a rate of 0.45–1.0 SREC per MWh, depending on system characteristics (the SREC factor); a system smaller than 25 kW authorized to interconnect before January 8, 2017, receives a factor of 1.0. Each unit receives SRECs for 40 calendar quarters and earns much less valuable Renewable Energy Certificates (RECs) thereafter.[4]

The ACP caps SREC prices, and the Massachusetts Department of Energy Resources (DOER) holds an annual fixed-price clearinghouse auction for unsold certificates, but the auction price is not a floor, because certificates that remain unsold can still be traded for two more compliance years. For Solar Carve-Out II, the ACP and auction price schedules were published at launch: the net clearinghouse price was \$285/MWh from 2014 to 2016 (a \$300 price less the 5% auction fee) and then fell by 5% per year. Publishing the schedules

[4]Certificates are issued by vintage, the calendar year of generation, and different vintages trade as distinct market segments. Owners typically sell through an aggregator, which charges small residential systems about 10% of the proceeds.

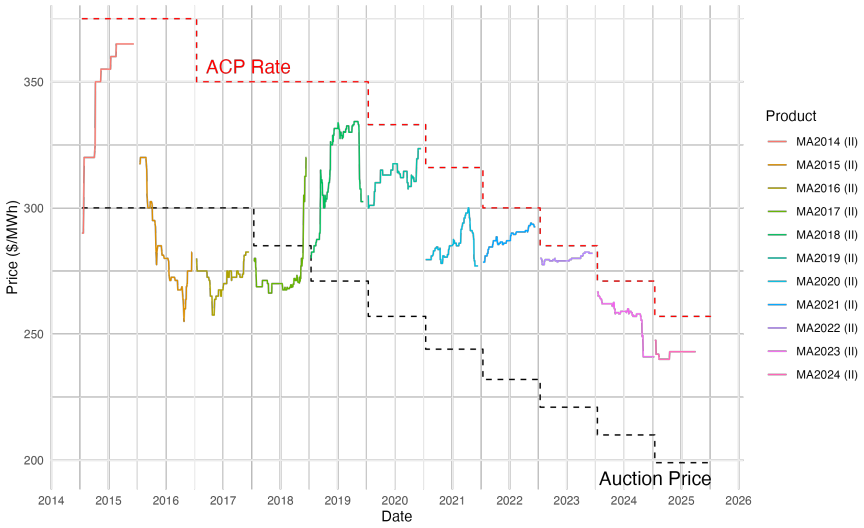


Figure 1: SREC II Market Prices with ACP and Auction Prices

The auction price is plotted gross of the 5% auction fee. An owner selling through the auction receives 95% of the plotted value ($285/MWh in 2014–2016). Each vintage is shown over its trading window, from July 15 of the vintage year to July 14 of the following year.

narrowed the range of likely SREC prices but did not remove price risk. Figure 1 shows that market prices were well below auction prices in the first few years because of oversupply and later followed the declining ACP.

DOER created the Solar Massachusetts Renewable Target (SMART) program to further encourage solar development in Massachusetts after the Solar Carve-Out II target was met. The program launched on November 26, 2018. Under SMART, the three investor-owned utilities (IOUs) in Massachusetts (Eversource, National Grid, and Unitil) compensate solar system owners for the electricity their systems generate, in exchange for the RECs attached to this generation.

The actual incentive rate for a given system is determined by four factors: the base compensation rate, household income status, whether the system includes a storage component, and the value of energy. The base compensation rate starts high and declines on a predetermined schedule as capacity blocks fill, and both block sizes and rate schedules vary by utility. The value of energy is set from three-year average retail electricity prices (including delivery and transmission charges) at the time of application. The SMART incentive rate equals the base compensation rate minus the value of energy. Small residential solar generation unit owners receive payments through the utility equal to this rate multiplied by their monthly generation, in addition to the bill savings their behind-the-meter systems already provide, and they can cash out a negative balance monthly. As electricity prices rise and capacity

blocks advance, the SMART incentive rate decreases until it reaches zero.[5]

Appendix Table 8 shows the base compensation rates for generation units of 25 kW AC or less, by utility and block.[6]

SMART was most favorable early in its life cycle, when the incentive rate was considerably higher than the price of RECs. In my sample, the SMART incentive rate is at most \$0.20/kWh, with a median of about \$0.066/kWh, while RECs traded at around \$0.04/kWh.[7] SMART also offers more predictable cash flows than the Solar Carve-Out program, because the rate is fixed and does not fluctuate with market conditions. The subsidy can therefore be calculated from modeled generation without an assumption about future prices. Another key difference between SMART and Solar Carve-Out is that SMART requires system interconnection to one of the three IOUs. As a result, households served by municipal utilities are ineligible for the program.

Table 1 summarizes the two programs. Four differences could affect how subsidy value is divided between buyers and installers: the form of payment (certificates sold at a market price capped by the ACP versus a fixed per-kWh rate); the timing of information (the SMART rate is fixed when the application is approved, so buyer and installer know the subsidy per kWh for ten years, whereas the SREC price an owner will receive is unknown at purchase); the size of the subsidy (about 45% of the average gross price under Solar Carve-Out and 15% under SMART, with almost no overlap between the two distributions); and coverage (all utility territories versus customers of the three IOUs).

These differences give reasons to expect pass-through to differ between the programs. A larger subsidy that arrives through a volatile market price is more likely to strain installers' capacity, which would lower pass-through. A small, certain subsidy is less likely to move installer capacity, which would raise pass-through. Because the programs also differ in years, technology costs, and eligible households, a comparison of the programs over time cannot test these predictions. Section 3.6 tests whether pass-through differs once the timing of the subsidy is absorbed.

Two rules within the programs change the subsidy discretely, and I use them both in Section 5. On January 8, 2017, the SREC factor for new small systems fell from 1.0 to 0.8, so a system that generates the same electricity receives 20% fewer certificates. Under SMART, the base rate of a utility steps down when one of its capacity blocks fills (Appendix

[5]System owners who do not wish to enroll in SMART can instead participate directly in the REC market, and DOER advises them to do so when the SMART rate becomes low. I do not observe how many owners opt out, because owners who do so do not appear in the SMART regulatory data.

[6]There are eight blocks for the three large utilities but only two and four for Nantucket and Fitchburg, respectively. Fitchburg is Unitil's operating company, and Mass Electric and Nantucket Electric are National Grid's operating companies in Massachusetts.

[7]Source: SRECTrade trading platform.

Table 1: Comparison of the Solar Carve-Out II and SMART Programs

	Solar Carve-Out II	SMART
Period in sample	Jan. 2014–Nov. 2018	Mar. 2019–Sep. 2023
Form of payment	Tradable SRECs	Fixed rate per kWh
Rate	Market price, capped by the ACP	Fixed at approval
Payment term	40 calendar quarters	10 years
Payer	Retail electricity suppliers	The three IOUs
Eligibility	All utility territories	IOU customers only
Mean subsidy (NPV, 8%)	$1,923 per kW	$566 per kW
Share of gross price	~45%	~15%

Subsidy is the present value per kW of DC capacity at an 8% discount rate with 0.5% annual degradation, from the analysis samples. Shares are ratios of the sample mean subsidy to the sample mean gross price per kW (systems of 25 kW or less).

Table 8).

2.2 Project and Price Data

I obtain the solar installation data from DOER. For SMART projects, the DOER data report compensation rates as well as project characteristics, including location, utility, size, installation cost, installer, and technical details. The technical details, which include azimuth, tilt, and inverter rating, are essential for modeling generation accurately. The compensation rate reported in the data is the actual SMART incentive rate that system owners receive, which reflects any adders and the value of energy.

For Solar Carve-Out II projects, DOER reports similar variables but without technical details. To model annual generation as accurately as for SMART, I merge the DOER data with the Tracking the Sun database published by Lawrence Berkeley National Laboratory (LBNL).[8] All projects in my data appear in Tracking the Sun, along with their technical details. I match projects across the two sources on operation date, size, and total cost, which gives an almost one-to-one match (Appendix A). SREC prices are not available in the project data. I construct an SREC price series from a major SREC trading platform, SRECTrade, and from the price support regulations published on DOER's website.

The Solar Carve-Out sample spans January 1, 2014, to November 26, 2018, the period covered by Solar Carve-Out II. The SMART sample begins on March 25, 2019, the first

[8]Tracking the Sun is LBNL's cleaned compilation of project records from state agencies, including DOER and the Massachusetts Clean Energy Center, and covers more than 60% of solar projects in the United States.

commercial operation date in the SMART data, and ends on September 20, 2023, by which time most projects received zero compensation and the number of applications had dropped sharply. For SMART projects, I date each project by its Commercial Operation Date (Appendix A).

The sample includes only host-owned, small-scale (≤25 kW DC) residential projects without battery storage. Small systems are exempt from the net metering capacity cap, so every system in the sample is on net metering; only host-owned projects have directly observed prices, because the price reported for third-party-owned projects is a "market" value rather than the cost to the owner (Pless and van Benthem, 2019); and battery storage qualifies for a storage adder and would distort the measured cost of the solar system. I also remove observations with a per-kW price above the 99th percentile or below the 1st percentile.[9] Appendix Figure 7 maps the projects. The main difference between the two samples is coverage: households in municipal utility territories are ineligible for SMART.[10] In both samples, the density of projects closely resembles the population distribution of Massachusetts.

2.3 Modeled Generation and the Subsidy

Measuring pass-through requires the value of the subsidy for each project, which the data do not report. Because both the SMART and Solar Carve-Out II programs are production-based subsidies, the total subsidy depends on how much electricity a system will generate and when the payments are made. SMART payments are made monthly through the utility, so the subsidy is the incentive rate times monthly generation. Under the SREC program, payment comes from selling certificates, which are issued every three months, so the subsidy is the market price of certificates times the number of certificates produced each quarter. For simplicity, I calculate the subsidy on an annual basis, which requires annual generation for each project.

I model annual generation using the PVWatts calculator, a tool developed by the National Renewable Energy Laboratory for modeling the energy production of grid-connected photovoltaic systems. The calculator combines solar irradiance and meteorological data for a project's location with the system's size, DC-to-AC ratio, module type, array type, azimuth, and tilt. The public DOER data contain only ZIP codes, which can cover large areas relative to the resolution of the solar resource data, so I obtained each project's street address from DOER and converted it to coordinates (Appendix A). I use the observed values of the other

[9] Winsorizing the price per kW at the 1st and 99th percentiles instead of dropping these projects leaves the estimates essentially unchanged.

[10] Excluding the 1,567 Solar Carve-Out projects in municipal utility territories changes the Solar Carve-Out estimates very little.

inputs where available, assume a fixed roof-mount array with standard modules, and keep the system loss parameter at the calculator default of 14.08%; Appendix A describes the handling of missing inputs, which affects fewer than 100 projects. The calculator applies the same default losses for shading, snow cover, soiling, and outages to every system, so its output is a model prediction rather than a measurement of what each system generates.

The output includes monthly and annual generation for each project. I summarize it using the capacity factor: a system's annual generation as a share of what it would produce if it ran at full rated capacity all year. Most projects have modeled capacity factors of 12.5–15%, and a few are much lower.[11]

What matters most for this study is how generation differs across roofs, and Figure 2 shows that these differences are large. The modeled capacity factor has a standard deviation of 0.012 around a mean of 0.143: more than half of projects (53%) differ from the statewide mean by more than 5%, and 18% by more than 10%. These differences are local. About 91% of the variance in the capacity factor is within towns, so it comes mostly from the orientation, tilt, and design of individual roofs rather than from regional differences in sunshine. Modeling each roof from its own address and design, rather than from its town or ZIP code, is what makes the comparison across roofs possible.

The average modeled capacity factor is 14.3%, slightly above the 13.1% that DOER reports from metered generation for the whole Massachusetts fleet in 2010–2023.[12] A difference of this kind, common to all projects, changes the scale of the subsidy rather than its variation across roofs, and Section 3.4 shows that it moves the pass-through estimates very little.

What matters more for the design is whether modeled generation ranks roofs correctly. The production estimates that installers report on SMART applications provide a second measure with which to check this. Section 4.3 shows that, within the cells the analysis compares, the two measures have a correlation of 0.57, so part of the remaining variation in modeled generation is error, and Section 3.4 shows that such error pulls the estimates toward one rather than toward zero. The analysis measures generation with PVWatts in both programs, so that the subsidy is built the same way for every roof, and uses the installer's estimate only as a sensitivity check.

[11] By entering the technical specifications into the calculator manually, I verified that the low values are correct. They come from suboptimal configurations: a north-facing system, for example, generates much less than a south-facing one, all else equal.

[12] Panel degradation closes roughly a fifth of this gap. Field evidence puts the loss at about 0.5% of output per year (Deline et al., 2024), applied here over the age distribution of the fleet. The larger part of the gap most plausibly reflects losses that PVWatts does not model. The default snow loss of zero is unrealistic for Massachusetts, and the default shading and soiling losses of 3% and 2% are likely too low given the long pollen season and restrictions on tree removal. The statewide fleet also includes lower-quality installations that predate my sample.

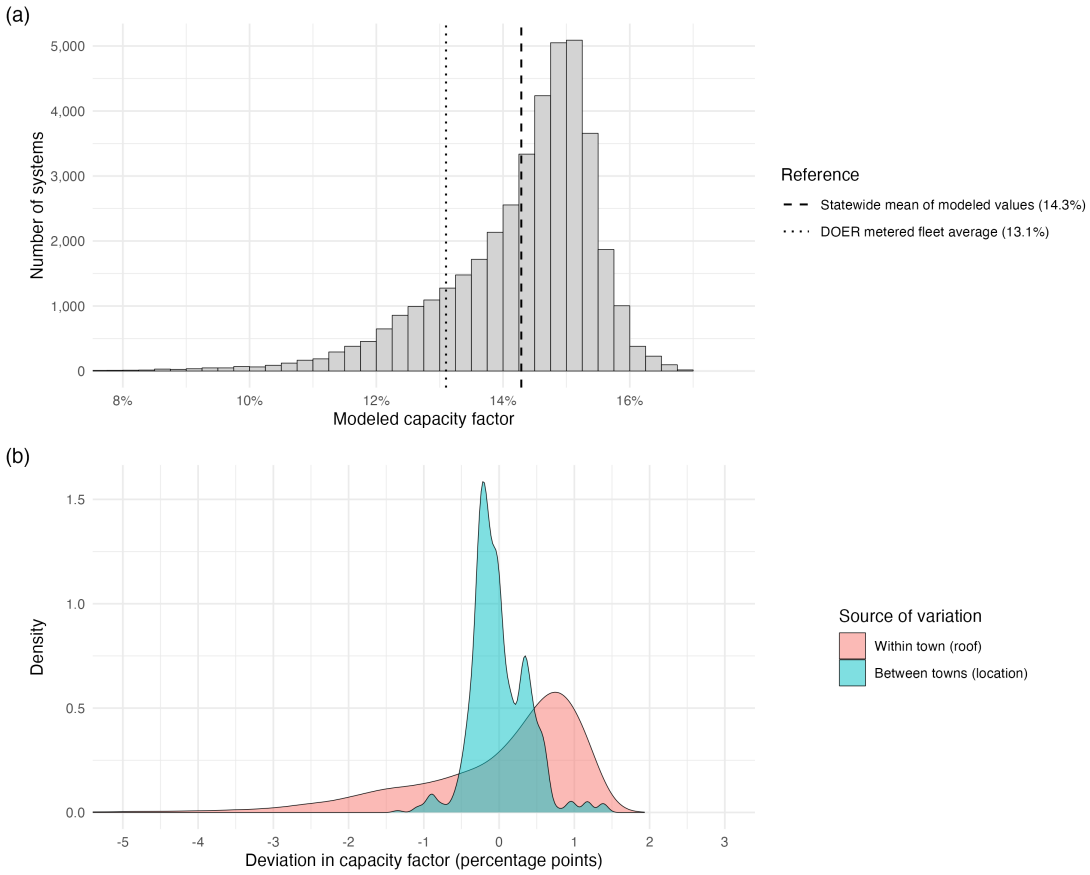


Figure 2: Variation in Modeled Capacity Factor Across Projects

Both analysis samples (39,808 projects). Panel (a) is the distribution of the modeled capacity factor, PVWatts annual generation as a share of a year's output at full rated capacity; the dashed line is the sample mean (14.3%) and the dotted line the fleet average that DOER reports from metered generation (13.1%). Panel (b) splits each project's deviation from the statewide mean, in percentage points, into the deviation of its town's mean from the statewide mean (between towns) and its own deviation from its town's mean (within town).

I first calculate the net present value of the subsidy for SMART projects, assuming a discount rate of 8%. Because the SMART compensation rate is locked in for ten years when the application is approved, the owner's compensation in a given year depends only on that year's generation. De Groote and Verboven (2019) find that system owners heavily discount the future cash flows from such subsidy programs, so I also report results for other discount rates. I follow a standard discounted cash flow model:

$$\text{NPV}_i = \sum_{t=1}^{10} \frac{r_i \times g_i (1-a)^{t-1}}{(1+d)^t}$$

where r_i is the SMART incentive rate, g_i is modeled annual generation, d is the discount rate, and a is the panel degradation rate. I assume annual degradation of 0.5% to account

for the performance loss from panel aging. Because r_i and g_i do not change over the ten years, the formula reduces to $\text{NPV}_i = r_i\, g_i\, A(d)$, where $A(d) = \sum_{t=1}^{10}(1-a)^{t-1}/(1+d)^t$ is the same for every project. The discount rate therefore rescales every SMART subsidy by the same factor, which is relevant for interpreting the robustness checks in Appendix B.1.

For Solar Carve-Out projects, I apply the same approach with certificate prices:

$$\text{NPV}_i = \sum_{t=1}^{10} \frac{g_i \times f_i \times 90\% \times P_t(1-a)^{t-1}}{(1+d)^t}$$

where g_i is modeled annual generation in MWh, f_i is the project's SREC factor, P_t is the SREC price in year t, and the 90% term reflects the 10% aggregator fee. Ten years of payments cover the whole entitlement of 40 calendar quarters.

Because SREC prices were set in a market, they are not in the Solar Carve-Out data, and the price an owner would receive was not known at the time of purchase. In the early years, market prices were well below the ACP and even below the auction price, so using either alone could misstate the subsidy. My main price path is the midpoint of the ACP and the auction price in each year, the price a buyer would expect without a view on where the market would settle between them. As an alternative, I use the annual average market price, which is what owners actually received, with the 2024 average for years not yet observed. Appendix B.1 compares these and other price paths.

2.4 Summary Statistics

Table 2 presents summary statistics for the two samples. The largest difference between the programs is in the subsidy per kW: \$1,923 on average under Solar Carve-Out, about 3.4 times the \$566 under SMART, or about 45% and 15% of the average gross price per kW. The two distributions have little mass in common. Every Solar Carve-Out system received at least \$757 per kW, while some SMART systems received nothing because the value of energy had caught up with the base rate, so a comparison of the two programs is a comparison of very different parts of the subsidy range. The systems themselves are similar: average size is 7.96 kW DC under Solar Carve-Out and 9.22 kW DC under SMART, and the mean capacity factor is about 14% in both samples.

Table 2: Summary Statistics by Program

	Solar Carve-Out		SMART	
	Mean	SD	Mean	SD
System size (kW DC)	7.96	3.20	9.22	3.53
Gross price per kW ($)	4,304	964	3,796	881
Subsidy per kW ($)	1,923	336	566	318
Net price per kW ($)	2,381	916	3,230	971
Subsidy as a share of the mean gross price	0.45		0.15	
Modeled capacity factor (%)	14.4	1.0	14.0	1.5
Tilt (degrees)	29.4	9.1	29.4	9.5
Azimuth (degrees)	181.0	41.3	179.9	56.9
SREC factor	0.92	0.10	—	
SMART rate ($/kWh)	—		0.070	0.038
Projects	25,491		14,317	
Installers	385		267	
Towns	348		303	

Means and standard deviations in the analysis samples. Subsidies are present values at an 8% discount rate with 0.5% annual panel degradation; Solar Carve-Out subsidies use the ACP–auction midpoint price path. The subsidy share is the mean subsidy per kW divided by the mean gross price per kW. The modeled capacity factor is PVWatts annual generation as a share of a year's output at full rated capacity. Tilt and azimuth are as recorded on the application; an azimuth of 180 degrees faces due south.

3 Empirical Strategy and Results

3.1 The Naive Model

I estimate the pass-through of solar subsidies to net installation prices with the regression

$$y_i = \alpha + \beta \text{Subsidy}_i + X_i\gamma + \delta_{u(i)} + \delta_{t(i)} + \varepsilon_i$$

where y_i is the net installation price per kW after subsidy (gross price minus subsidy, per kW of DC capacity), Subsidy_i is the net present value per kW of the subsidy accumulated over ten years, X_i is a set of controls, $\delta_{u(i)}$ is the fixed effect for the utility $u(i)$ that serves project i, and $\delta_{t(i)}$ is the fixed effect for the year $t(i)$ in which it began operating. Because y is net of the subsidy, the coefficient β equals the negative of the pass-through rate ρ: a coefficient of −1 means that a dollar of subsidy lowers the net price by a dollar (the gross price does not change), and a coefficient of −0.76 means the gross price rises by 24 cents.

The controls X_i include fixed effects for panel manufacturer, inverter manufacturer, installer, and city, all of which affect system prices. Panel and inverter manufacturer fixed effects control for equipment quality, installer fixed effects control for labor costs and in-

staller quality, and city fixed effects control for demographic characteristics and local price levels. City here means the municipality, and I use "city" and "town" interchangeably. Because SMART rates vary across utility territories, utility fixed effects allow me to identify pass-through from changes in the subsidy within a utility territory, where the compensation structure is the same. Year fixed effects capture the effect of technology and learning by doing on installation costs that is common to a year.

I estimate the model separately for the SMART and Solar Carve-Out samples, using the net present value of the subsidy at a discount rate of 8%. The programs differ in period, in how the subsidy is measured, and in the fixed effects the policy requires, so I let every coefficient differ by program and compare the estimates directly (Appendix B.4 compares this with pooled estimation).[13] All prices and subsidies are in nominal dollars.[14]

I call this model the naive specification because it uses all of the variation in the subsidy, including its movement within years, without separating the part that comes from the timing of the subsidy from the part that comes from differences across roofs. Section 3.2 explains why the timing variation is confounded with installation costs, and Section 5 shows the consequences for the estimates.

3.2 Two Sources of Variation and the Timing Confound

The subsidy per kW varies for two reasons, which differ in how plausibly they are independent of installation costs. Under Solar Carve-Out, the subsidy varies across installation dates through the SREC price path and, in January 2017, through the cut in the SREC factor; year fixed effects absorb a level shift common to all projects in a year, but not the movement of the price path within a year or the interaction between the annual price path and the timing of payments (a project that begins operating in March sells more certificates at that year's vintage price than one that begins in November, and the January 2017 factor cut falls inside a year). Under SMART, the base rate falls as capacity blocks fill within a utility territory, and the value of energy differs by utility and over time. The second source is generation across roofs: two otherwise identical projects in different locations or with different orientations face different subsidies, in both programs, and this variation does not depend on the policy calendar, so it cannot move together with installation costs over time.

Write the gross price of project i as $P_i = c_i + (1-\rho)S_i$, where c_i is the cost-based price, the

[13] I do not include utility fixed effects for the Solar Carve-Out sample because SREC prices do not depend on the utility. Adding them leaves the Solar Carve-Out estimate essentially unchanged (Appendix Table 13).

[14] Both are nominal contracts, and pass-through is a within-period object, the effect of a dollar of subsidy on the price paid on the same date, so nominal dollars are the natural unit, and time fixed effects absorb the economy-wide price level. The choice does affect the naive SMART estimate, because consumer prices rose by about 22% over the SMART sample while rates were stepping down (Section 5.1 and Appendix C.4).

price the installer would charge without the subsidy, and S_i is the subsidy. The dependent variable is the net price $y_i = P_i - S_i = c_i - \rho S_i$. After the fixed effects and controls are partialled out, the regression coefficient converges to

$$\hat{\rho} \;\to\; \rho - \frac{\mathrm{Cov}(\tilde{c}, \tilde{S})}{\mathrm{Var}(\tilde{S})}, \tag{1}$$

where tildes denote residuals. Any component of cost that moves with the residual subsidy biases the estimate. When the identifying variation comes from timing, $\tilde{c}$ contains the part of the cost trend that the time controls do not absorb. Costs that fall as the subsidy falls ($\mathrm{Cov} > 0$) bias $\hat{\rho}$ down, and costs that rise as the subsidy falls ($\mathrm{Cov} < 0$) bias it up.

The two programs give opposite signs. Under Solar Carve-Out, the SREC price path and installation prices fell together from 2014 to 2018. Under SMART, the rate fell throughout while gross prices rose after 2020 (Figure 3). Equation (1) therefore predicts the sign of the bias for every design that relies on timing: downward under Solar Carve-Out and upward under SMART. Narrowing the window around a rule-based change in the subsidy shrinks the cost trend within the window but does not remove it, and windows narrow enough to remove it contain few projects and are exposed to bunching before the change. Demand shifts point the same way under SMART: blocks fill faster when demand is strong, so strong demand raises the gross price while moving buyers to a lower block, which also biases the estimate up.

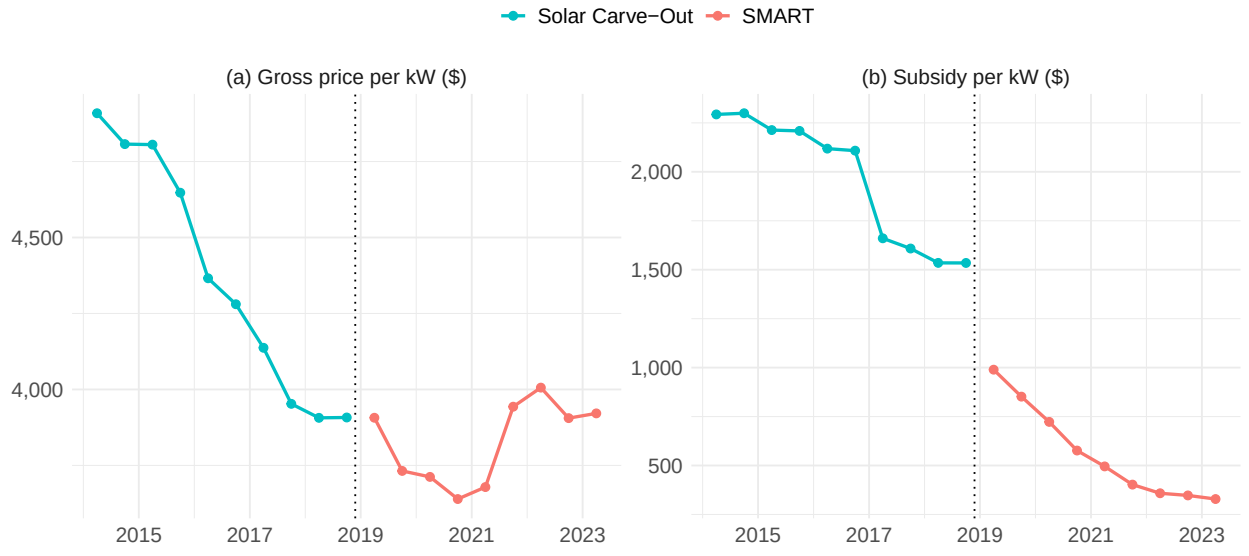


Figure 3: Gross Price and Subsidy per kW over Time

Half-year means of the gross price per kW and the subsidy per kW (present value at an 8% discount rate) by program, without controls. Half-years with fewer than 100 projects are omitted. The dotted line marks the start of SMART.

3.3 The Cross-Roof Design

The preferred specification removes the timing of the subsidy. It replaces year fixed effects with year-month fixed effects (utility-by-year-month for SMART), which absorb the SREC price path and the timing of SMART block changes, adds installer-by-year and city-by-year fixed effects, and controls for the log of system size and five bins of the DC-to-AC ratio. The fixed effects are additive, so pass-through is identified from differences in the subsidy that remain after removing installer-by-year, town-by-year, and month effects, which come mainly from differences in modeled generation. A stricter version replaces installer-by-year with installer-by-month fixed effects, so that only roofs equipped by the same installer in the same month are compared. More than 90% of projects share an installer and month with another project, and the estimates are nearly identical (Section 3.7). Under SMART, part of the remaining variation also comes from differences in incentive rates among projects that began operating in the same month, because the rate is locked at approval, which is usually months before operation. A further specification instruments the subsidy with the project's modeled capacity factor, so that only the variation that comes from the roof's orientation, tilt, and location identifies pass-through.

Applied to this variation, equation (1) shows what remains a concern. With the timing of the subsidy absorbed, $\tilde{c}$ contains only cost differences across roofs that are correlated with generation, such as a flat roof that needs tilted racking. These differences are the threat to the cross-roof design, and Section 4 examines them.

The two designs estimate different objects. The naive specification uses changes in the subsidy common to all buyers, whereas the cross-roof design uses differences across buyers, and the pass-through of a market-wide change can differ from that of a change specific to one buyer (Muehlegger and Sweeney, 2022). Section 6.1 shows that, with log-concave demand and marginal cost that does not fall with volume, market-wide pass-through cannot exceed the pass-through of a roof-specific subsidy, provided that conduct and the curvature of demand are the same for the two kinds of change. Two further issues remain. The discount rate affects both designs: it enters both the dependent variable and the regressor, so the coefficient depends on it mechanically, and I report a range of rates (Appendix B.1). Anticipation affects the timing-based designs: because installers know the subsidy schedule, such a design cannot separate the response of prices to a subsidy cut from pricing that anticipated the cut.

Inference. I cluster standard errors at the installer level. Installers set prices, the same installer appears in many projects and years, and an installer's pricing errors are likely to be correlated across its projects. There are 385 installers in the Solar Carve-Out sample

and 267 in the SMART sample, which is enough for standard asymptotic inference.[15] To compare the programs, I estimate both in one stacked regression in which the slope, the controls, and every fixed effect are specific to each program, and I cluster by installer across programs. The point estimates equal those from separate regressions, and the standard error of the difference accounts for the 83 installers that operate in both programs, who installed 56% of Solar Carve-Out projects and 69% of SMART projects.

3.4 Measurement Error in the Subsidy

The subsidy is built from modeled generation, so it is measured with error, and the error enters both sides of the regression, because the dependent variable subtracts the measured subsidy. Write the measured subsidy as $\hat{S}_i = S_i + u_i$. With classical error, $\hat{\rho} \to \lambda\rho + (1 - \lambda)$, where $\lambda = \text{Var}(S)/[\text{Var}(S) + \text{Var}(u)]$, so error pulls the estimate toward one rather than toward zero.[16] Errors in the inputs, such as a misrecorded tilt or azimuth or a geocoding error, are of this kind. Losses that PVWatts does not model and that are unrelated to its inputs, such as shading from a neighbor's tree, are not: if true generation equals modeled generation plus an independent error, the estimate is unbiased. An error common to all projects matters little. If PVWatts overstated every project's generation by a factor k, the estimate would converge to $1 - (1 - \rho)/k$, and even the full statewide gap of $k \approx 1.09$ would turn a true rate of 0.84 into 0.85.

The pull toward one is stronger in the cross-roof design than in the naive design, because roofs within the same cells differ less in true generation, so error makes up a larger share of the remaining variation. Error can move an estimate toward one but not past it, whereas the upward timing bias under SMART can. Under Solar Carve-Out, removing the downward timing bias and the pull of error both raise the cross-roof estimate above the naive one, so the gap between the two cannot be attributed to either alone. Section 5.1 uses these directions to read the gap between the naive and cross-roof estimates, and Section 4.3 uses the installers' production estimates to gauge the size of the error.

3.5 Cross-Roof Estimates

Table 3 reports pass-through estimates from specifications that progressively remove the variation in the subsidy that comes from its timing. Column (1) is the naive specification. Column (2) replaces year fixed effects with year-month fixed effects, interacted with utility

[15] Appendix Table 12 reports the standard errors under alternative clustering; utility clustering has too few clusters in the SMART sample, and the conclusions do not change.

[16] Equivalently, in a regression of the gross price on the measured subsidy, the coefficient $1 - \rho$ is attenuated toward zero.

for SMART. Column (3) adds installer-by-year and city-by-year fixed effects. Column (4) adds the log of system size in kW DC and indicators for five bins of the DC-to-AC ratio and is the preferred specification, because the price per kW falls with system size and depends on how much inverter capacity each kW of panels needs, and both can differ across roofs that generate different amounts. Column (5) instruments the subsidy with the project's modeled capacity factor in the specification of column (4).[17]

Table 3: Pass-Through from Cross-Roof Variation in Generation

	Dependent variable: net price per kW				
	(1)	(2)	(3)	(4)	(5)
Panel A: Solar Carve-Out					
Subsidy per kW	−0.756***	−0.799***	−0.826***	−0.863***	−0.856***
	(0.044)	(0.046)	(0.044)	(0.044)	(0.050)
Panel B: SMART					
Subsidy per kW	−1.017***	−0.887***	−0.839***	−0.847***	−0.894***
	(0.065)	(0.062)	(0.070)	(0.073)	(0.115)
Difference (SMART − Solar Carve-Out)	0.260***	0.087	0.013	−0.016	0.038
	(0.090)	(0.080)	(0.080)	(0.085)	(0.114)
Time fixed effects	Year	Year-month	Year-month	Year-month	Year-month
Installer × year, city × year			✓	✓	✓
System size, DC/AC ratio				✓	✓
Subsidy instrumented with capacity factor					✓

Panels A and B report the coefficient on subsidy per kW. Pass-through is minus this coefficient. The row below the panels reports the difference in pass-through between the programs. All columns include installer, inverter, manufacturer, and city fixed effects (and utility fixed effects for SMART), or their interactions as listed. Discount rate 8%. Observations: 25,491 Solar Carve-Out (23,503 in columns 4 and 5) and 14,317 SMART. Standard errors clustered by installer in parentheses. The difference and its standard error come from one stacked regression with program-specific slopes, controls, and fixed effects, clustered by installer across programs. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

In the preferred specification, column (4), pass-through is 0.86 under Solar Carve-Out and 0.85 under SMART. The difference between the programs is −0.02 with a standard error of 0.085. Both estimates are high but significantly below one. With modeled generation, buyers capture about 85 cents of each dollar of subsidy, whether the subsidy is paid through certificates sold at an uncertain market price or through a fixed rate known at approval. Columns (3) and (5) give the same picture, with differences of 0.01 and 0.04. Instrumenting with the capacity factor, which removes the residual variation in SMART rates, gives 0.89 under SMART. Measurement error in modeled generation pulls both levels toward one, so true pass-through is, if anything, lower, and Section 4.3 gauges by how much.

[17]The first-stage Wald statistics are 3,661 (Solar Carve-Out) and 1,298 (SMART). Columns (4) and (5) drop 1,988 Solar Carve-Out projects with missing AC capacity.

The estimates are in the range of prior estimates for solar rebate programs, 78% to close to 100% for host-owned systems (Pless and van Benthem, 2019; Gillingham and Tsvetanov, 2019; Dong, Wiser and Rai, 2018). Those estimates rely on the timing of the subsidy, and Section 5 shows that in Massachusetts such estimates move by more than 0.1 in either direction depending on how costs moved relative to the subsidy.

3.6 What the Comparison Rules Out

The comparison excludes large differences between the programs. In the preferred specification, the 95% confidence interval for the difference between the programs is −0.18 to 0.15. It excludes the naive gap of 0.26, so the cross-roof estimates reject the difference that a comparison of the two programs over time would suggest. Column (3), the specification with tilt bins (Table 4), and the specification with installer-by-month cells give similar intervals. The instrumental variables estimates in column (5) and the estimate identified from azimuth alone give point differences of 0.04 and 0.09, but they are less precise, and their intervals reach the naive gap.[18]

The programs differ in the form of the payment, in its predictability, and in its size (Section 2). Each difference could change pass-through, and the data detect no change from any of them. A buyer receives about the same share of a dollar of subsidy whether it arrives through a certificate market or through a fixed rate, and whether the subsidy covers 45% or 15% of the price. Because the two estimates are nearly identical, a single pooled slope summarizes both programs at 0.86. I estimate them separately because pooling would impose the equality that the comparison tests (Appendix B.4).

3.7 Robustness

Appendix B reports further checks, all in the preferred specification. Comparing only roofs equipped by the same installer in the same month, with installer-by-month in place of installer-by-year fixed effects, gives 0.860 under Solar Carve-Out and 0.859 under SMART, a difference of −0.002, and installer-by-town-by-year fixed effects give 0.853 and 0.848 (Appendix Table 14). The Solar Carve-Out estimate barely moves with the assumed SREC price path, from 0.844 with the auction price to 0.878 with the ACP. Both estimates fall as the discount rate rises from 3% to 15%, from 0.890 to 0.821 under Solar Carve-Out and from 0.879 to 0.796 under SMART, and both stay significantly below one at every rate.

[18] Accounting for installers that operate in both programs barely changes the standard error of the difference in the preferred specification (0.085, compared with 0.084 if the programs were independent) but raises it in the naive specification, from 0.078 to 0.090. The naive gap remains significant ($p = 0.004$).

The change is a mechanical rescaling of the subsidy. Adding utility fixed effects to the Solar Carve-Out model leaves the estimate essentially unchanged, and clustering by town or utility instead of by installer does not change the conclusions. SMART pass-through does not vary with ZIP-code income or installer concentration (Appendix B.7).

4 Threats to the Cross-Roof Design

The cross-roof design rests on the assumption that, within an installer, town, and period, generation affects the gross price only through the subsidy. This section examines four threats. Two concern the assumption itself: roofs that generate more could cost more or less to install, and a productive roof could be worth more to the buyer for reasons other than the subsidy, so that installers price the roof's value rather than its subsidy. The other two leave the assumption intact and affect what the estimates measure: error in modeled generation and, under SMART, variation in rates within a month.

4.1 Roof Traits and Installation Cost

Roofs that generate more could cost more, or less, to install, in ways that system size and the DC-to-AC ratio do not capture, such as pitch, access, or racking. If installation cost changes by c dollars per kW for each percentage point of capacity factor, and the subsidy rises by s, equation (1) gives $\hat{\rho} = \rho - c/s$. Within the cells of the preferred specification, the subsidy rises by \$131 per kW per percentage point under Solar Carve-Out and by \$40 under SMART, so a cost difference common to both programs would move the SMART estimate about 3.3 times as much as the Solar Carve-Out estimate and would open a gap between them.

Within cells, the capacity factor varies mainly with the direction the roof faces. Azimuth bins explain 83% (Solar Carve-Out) and 89% (SMART) of the residual variance of the capacity factor, and tilt bins 9% and 3%. Azimuth determines where the panels point, not how much work the installation takes. Tilt can change cost: flat roofs need tilted racking and steep roofs need more safety work. Appendix Table 16 reports the gross price by tilt bin, holding the subsidy fixed. The only sizable difference is for nearly flat roofs under Solar Carve-Out, which cost an average \$372 more per kW than roofs pitched at 25–30 degrees and generate less. Because those roofs cost more and generate less, they bias the estimate up.

Table 4 removes these differences. Holding tilt fixed with five-degree bins lowers the Solar Carve-Out estimate from 0.86 to 0.82, as the flat-roof premium predicts, and leaves

SMART at 0.85. The difference between the programs becomes 0.03, with a 95% interval of −0.13 to 0.18. Identifying pass-through only from azimuth, by instrumenting the subsidy with azimuth bins, gives 0.76 under Solar Carve-Out and 0.86 under SMART, a difference of 0.09 that is not significant. I keep column (4) as the preferred specification because it is the same specification for both programs and the most precise. Under SMART, pass-through is 0.85–0.86 across specifications. Under Solar Carve-Out, the estimates range from 0.76 to 0.86 depending on which roof traits identify pass-through. The flat-roof premium accounts for about 0.04 of the column (4) estimate, and the azimuth-only value is what one would expect if part of the remaining variation in modeled generation were error in the other inputs, such as tilt and location: azimuth is recorded directly on the application, and instrumenting with it removes error in the other inputs that is unrelated to azimuth. The conclusion that the two programs pass through the same share holds in every specification, but the Solar Carve-Out level is less certain than its standard error suggests. Section 4.3 shows that error of this size is plausible.

Table 4: Cross-Roof Pass-Through and Roof Geometry

	Dependent variable: net price per kW		
	Cross-roof (preferred) (1)	Tilt bins added (2)	Azimuth only (IV) (3)
Panel A: Solar Carve-Out			
Subsidy per kW	−0.863***	−0.820***	−0.762***
	(0.044)	(0.035)	(0.039)
First-stage F statistic			2942.6
Observations	23,503	23,503	23,503
Panel B: SMART			
Subsidy per kW	−0.847***	−0.845***	−0.855***
	(0.072)	(0.072)	(0.121)
First-stage F statistic			253.6
Observations	14,317	14,317	14,317
Difference in pass-through (SMART − SCO)	−0.016	0.026	0.093
	(0.085)	(0.080)	(0.116)

Column (1) is column (4) of Table 3. Column (2) adds five-degree tilt bins, specific to each program. Column (3) instruments the subsidy with fifteen-degree azimuth bins in the specification of column (2), so that only azimuth identifies pass-through. All columns come from one stacked regression with program-specific slopes, controls, and fixed effects, clustered by installer across programs. The first-stage F statistics are computed under homoskedasticity. Discount rate 8%. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

4.2 Pricing the Roof

A roof that generates more earns more than a larger subsidy. It also saves more on the electricity bill, because net metering credits each kWh at the retail rate, and these savings are proportional to the same modeled generation. Over a system's life they are the larger part of the roof's value. At the discount rate used here, ten years of the median SMART rate are worth about \$0.43 per kWh of annual output, whereas twenty-five years of retail savings at \$0.22–0.30 per kWh, the range of Massachusetts residential rates over 2019–2023, are worth about \$2.20–3.00, so bill savings are several times the SMART subsidy per kWh and comparable to the Solar Carve-Out subsidy.[19] Installers build quotes around the payback on total savings. If they priced the roof's total value, the gross price would rise with generation for a reason unrelated to the subsidy, and the regression would attribute the whole gradient to the subsidy. A second channel points the same way: households that install on an east- or west-facing roof are a selected group, and if they differ in price sensitivity from buyers with south-facing roofs, the gradient of the price in generation mixes pricing of the subsidy with differences in markups across buyer types.

Both channels predict a gross-price gradient of similar dollar size in the two programs, because bill savings per kWh were similar in the two periods, or higher under SMART as retail rates rose, and buyer selection on roof orientation does not depend on the program. Pricing of the subsidy predicts instead a gradient proportional to the subsidy gradient, which is 3.3 times larger under Solar Carve-Out. Table 5 gives the test. Per percentage point of capacity factor, the gross price rises by about \$19 under Solar Carve-Out and \$4 under SMART, with standard errors of about \$7 and \$5, while the subsidy rises by \$131 and \$40. The SMART gradient is a fifth of the Solar Carve-Out gradient, in line with the ratio of the subsidy gradients, whereas pricing of bill savings or selection on roof type would make it at least as large. The difference between the two gradients is about \$15 per percentage point with a standard error of about \$8, so equal gradients are rejected at the 10% level in a two-sided test and at the 5% level against the one-sided alternative that these channels imply. The evidence is that installers price the roof's subsidy, and only a little, rather than its value. The test assumes that installers priced roof value the same way in both periods, and the standard error of the SMART gradient (about \$5) limits its power.

[19] Residential rates are the U.S. Energy Information Administration's annual averages for Massachusetts. The present values use the 8% discount rate and 0.5% degradation of Section 2.3 and a 25-year system life. Under Solar Carve-Out, the mean subsidy of \$1,923 per kW is about \$1.50 per kWh of annual output at the mean capacity factor, against retail savings of about \$1.70–2.20 at the \$0.17–0.22 rates of 2014–2018.

Table 5: Price Dispersion and Pricing on Roof Productivity

	Solar Carve-Out	SMART
Mean gross price per kW ($)	4,270	3,796
Residual SD of gross price per kW ($)	485	521
within installer and year-month ($)	447	488
SD of modeled capacity factor (pp)	1.04	1.45
Subsidy gap for a 1-SD more productive roof ($/kW)	137	59
Gross-price gap for a 1-SD more productive roof ($/kW)	20	6
	(7)	(7)
Gross-price gap as % of residual SD	4.1%	1.2%

Residual standard deviations of the gross price per kW after installer-by-year, city-by-year, year-month (utility-by-year-month for SMART), inverter, and manufacturer fixed effects, log system size, and DC-to-AC ratio bins. The within installer and year-month row uses projects that share an installer and month with at least one other project. The gaps are the coefficients on the modeled capacity factor (percentage points) in regressions of the subsidy per kW and of the gross price per kW with the same controls, multiplied by the standard deviation of the capacity factor. Discount rate 8%. Standard errors clustered by installer in parentheses.

4.3 Measurement Error in Modeled Generation

As Section 3.4 shows, error in modeled generation pulls the estimates toward one, and more so in the cross-roof design. Instrumenting the subsidy with the modeled capacity factor (column 5 of Table 3) does not remove this, because the instrument comes from the same PVWatts run.

A second measure of generation gives a sense of how large this pull is. Each SMART application reports an estimated annual production for the system, which the applicant, usually the installer, computes with its own design tools; I call it the installer's production estimate. I drop the 1% of projects whose estimate is less than half or more than one and a half times the PVWatts value, which are clear reporting errors such as an annual production of 6 kWh for a 15 kW system; for the remaining projects, the median estimate is 6% below the PVWatts value. Across all projects, the subsidies per kW computed from the two measures have a correlation of 0.97, mainly because both use the SMART rate. Within the cells of the preferred specification, however, the two capacity factors have a correlation of only 0.57 (Table 6). If the errors in the two measures are independent, the share of the within-cell variation in the modeled capacity factor that is true variation is 0.57, and for the modeled subsidy, which also varies with the SMART rate, it is 0.80; if the errors are positively correlated, for example because both measures use the same recorded tilt and azimuth, these shares are upper bounds. Under either assumption, at least two-fifths of the remaining variation in the modeled capacity factor, and at least a fifth of that in the modeled

subsidy, is therefore error. Computing the subsidy from the installer's estimate instead of PVWatts gives SMART pass-through of 0.77, compared with 0.85 with PVWatts on the same projects, and instrumenting the PVWatts subsidy with the installer-estimate subsidy gives 0.74. Complete pass-through is rejected in every specification.

PVWatts remains the main measure, for three reasons. First, the installer's estimate is prepared for the sale, with tools and loss assumptions that differ across installers, and its error can be correlated with the price, which would bias both estimates. An installer that overstates output to make a system look attractive may also charge more for it, which would push both estimates down. Second, the instrumental variables estimate requires the errors of the two measures to be independent, which fails if both rely on the same recorded tilt and azimuth, and it requires the installer's error to be unrelated to the price. Third, Solar Carve-Out records include no production estimate, so the alternative measure cannot be built consistently for both programs. PVWatts applies one model and one set of loss assumptions to every roof in both programs, which is what the comparison between them requires.

The sensitivity check therefore applies to the level of pass-through, not to the comparison between the programs. Because this error pulls the estimates toward one, true pass-through is, if anything, below 0.85, which strengthens the conclusion that it is incomplete. Because both subsidies are built from the same model, the error has the same source in both programs. Under Solar Carve-Out, the estimate identified from azimuth alone, which is less exposed to error in the other inputs, is 0.76, as reported in Table 4, close to the SMART values with the installer's estimate. Appendix D.4 shows how the bound on market power changes with the amount of error.

4.4 Rates Locked at Approval

Under SMART, the rate is locked when the application is approved, while the data record the month the system began operating. Projects that began operating in the same month in the same utility can therefore carry different rates, so part of the SMART variation in column (4) comes from rates rather than generation. Within the cells of column (4), about two-thirds of the remaining variance of the SMART subsidy comes from differences in rates and about one-third from differences in modeled generation.[20] These rate differences reflect when a project applied, which could be correlated with its cost. Column (5) of Table 3 uses

[20]Splitting the residual subsidy into a rate component (the project's rate at the mean generation per kW) and a generation component (the mean rate at the project's generation per kW) gives variance shares of 0.64 and 0.28, with a covariance term near zero. Adding utility-by-rate fixed effects to column (4) leaves 35% of the residual variance. Low-income rate adders are not a source of this variation: about 40 of the 45,797 records in the SMART qualified-units file, which covers all sectors, carry a low-income off-taker adder.

Table 6: Modeled Generation and Installers' Production Estimates under SMART

Panel A: Agreement of the two measures (SMART)	Capacity factor	Subsidy per kW
Correlation, all projects	0.472	0.965
Correlation within cells of column (4)	0.565	0.831
Reliability of PVWatts within cells, λ	0.570	0.795
Panel B: SMART pass-through by measure of generation		
	$\hat{\rho}$	SE
PVWatts (column 4, validation sample)	0.846	(0.076)
Installer's production estimate	0.773	(0.069)
PVWatts, instrumented with the installer's estimate	0.739	(0.079)

Panel A compares PVWatts output with the estimated annual production reported on each SMART application, for the 14,196 projects (99%) whose estimate lies within 50% of the PVWatts value. Within-cell values use residuals after the fixed effects and controls of column (4) of Table 3. $\lambda = \text{Cov}(r_{PV}, r_I)/\text{Var}(r_{PV})$, where r_{PV} and r_I are the residuals of the PVWatts and installer measures, is the share of the within-cell variation in the PVWatts measure that is true variation if the errors of the two measures are independent. The subsidy per kW varies with the SMART rate as well as with generation. In Panel B, the installer-estimate subsidy uses the same rate and present-value factor with the installer's estimate of generation. The instrumental variables estimate instruments the PVWatts subsidy with it (first-stage Wald statistic 1,371). Both alternative estimates are sensitivity checks: the installer's estimate is prepared for the sale and its error may be correlated with the price. Standard errors clustered by installer in parentheses. Discount rate 8%.

only the variation in the subsidy that comes from the modeled capacity factor and gives 0.89, close to column (4), so the point estimate does not depend on the rate variation. Part of the precision does. Column (5) does not reject complete pass-through, and the gross-price gap for a more productive roof in Table 5 is \$6 with a standard error of \$7. Under SMART the paper therefore has one estimate identified from roofs alone, 0.89, and one more precise estimate, 0.85, part of whose precision comes from rate variation within a month; the rejection of complete pass-through at conventional levels rests on the latter.

5 Subsidy Timing and Cost Trends

Equation (1) predicts that any estimate identified from the timing of the subsidy is biased down under Solar Carve-Out and up under SMART. This section tests the prediction with every timing-based estimate the data allow: the naive specification, a regression discontinuity at the cut in the SREC factor, a stacked event study around the step-downs in SMART rates, and alternative treatments of the price level. Appendix Table 20 collects them, and Figure 4 plots them against the length of the period over which each design compares projects. The shaded band is the range of cross-roof estimates across the specifications of Section 4. Under

SMART, the naive estimate and the event study lie above the band and at or above one. Under Solar Carve-Out, the estimates scatter on both sides of the band, and the regression discontinuity falls as its window widens.

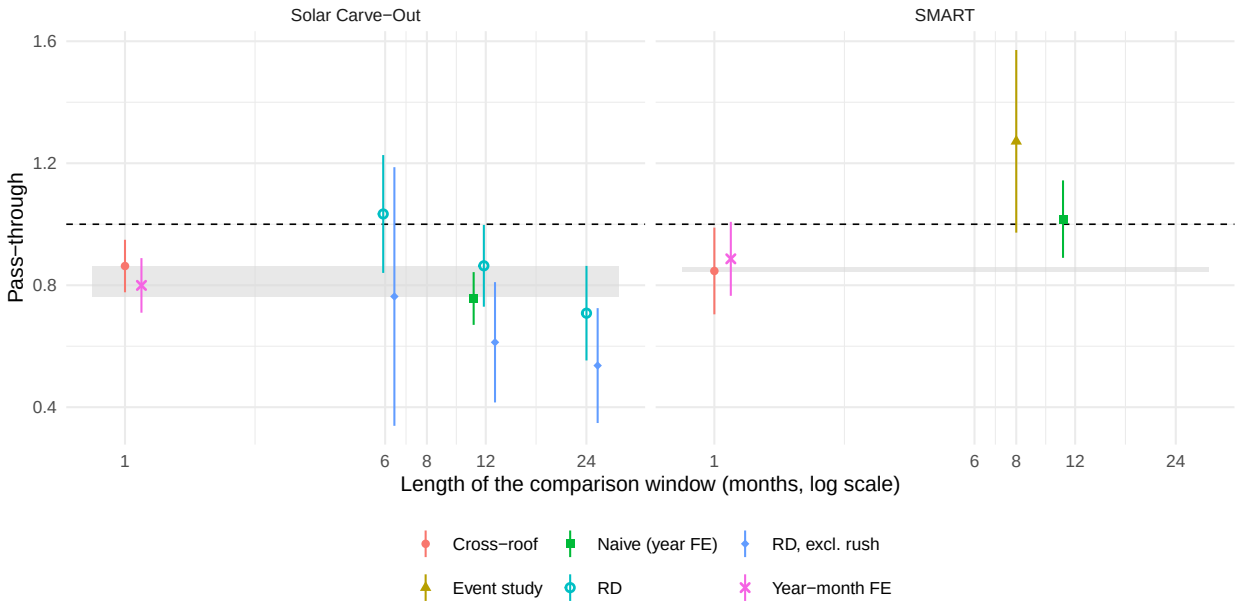


Figure 4: Pass-Through and the Length of the Comparison Window

Each point is an estimate from Appendix Table 20, plotted against the length of the period over which the design compares projects (log scale), with a 95% confidence interval. Cross-roof and year-month estimates compare projects within a month, the event study within eight months, the regression discontinuity within ±90, ±180, and ±365 days, and the naive model within a year. The shaded band is the range of cross-roof estimates: column (4) of Table 3, tilt bins, azimuth only, and installer-by-month cells. The dashed line marks complete pass-through.

5.1 The Naive Gap

Column (1) of Table 3 gives pass-through of 0.76 under Solar Carve-Out and 1.02 under SMART, a gap of 0.26. Complete pass-through is rejected for Solar Carve-Out but not for SMART. The Solar Carve-Out estimate is insensitive to controls, while the SMART estimate rises from 0.91 to 1.02 as installer and city fixed effects are added (Appendix Table 21). Taken at face value, the naive estimates would suggest that a fixed-rate subsidy is passed through more fully than a certificate subsidy.

The SMART base rate is not tied to installation costs: it steps down on a fixed schedule as cumulative capacity fills the blocks, and the incentive also falls when the value of energy rises. Because blocks fill over time, however, the subsidy changes during the same period in which installation costs change, and the SREC price path co-moves with market conditions.

As Section 3.2 notes, the Solar Carve-Out subsidy and gross price fell together from 2014 to 2018 (Figure 3). Under SMART, the subsidy fell from $989 per kW in the first half of 2019 to $329 in the first half of 2023, while the gross price fell until 2020 and then rose from $3,639 to $4,006 per kW by the first half of 2022, as consumer prices rose (half-year means in Figure 3). The same pattern holds within years. After removing year means, the monthly mean gross price and subsidy are positively correlated under Solar Carve-Out and negatively correlated under SMART. These are the signs of $\mathrm{Cov}(\tilde{c}, \tilde{S})$ in equation (1) behind the predictions of Section 3.2, and they match the directions in which the naive estimates differ from the cross-roof estimates.

Year-month fixed effects alone do most of the work: the gap is already down to 0.09 in column (2), before any of the other controls are added. As the timing of the subsidy is absorbed, the Solar Carve-Out estimate rises from 0.76 to 0.86 and the SMART estimate falls from 1.02 to 0.84–0.89. Most of the naive gap therefore reflects how each subsidy moved over time relative to costs, not how the two designs divide the subsidy. The SMART movement is the more informative. The fall in the SMART estimate goes against the pull of measurement error, so the naive SMART estimate is biased up by at least 0.17, and by more if error inflates the cross-roof estimate. The rise in the Solar Carve-Out estimate goes in the direction measurement error would also produce, and the cross-roof estimate identified from azimuth alone, 0.76, equals the naive estimate. For Solar Carve-Out, the comparison therefore does not separate the timing bias from the pull of error, even though costs fell with the subsidy. The naive gap between the programs comes mainly from the SMART estimate.

Three further results point the same way for SMART. The naive SMART estimate is 0.91 with the CPI as a control, 0.93 with year-by-month fixed effects, and 0.85 in CPI-deflated dollars, while the Solar Carve-Out estimate stays between 0.73 and 0.80 (Appendix C.4). Each alternative removes part of the co-movement between the subsidy and costs, and each moves the SMART estimate toward the cross-roof estimate. By quartile of the subsidy, the naive SMART estimates are 1.55 and 1.32 in the two lowest quartiles, where the subsidy is smallest and the projects are latest (Appendix Table 25). And in the naive specification SMART pass-through falls with ZIP-code income, from 1.09 to 0.91 ($p = 0.011$), a gradient that arises if adoption timing differs by income and that disappears once the timing of the subsidy is absorbed (Appendix B.7).

5.2 Regression Discontinuity at the SREC Factor Cut

On January 8, 2017, the SREC factor for new small systems fell from 1.0 to 0.8 (Section 2), so the subsidy per kW falls discretely without any change in the installer's costs: from

about $2,110 in December 2016 to about $1,610 from February 2017 (Figure 5), and every project operating after the cutoff has a factor of 0.8, with a single exception. A regression discontinuity in time at this cut is the design one would use to avoid trends. I use the commercial operation date as the running variable, with a linear trend on each side of the cutoff and installer, inverter, manufacturer, and city fixed effects, and I instrument the subsidy with the post-cutoff indicator. Two features of the raw data complicate the comparison. First, gross prices fall throughout 2016–2017, by roughly $230 per kW per year. Second, the number of projects operating in December 2016 (879) is well above that in each of the eleven preceding months (358–647), which suggests a rush to complete before the deadline, and December projects have unusually low gross prices (about $4,180 per kW, compared with $4,240–$4,380 in the preceding months). I therefore report estimates with and without the two months around the cutoff.

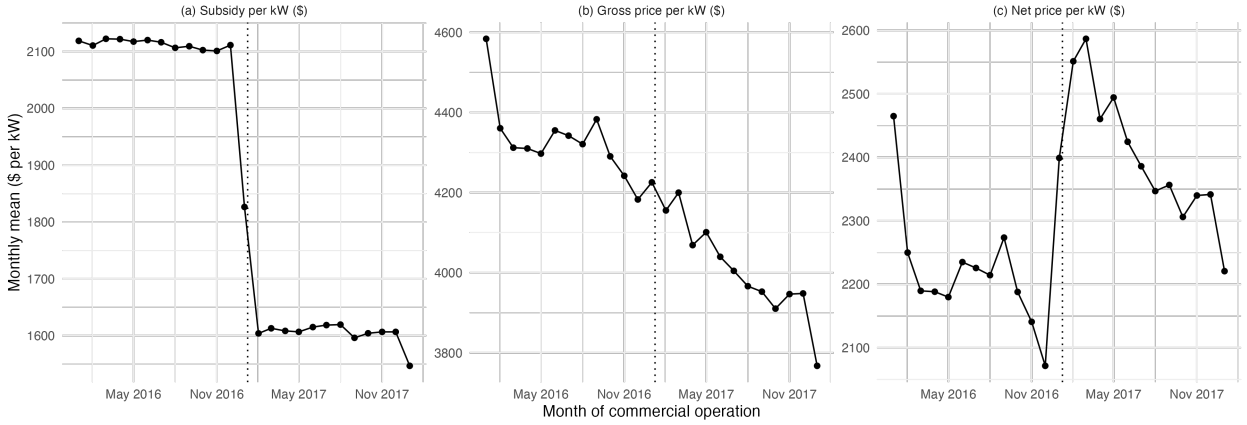


Figure 5: Subsidy, Gross Price, and Net Price by Month of Operation around the SREC Factor Cut

Solar Carve-Out projects within a year of the cut. Each point is a monthly mean in dollars per kW of DC capacity, by month of commercial operation: the subsidy (present value at an 8% discount rate, midpoint-price measure), the gross price, and the net price. The dotted line marks the January 8, 2017 cut in the SREC factor from 1.0 to 0.8.

Appendix Table 22 reports the estimates. The subsidy falls by $465–$500 per kW in the midpoint-price measure and $420–$435 in the factor-only measure, which holds the SREC price path fixed and isolates the change caused by the factor. The estimate depends on the window. With all projects, pass-through falls from 1.03 at ±90 days to 0.86 at ±180 days and 0.71 at ±365 days. Excluding the two months around the cutoff, it falls from 0.76 to 0.61 and 0.54. The design fits a linear trend on each side of the cutoff, so widening the window changes the estimate only if the cost trend is not linear within the window or if the bunching in December 2016 distorts the fit near the cutoff. As the window widens, the estimates fall through and then below the naive 0.76, which is the direction equation (1) predicts when

the linear fit does not absorb the whole cost trend. The narrowest windows are imprecise and cannot reject the cross-roof estimates of 0.76–0.86, and the all-project estimate at ±90 days is pushed up by the low December prices. The discontinuity shows that a design built to avoid trends still gives an answer that moves with the cost trend in its window. It leaves the sign of the Solar Carve-Out bias open.

5.3 Event Study around SMART Step-Downs

The same logic applies to SMART: when a capacity block fills, the base rate for the utility falls by a known amount. I identify step-downs from the monthly modal SMART rate, a step being a month in which the modal rate falls by more than \$0.004 per kWh. Reading the rate this way requires many projects in each month, so I use the two utilities that account for about 92% of SMART projects, Eversource East (7,400 projects, 51 months with at least 15 projects) and National Grid (5,747 projects, 46 such months). I keep the nine steps that are at least four months from any other step, each of which lowers the rate by \$0.011–\$0.020 per kWh, and stack them to estimate the change in subsidy, gross price, and net price from four months before to three months after each step, relative to the month before it.[21] Because the rate is fixed at approval, usually months before operation, the subsidy does not jump at the step but declines across the window, from about \$51 per kW above its month −1 level four months before to \$88 below it three months after (Figure 6, panel a). If pass-through is complete, the gross price should not respond, and the net price should rise by the amount the subsidy falls. Panel (b) shows that the gross price does not break at the step: its event-time coefficients lie between −\$40 and +\$26 per kW, with confidence intervals of roughly ±\$40. Panel (c) shows that the net price rises as the subsidy falls.

Pooling the post-step periods, the subsidy is \$85 per kW lower after a step, the gross price \$23 higher, and the net price \$108 higher, which implies pass-through of 1.27 when the post indicator instruments the subsidy (Appendix Table 23). The estimate is above one, as the naive SMART estimate is, and in the direction equation (1) predicts for a period in which gross prices drifted up as the subsidy fell. Adding a linear trend in event time gives 1.65 with a standard error of 1.43, so once the gradual decline in the subsidy is absorbed by the trend, the design does not identify pass-through. The event study compares projects a few months apart rather than across a sharp break, rests on nine events in two utilities, and cannot rule out that some step-downs coincide with updates to the value of energy.

[21] Adding Eversource West, whose modal rate is noisy in months with few installations, gives twelve steps and an estimate of 1.26.

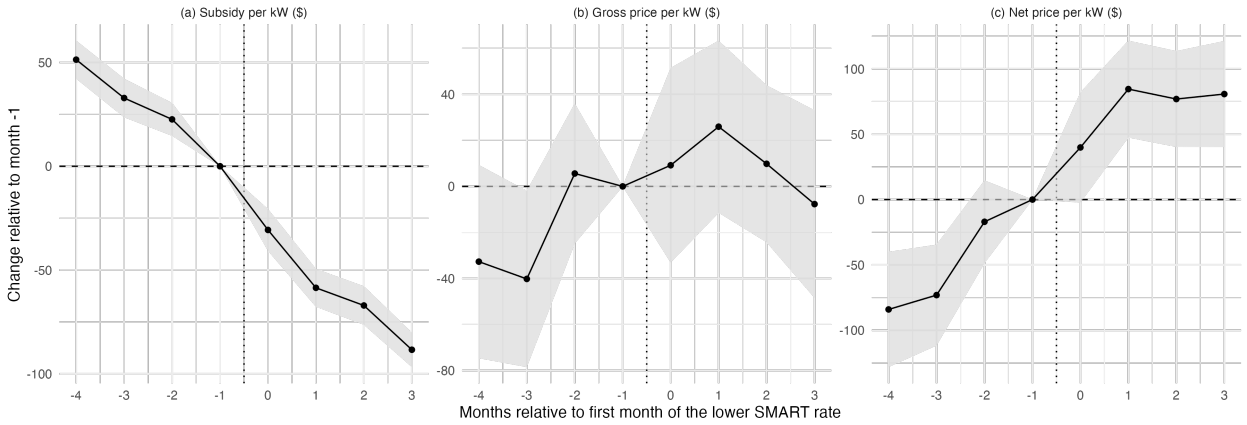


Figure 6: Event Study around SMART Rate Step-Downs

SMART projects in the stacked sample of nine step-downs (Appendix Table 23). Each point is the coefficient on an event-time indicator, in dollars per kW, relative to the month before the first month of the lower rate (event time −1), from regressions with event, installer, inverter, manufacturer, and city fixed effects; shaded areas are 95% confidence intervals with standard errors clustered by installer. The subsidy declines across the whole window rather than jumping at the step because the rate is locked at approval, months before operation. Discount rate 8%.

5.4 Implications for Existing Estimates

The timing-based designs and the cross-roof design estimate different objects, so a difference between them need not reflect bias. Under SMART, the timing-based estimates exceed the cross-roof estimate, and with log-concave demand and marginal cost that does not fall with volume, market-wide pass-through cannot exceed the pass-through of a roof-specific subsidy (Section 6.1). Estimates of 1.02 and 1.27 therefore cannot come from the difference between market-wide and roof-specific pass-through. Under Solar Carve-Out, the naive estimate lies within the range of cross-roof estimates (0.76–0.86), and the regression discontinuity estimates scatter on both sides of it, so there is no gap for this difference to explain. Appendix C.5 considers three further explanations for the naive gap, the size of the subsidy, its predictability, and competition, and finds that none of them can be separated from the timing of the two programs.

The solar pass-through estimates in the literature use the kind of variation that this section shows to be confounded in Massachusetts. Pless and van Benthem (2019) identify pass-through from jumps in CSI rebate levels relative to smooth time trends, with the steps timed differently across three utilities; Dong, Wiser and Rai (2018), Gillingham and Tsvetanov (2019), and Podolefsky (2013) identify from rebate step-downs or a tax-credit change with year effects or a polynomial trend. A trend absorbs the average movement of costs over the window, not the cost path at the dates when the subsidy changes, and three features of these programs put the two at odds. The steps are triggered by cumulative capacity, which

fills fastest when costs fall, so they land in the periods a global trend fits worst. Costs move in episodes, such as collapses in module prices, tariffs, and the 2021–22 inflation, rather than along a slope. And because installers watch the capacity counters, quotes before a step can already embed the lower subsidy, and the rush to beat a deadline changes who installs and at what price.

The Massachusetts estimates show what follows. The regression discontinuity fits a linear trend on each side of the SREC factor cut and still falls from 1.03 to 0.71 as its window widens, so the trend absorbs less of the cost path the more of it the window contains. The event study loses identification once a trend in event time is added. And the naive SMART estimate ranges from 0.85 to 1.02 depending only on how the price level is treated. A timing-based design can therefore obtain a wide range of answers from its choice of trend, and the data do not tell which is right. The same pattern appears in Pless and van Benthem (2019), whose host-owned estimate moves from 0.78 to 1.01 across the trend specifications, with the largest change from utility-specific trends, which absorb each utility's own cost and demand path.

The direction of any bias in their estimates is predictable. Under the California schedule, which fell as costs fell, equation (1) predicts that a trend that under-absorbs the cost decline biases the estimate down, so the true host-owned rate would, if anything, be higher than 0.78. Their identifying variation, unlike variation across roofs at the same time, does not separate how the subsidy is shared from how costs moved, and no choice of trend can be shown to do so from the data these designs use.

6 Implications for Market Power

High pass-through often serve as evidence of competition. This section asks what the cross-roof estimates reveal about installers' market power. The answer is that the market power the subsidy reveals is a limited amount of price discrimination, which under SMART bounds the markup installers earn by pricing the roof at 2–7% of price with modeled generation, and below 10% in 15 of the 18 combinations of pass-through and demand estimates in Table 7, and that the other sources of market power the data can measure leave little trace.

6.1 Framework

Weyl and Fabinger (2013) provide a framework for interpreting the size of pass-through. Consider a symmetric market in which installers' conduct is summarized by a parameter θ, which equals 0 under perfect competition and 1 under monopoly and does not change with

the subsidy. The pass-through rate of a per-unit subsidy that changes for all buyers at once is

$$\frac{1}{\rho_M} - 1 = r + \frac{\theta}{\varepsilon_{ms}}, \tag{2}$$

where $r \geq 0$ is the slope of marginal cost relative to the slope of inverse demand and ε_{ms} is the elasticity of the marginal consumer surplus, which measures the curvature of demand. The term $1/\varepsilon_{ms}$ is zero when demand is log-linear, positive when demand is log-concave (for example, linear), and negative when it is log-convex (for example, constant elasticity).[22] With log-concave demand, market power pushes pass-through below one, so a market-wide pass-through rate above one requires marginal cost that falls with volume.

The cross-roof estimates use a different kind of variation: differences in the subsidy across buyers of the same installer in the same town and month. A higher subsidy for one buyer does not change the installer's volume appreciably, so the supply term drops out, and the pass-through of a roof-specific subsidy is

$$\frac{1}{\rho_B} - 1 = \frac{\theta_B}{\varepsilon^B_{ms}},$$

where θ_B measures the installer's pricing power over individual buyers and ε^B_{ms} is the curvature of demand among buyers with similar roofs (Appendix E derives this expression from a pricing condition for each roof type). With log-concave demand, $r \geq 0$, and the same conduct and curvature for the two kinds of change, $\rho_M \leq \rho_B$. If installers charged one price per kW regardless of the roof, then $\rho_B = 1$, however large the installer's markup, and under perfect competition $\rho_B = 1$ with no markup at all. Pass-through below one therefore requires that the gross price rise with the roof's subsidy: the shortfall $1 - \rho_B$ is the share of a roof's subsidy that the installer captures by pricing the roof. Unlike a cost shock specific to one firm (Muehlegger and Sweeney, 2022), a roof-specific subsidy is faced by every installer that quotes the roof, so competition for the job already reflects it.

6.2 Price Discrimination

Pricing the roof is a form of price discrimination, and the conditions for it hold particularly well in this market. Installers quote each project individually, and they compute the roof's expected output themselves: every SMART application reports the installer's esti-

[22] Pless and van Benthem (2019) use the same expression with an additional term for conduct that changes with the subsidy. In California, pass-through of 153% to lessees together with convex demand revealed market power in the third-party-owned segment, while host-owned buyers received 78%; log-concave demand in Massachusetts turns pass-through below one into an upper bound on market power rather than evidence of it.

mate of the system's annual production (Section 4.3), so the installer knows the subsidy the buyer will receive. An installation cannot be resold. A production-based subsidy also makes such pricing easier than an upfront rebate does, because it varies with a characteristic the installer observes. An installer with market power would therefore have both the means and the incentive to charge more for roofs that earn more. Installers do set individual prices, but they price the roof only a little. After the controls of column (4) of Table 3, the gross price per kW varies with a standard deviation of \$447 under Solar Carve-Out and \$488 under SMART among buyers of the same installer in the same month, so prices are not set from a uniform schedule, yet a roof that is one standard deviation more productive, which earns \$137 and \$59 more subsidy per kW, is charged only \$20 and \$6 more, about 4% and 1% of that dispersion (Table 5); the ratio of the two gaps is $1 - \rho$ in column (5) of Table 3. Pass-through is also nearly the same on more and less productive roofs (Appendix Table 17), and Appendix F discusses the welfare effects of pricing by roof.

6.3 Bound Market Power under SMART

Turning the shortfall into a bound on conduct requires the curvature of demand. Following Pless and van Benthem (2019), I estimate demand at the ZIP-code and year level, with installations per 1,000 residents quadratic in the mean net price per watt, ZIP characteristics, and county and year fixed effects, and I instrument the price with the mean subsidy per watt and the county wage in building-equipment contracting (Appendix D.1). Under SMART, demand slopes down, with an elasticity at the mean price of −1.57 in the quadratic form and −2.16 in the log-quadratic form, and it is log-concave: the curvature $1/\varepsilon_{ms}$ is 5.38 with the quadratic in levels and 2.14 with the log-quadratic form, which rejects log-linear demand at the 10% and 5% levels. The bound below depends on the sign of the curvature. The subsidy instrument can be correlated with demand shocks, because blocks fill faster where demand is strong; county-by-year fixed effects leave the curvature almost unchanged, and an instrument built from roof geometry alone is too weak to be informative (Appendix D.2), so the bound is conditional on the demand estimates. Under Solar Carve-Out, demand is not identified, because the subsidy is driven by the statewide certificate price and varies little across ZIP codes within a county and year. I therefore draw conclusions only for SMART.

Because r is absent, $\theta_B = (1/\rho_B - 1)\,\varepsilon_{ms}$ holds with equality if all of the shortfall below one reflects pricing. It is an upper bound if any remaining cost differences make costlier roofs also more productive; costs that run the other way, such as the flat-roof premium of Section 4.1, and error in modeled generation (Appendix D.4) would loosen it. The Lerner index equals θ/ε_D. Table 7 reports the bound on markups for three values of ρ_B and six

demand estimates. With modeled generation ($\rho_B = 0.847$, Table 3, column 4), θ is at most 0.034 with the quadratic demand and 0.085 with the log-quadratic demand, so the markup from pricing the roof is at most 2.1% and 3.9% of price, and 4.1–7.0% with the flexible forms of Appendix D.3. If pass-through is instead measured with the installer's production estimate, at 0.77 or 0.74 (Section 4.3), the shortfall below one is half again as large or more, and the bound rises to 3.5–11.3% and 4.2–13.7% of price. The bound stays below 10% of price in 15 of the 18 combinations. The calculation uses the curvature of aggregate ZIP-level demand in place of the curvature among buyers with similar roofs. Sampling uncertainty widens the upper end: drawing ρ and $1/\varepsilon_{ms}$ from their estimated distributions at $\rho_B = 0.847$, the 95th percentile of the bound on θ is 0.15 with the quadratic demand and 0.35 with the log-quadratic demand, which correspond to markups of about 10% and 16% of price. The bound also assumes that conduct does not change with the subsidy and that marginal cost does not fall with volume; economies of scale would loosen it. Appendix D.4 shows how the bound changes with the share of error in modeled generation and gives the implied division of surplus.

Table 7: Upper Bound on Installer Markups under SMART

	Roof-specific pass-through ρ_B		
	0.847	0.773	0.739
Demand estimate	Modeled generation	Installers' estimates	
Quadratic, 2SLS	2.1	3.5	4.2
Log-quadratic, 2SLS	3.9	6.3	7.7
Quadratic, control function	5.5	8.8	10.7
Natural spline, 3 df	4.6	7.4	8.9
Natural spline, 4 df	4.1	6.6	8.0
Natural spline, 5 df	7.0	11.3	13.7

Entries are the upper bound on the markup in percent of price, $(1/\rho_B - 1)/(1/\varepsilon_{ms})$ divided by the absolute elasticity of demand, at the point estimates of the curvature and elasticity at the mean SMART price (Tables 26 and 28). ρ_B is column (4) of Table 3, with modeled generation, and the two sensitivity values of Table 6, with the installers' production estimates (OLS and instrumental variables).

6.4 Other Sources of Market Power

Price discrimination is one way installers could earn a markup. Others would leave different traces, and Appendix Table 19 tests their predictions within the preferred specification. If installers' pricing power over individual buyers, θ_B, were sustained by few competitors, cross-roof pass-through would be lower where competition is weaker and gross prices higher. If installers priced roofs more aggressively when busy, pass-through would fall with workload.

The discrimination on willingness to pay rather than on the roof would make pass-through vary with buyer income. None of these patterns appears. Pass-through does not fall when an installer is busy, although under SMART it is lower by 0.073 per standard deviation of statewide volume, so installers price the roof somewhat more when demand across the state is strong. SMART pass-through rises with town concentration, by 0.099 per standard deviation, the opposite of what pricing power sustained by concentration would predict, and does not change with concentration under Solar Carve-Out. SMART gross prices are lower where more installers are active, by $66 per kW per log point, or 1.7% of the mean price, which is small and may reflect demand, since installers enter towns where demand is growing. Across terciles of ZIP-level concentration and of ZIP-code income, SMART pass-through does not differ (Appendix B.7). These are correlations rather than causal estimates, because workload and the number of competitors respond to demand, and the concentration measures vary within a range of low concentration (Appendix C.5.3).

One source of market power pass-through can not detect is a markup common to all buyers in a market, such as rents that free entry dissipates through customer-acquisition costs. Under constant marginal cost and log-linear demand, such a markup is a fixed dollar amount that does not respond to the subsidy, and pass-through is one even though installers gain from the subsidy because they sell more systems at a positive markup. The evidence above shows that the market power the subsidy reveals is limited. Whether installers earn a markup common to all buyers is a question for cost data or a structural model of installer pricing.

The bound of Section 6.3 is derived for SMART. The two programs pass through the same share across roofs, and 83 installers, accounting for 56% of Solar Carve-Out projects, operate in both, so if conduct and the curvature of demand were similar in the two periods the bound would apply to Solar Carve-Out as well. Unfortunately the demand under Solar Carve-Out is not identified. Under Solar Carve-Out, prices are lower in concentrated town-years, which likely reflects the Solarize Mass campaigns, in which a town selected a single installer and offered tiered group prices that fell as more households signed up (Clean Energy States Alliance, 2015).

7 Conclusion

This paper measures the pass-through of production-based solar subsidies to the prices homeowners pay. The inherent difficulty is that subsidy schedules are built to fall as deployment grows, so the timing of the subsidy moves with installation costs which is rarely observed, and a regression on that timing cannot separate the two. I address this problem

by comparing roofs that the same installer equips in the same town and month, which earn different subsidies because the payment depends on each roof's generation. The two Massachusetts programs also provide a contrast as the subsidy fell together with installation costs under one and against them under the other.

I find that buyers receive about 85% of the subsidy under both the certificate-based and the fixed-rate program, although the two subsidies differ in form, in predictability, and by a factor of three in size. A comparison of the two programs over time shows a gap of 0.26 between them. That gap comes from the co-movement of each subsidy with installation costs, and in each program it has the sign that the co-movement predicts. Error in modeled generation pulls the estimates toward one, so the result is likely an overestimate. The lowest cross-roof estimate, with installers' production estimates, is 0.74.

Installers quote each job individually and have information on each roof's subsidy, yet they price the roof only a little. With log-concave demand, this bounds the markup that installers earn from price discrimination at 2–7% of price. Local concentration, the number of competitors, and buyer income leave little trace in pass-through or prices. A markup charged to every buyer alike is not captured by the estimates. Whether such markups make residential solar expensive is a question for data on installers' costs, and a natural direction for future work.

I focused on Massachusetts because its two programs offer the rare case of two subsidy designs in one market, with a payment that varies across buyers at a point in time. The approach extends beyond solar. Subsidy tied to a characteristic of the buyer, such as a rebate that varies with income or with the attributes of a home, provides differences across buyers at the same time. And any subsidy scheduled to fall as deployment grows, such as those for electric vehicles and batteries, exposes a timing-based estimate to the cost trend this paper documents.

References

Aguirre, Iñaki, Simon Cowan, and John Vickers. 2010. "Monopoly price discrimination and demand curvature." *American Economic Review*, 100(4): 1601–1615.

Bagnoli, Mark, and Ted Bergstrom. 2005. "Log-Concave Probability and Its Applications." *Economic Theory*, 26(2): 445–469.

Barwick, Panle Jia, Hyuk-soo Kwon, Binglin Wang, and Nahim Bin Zahur. 2023. "Pass-through of electric vehicle subsidies: A global analysis." *AEA Papers and Proceedings*, 113: 323–328.

Bistline, John, Neil R. Mehrotra, and Catherine Wolfram. 2023. "Economic implications of the climate provisions of the Inflation Reduction Act." *Brookings Papers on Economic Activity*, 2023(1): 77–182.

Borenstein, Severin. 2017. "Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates." *Journal of the Association of Environmental and Resource Economists*, 4(S1): S85–S122.

Burr, Chrystie. 2016. "Subsidies and investments in the solar power market." *University of Colorado at Boulder Working Paper*.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer. 2006. "$1,000 cash back: The pass-through of auto manufacturer promotions." *American Economic Review*, 96(4): 1253–1270.

Butters, R Andrew, Daniel W Sacks, and Boyoung Seo. 2022. "How do national firms respond to local cost shocks?" *American Economic Review*, 112(5): 1737–1772.

Cabral, Marika, Michael Geruso, and Neale Mahoney. 2018. "Do larger health insurance subsidies benefit patients or producers? Evidence from Medicare Advantage." *American Economic Review*, 108(8): 2048–2087.

Clean Energy States Alliance. 2015. "States Spur Solarize Successes." *https://www.cesa.org/states-spur-solarize-successes/*.

Colas, Mark, and Emmett Reynier. 2026. "Optimal Subsidies for Residential Solar." *Journal of Political Economy Microeconomics*, 4(1): 213–264.

Crago, Christine Lasco, and Ilya Chernyakhovskiy. 2017. "Are policy incentives for solar power effective? Evidence from residential installations in the Northeast." *Journal of Environmental Economics and Management*, 81: 132–151.

De Groote, Olivier, and Frank Verboven. 2019. "Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems." *American Economic Review*, 109(6): 2137–2172.

Deline, Chris, Dirk Jordan, Bill Sekulic, Josh Parker, Byron McDanold, and Allan Anderberg. 2024. "PV Lifetime Project-2024 NREL Annual Report." National Renewable Energy Laboratory (NREL), Golden, CO (United States).

Dong, Changgui, Ryan Wiser, and Varun Rai. 2018. "Incentive pass-through for residential solar systems in California." *Energy Economics*, 72: 154–165.

Fabinger, Michal, and E. Glen Weyl. 2014. "A Tractable Approach to Pass-Through Patterns." Working paper.

Fabra, Natalia, and Mar Reguant. 2014. "Pass-through of emissions costs in electricity markets." *American Economic Review*, 104(9): 2872–2899.

Fan, Ying, and Ge Zhang. 2022. "The welfare effect of a consumer subsidy with price ceilings: the case of Chinese cell phones." *The RAND Journal of Economics*, 53(2): 429–449.

Ganapati, Sharat, Joseph S Shapiro, and Reed Walker. 2020. "Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes." *American Economic Journal: Applied Economics*, 12(2): 303–342.

Genakos, Christos, and Mario Pagliero. 2022. "Competition and pass-through: evidence from isolated markets." *American Economic Journal: Applied Economics*, 14(4): 35–57.

Gerarden, Todd D. 2023. "Demanding innovation: The impact of consumer subsidies on solar panel production costs." *Management Science*, 69(12): 7799–7820.

Gillingham, Kenneth, and Tsvetan Tsvetanov. 2019. "Hurdles and steps: Estimating demand for solar photovoltaics." *Quantitative Economics*, 10(1): 275–310.

Gulati, Sumeet, Carol McAusland, and James M Sallee. 2017. "Tax incidence with endogenous quality and costly bargaining: Theory and evidence from hybrid vehicle subsidies." *Journal of Public Economics*, 155: 93–107.

Hughes, Jonathan E, and Molly Podolefsky. 2015. "Getting green with solar subsidies: evidence from the California solar initiative." *Journal of the Association of Environmental and Resource Economists*, 2(2): 235–275.

Hu, Yunyi, Haitao Yin, and Li Zhao. 2025. "Subsidy phase-out and consumer demand dynamics: Evidence from the battery electric vehicle market in China." *Review of Economics and Statistics*, 107(2): 458–475.

Kenkel, Donald S. 2005. "Are alcohol tax hikes fully passed through to prices? Evidence from Alaska." *American Economic Review*, 95(2): 273–277.

Kirwan, Barrett E. 2009. "The incidence of US agricultural subsidies on farmland rental rates." *Journal of political economy*, 117(1): 138–164.

Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod. 2016. "Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes." *American Economic Journal: Economic Policy*, 8(2): 251–286.

Kosonen, Tuomas. 2015. "More and Cheaper Haircuts after VAT Cut? On the Efficiency and Incidence of Service Sector Consumption Taxes." *Journal of Public Economics*, 131: 87–100.

Lade, Gabriel E, and James Bushnell. 2019. "Fuel subsidy pass-through and market structure: Evidence from the renewable fuel standard." *Journal of the Association of Environmental and Resource Economists*, 6(3): 563–592.

Matisoff, Daniel C, and Erik P Johnson. 2017. "The comparative effectiveness of residential solar incentives." *Energy Policy*, 108: 44–54.

Miklós-Thal, Jeanine, and Greg Shaffer. 2021. "Pass-through as an economic tool: On exogenous competition, social incidence, and price discrimination." *Journal of Political Economy*, 129(1): 323–335.

Muehlegger, Erich, and David S Rapson. 2018. "Subsidizing mass adoption of electric vehicles: Quasi-experimental evidence from California." National Bureau of Economic Research NBER Working Paper 25359.

Muehlegger, Erich, and Richard L. Sweeney. 2022. "Pass-Through of Own and Rival Cost Shocks: Evidence from the U.S. Fracking Boom." *Review of Economics and Statistics*, 104(6): 1361–1369.

Nakamura, Emi, and Dawit Zerom. 2010. "Accounting for incomplete pass-through." *The review of economic studies*, 77(3): 1192–1230.

Newey, Whitney K., James L. Powell, and Francis Vella. 1999. "Nonparametric Estimation of Triangular Simultaneous Equations Models." *Econometrica*, 67(3): 565–603.

Pless, Jacquelyn, and Arthur A. van Benthem. 2019. "Pass-through as a test for market power: An application to solar subsidies." *American Economic Journal: Applied Economics*, 11(4): 367–401.

Podolefsky, Molly. 2013. "Tax evasion and subsidy pass-through under the solar investment tax credit." *Working Paper*.

Ros, Agustin J, and Sai Shetty. 2023. "Residential rooftop solar demand in the US and the impact of net energy metering and electricity prices." *Energy Economics*, 118: 106491.

Sallee, James M. 2011. "The surprising incidence of tax credits for the Toyota Prius." *American Economic Journal: Economic Policy*, 3(2): 189–219.

U.S. Energy Information Administration. 2023. "Federal Financial Interventions and Subsidies in Energy in Fiscal Years 2016–2022." *U.S. Department of Energy, Washington, DC.* https://www.eia.gov/analysis/requests/subsidy/.

Weyl, E Glen, and Michal Fabinger. 2013. "Pass-through as an economic tool: Principles of incidence under imperfect competition." *Journal of political economy*, 121(3): 528–583.

A Program Details and Data Construction

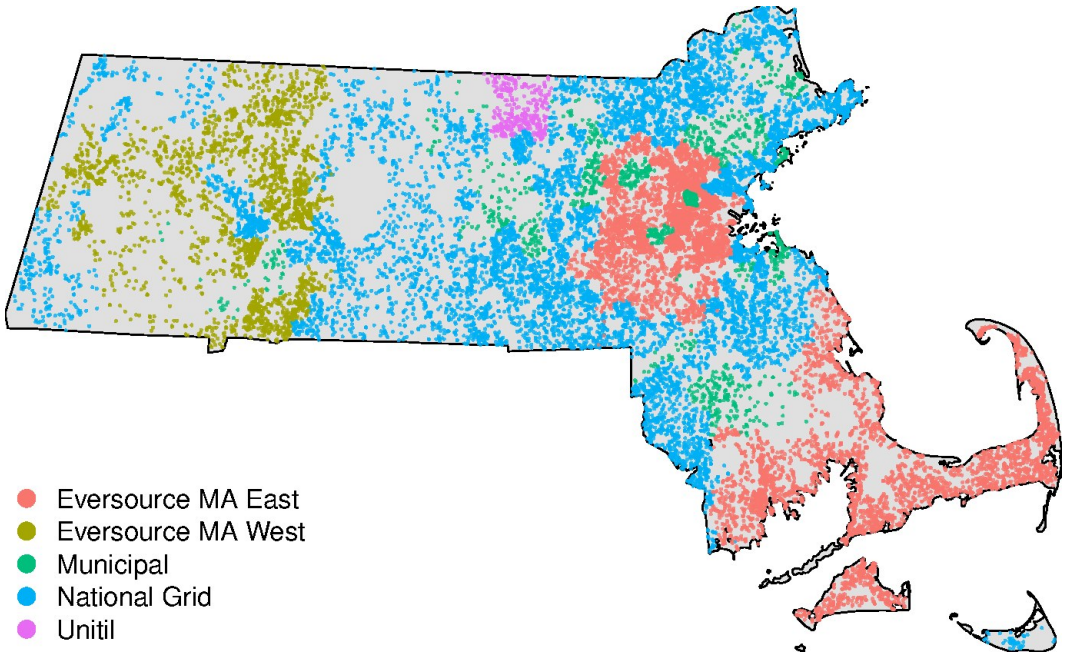
Merging DOER and Tracking the Sun. Projects are identified differently in the two sources. I match them on operation date, size, and total cost simultaneously. This gives an almost perfect one-to-one match, with only a few projects matched to the same Tracking the Sun record, and I verified manually that these projects have the same technical details.

Dates. For the SMART sample, the DOER data list two dates for each project: the Commercial Operation Date and the Incentive Payment Effective Date. I use the Commercial Operation Date, the earliest date in the data. The SMART rate is fixed earlier, when the application is approved, but the data do not record the approval date (Section 4.4).

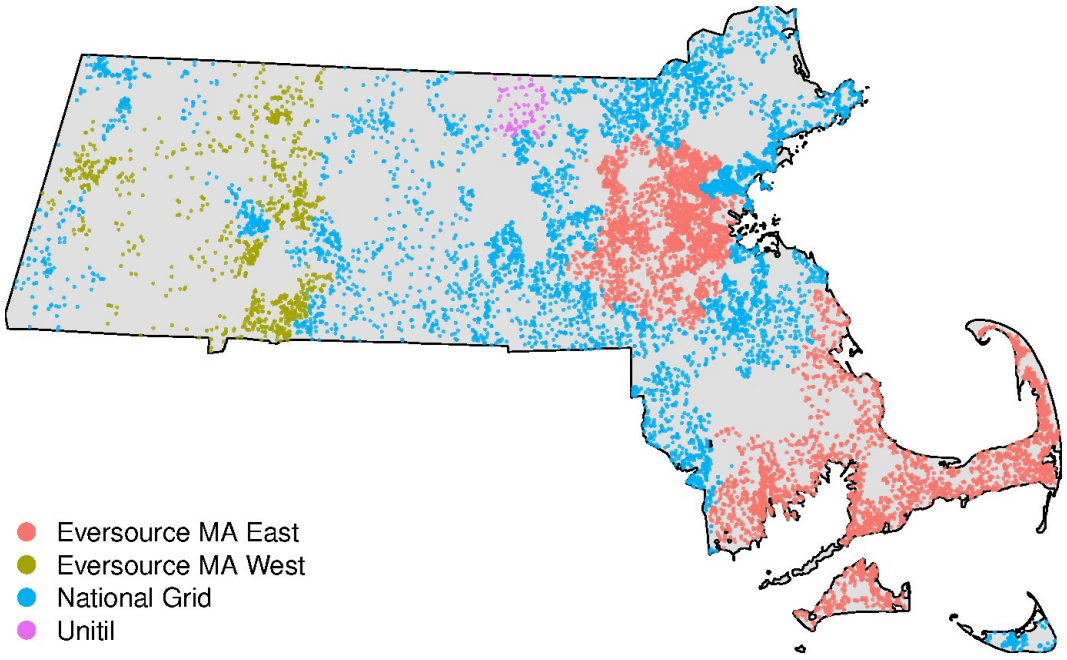
Geocoding. The PVWatts location input accepts street addresses, ZIP codes, or coordinates. The public DOER data contain only ZIP codes, which can cover large areas relative to the resolution of the solar resource data, so I obtained detailed addresses for each project through a separate data request to DOER. The calculator returned unreliable results for street addresses entered directly, so I converted the addresses to coordinates with the U.S. Census geocoder and used the OpenStreetMap geocoding service for addresses that the Census geocoder could not match. All addresses were geocoded by one of the two services, and I verified that all coordinates lie in Massachusetts.

PVWatts inputs. I enter the coordinates as the location input. System size is available for every project, but other parameters are missing for some projects. For a missing parameter, I calculate an implied value; for example, I compute the DC-to-AC ratio for SMART projects from their DC and AC ratings. When no implied value is available, I use the calculator's default values, which represent typical rooftop systems. Fewer than 100 projects require default values.

SMART rate schedule and project locations. Table 8 reports the SMART base compensation rates by utility and capacity block, and Figure 7 maps the projects in the two samples.



(a) Solar Carve-Out II



(b) SMART

Figure 7: Project Locations by Utility Territory

Each dot is one project in the analysis sample. Municipal utility customers are ineligible for SMART.

Table 8: SMART Base Compensation Rates for Generation Units of 25 kW AC or Less

Utility	Low income	Block 1	2	3	4	5	6	7	8
Fitchburg	Yes	$0.358	$0.326	$0.298	$0.272	—	—	—	—
	No	$0.311	$0.284	$0.259	$0.236	—	—	—	—
Mass Electric	Yes	$0.358	$0.344	$0.330	$0.317	$0.304	$0.292	$0.280	$0.269
	No	$0.311	$0.299	$0.287	$0.275	$0.264	$0.254	$0.244	$0.234
Nantucket	Yes	$0.391	$0.328	—	—	—	—	—	—
	No	$0.340	$0.286	—	—	—	—	—	—
Eversource E.	Yes	$0.391	$0.375	$0.360	$0.346	$0.332	$0.319	$0.306	$0.294
	No	$0.340	$0.326	$0.313	$0.301	$0.289	$0.277	$0.266	$0.255
Eversource W.	Yes	$0.329	$0.315	$0.303	$0.291	$0.279	$0.268	$0.257	$0.247
	No	$0.286	$0.274	$0.263	$0.253	$0.243	$0.233	$0.224	$0.215

Base compensation rates in dollars per kWh by utility and capacity block. "Low income" indicates the rates for low-income customers.

B Cross-Roof Estimates: Additional Results

B.1 SREC Prices, Discount Rates, Clustering, and Tests of Complete Pass-Through

SREC prices. I perform a further robustness check specific to Solar Carve-Out using alternative assumptions about system owners' expectations of SREC value. An owner with perfect foresight would know future SREC prices and sell SRECs once a year at that year's mean price. However, it is hard for an individual to predict the actual price of SRECs, especially in the initial years after the launch of Solar Carve-Out II. The clearinghouse auction price was intended to support SREC prices, but for a period SRECs consistently traded below it because of oversupply (newly installed capacity quickly surpassed the annual target set by the RPS). The common practice of contracting aggregators or brokers to handle SREC trades makes the actual SREC sale price even more difficult to predict. I therefore consider three ways of predicting SREC prices in addition to actual market prices. Optimistic customers use the ACP—the theoretical upper bound—as the predicted SREC price. Pessimistic customers use the clearinghouse auction price. Neutral customers use the average of the ACP and the clearinghouse auction price, which is the main assumption. A standard 10% transaction fee applies to both optimistic and neutral customers. Pessimistic customers are not subject to any other fees, because the transaction fee is already embedded in the clearinghouse auction price.

Table 9 presents pass-through estimates for Solar Carve-Out projects under different SREC price assumptions, in the preferred specification (column (4) of Table 3). Only the Solar Carve-Out sample is shown, and the results are again very close to each other. The highest estimated pass-through rate comes from using the alternative compliance payment as the SREC price, and the lowest from using auction prices. The ordering is as expected, because the ACP is the upper bound of SREC prices and yields the largest subsidy, while the auction price is the lowest assumed path and yields the smallest. The estimates range from 0.844 to 0.878. The estimate using actual market prices, 0.857, is close to the midpoint estimate of 0.863. Whatever owners expected the SRECs to earn, the Solar Carve-Out estimate is robust to the assumed price path.

Discount rates. Because the subsidy is a discounted sum of future cash flows, the discount rate affects both the dependent variable (y_i) and the regressor (Subsidy_i). De Groote and Verboven (2019) find that when subsidies are delivered through electricity production, consumers apply high discount rates to future benefits. They estimate an average rate of 15% in Belgium, well above my main assumption and above the contemporaneous risk-free

Table 9: Pass-Through under Alternative SREC Price Paths

Dependent Variable:	Net Price per kW			
Model:	(1) ACP	(2) Auction	(3) Actual price	(4) Midpoint
Variables				
Subsidy per kW	−0.878***	−0.844***	−0.857***	−0.863***
	(0.039)	(0.050)	(0.046)	(0.044)
Installer × year	✓	✓	✓	✓
City × year	✓	✓	✓	✓
Year-month	✓	✓	✓	✓
Inverter, manufacturer	✓	✓	✓	✓
Log size, DC/AC ratio bins	✓	✓	✓	✓
Observations	23,503	23,503	23,503	23,503

Preferred specification: column (4) of Table 3. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1. Subsidies are discounted at 8% with 0.5% annual panel degradation.

euro interest rate, which remained below 1% during their study period. I therefore also use rates of 3%, 5%, and 15%.

Table 10 presents pass-through estimates under varying discount rate assumptions, in the preferred specification. In Panel A, the estimated pass-through rate declines from 0.890 at a 3% discount rate to 0.821 at 15%, and in Panel B from 0.879 to 0.796. The discount rate mechanically changes the size of the regressor and the dependent variable. For SMART the change is purely mechanical. The subsidy is $r_i g_i A(d)$ (Section 2.3), so changing d multiplies every subsidy by the same factor. The coefficient on the subsidy in the gross price, $1 - \rho$, is divided by that factor, and $(\hat{\rho}_d - 1)A(d)$ is the same at every discount rate. With $A(8\%)/A(3\%) \approx 0.79$, the preferred estimate implies $\hat{\rho}_{3\%} - 1 \approx -0.153 \times 0.79 \approx -0.121$, which is the estimate in the table, and the test of $\rho = 1$ gives the same p-value (about 0.035) at every rate by construction. The SMART results therefore carry no information about how buyers discount the subsidy. Under Solar Carve-Out the SREC price path differs across cohorts, but the year-month fixed effects absorb these differences, so within a month the discount rate scales every subsidy by nearly the same factor and the result is close to mechanical as well. At every discount rate, both estimates are significantly below one (Table 11).

Table 11 reports the corresponding tests of complete pass-through. For SMART, the standard error times $A(d)$ is also the same at every rate, and so is the p-value. The Solar Carve-Out p-values are nearly constant for the same reason.

Clustering. Table 12 reports the standard errors of the preferred estimates under alter-

Table 10: Pass-Through at Alternative Discount Rates

	Panel A: Solar Carve-Out				Panel B: SMART			
	Dependent Variable: Net Price per kW				Dependent Variable: Net Price per kW			
	(1) d=3%	(2) d=5%	(3) d=8%	(4) d=15%	(5) d=3%	(6) d=5%	(7) d=8%	(8) d=15%
Variables								
Subsidy per kW	−0.890***	−0.880***	−0.863***	−0.821***	−0.879***	−0.867***	−0.847***	−0.796***
	(0.035)	(0.039)	(0.044)	(0.057)	(0.057)	(0.063)	(0.073)	(0.097)
Installer × year, city × year	✓	✓	✓	✓	✓	✓	✓	✓
Inverter, manufacturer	✓	✓	✓	✓	✓	✓	✓	✓
Log size, DC/AC ratio bins	✓	✓	✓	✓	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓				
Utility × year-month FE					✓	✓	✓	✓
Observations	23,503	23,503	23,503	23,503	14,317	14,317	14,317	14,317

Preferred specification: column (4) of Table 3. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1. 0.5% annual panel degradation. SREC subsidies use the ACP–auction midpoint price path.

Table 11: Test of $\rho = 1$ by Program and Discount Rate

	d=3%	d=5%	d=8%	d=15%
Panel A: Solar Carve-Out				
Subsidy per kW	−0.890***	−0.880***	−0.863***	−0.821***
	(0.035)	(0.039)	(0.044)	(0.057)
p-value, $\rho = 1$	0.002	0.002	0.002	0.002
Panel B: SMART				
Subsidy per kW	−0.879***	−0.867***	−0.847***	−0.796***
	(0.057)	(0.063)	(0.073)	(0.097)
p-value, $\rho = 1$	0.035	0.035	0.035	0.035

Coefficient on subsidy per kW from the preferred specifications, column (4) of Table 3. The dependent variable is the net price per kW, and pass-through ρ is minus the coefficient. The p-values test complete pass-through ($\rho = 1$, a coefficient of −1), using installer-clustered standard errors. Stars test the coefficient against zero. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

native clustering choices.

B.2 Solar Carve-Out with Utility Fixed Effects

Table 13 adds utility fixed effects to the preferred Solar Carve-Out specification, column (4) of Table 3. SREC prices do not depend on the utility, so the estimate is essentially unchanged.

Table 12: Standard Errors under Alternative Clustering

Clustered by	Installer	City	Utility
Panel A: Solar Carve-Out			
Subsidy per kW	−0.863***	−0.863***	−0.863***
	(0.044)	(0.028)	(0.068)
Panel B: SMART			
Subsidy per kW	−0.847***	−0.847***	−0.847***
	(0.073)	(0.061)	(0.071)

Coefficient on subsidy per kW from the preferred specifications, column (4) of Table 3. The dependent variable is the net price per kW, and pass-through is minus the coefficient. Discount rate 8%. Standard errors in parentheses under each clustering. Utility clustering has only four clusters in the SMART sample and is shown for comparison only. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 13: Solar Carve-Out Pass-Through with and without Utility Fixed Effects

	Preferred	With Utility FE
Subsidy per kW	−0.863***	−0.862***
	(0.044)	(0.044)
Installer × year, city × year FE	✓	✓
Year-month, inverter, manufacturer FE	✓	✓
Log system size, DC/AC ratio bins	✓	✓
Utility FE		✓
Observations	23,503	23,503

Preferred specification: column (4) of Table 3. Standard errors clustered by installer in parentheses. Discount rate 8%. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

B.3 Comparisons within Finer Cells

The preferred specification includes installer-by-year, city-by-year, and month fixed effects additively, so it does not require two compared roofs to share an installer, a town, and a month at once. Table 14 replaces installer-by-year with installer-by-month fixed effects, which compare only roofs that the same installer equipped in the same month, and alternatively with installer-by-town-by-year fixed effects. In the installer-by-month specification, 94% of Solar Carve-Out and 92% of SMART projects share an installer and month with another project. The estimates barely move, and the difference between the programs is −0.002 (0.087).

Table 14: Cross-Roof Pass-Through within Finer Cells

	Solar Carve-Out	SMART
Installer × year (column 4)	0.863	0.847
	(0.044)	(0.073)
Installer × year-month	0.860	0.859
	(0.046)	(0.077)
Installer × city × year	0.853	0.848
	(0.050)	(0.081)
Difference (SMART − SCO), installer × year-month	−0.002 (0.087)	

Coefficients are pass-through (minus the coefficient on subsidy per kW). Each row replaces the installer-by-year fixed effects of column (4) of Table 3 with the stated fixed effects and keeps the other fixed effects and controls. The difference comes from one stacked regression with program-specific slopes, controls, and fixed effects, clustered by installer across programs. Discount rate 8%. Standard errors clustered by installer in parentheses.

B.4 Pooled versus Separate Estimation

Table 15 compares the separate estimates with pooled ones. Columns (1) and (2) come from one stacked regression in which the slope, the controls, and every fixed effect are specific to each program, clustered by installer across programs. Its estimates equal those from separate regressions, and the difference between the programs, reported below the two panels, has a standard error that accounts for installers active in both. Columns (3) to (6) impose one slope. With program-specific fixed effects (columns 3 and 4), the pooled slope is a weighted average of the two program slopes, with weights proportional to the residual variance of the subsidy in each program. Solar Carve-Out receives two-thirds to three-quarters of the weight because its sample is larger. With fixed effects shared across programs (columns 5 and 6), installer, equipment, and town effects are also assumed to be constant from 2014 to 2023.

Table 15: Pooled versus Separate Estimation

	Dependent variable: net price per kW					
	Slope by program		One slope, program FE		One slope, shared FE	
	Naive	Cross-roof	Naive	Cross-roof	Naive	Cross-roof
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Solar Carve-Out						
Subsidy per kW	−0.756***	−0.863***				
	(0.044)	(0.044)				
Panel B: SMART						
Subsidy per kW	−1.017***	−0.847***				
	(0.065)	(0.072)				
Difference in pass-through (SMART − SCO)	0.260***	−0.016				
	(0.090)	(0.085)				
Standard error if programs independent	(0.078)	(0.084)				
Subsidy per kW, both programs			−0.843***	−0.859***	−0.897***	−0.858***
			(0.031)	(0.037)	(0.038)	(0.037)
Weight on Solar Carve-Out			0.666	0.736		

Entries are coefficients on subsidy per kW, except the difference row, which reports the difference in pass-through. Pass-through is minus the coefficient. Naive: the specification of column (1) of Table 3. Cross-roof: column (4). Standard errors clustered by installer across programs in parentheses. The weight on Solar Carve-Out is its share of the residual variance of the subsidy after the fixed effects and controls. Discount rate 8%. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Pless and van Benthem (2019) pool host-owned and third-party systems, which share one program, period, subsidy schedule, and set of fixed effects. The two programs here differ in all four, so I let every coefficient differ by program: pooling imposes equal pass-through, which the naive specification rejects (p = 0.004), so that the pooled naive rate of 0.84 describes neither program; the SMART subsidy is a fixed contract whereas the Solar Carve-Out subsidy depends on an assumed SREC price path, which pooling would let affect

the SMART estimate; and the fixed effects must differ by program anyway, since SMART needs utility-by-time effects and installer, town, and equipment effects plausibly changed between the two periods. Sharing the fixed effects raises the naive pooled rate from 0.84 to 0.90, because it lets comparisons across programs, in which the subsidy fell by a factor of 3.4 while costs were falling, enter the estimate.

In the preferred specification, the two program estimates are nearly identical, and every specification gives about 0.86. Pooling is then harmless, and 0.86 (0.04) is a fair summary of both programs.

B.5 Roof Tilt and Installation Cost

Table 16 reports the gross price per kW by five-degree tilt bin, relative to roofs pitched at 25–30 degrees, in the preferred specification with the subsidy per kW held fixed. Section 4.1 uses it.

Table 16: Gross Price per kW by Roof Tilt

	Panel A: Solar Carve-Out			Panel B: SMART		
Tilt (degrees)	$/kW	SE	Projects	$/kW	SE	Projects
0–5	372***	(122)	333	88	(67)	208
5–10	86**	(37)	407	25	(31)	322
10–15	6	(19)	773	−28	(24)	556
15–20	−19	(18)	2,598	13	(21)	1,556
20–25	−4	(16)	4,210	−6	(18)	2,431
30–35	42*	(21)	4,804	6	(17)	2,436
35–40	−12	(19)	3,518	−21	(17)	2,580
40–45	48*	(27)	2,045	17	(24)	1,292
45–50	−23	(55)	131	20	(69)	113
50–55	−165	(140)	16	137	(152)	18
55–60	162	(186)	11	−61	(193)	4
60+	199*	(118)	10	86	(251)	6
25–30 (reference)			4,647			2,795

Coefficients on tilt-bin indicators in a regression of the gross price per kW on the subsidy per kW, the tilt bins, and the controls and fixed effects of column (4) of Table 3. The reference bin is 25–30 degrees. Discount rate 8%. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

B.6 Pass-Through by Roof Productivity

This test estimates cross-roof pass-through by tercile of modeled capacity factor, using the specification in column (4) of Table 3, with terciles defined within each program. In the terminology of Appendix F, roofs in the low tercile are the weak market, where a discriminating installer charges less, and roofs in the high tercile the strong market. In one regression per program, with a separate slope on the subsidy and a separate intercept for each tercile and common fixed effects, pass-through is nearly the same in all three terciles: 0.98, 0.95, and 0.97 under Solar Carve-Out and 0.87, 0.83, and 0.87 under SMART. Equality is not rejected ($p = 0.66$ and $p = 0.46$), and the difference between the high and the low tercile is −0.01 (0.04) in both programs. Estimating column (4) separately on each tercile's projects gives point estimates that fall from the low to the high tercile, from 0.95 to 0.81 under Solar Carve-Out and from 0.96 to 0.79 under SMART, but the standard errors are 0.07 to 0.20

and the differences are not significant. The estimates within terciles are higher than those in column (4) because the tercile intercepts remove the variation in generation between terciles. Measurement error in generation, which pulls the estimate toward one (Section 3.4), has a larger effect on the variation that remains. The comparison with installers' production estimates in Section 4.3 suggests that this effect can be large: within the cells of column (4), the two measures of the capacity factor have a correlation of only 0.57, and within terciles the true variation is smaller still. The data therefore show no significant difference in pass-through between the weak and the strong market.

Table 17: Pass-Through by Tercile of Modeled Capacity Factor

	Dependent variable: net price per kW					
	One regression			Separate regressions		
	Low CF	Middle CF	High CF	Low CF	Middle CF	High CF
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Solar Carve-Out						
Subsidy per kW	−0.983***	−0.952***	−0.972***	−0.947***	−0.806***	−0.809***
	(0.069)	(0.059)	(0.058)	(0.075)	(0.162)	(0.198)
p-value, equal across terciles		0.658				
Panel B: SMART						
Subsidy per kW	−0.873***	−0.825***	−0.867***	−0.963***	−0.830***	−0.787***
	(0.091)	(0.095)	(0.086)	(0.140)	(0.147)	(0.160)
p-value, equal across terciles		0.459				

Coefficient on subsidy per kW by tercile of modeled capacity factor (CF), defined within each program, in the specification of column (4) of Table 3. Pass-through is minus the coefficient. Columns (1)–(3) come from one regression that interacts the subsidy with tercile indicators and includes tercile intercepts and common fixed effects, and the *p*-value is from a Wald test that the three slopes are equal. Columns (4)–(6) estimate column (4) of Table 3 separately on each tercile's projects. Discount rate 8%. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

B.7 Heterogeneity by Income and Installer Concentration

I examine whether SMART pass-through differs across ZIP codes by median household income and by local installer concentration. Income is the ZIP-level median household income from the American Community Survey (2016–2020 five-year estimates; the demand estimation in Appendix D.1 uses the 2017–2021 release). Concentration is the Herfindahl–Hirschman index (HHI) of installers' shares of SMART projects in the ZIP code. I divide ZIP codes into terciles on each measure. For each measure, I estimate a single regression that interacts the subsidy with tercile indicators and keeps the fixed effects common across groups, and then test whether the three coefficients are equal. Table 18 reports the results.

Table 18: Heterogeneity in SMART Pass-Through

	Dependent variable: net price per kW							
	ZIP median income tercile				Installer HHI tercile			
Discount rate	3%	5%	8%	15%	3%	5%	8%	15%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Subsidy per kW × tercile 1	−0.877***	−0.865***	−0.844***	−0.793***	−0.841***	−0.824***	−0.798***	−0.730***
	(0.058)	(0.064)	(0.073)	(0.098)	(0.067)	(0.074)	(0.085)	(0.114)
Subsidy per kW × tercile 2	−0.886***	−0.874***	−0.855***	−0.806***	−0.904***	−0.894***	−0.878***	−0.838***
	(0.063)	(0.070)	(0.080)	(0.107)	(0.059)	(0.065)	(0.074)	(0.099)
Subsidy per kW × tercile 3	−0.864***	−0.849***	−0.827***	−0.769***	−0.888***	−0.876***	−0.857***	−0.810***
	(0.068)	(0.075)	(0.086)	(0.115)	(0.059)	(0.065)	(0.075)	(0.100)
p-value, equal pass-through	0.845	0.845	0.845	0.845	0.216	0.216	0.216	0.216
Observations	14,195	14,195	14,195	14,195	14,195	14,195	14,195	14,195

SMART projects. Each column is one regression, with the subsidy per kW interacted with tercile indicators. Tercile 1 is the lowest-income tercile in columns (1)–(4) and the least concentrated tercile in columns (5)–(8). Entries are coefficients, so pass-through is minus each entry. All regressions use the preferred specification, column (4) of Table 3: utility-by-year-month, installer-by-year, city-by-year, inverter, and manufacturer fixed effects, log system size, and DC-to-AC ratio bins. Standard errors clustered by installer in parentheses. The p-value is from a joint Wald test that the three interaction coefficients are equal. It is the same at every discount rate because the discount rate rescales every SMART subsidy by the same factor (Appendix B.1). The sample has the 14,195 projects with ZIP-level income data, compared with 14,317 in the full SMART sample. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

Table 18 uses the preferred specification. Pass-through does not vary with income. At the 8% discount rate, the estimates are 0.84 in the lowest tercile, 0.85 in the middle tercile, and 0.83 in the highest tercile, and the hypothesis that they are equal is not rejected (p = 0.84). The pattern is mechanically the same at every discount rate (Appendix B.1). By concentration, the estimates are 0.80, 0.88, and 0.86 from the least to the most concentrated tercile, so pass-through is lowest where installers are least concentrated, the opposite of what market power would predict, and the equality test does not reject (p = 0.22). In the naive specification (not shown in the table), pass-through falls with income, from 1.09 to 0.91 (p = 0.011), but this gradient disappears once the timing of the subsidy is absorbed.

Two cautions apply. First, the pattern depends on how the fixed effects are estimated. If I estimate the model separately within each tercile, allowing every fixed effect to differ by group, the estimates are 0.74, 0.82, and 1.00 from the lowest to the highest income tercile, and 0.85, 0.94, and 0.79 across concentration terciles, with no monotone pattern by concentration (Figures 8 and 9). The split samples are a third of the full sample with installer-by-year and city-by-year effects estimated within each, so these estimates are imprecise, and I do not put weight on either set of gradients. Second, income and concentration are measured at the ZIP level from SMART records only. Income is an ACS estimate for the ZIP, not the buyer's own income, and the HHI is computed from the same projects used in the regression, so ZIPs with few projects have mechanically high concentration. Terciles are equal-sized groups of projects and not of ZIP codes.

Taken together, the results give no evidence that pass-through differs with local income or that it is lower where installers are more concentrated. The second finding matters for the interpretation of the main results: if pass-through under SMART reflected market power, one would expect it to vary with local competition and, since demand is log-concave (Section D.1), to be lower where competition is weaker. The flat gradient is evidence against market power as the main determinant of pass-through in this market. The test has limited power, because the HHI varies only within a range of low concentration statewide (Appendix C.5.3).

Figures 8 and 9 report SMART pass-through when the model is estimated separately within income and concentration terciles, so that every fixed effect can differ by group.

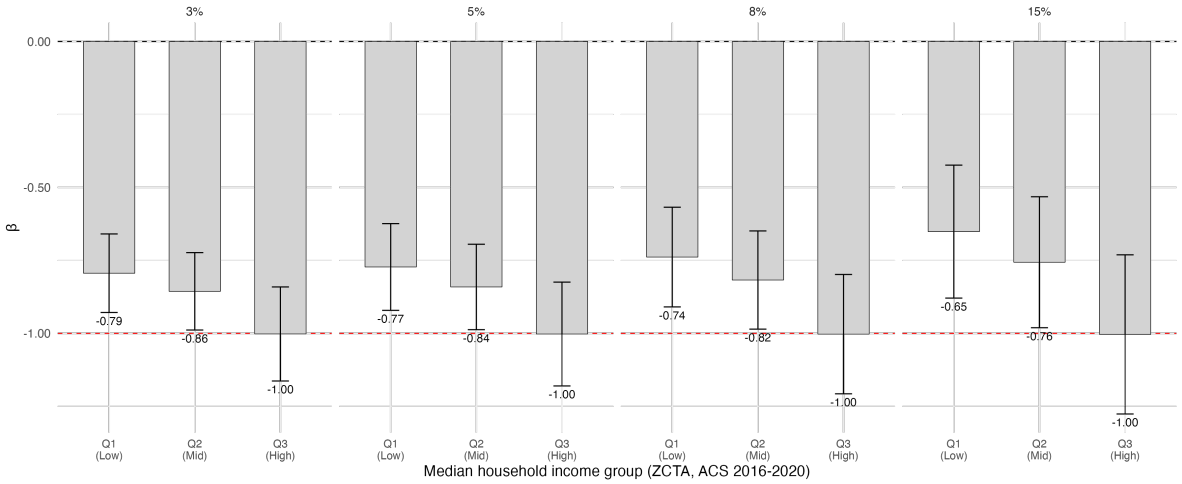


Figure 8: SMART Pass-Through by ZIP Income Tercile, Separate Regressions

SMART projects. Bars are the coefficient on subsidy per kW, so pass-through is minus the bar, from the preferred specification estimated separately within each tercile of ZIP median household income (ACS 2016–2020), at discount rates of 3%, 5%, 8%, and 15%, with 95% confidence intervals from installer-clustered standard errors. The red dashed line marks complete pass-through.

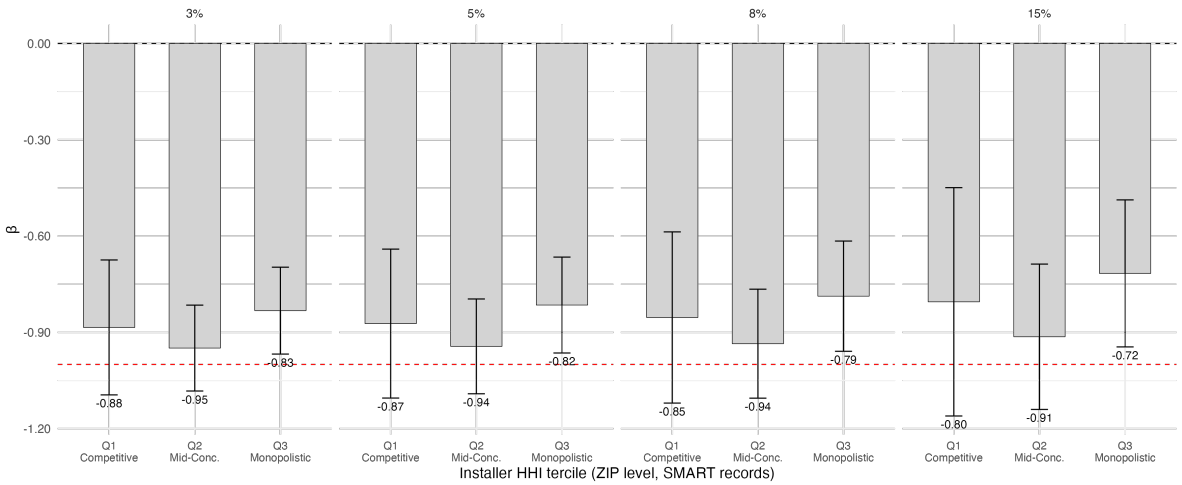


Figure 9: SMART Pass-Through by Installer Concentration Tercile, Separate Regressions

SMART projects. Bars are the coefficient on subsidy per kW, so pass-through is minus the bar, from the preferred specification estimated separately within each tercile of the ZIP-level installer HHI computed from SMART records, at discount rates of 3%, 5%, 8%, and 15%, with 95% confidence intervals from installer-clustered standard errors. The red dashed line marks complete pass-through.

B.8 Workload, Competition, and Pass-Through

Table 19 reports the interactions and town-level regressions discussed in Section 6.4. I measure an installer's workload as the log of its installations in the month of operation and the month before, relative to the installer's own mean, and statewide workload as the log of monthly installations relative to the mean of that year, and I interact each, standardized, with the subsidy. For town competition, I regress the gross price per kW on the log number of installers active in the town and year and, separately, on the town-year HHI, for towns with at least ten installations; because city fixed effects are included, these comparisons are within a town over time.

Table 19: Workload, Competition, and Pass-Through

	Net price per kW			Gross price per kW	
	Installer workload (1)	Statewide workload (2)	Town HHI (3)	Town competitors (4)	Town HHI (5)
Panel A: Solar Carve-Out					
Subsidy per kW	−0.863***	−0.857***	−0.872***		
	(0.043)	(0.044)	(0.040)		
Subsidy per kW × installer workload	−0.036*				
	(0.022)				
Subsidy per kW × statewide workload		−0.014			
		(0.024)			
Subsidy per kW × town HHI			0.008		
			(0.030)		
Log number of installers in town-year				42.5	
				(31.3)	
Town HHI					−95.7***
					(14.4)
Observations	23,503	23,503	20,086	21,783	21,783
Panel B: SMART					
Subsidy per kW	−0.847***	−0.881***	−0.883***		
	(0.072)	(0.083)	(0.069)		
Subsidy per kW × installer workload	−0.039				
	(0.025)				
Subsidy per kW × statewide workload		0.073**			
		(0.034)			
Subsidy per kW × town HHI			−0.099**		
			(0.047)		
Log number of installers in town-year				−66.1*	
				(34.0)	
Town HHI					−0.8
					(16.3)
Observations	14,317	14,317	11,346	11,346	11,346

The dependent variable is the net price per kW in columns (1)–(3) and the gross price per kW in columns (4) and (5). Columns (1)–(3) use the preferred specification, column (4) of Table 3, so the interactions are identified from differences across roofs within its cells. The main effects of statewide workload and town HHI are absorbed by its fixed effects. Columns (4) and (5) relate prices to the number of installers and the HHI of a town-year, which vary only at the town-year level that city-by-year fixed effects absorb, so they use the naive fixed effects (installer, inverter, manufacturer, city, and year, and utility for SMART) and compare a town with itself over time. Installer workload is the log of the installer's installations in the month of operation and the month before, minus the installer's mean. Statewide workload is the log of all installations in the month, minus the mean for that year. Town competition is measured in town-years with at least ten installations, and the HHI is on a 0–10,000 scale. Workload and HHI are standardized within each program. The coefficient on subsidy per kW is minus pass-through at the mean, and the coefficient on an interaction is minus the change in pass-through per standard deviation. Gross-price regressions also control for subsidy per kW. Discount rate 8%. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

C Timing-Based Designs: Additional Results

Table 20 collects the estimates discussed in Section 5, ordered by the length of the period over which each design compares projects; Figure 4 plots them.

Table 20: Pass-Through by Source of Variation and Length of the Comparison Window

Design	Comparison window	Solar Carve-Out	SMART
Cross-roof, preferred	1 month	0.86 (0.04)	0.85 (0.07)
Cross-roof, other versions	1 month	0.76–0.86	0.85–0.86
Year-month fixed effects	1 month	0.80 (0.05)	0.89 (0.06)
Event study, SMART step-downs	8 months		1.27 (0.15)
Regression discontinuity,	±90 days	1.03 (0.10)	
SREC factor cut	±180 days	0.86 (0.07)	
	±365 days	0.71 (0.08)	
excluding	±90 days	0.76 (0.21)	
Dec. 2016–Jan. 2017	±180 days	0.61 (0.10)	
	±365 days	0.54 (0.10)	
Naive, year fixed effects	1 year	0.76 (0.04)	1.02 (0.07)
Predicted sign of the timing bias		−	+

Pass-through estimates with installer-clustered standard errors in parentheses. The comparison window is the period over which a design compares projects. Cross-roof designs compare projects in the same month within installer-by-year (or installer-by-month) and town-by-year cells; the event study compares projects from four months before to three months after a step-down. Sources: Table 3, columns (4), (2), and (1); Tables 4 and 14 for the other cross-roof specifications (tilt bins, azimuth only, installer-by-month cells); Appendix Table 23 for the event study; Appendix Table 22 for the regression discontinuity, which uses the midpoint-price subsidy measure. The predicted sign follows from equation (1) and the co-movement of costs and the subsidy in Figure 3: costs fell with the subsidy under Solar Carve-Out and rose as it fell under SMART. Discount rate 8%.

C.1 Naive Specifications

Table 21 reports the naive model as controls are added. Columns (1) and (5) control only for year fixed effects in Panel A (Solar Carve-Out) and for utility and year fixed effects in Panel B (SMART). Columns (2) and (6) add inverter and panel manufacturer fixed effects. Columns (3) and (7) add installer and city fixed effects. Columns (4) and (8) include all controls and are column (1) of Table 3.

Across specifications in Panel A, the naive estimates with few or many controls are all close to the full model, so the estimate is not driven by variation in the cost factors that the controls capture. Specifications in Panel B are less similar to one another. Adding

Table 21: Naive Pass-Through Estimates: Solar Carve-Out and SMART

	Panel A: Solar Carve-Out				Panel B: SMART			
	Dependent Variable: Net Price per kW				Dependent Variable: Net Price per kW			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
Subsidy per kW	−0.709***	−0.729***	−0.751***	−0.756***	−0.908***	−0.888***	−1.005***	−1.017***
	(0.109)	(0.073)	(0.042)	(0.044)	(0.138)	(0.089)	(0.061)	(0.065)
Controls								
Inverter		✓		✓		✓		✓
Manufacturer		✓		✓		✓		✓
Installer			✓	✓			✓	✓
City			✓	✓			✓	✓
Utility FE					✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	25,491	25,491	25,491	25,491	14,317	14,317	14,317	14,317

Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1. Subsidies are discounted at 8% with 0.5% annual panel degradation. SREC subsidies use the midpoint of the ACP and auction prices.

manufacturer and inverter controls alone lowers the estimate slightly (0.888), while adding installer and city fixed effects raises it to 1.005, very close to the full model (1.017). The standard errors fall as controls are added, from 0.138 to 0.065 in Panel B, which is consistent with the controls removing cost variation that is unrelated to the subsidy.

C.2 Regression Discontinuity at the SREC Factor Cut

Table 22 reports the regression discontinuity estimates of Section 5.2 for both subsidy measures and all three windows, with and without the two months around the cutoff. The low December 2016 prices are why the all-project estimate at ±90 days shows a larger apparent jump in net prices.

Table 22: Discontinuity at the January 2017 SREC Factor Cut

	All projects			Excluding Dec. 2016–Jan. 2017		
Bandwidth (days)	±90	±180	±365	±90	±180	±365
	(1)	(2)	(3)	(4)	(5)	(6)
Subsidy per kW, midpoint-price measure (first stage)	−466***	−478***	−478***	−495***	−497***	−489***
	(12)	(9)	(8)	(34)	(13)	(10)
Subsidy per kW, factor-only measure (first stage)	−434***	−430***	−424***	−425***	−427***	−422***
	(11)	(9)	(8)	(34)	(13)	(11)
Gross price per kW	16	−65**	−139***	−117	−192***	−227***
	(46)	(33)	(37)	(104)	(50)	(46)
Net price per kW	482***	413***	339***	378***	305***	263***
	(48)	(34)	(41)	(114)	(51)	(49)
Pass-through $\hat{\rho}$, midpoint-price measure	1.034***	0.864***	0.708***	0.763***	0.613***	0.537***
	(0.098)	(0.068)	(0.079)	(0.213)	(0.100)	(0.096)
Pass-through $\hat{\rho}$, factor-only measure	1.110***	0.960***	0.799***	0.888***	0.713***	0.622***
	(0.105)	(0.076)	(0.089)	(0.246)	(0.117)	(0.110)
Observations	3,206	5,995	11,795	1,722	4,527	10,328

Solar Carve-Out projects. Each column is a set of regressions on projects within the stated window around January 8, 2017, using commercial operation date as the running variable, with a linear trend on each side of the cutoff and installer, inverter, manufacturer, and city fixed effects. Entries for the first four rows are the estimated jump at the cutoff in dollars per kW. Pass-through is estimated by two-stage least squares with the post-cutoff indicator as the instrument for the subsidy. The factor-only measure holds the present-value factor of the SREC price path at its 2016 value, so that the subsidy varies only through the SREC factor and generation. Standard errors clustered by installer in parentheses. Discount rate 8%. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

C.3 Event Study around SMART Step-Downs

Table 23 reports the pooled estimates of the event study in Section 5.3, without and with a linear trend in event time. Of the two utilities not used, Unitil has 88 projects and never reaches 15 in a month, and Eversource West has 1,082 projects and a noisy modal rate in months with few installations.

Table 23: Pooled Event Study around SMART Rate Step-Downs

	No trend	Linear event-time trend
	(1)	(2)
Subsidy per kW, post vs. pre ($)	−85.2***	−18.2***
	(2.4)	(5.8)
Gross price per kW, post vs. pre ($)	23.2*	11.8
	(12.9)	(25.7)
Net price per kW, post vs. pre ($)	108.4***	30.0
	(12.8)	(26.1)
Pass-through, IV	1.272***	1.645
	(0.153)	(1.435)
Observations	10,851	10,851

SMART projects. Stacked sample of nine step-downs in the modal SMART rate (two National Grid, seven Eversource East), with projects from four months before to three months after the first month of the lower rate. Post is the indicator for event time $k \geq 0$, and column (2) adds a linear trend in event time. Regressions include event, installer, inverter, manufacturer, and city fixed effects. Pass-through is estimated by two-stage least squares with the post indicator as the instrument for the subsidy. Standard errors clustered by installer in parentheses. Discount rate 8%. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

C.4 Price Level

The naive specification is in nominal dollars, for the reasons given in Section 3.1. The SMART estimate is, however, sensitive to how price-level changes are treated: it is 1.02 in nominal dollars, 0.91 when the CPI is added as a control, 0.93 with year-by-month fixed effects, and 0.85 in CPI-deflated dollars (Table 24), while the Solar Carve-Out estimate stays between 0.73 and 0.80. Consumer prices rose by about 22% over the SMART sample while rates were stepping down, so each of these alternatives also removes part of the co-movement between the subsidy and costs discussed in Section 5.1. Because the five specifications use the same projects, I test whether they differ with one stacked regression that gives each specification its own slope, controls, and fixed effects and clusters standard errors by installer across the stacked copies. The SMART estimates differ significantly: every alternative is below the naive estimate, by 0.11 (standard error 0.03) with the CPI control, 0.09 (0.04) with year-by-month fixed effects, 0.17 (0.01) in CPI-deflated dollars, and 0.13 (0.05) in CPI-deflated dollars with year-by-month fixed effects, and equality of all five is rejected ($p < 0.001$). Under Solar Carve-Out the differences are small, from −0.04 to 0.03, but precisely estimated, and equality is also rejected. That the SMART estimate falls whenever part of the rise in the price level is removed is consistent with the naive estimate being pushed up by the timing of the subsidy.

Table 24: Pass-Through under Alternative Treatments of the Price Level

	Dependent variable: net price per kW				
	Nominal dollars			CPI-deflated dollars	
	Year FE (naive) (1)	Year FE, CPI control (2)	Year-by-month FE (3)	Year FE (4)	Year-by-month FE (5)
Panel A: Solar Carve-Out					
Subsidy per kW	−0.756***	−0.760***	−0.799***	−0.729***	−0.797***
	(0.044)	(0.046)	(0.046)	(0.042)	(0.045)
Panel B: SMART					
Subsidy per kW	−1.017***	−0.906***	−0.929***	−0.846***	−0.883***
	(0.065)	(0.063)	(0.062)	(0.062)	(0.062)

Full specifications with installer, inverter, manufacturer, and city fixed effects (and utility fixed effects for SMART). Pass-through is minus the coefficient on subsidy per kW. Discount rate 8%. Deflated values use the monthly CPI-U and are expressed in dollars of the first half of 2019 (index 254.412). The CPI ranges from 234 to 253 over the Solar Carve-Out sample and from 248 to 302 over the SMART sample. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

C.5 Other Explanations for the Naive Gap

The naive estimates differ between the programs: 0.76 for Solar Carve-Out and about one for SMART. Section 5 shows that most of this gap disappears once the timing of the subsidy is absorbed, so the main explanation is how each subsidy moved over time relative to installation costs. This section examines three further candidate explanations. Because the programs differ in size, certainty, period, and eligible population at the same time, the evidence on each is descriptive. The equal cross-roof estimates suggest that none of them is large.

C.5.1 Size

The first explanation is that SMART subsidies are small enough to be absorbed without a supply response, whereas Solar Carve-Out subsidies, worth about 45% of the gross price, strain installers' capacity. The two subsidy distributions have little mass in common (Section 2), so a direct comparison at the same subsidy level is not possible; Table 25 instead reports naive pass-through by quartile of subsidy per kW within each program. Under Solar Carve-Out the estimate is flat, between 0.77 and 0.78 in every quartile. Under SMART it falls from 1.55 in the lowest quartile to 1.09 in the highest, the direction the size explanation predicts, but the lowest quartile has a median subsidy of $127 per kW and a standard error of 0.23, the gradient may also reflect rising costs in the later SMART years, when the subsidy was smallest, and estimates above one are hard to reconcile with log-concave demand (Section 5.4). The within-program evidence therefore does not support the size explanation under Solar Carve-Out and is confounded under SMART.

Table 25: Pass-Through by Quartile of Subsidy per kW

	Q1	Q2	Q3	Q4
Panel A: Solar Carve-Out				
Median subsidy ($/kW)	1,528	1,698	2,151	2,305
Subsidy per kW	−0.779***	−0.776***	−0.780***	−0.770***
	(0.073)	(0.068)	(0.062)	(0.059)
Panel B: SMART				
Median subsidy ($/kW)	127	458	669	947
Subsidy per kW	−1.552***	−1.317***	−1.189***	−1.091***
	(0.228)	(0.131)	(0.091)	(0.076)

Coefficient on subsidy per kW in each quartile, from a regression of the net price per kW in which subsidy per kW is interacted with quartile indicators of the same variable, with the full set of fixed effects. Pass-through is minus the coefficient. Discount rate 8%. Standard errors clustered by installer in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

C.5.2 Predictability

The second explanation is that fixed SMART rates reduce uncertainty: buyers who cannot forecast SREC revenue may value it less than its expected present value, so installers cannot raise prices as far. Two observations weigh against it. The Solar Carve-Out estimate is not sensitive to how one forecasts prices: in the naive specification (not tabulated) it ranges from 0.723 (auction price) to 0.782 (ACP) and is 0.747 with actual market prices, and in the preferred specification from 0.844 to 0.878 (Table 9); if uncertainty about the price path were the main driver, the estimate should be more sensitive to the assumed path. Naive pass-through under Solar Carve-Out does fall with the discount rate (0.805 at 3% versus 0.682 at 15%), whereas the SMART estimate cannot respond to it by construction (Appendix B.1), which is consistent with buyers applying high discount rates to uncertain, back-loaded certificate revenue but is also partly the mechanical effect of a larger regressor.

C.5.3 Market Competition

The third explanation concerns competition. Figure 10 shows the concentration of installers over the sample. At the state level the market is unconcentrated throughout: the statewide Herfindahl–Hirschman index (HHI) is 603 in 2015 and falls to about 300 by 2017, where it stays through 2022; the top four installers' combined share falls from about 40% in 2015 to 23–27% from 2017 on; and 128–216 installers are active in each year. Town-level concentration, which is more relevant for a household choosing among installers who quote locally, is higher: the project-weighted town HHI falls from about 2,400 in 2015 to about 1,400 in 2018 and is 1,500–1,900 in the SMART years with no downward trend.[23] Concentration did not change in a way that could explain the jump in pass-through: most of the fall in concentration happened before SMART started, when Solar Carve-Out pass-through was already about 0.76, and within each program pass-through does not fall with local concentration (Section 6.4 and Appendix B.7). Market power is therefore not the main reason for the gap between the two programs.

[23] The figure starts in 2015 because Solar Carve-Out II opened in April 2014 and, in that partial year, only 28 towns have at least 20 installations.

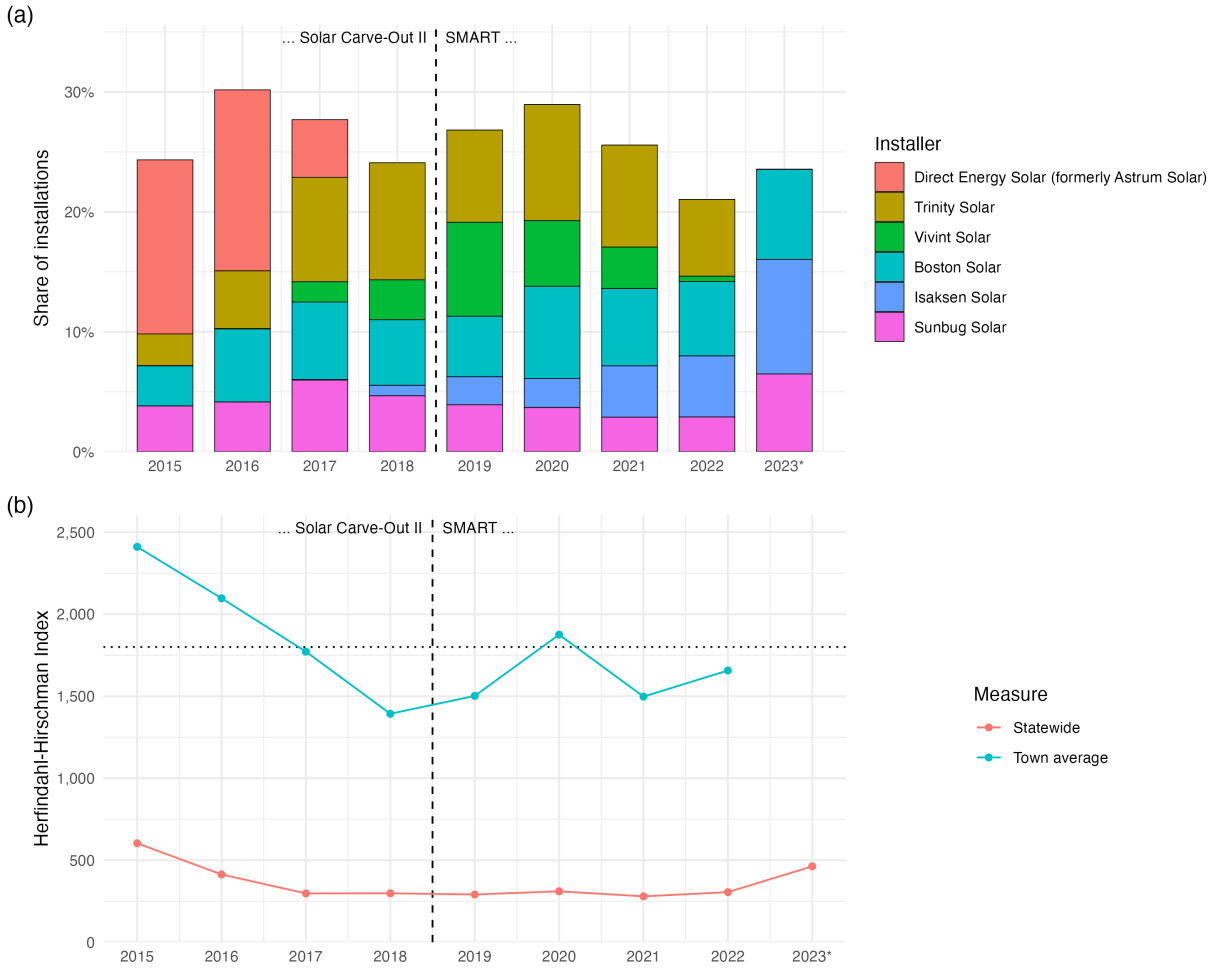


Figure 10: Installer Market Share and Concentration

Both analysis samples, Solar Carve-Out through 2018 and SMART from 2019, by year of operation; 2023 is a partial year. Panel (a): shares of installations of the six largest installers over the sample period. Panel (b): Herfindahl–Hirschman index of installers' shares on a 0–10,000 scale, statewide and as the project-weighted average across towns with at least 20 installations in the year; the dotted line marks 1,800.

D Demand Estimation

D.1 Main Estimates

The bound in Section 6.3 requires the curvature of demand. I estimate demand following Pless and van Benthem (2019). The unit is a ZIP code and year, and the outcome is installations per 1,000 residents. Demand is quadratic in the mean net price per watt,

$$q_{zt} = \beta_1 p_{zt} + \beta_2 p_{zt}^2 + X_z\gamma + \delta_{c(z)} + \tau_t + e_{zt}, \tag{3}$$

so that β_2 measures convexity. The vector X_z contains ZIP characteristics from the 2017–2021 American Community Survey (log median household income, the share of adults with a bachelor's degree, median age, log median home value, and the share of owner-occupied homes), $\delta_{c(z)}$ are county fixed effects, and τ_t are year fixed effects. The net price responds to demand shocks, so I instrument p and p^2 with the mean subsidy per watt and its square and with the average weekly wage in the county's building-equipment contracting sector (NAICS 2382, from the Quarterly Census of Employment and Wages) and its square. The subsidy shifts the price buyers pay, and the wage shifts installation costs. Neither should affect adoption except through price. This is the functional form that Pless and van Benthem (2019) use, with the same types of instruments, so a difference in curvature between California and Massachusetts is not an artifact of the specification. Standard functional forms impose the curvature of demand: linear demand is log-concave, exponential demand is log-linear, and constant-elasticity demand is log-convex (Fabinger and Weyl, 2014). I use a quadratic because it is the lowest-order form that estimates curvature rather than imposing it, and I evaluate curvature at the mean price because demand curves need not have constant curvature. Log-concavity is also what one would expect for a discrete adoption decision when willingness to pay has a log-concave density (Bagnoli and Bergstrom, 2005). Appendix D.3 replaces the quadratic with a natural spline. From the estimates, I compute the price elasticity and the curvature $1/\varepsilon_{ms} = 1 - qq''/q'^2$ at the mean price, with standard errors by the delta method. Because the quadratic in levels imposes a particular shape, I also estimate the same equation with $\log q_{zt}$ as the outcome. In that form the coefficient on p^2 is zero exactly when demand is log-linear, so it is a direct test of log-linear demand, and $1/\varepsilon_{ms} = -2\beta_2/(\beta_1 + 2\beta_2 p)^2$. The sample includes ZIP-years with at least one installation and at least 500 residents.

Table 26 reports the estimates. Under SMART, demand slopes down with an elasticity of −1.57 at the mean price of \$3.35 per watt. This is nearly twice the −0.85 that Pless and van Benthem (2019) report for host-owned systems in California at a mean price of \$3.70 per

Table 26: Demand Estimates

	Panel A: Solar Carve-Out		Panel B: SMART	
Dependent variable	Quadratic q (1)	Log-quadratic $\log q$ (2)	Quadratic q (3)	Log-quadratic $\log q$ (4)
Net price	0.777	2.110	2.516***	2.343***
	(4.257)	(2.690)	(0.605)	(0.563)
Net price squared	−0.124	−0.362	−0.439***	−0.446***
	(0.716)	(0.459)	(0.082)	(0.075)
Price elasticity at the mean	0.27	0.79	−1.57***	−2.16***
	(1.21)	(1.09)	(0.42)	(0.42)
Curvature $1/\varepsilon_{ms}$ at the mean	14.86	7.05	5.38*	2.14**
	(56.01)	(11.35)	(2.77)	(1.02)
First-stage Wald statistic, p / p^2	20.3 / 16.3	20.3 / 16.3	61.2 / 56.3	61.2 / 56.3
Mean net price ($/W)	2.47	2.47	3.35	3.35
Mean installations per 1,000 residents	1.54	1.54	0.91	0.91
ZIP-years	2,171	2,171	1,715	1,715

Two-stage least squares estimates. Columns (1) and (3): equation (3), with q the installations per 1,000 residents. Columns (2) and (4): the same equation with $\log q$ as the outcome, so that the coefficient on the squared price is zero under log-linear demand. Instruments: mean subsidy per watt and its square, and the county building-equipment contractor wage and its square. Controls: ACS 2017–2021 ZIP characteristics, county and year fixed effects. $1/\varepsilon_{ms} = 0$ for log-linear demand, > 0 for log-concave demand, and < 0 for log-convex demand. Standard errors clustered by ZIP code in parentheses. Elasticity and curvature standard errors by the delta method. Appendix D.2 reports alternative specifications. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

watt.[24] The squared term is negative and precisely estimated, so demand is concave, and the curvature measure is 5.38. The log-quadratic form gives the same answer more sharply: the coefficient on the squared price is −0.446, which rejects log-linear demand, and the curvature is 2.14. Demand under SMART is therefore log-concave, not log-linear and not log-convex, whereas Pless and van Benthem (2019) find convex demand in California with the same functional form.[25] The result holds across specifications. The elasticity ranges from −1.14 to −1.92 and the curvature from 3.4 to 9.8 when I add utility fixed effects, replace county and year fixed effects with county-by-year fixed effects, weight by population, drop ZIP-years with fewer than three installations, or cluster by county (Appendix D.2). Under Solar

[24] Two differences between the settings could explain the gap. Buyers in the later period had more alternatives to buying a system outright, such as leases and power purchase agreements, which are outside my sample, and later adopters are likely more price-responsive than early adopters, who install at almost any price. The samples and instruments also differ.

[25] Both forms bend back at low prices: the level quadratic slopes upward below $2.87 per watt (a quarter of ZIP-years) and the log-quadratic below $2.62 (13%), so the estimates describe demand around the mean rather than over the whole range.

Carve-Out, demand is not identified. The subsidy is driven by the statewide certificate price and varies little across ZIP codes within a county and year, and the estimates change sign across specifications. I therefore draw structural conclusions only for SMART.

Two further limitations qualify the SMART result. First, the bound on market power in Section 6.3 depends on the sign of the curvature, and the t-statistics are about 1.9 for the quadratic estimate of 5.38 and 2.1 for the log-quadratic estimate of 2.14. Second, the sample excludes ZIP-years with no installations, which truncates demand at low quantities, and truncation affects the estimated curvature directly.

D.2 Alternative Specifications

Table 27 reports the quadratic demand estimates of Section D.1 under alternative specifications. With county-by-year fixed effects and only the subsidy instruments, the log-quadratic SMART estimates are almost unchanged: the coefficient on the squared price is −0.425 (0.103) and the curvature is 2.13 (1.12). Under SMART, the price coefficients and the curvature keep their sign in every specification. Under Solar Carve-Out, the estimates are imprecise and change sign, so demand is not identified. I also replaced the subsidy instruments with the mean modeled capacity factor of each ZIP code's projects and its square, which shift the subsidy only through roof geometry. With the wage instruments, the first-stage Wald statistics are 7.6 and 8.6, and the curvature is positive but has standard errors of 13 (log-quadratic) and more than 1,000 (quadratic). Without them, the first stage has no power. These instruments are too weak to test the subsidy instrument.

Table 27: Demand Estimates under Alternative Specifications

	Baseline	Utility FE added	County × year FE, subsidy IV	Population-weighted	ZIP-years with 3+ installations	Clustered by county
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Solar Carve-Out						
Price elasticity at the mean	0.27	0.30	−2.51	2.24	−4.93**	0.27
	(1.21)	(1.29)	(1.61)	(1.40)	(2.46)	(1.12)
Net price squared (β_2)	−0.124	−0.216	−0.202	−1.001	5.026**	−0.124
	(0.716)	(0.765)	(0.909)	(0.857)	(2.132)	(0.768)
Curvature $1/\varepsilon_{ms}$ at the mean	14.86	19.58	1.25	2.59***	−0.42	14.86
	(56.01)	(97.34)	(1.43)	(0.78)	(0.91)	(58.00)
ZIP-years	2,171	2,171	2,171	2,171	1,809	2,171
Panel B: SMART						
Price elasticity at the mean	−1.57***	−1.92**	−1.62***	−1.21***	−1.14**	−1.57***
	(0.42)	(0.77)	(0.43)	(0.37)	(0.47)	(0.51)
Net price squared (β_2)	−0.439***	−0.417***	−0.251***	−0.216***	−0.570***	−0.439***
	(0.082)	(0.092)	(0.082)	(0.062)	(0.116)	(0.122)
Curvature $1/\varepsilon_{ms}$ at the mean	5.38*	3.80	3.35*	4.64	9.78	5.38
	(2.77)	(2.65)	(1.73)	(2.97)	(8.39)	(3.55)
ZIP-years	1,715	1,715	1,715	1,715	1,237	1,715

Two-stage least squares estimates of equation (3). Elasticity and curvature at the mean net price and installation rate of each sample. Baseline: county and year fixed effects, with the subsidy per watt, the county building-equipment contractor wage, and their squares as instruments. Column (3) uses only the subsidy instruments, because the wage varies at the county-year level. Standard errors clustered by ZIP code (by county in column 6) in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

D.3 Flexible Functional Form

The quadratic in Section D.1 estimates curvature from the data but restricts how curvature varies with price. To check whether this restriction matters, I replace the quadratic with a natural cubic spline in the net price, with three to five degrees of freedom and interior knots at quantiles of the price distribution. A spline enters the endogenous price through several nonlinear terms, and two-stage least squares with polynomials of the instruments identifies them only weakly: with nine instruments, the curvature at the mean ranges from −87 to 1.9 across three to five degrees of freedom, with standard errors of up to 389. I therefore use a control function (Newey, Powell and Vella, 1999). The first stage regresses the price on the instruments, controls, and fixed effects of equation (3), and the demand equation adds the first-stage residual to the spline. Standard errors come from 499 bootstrap draws that resample ZIP codes and re-estimate both stages. Under Solar Carve-Out, the curvature estimates have standard errors in the thousands under every form, consistent with demand not being identified (Section D.1), so Table 28 reports SMART only.

Three findings stand out. First, curvature is positive in every specification, between 0.73 and 1.51 at the mean price, and 97% to 100% of bootstrap draws imply log-concave demand. The conclusion that demand under SMART is log-concave therefore does not depend on the quadratic. Second, with the estimator held fixed, the quadratic and the splines give similar curvature (1.31 for the quadratic and 0.73 to 1.51 for the splines), so the quadratic is flexible enough near the mean. Third, the baseline curvature of 5.38 is larger than under any flexible form. Figure 11 shows why. The baseline quadratic turns upward at \$2.87 per watt, not far below the mean of \$3.35, and the curvature of a quadratic rises steeply as its slope approaches zero. From about \$3.80 per watt upward, all three curves in the figure lie between about one and two.

Smaller curvature loosens the bound. At the cross-roof pass-through estimate with modeled generation, the implied markup is at most 4.1% to 7.0% of price under the control-function estimates, compared with 2.1% for the baseline quadratic and 3.9% for the log-quadratic form. With the sensitivity values of pass-through from the installer's estimate, 0.77 and 0.74, the control-function bounds are 6.6–11.3% and 8.0–13.7% (Table 7). The baseline quadratic gives the tightest bound.

Table 28: SMART Demand Curvature under Flexible Functional Forms

	Two-stage least squares		Control function			
	Quadratic (baseline)	Log-quadratic	Quadratic	Natural spline		
				3 df	4 df	5 df
	(1)	(2)	(3)	(4)	(5)	(6)
Elasticity at the mean	−1.57***	−2.16***	−2.54***	−3.01***	−2.94***	−3.52***
	(0.42)	(0.42)	(0.36)	(0.41)	(0.39)	(0.76)
Curvature $1/\varepsilon_{ms}$ at the mean	5.38*	2.14**	1.31***	1.32***	1.51***	0.73
	(2.77)	(1.02)	(0.17)	(0.14)	(0.56)	(0.98)
Share of bootstrap draws log-concave			1.00	1.00	1.00	0.97
Markup bound (% of price)	2.1	3.9	5.5	4.6	4.1	7.0

SMART ZIP-years (1,715). Elasticity and curvature at the mean net price and installation rate. Columns (1) and (2): two-stage least squares from Table 26, with standard errors clustered by ZIP code. Columns (3)–(6): control function, in which the demand equation includes the first-stage residual of the price. Natural splines have interior knots at quantiles of the net price. Standard errors in columns (3)–(6) and the share of log-concave draws come from 499 bootstrap draws that resample ZIP codes and re-estimate both stages. The markup bound is $(1/\rho_B - 1)/(1/\varepsilon_{ms})$ divided by the absolute elasticity, at $\rho_B = 0.847$ (Table 3, column 4). All specifications include the ZIP controls and county and year fixed effects of equation (3). Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

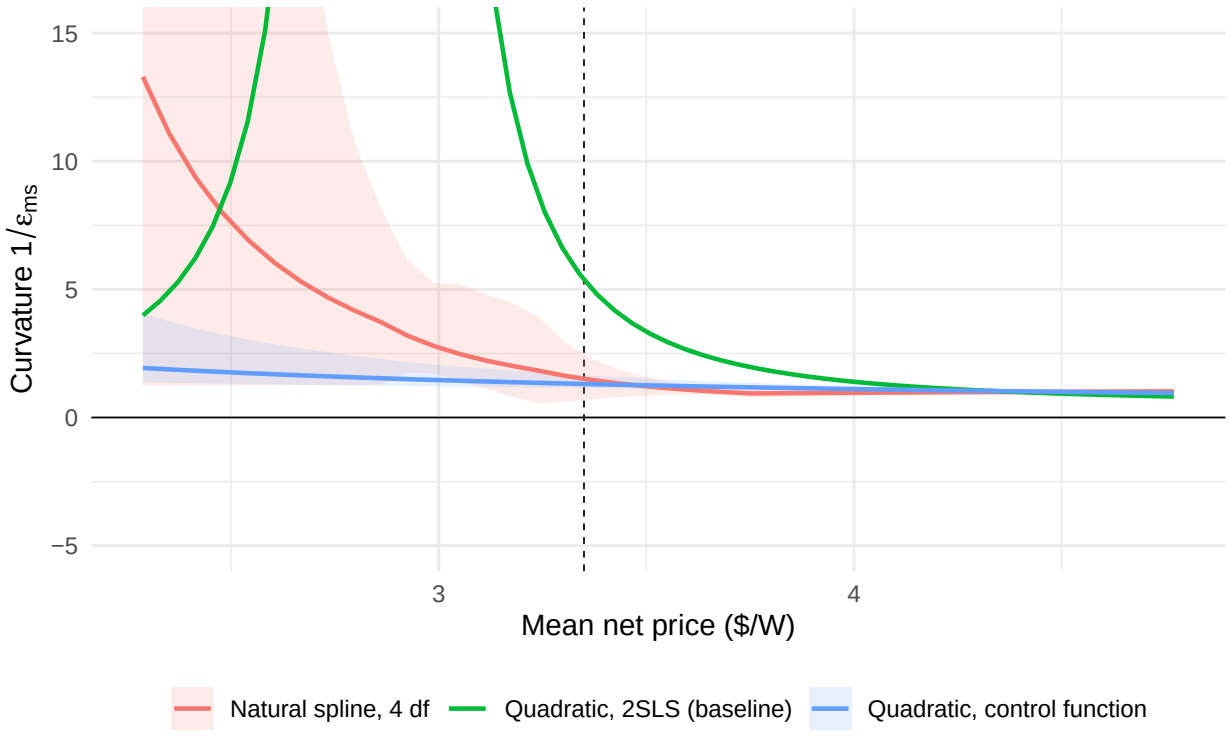


Figure 11: Curvature of SMART Demand across Prices

Curvature $1/\varepsilon_{ms}$ evaluated along each fitted demand curve between the 5th and 95th percentiles of the net price. Shaded areas are 90% bootstrap intervals for the control-function estimates. The dashed line marks the mean net price.

D.4 The Bound under Measurement Error and the Division of Surplus

Table 29 shows the bound of Section 6.3 as a function of the share λ of the remaining variation in the measured subsidy that is true variation. If λ is known, $\rho_B = 1 - (1 - \hat{\rho}_B)/\lambda$. Markups exceed 5% of price only if λ is below 0.82 with the log-quadratic demand or 0.52 with the quadratic demand, and 10% only if it is below 0.48 or 0.33. Taken at face value, the installers' estimates in Section 4.3 would put λ at about 0.6–0.8, where markups are 3–7% of price.

Table 29: Bound on Market Power under Measurement Error in Generation

		Quadratic demand		Log-quadratic demand	
λ	ρ_B	θ bound	Markup (% of price)	θ bound	Markup (% of price)
1.0	0.847	0.034	2.1	0.085	3.9
0.9	0.830	0.038	2.4	0.096	4.4
0.8	0.808	0.044	2.8	0.111	5.1
0.7	0.781	0.052	3.3	0.131	6.1
0.6	0.745	0.064	4.1	0.161	7.4
0.5	0.693	0.082	5.2	0.207	9.6
λ at which the markup reaches 5%: 0.52 (quadratic), 0.82 (log-quadratic)					
λ at which the markup reaches 10%: 0.33 (quadratic), 0.48 (log-quadratic)					

λ is the share of the variation in the measured subsidy, within the cells of the preferred specification, that reflects true differences in generation. $\rho_B = 1 - (1 - \hat{\rho}_B)/\lambda$ with $\hat{\rho}_B = 0.847$ (Table 3, column 4). The bound is $\theta \le (1/\rho_B - 1)\,\varepsilon_{ms}$ and the markup is θ/ε_D, at the point estimates of the SMART demand curvature and elasticity in Table 26.

The same framework gives the division of surplus. The ratio of the gain in consumer surplus to the gain in producer surplus is $I = \rho/[1-(1-\theta)\rho]$, so the buyers' share of the surplus that the subsidy creates is $I/(1+I) = \rho/(1+\theta\rho)$, which equals ρ under perfect competition. The formula requires market-wide pass-through, and because $\rho_M \le \rho_B$, evaluating it at ρ_B gives an upper bound on the buyers' share. At $\rho_B = 0.847$, buyers receive at most 85% of the surplus if the market is perfectly competitive and at most 82% and 79% at the bounds on θ of 0.034 and 0.085. Market power at the bound therefore shifts at most a few percentage points of the surplus from buyers to installers, beyond what the shortfall in pass-through already shows.

E Pass-Through of a Roof-Specific Subsidy

This appendix derives the expression for roof-specific pass-through in Section 6.1. Group buyers by roof type k, which determines the subsidy s_k a buyer receives, and let $D_k(p)$ be the probability that a buyer of type k installs at net price $p = P_k - s_k$, where P_k is the gross price the installer charges. The installer observes k, because it computes the roof's expected output itself, and its cost of installing roof k is c_k, which does not depend on the subsidy. Following Weyl and Fabinger (2013), I summarize conduct by a parameter $\theta_B \in [0,1]$ in the pricing condition

$$P_k - c_k = \theta_B \, m_k(P_k - s_k), \qquad m_k(p) \equiv -\frac{D_k(p)}{D_k'(p)},$$

which is marginal-cost pricing when $\theta_B = 0$ and monopoly pricing of roof k when $\theta_B = 1$. Intermediate values describe competition among the installers that quote the roof. Because one roof is a negligible share of an installer's volume, a change in s_k does not move the installer along its marginal cost curve, so c_k is fixed.

Differentiating with respect to s_k,

$$\frac{dP_k}{ds_k} = \theta_B \, m_k'(p)\left(\frac{dP_k}{ds_k} - 1\right) \quad \Longrightarrow \quad \frac{dP_k}{ds_k} = \frac{-\theta_B m_k'}{1 - \theta_B m_k'}.$$

Since $m_k' = -(1 - D_k D_k''/D_k'^2) = -1/\varepsilon_{ms}^B$, the pass-through of the roof's subsidy to the buyer's net price is

$$\rho_B = 1 - \frac{dP_k}{ds_k} = \frac{1}{1 + \theta_B/\varepsilon_{ms}^B}, \qquad \text{so} \qquad \frac{1}{\rho_B} - 1 = \frac{\theta_B}{\varepsilon_{ms}^B}.$$

For a subsidy that changes for all buyers at once, the installer's total volume changes, and the slope of marginal cost adds the term r of equation (2). With $r \geq 0$ and the same conduct and curvature in the two expressions, the market-wide rate is therefore no higher than the roof-specific rate.

Three cases give $\rho_B = 1$ regardless of the installer's markup. Under perfect competition, $\theta_B = 0$. If the installer charges one price per kW for every roof, P_k cannot depend on s_k. And if demand is log-linear, $1/\varepsilon_{ms}^B = 0$, and the markup m_k is a constant number of dollars. Pass-through below one therefore requires both pricing by roof and log-concave demand among buyers with similar roofs.

Two qualifications apply to the bound. First, if installation cost varies with the subsidy across roofs, $c_k = c + \kappa s_k$, then dP_k/ds_k gains the term $\kappa/(1 - \theta_B m_k')$, and the estimated shortfall $1 - \hat{\rho}_B$ overstates the shortfall due to conduct when $\kappa > 0$ and understates it when $\kappa < 0$. Section 4.1 shows that, holding tilt fixed, cost differences across roofs leave the SMART estimate unchanged. Second, the bound uses the curvature of aggregate ZIP-level

demand in place of ε^B_{ms}, the curvature among buyers with similar roofs, which the data do not identify separately.

F Welfare Effects of Pricing by Roof

If installers do price by roof, a natural question is whether buyers as a group gain or lose relative to a uniform price. Aguirre, Cowan and Vickers (2010) and Miklós-Thal and Shaffer (2021) show that the answer depends on pass-through and on the curvature of demand in each market. Treat roofs with high generation as the strong market, where the discriminatory price is higher, and roofs with low generation as the weak market. Relative to uniform pricing, discrimination raises output if pass-through is higher in the weak market ($\rho_W > \rho_S$, averaged over the relevant range) and inverse demand is convex in both markets. It lowers output if $\rho_W < \rho_S$ and inverse demand is concave in both markets. The same conditions, with pass-through weighted by the markup, determine the welfare effect. Miklós-Thal and Shaffer (2021) show that the ranking of pass-through rates alone is not sufficient, so the sign of curvature is needed. Demand under SMART is concave (Table 26). If pass-through were lower on roofs with low generation, pricing by roof would therefore reduce adoption relative to a uniform price per kW. The test in Appendix B.6 finds no such difference, so the data do not sign this effect. Miklós-Thal and Shaffer (2021) also note that formulas based on pass-through hold exactly only for small changes unless demand is linear, so the division of surplus in Appendix D.4, which is evaluated at the mean, is an approximation for a subsidy of this size.