

Carattini, Stefano; Davis, Wade; Figge, Béla; Heimerdinger, Anton

Working Paper

Solar Adoption by Mandates

CESifo Working Paper, No. 12359

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Carattini, Stefano; Davis, Wade; Figge, Béla; Heimerdinger, Anton (2025) : Solar Adoption by Mandates, CESifo Working Paper, No. 12359, Munich Society for the Promotion of Economic Research - CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/336059>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CES ifo

**12359
2025**

December 2025

Working Papers

Solar Adoption by Mandates

Stefano Carattini, Wade Davis, Béla Figge, Anton Heimerdinger

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

Solar Adoption by Mandates*

Stefano Carattini[†], Wade Davis[‡], Béla Figge[§], Anton Heimerdinger[¶]

December 21, 2025

Abstract

Rooftop solar photovoltaic mandates are becoming a popular policy across Europe and the United States. In this paper, we leverage the frontrunner experience of California to examine their economics. First, we evaluate the private payoffs of solar adoption for residential new construction. To do so, we assemble a rich dataset of new construction building projects and parameterize an engineering model to provide estimates of private payoffs across our sample. Second, we evaluate the hypothesis of what we call a “solar gap” by comparing the cost-effectiveness estimates from the engineering model to observed builder decisions. We find substantial variation in the cost-effectiveness of solar across building locations and characteristics, though the estimated private payoffs are generally positive across a robust variety of model parameterizations and financial assumptions. We observe that the majority of buildings in our data do not adopt solar despite engineering estimates suggesting opportunities for positive payoffs. Relatedly, we find that payoffs explain little of the variation in solar adoption decisions. Lastly, we estimate the effectiveness of both San Francisco’s citywide solar mandate and California’s statewide mandate. Across a variety of empirical approaches, we find that both the citywide and statewide mandates increased solar adoption. However, new construction solar adoption remains below 100 percent. We discuss compliance accordingly.

JEL Codes: H70; Q42; Q48

Keywords: Solar Photovoltaic, Building Codes, Regulation, Engineering Models

*First version: August 2024. We thank Kenneth Gillingham, Geoffrey Heal, Javier López Prol, Matthew Kotchen, and Joseph Shapiro for useful comments as well as participants to the 2024 annual conference of the Association of Environmental and Resource Economists and 2025 conference of the European Association of Environmental and Resource Economists, among others. We also thank Asmin Acharya, Iman Ahmadi Naeeni, Lucie Gayelle Assonfack Kana, Raka Datta, Josselyn Jimenez, Pranay Kesari, Xiangyu Meng, Byron Owusu-Ansah, Wes Parrish, Arnaud Schuele, and Denise Wong for excellent research assistance. We are grateful as well to Barry Hooper from SF Environment, City and County of San Francisco, for helpful insights about San Francisco’s policy. Carattini acknowledges support from the Department of Energy, grant number DE-EE0009363, and from the Swiss National Science Foundation, grant number PZ00P1.180006/1. This paper uses restricted-access data, for which Carattini and Davis first received approval from the Institutional Review Board of Yale University in 2017.

[†]Corresponding author: scarattini@gsu.edu. Georgia State University, CEPR, and CESifo.

[‡]U.S. Environmental Protection Agency. The views expressed in this paper are the authors’ own and do not necessarily represent those of the U.S. Environmental Protection Agency (EPA). In addition, the research described in this paper has not been subjected to the Agency’s required peer and policy review. No official Agency endorsement should be inferred. All errors and omissions are the authors’ own. Wade Davis did much of his work on this project during his time as a graduate student at Yale University, including aspects of the project related to procurement of restricted-access data.

[§]Analysis Group.

[¶]University of St. Gallen.

1 Introduction

In the absence of a first-best international or even national climate policy, many states and local jurisdictions are adopting their own second- and n^{th} -best policies to reduce greenhouse gas emissions, including various building codes and technology mandates. In recent years, proposed building policies have included bans on natural gas (e.g., Davis 2025), requirements for electric appliances (e.g., Gould et al. 2023), and various energy efficiency standards (e.g., Jacobsen and Kotchen 2013, Levinson 2016). One increasingly common policy is introduced by municipalities mandating rooftop solar photovoltaic (PV) panels on new construction. Municipalities in California, including the prominent example of San Francisco, have been frontrunners in this respect. Importantly, in January 2020, the state of California began requiring that all newly constructed low-rise residential buildings in the state install rooftop solar PV, thus imposing a solar mandate statewide in the world’s 5th largest economy. Our paper examines the economics of solar mandates in California, leveraging their frontrunner role for informing policy there as well as in the many jurisdictions that recently adopted or are currently considering solar mandates, including the European Union and Switzerland.

In this paper, we address the following questions. First, we assess solar adoption at baseline in the absence of mandates. Many of the solar mandates were preceded by reports suggesting that rooftop solar photovoltaics were cost-effective for all new construction (Davis et al. 2014; Energy and Environmental Economics Inc. 2017). If that were the case, however, we would expect adoption to be ubiquitous, and mandates would be unnecessary and nonbinding. However, we find that many builders do not adopt solar on new construction in the absence of mandates.

Second, we examine the rationale for solar mandates. One justification for solar mandates has been the claim that many builders forego cost-effective solar investments, leaving money on the table. That is, the finding that most new construction projects have not adopted solar could be evidence of what we call a “solar gap” where builders fail to adopt solar even when it would be cost-effective to do so. Alternatively, the modest adoption rates could reflect some misspecification or missing costs known to builders but not accounted for in the engineering models that have found ubiquitous cost-effectiveness. In other contexts, the literature has explored the hypothesis of the “energy efficiency gap” where observed adoption of energy efficiency technologies have not

aligned with engineering estimates of cost-effectiveness, in part because of principal-agent issues and behavioral failures (see Allcott and Greenstone 2012; Gillingham and Palmer 2014; Gerarden et al. 2017 for reviews of the literature), but also because of overly optimistic engineering estimates (Metcalf and Hassett 1999; Davis et al. 2014; Grimes et al. 2016; Levinson 2016; Fowle et al. 2018).

We address this second question on the hypothesis of a “solar gap” as follows. First, we evaluate the private payoffs of solar adoption for residential new construction. Using detailed building-level data, we parameterize an engineering model to provide estimates of solar cost-effectiveness in our sample (i.e., profitability based on private costs and benefits). We compile our novel building project-level data from a variety of sources including municipal building permits records, fine-scale climate and solar irradiance data, rooftop-level shading estimates from Google Project Sunroof, and restricted-access electricity consumption data for a sample of new construction building projects that we obtained in 2018. This portion of our analysis focuses on the San Francisco Bay Area because several municipalities in the region adopted their own municipal mandates in advance of the statewide policy, making this region a good context for examining baseline adoption. Cost-effectiveness reports behind those municipal mandates as well as the eventual statewide mandate pointed to ubiquitous cost-effectiveness of solar PV. For our own analysis of solar system design, performance, and cost-effectiveness, we use the System Advisor Model (SAM) from the U.S. National Renewable Energy Laboratory (NREL). The estimates from the engineering model provide a distribution of project payoffs across project characteristics and allow us to investigate sensitivity to model parameterization. NREL’s System Advisor Model is actually the same model used in many of the abovementioned cost-effectiveness reports, but we find much greater heterogeneity in project cost-effectiveness in parameterizing the model across our dataset of observed building projects and their detailed characteristics. Although estimated payoffs are generally positive, the payoffs are typically small especially compared to overall new construction costs.

However, even the absence of ubiquitous cost-effectiveness may not necessarily imply absence of a “solar gap.” Along the distribution of payoffs, there may still be builders leaving money on the table, even in presence of non-monetary costs. Hence, after parameterizing and estimating the engineering model across our building projects data, we compare payoff estimates with actual solar adoption decisions, at the building level. We observe that positive expected payoffs explain little of the observed variation in adoption decisions. In fact, we find little evidence of any strong

relationship between solar adoption and the level of estimated payoffs, and builders forego solar across the entire support of payoff estimates. We conclude that these findings do suggest the existence of a “solar gap” where some builders forego solar even among the most cost-effective projects.

Lastly, we take an ex-post approach and estimate the policy impact on solar adoption relative to a no-policy counterfactual using synthetic control methods. These exercises estimate the extent to which mandates increase solar adoption, accounting for compliance and uptake of exemptions. Even with limited compliance and exemptions, solar mandates can still effectively increase adoption. We find that San Francisco’s citywide mandate increased solar photovoltaic adoption by 33 percentage points, approximately doubling the pre-policy adoption rate. Post-mandate adoption in San Francisco remains below 100 percent, consistent with the ordinance’s exemptions and alternative compliance pathways. We investigate compliance patterns accordingly. For California’s statewide policy, we find an increase of 5.8 to 8.4 kW (between 30% to 43%) in solar photovoltaic capacity, which is not far from the median capacity of newly installed residential photovoltaic systems in the United States (Barbose et al. 2024).

To our knowledge, ours is the first study to investigate the economics of solar mandates, a policy that is currently being adopted by a growing number of jurisdictions, aside from the ex-ante cost-effectiveness reports. In our analysis of the economics of these mandates, we contribute to several strands of literature. One relevant stream of research is the literature on building codes and restrictions at the local level in a variety of contexts (e.g., Pollakowski and Wachter 1990, Glaeser and Gyourko 2003, Hsieh and Moretti 2019), including building codes motivated to achieve environmental outcomes (e.g., Aroonruengsawat et al. 2012, Jacobsen and Kotchen 2013, Levinson 2016, Bruegge et al. 2019, Novan et al. 2022, and Carattini et al. 2024). We also contribute to the literature on distributed solar policies and solar adoption decisions generally (e.g., Darghouth et al. 2011, Borenstein 2012, Rai and Sigrin 2013, Hughes and Podolefsky 2015, Crago and Chernyakhovskiy 2017). Further, our exploration of a possible “solar gap” presents an additional context for studying the mechanisms identified in the extensive literature on the energy efficiency gap, going back as far as Hausman (1979), among others (see Allcott and Greenstone 2012, Gillingham and Palmer 2014, and Gerarden et al. 2017 for reviews). Ex-post comparisons of engineering estimates to observed outcomes include Davis et al. (2014) on household cooling appliances and Fowlie et al. (2018) on

building weatherization.

Our paper proceeds as follows. Section 2 reviews related literature and provides institutional background and design details about the mandates that we examine in this paper; Section 3 describes our estimates of private payoffs, including our data sources and engineering model parameterization, and evaluates evidence of a “solar gap”; Section 4 presents our policy evaluation for both San Francisco’s citywide mandate and California’s statewide mandate; and Section 5 concludes.

2 Background

2.1 The possible case for solar mandates

Canonical environmental economics suggests that standards and mandates are less efficient policies than market-based instruments. However, standards and mandates may for instance be rationalized under the circumstances of an “energy-efficiency gap” where individuals or firms fail to take seemingly cost-effective actions for a variety of behavioral or institutional reasons, which additional economic incentives from market-based instruments may not be able to address. Across a variety of technologies and institutional settings, the literature in economics has explored the existence of an “energy-efficiency gap,” and identified various potential mechanisms for its emergence (starting with Hausman 1979 and Jaffe and Stavins 1994; see also the reviews by Allcott and Greenstone 2012, Gillingham and Palmer 2014, and Gerarden et al. 2017). Similarly, the existence of a “solar gap” could provide support for solar mandates, and many of the explanations for the energy-efficiency gap in other settings may be applicable in the failure of homeowners and builders to undertake profitable rooftop solar investments.

Prominent energy-efficiency gap explanations for under-adoption include market failures, such as liquidity constraints, imperfect information, and principal-agent problems. Behavioral explanations include myopia and inattention. For example, Tsvetanov and Segerson (2013) present a behavioral model of consumer purchase decisions wherein consumers can be “tempted” by the low purchase price of less energy-efficient goods. Under this model of welfare outcomes, the authors demonstrate that energy-efficiency standards can dominate a tax in the presence of this behavioral distortion. In the rooftop solar context, payoffs (or electricity bill savings) are realized over long horizons, often 20 to 30 years. Also, in the context of new construction, builders may be inattentive to

the value of the solar investment because it is small relative to the overall building project, which may cost several hundreds of thousands or even millions of dollars. Levinson and Niemann (2004) explore principal-agent problems and information asymmetries in energy-efficient investments in rental buildings, suggesting that landlords may pay tenants' utilities to internalize the value of their energy efficiency investments, which is an option that however also comes with downsides, including that it may not be salient to prospective tenants during their rental search. It is possible that builders face similar challenges in capitalizing the value of solar in the ultimate sale or rental of the building. Through rational expectations, builders that do not plan to keep the property over the long run may then decide against the installation of solar panels.

Some strands of the energy-efficiency gap literature explore settings where a purported gap may not exist at all. For example, Myers (2019) exploits the differences in home sales prices between oil-versus gas-heated homes in Massachusetts, finding that changes in the relative prices of these fuels are rationally capitalized. Further, a routine finding in the literature is that engineering estimates overstate the savings from energy efficiency investments (Metcalf and Hassett 1999, Davis et al. 2014, Grimes et al. 2016, Levinson 2016, and Fowlie et al. 2018). In the context of California building codes, Levinson (2016) finds that realized savings are substantially lower than was forecast by regulators. Based on a literature review of engineering estimates finding that the promised energy savings do not actually materialize, Fowlie et al. (2018) suggest that many ex-ante engineering analyses may have systematic misspecifications. They also provide evidence from a field experiment on the gap between engineering estimates and realized savings, even accounting for a potential rebound effect.

The literature has identified a variety of explanations for overestimation of technology performance and payoffs in engineering studies. Engineering studies may fail to account for behavioral factors or assume idealized conditions, including optimized operations and maintenance. Davis et al. (2014) explore circumstances where engineering models are overly pessimistic about the inefficiency of the old technologies being replaced. Another source of engineering model misspecification could be the presence of hidden costs, which are not accounted for in the engineering model (Gillingham and Palmer 2014). Hidden costs could include search frictions and administrative costs, such as the costs of finding a suitable solar installer and the administrative burden of managing the contractor, interconnecting to the electricity grid, and applying for government subsidies. We are going to refer

to these costs with the generic term of “hassle.”

As a counterpoint to the critiques of the engineering literature, some studies have identified settings where the engineering estimates are validated. Jacobsen and Kotchen (2013) find that building codes implemented in Florida in 2002 reduce electricity consumption by 4 percent and natural gas by 6 percent, in line with the ex-ante predictions. In a refinement of this work over an extended time horizon, Kotchen (2017) finds that the energy savings from the building code are persistent for natural gas consumption. However, he finds that electricity consumption appears to revert to pre-policy levels after a few years, suggesting that the effectiveness of the measures degrades over time.

Our analysis focuses on a new and fast-growing policy, which has, however, not been studied yet in the literature. In doing so, it extends these strands of literature comparing ex-post outcomes to ex-ante engineering estimates with an emphasis on the importance of heterogeneity in potential ex-ante misspecification.

2.2 Rooftop solar mandates

Originating in California, rooftop solar mandates are an increasingly common variety of building code mandates for new construction projects around the world. In January 2017, San Francisco became the first major city to implement a rooftop solar mandate, and the example was quickly followed by other municipalities in the San Francisco Bay Area and ultimately the state of California.¹ Prior to California’s statewide implementation of the rooftop solar mandate in January 2020, more than two dozen California municipalities adopted solar photovoltaic requirements between 2017 and 2020, including Brisbane, Culver City, Davis, Fremont, Lancaster, San Francisco, San Mateo, Santa Monica, and Sebastopol (California Energy Commission 2016; California Statewide Codes & Standards Program 2022). Outside of California, jurisdictions with rooftop solar mandates include South Miami, Tokyo, and several German municipalities and German states, such as Baden-Württemberg, Berlin, and Hamburg, which were then followed by larger jurisdictions such as the European Union and Switzerland. While these policies share in common the requirement of rooftop solar on new construction, there is substantial variation in the specific details across

¹See <https://sfplanning.org/project/better-roofs> for the mandate in the city of San Francisco (last accessed, May 9, 2024) and <https://www.energy.ca.gov/programs-and-topics/programs/building-energy-efficiency-standards/online-resource-center/solar> for the statewide mandate (last accessed, May 9, 2024).

jurisdictions.

In support of solar mandates, officials in several jurisdictions have cited ex-ante cost-effectiveness reports that estimate builders and building occupants achieve positive payoffs from the mandates (e.g., Davis Energy Group, Inc. et al. 2016 and Energy and Environmental Economics, Inc. 2017). These cost-effectiveness reports apply engineering models to a small number of hypothetical new buildings with stylized characteristics. The suggestion that rooftop solar is cost-effective on all new construction must be reconciled with the evidence from the data that most new construction buildings do not adopt solar at baseline, as shown in Figure 1 for San Francisco and other Bay Area municipalities. In our review of building permits data and solar adoption behavior, we do not find any jurisdictions with ubiquitous solar adoption on new construction in the absence of mandates.

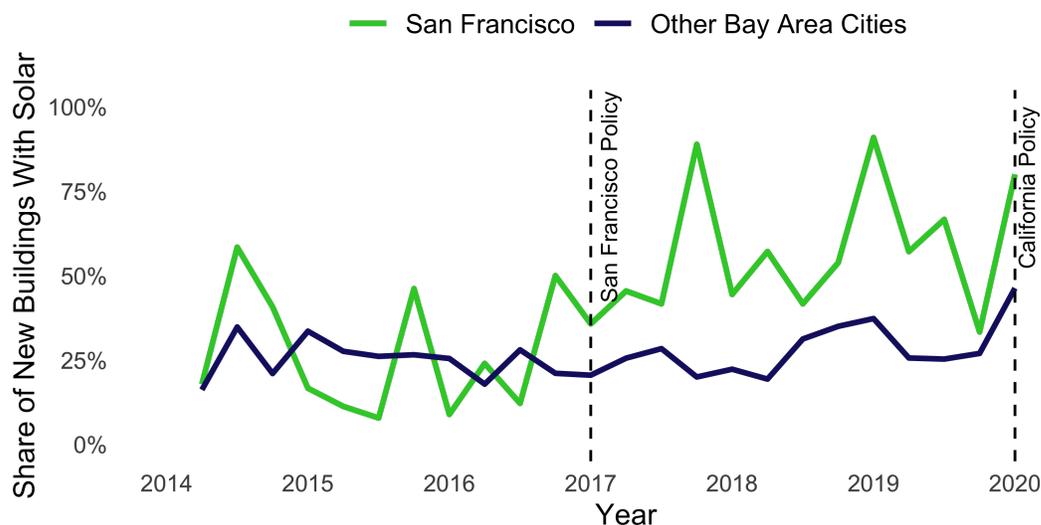
In official statements, the motivations for solar mandates have included reducing electricity production from fossil fuels and electricity system cost-savings and resiliency associated with distributed generation. In regions with cap-and-trade programs, such as California and much of Europe, the extent to which solar mandates may reduce greenhouse gas emissions depends on the exogeneity of the cap. If solar mandates increase the adoption of solar energy, which represents our research question, they then contribute to lower the price of allowances, potentially helping cap-and-trade programs withstand political pressure.

In any case, solar mandates are for some jurisdictions part of the pursuit of more specific outcomes, such as various “net-zero” emissions targets, including California’s goal of zero net energy (ZNE) buildings. Also, informed by the finding of ubiquitous cost-effectiveness in the ex-ante engineering studies, some officials may have had interest in using the mandates to help more of their constituents fully capture the private benefits of solar adoption, including taking maximum advantage of national and regional incentive programs.

The exact requirements and exemptions of solar mandates vary by jurisdiction, so here, we summarize the main features of the San Francisco citywide and California statewide policies with more detailed descriptions available in Appendices A.1 and A.2, respectively. In 2017, San Francisco implemented the Better Roofs Ordinance mandating rooftop solar installations on new construction residential, commercial, and municipal buildings. The policy went into effect for any new construction project that applied for building permits on or after January 1, 2017, with exemptions for buildings over ten stories, data centers, laboratory buildings, and non-residential buildings with

gross floor area less than 2,000 square feet. Minimum size rules required that the solar installations cover at least 15 percent of the total roof area of multifamily and non-residential buildings, or 250 square feet for single-family homes. Aside from rooftop solar PV, the Better Roofs Ordinance allowed two alternative compliance pathways, which were either rooftop solar thermal or building a living roof (also known as green roof) engineered for stormwater management. San Francisco enforced these and other building codes through the building permitting and inspection process, with fines up to \$500 per day for a building code infraction. The Better Roofs Ordinance was superseded in January 2020 when the California Energy Commission introduced a statewide solar requirement in the 2019 Building Energy Efficiency Standards. The statewide standards required all newly constructed low-rise residential buildings (defined as residential buildings up to three stories) to install solar PV. The statewide standards did not apply to commercial and municipal buildings, nor include alternative compliance pathways.

Figure 1: Share of Newly Completed Residential Buildings Adopting Solar



Notes: Solar adoption shares are for completed residential new construction projects. (1) This sample of other Bay Area municipalities omits municipalities that adopted their own solar mandate prior to the 2020 California policy. (2) Incomplete and ongoing building projects are omitted. (3) Dates are based on the year a project permit was first filed, which is also the applicability date for building codes and solar mandates. (4) We omit data after 2020 given the implementation of a statewide mandate in California. (5) See Sections 3.2 and 3.4 for data details and related data summaries.

3 Private payoffs and “solar gap”

3.1 Setup

We begin with a review of time series of new residential construction and solar adoption in municipalities of the San Francisco Bay Area. Sections 3.2 and 3.3 describe our data sources in detail. After reviewing the data, we pursue two distinct sets of analyses parameterizing an engineering model to estimate the private payoffs of a solar installation. First, we parameterize the engineering model to estimate the heterogeneity in the private payoffs of solar adoption, using our project-level data. Understanding the heterogeneity in estimated payoffs and the sensitivity of the model parameterization allows us to explore the importance of accounting for project-level characteristics, the possibility of model misspecification and bias, and the extent that unobserved costs might affect payoffs (including administrative and other “hassle” costs as introduced in Section 2). Second, we compare the engineering estimates to observed solar adoption decisions and test whether variation in project characteristics and payoffs rationalizes solar adoption decisions. Whether or not expected payoffs rationalize solar adoption speaks to the question of whether there is a “solar gap” where builders fail to adopt solar even when it would be cost-effective to do so. Engineering estimates of the private payoffs of a solar installation are not a welfare measure and do not account for social benefits and costs of each solar installation. Expected private payoffs are the measure of interest in the behavior of the private decision-maker, such as the home builder.

For the engineering estimates of private payoffs from a solar installation, we apply the System Advisor Model (SAM) from the U.S. National Renewable Energy Laboratory (NREL). SAM provides a rich representation of a rooftop solar project including geography-specific variation in climate and solar irradiance; installation materials, labor, and financing costs; taxes and subsidies; future electricity consumption and tariff structure; technological specifications of the solar module and inverter; and shading and other system losses. In fact, SAM accommodates dozens of possible user inputs and system configurations. To estimate the private payoff of the solar installation relative to a counterfactual without the solar panels, the SAM outputs include total costs over the system life, electricity production, electricity bill savings, compensation for excess electricity generation, taxes, and subsidies. Principally, SAM estimates the net present value (NPV) of the installation. In Sections 3.2 and 3.3, we detail the specific data sources we use to parameterize the

model, and Appendix Section B.2 provides a more detailed model description. Having estimated the distribution of solar payoffs across all building projects, we are able to illustrate the heterogeneity in project payoffs. We also compare these distributions to payoffs estimated by previous cost-effectiveness reports, which generally only considered a few stylized sets of building characteristics. The comparison to prior engineering studies as well as our sensitivities with alternative SAM parameterizations allow us to assess the importance of using such granular project-level data and accounting for differences in project characteristics. We also compare our results to the finding of earlier cost-effectiveness studies that California solar projects have strictly positive payoffs across project types.

3.2 Building permits data

For our core data, we create a novel address-level dataset of new construction building projects in Bay Area municipalities, which we collect from individual permitting offices of individual municipalities. In addition to new construction permits, we also collect solar and electrical permits.² In matching the solar electrical permits to building permits, we are able to identify the new construction projects that adopt rooftop solar.³ Specifically, we match the new construction building permit addresses to solar permits in the corresponding jurisdiction’s electrical permits. We are then able to validate these solar adoption statistics from individual jurisdictions’ electrical and solar permitting databases against California’s solar interconnection data.⁴ The building permits data from individual municipalities vary in the level of data availability and details of each project’s characteristics. However, the general structure of the data tracks each project’s status from initial filing date through various approvals and inspections to final completion. We restrict the data to only completed projects, for which we are able to observe the final status of a building’s solar adoption decision.

For this analysis and also the San Francisco municipal mandate policy evaluation in Section 4, we sorted municipalities in the San Francisco Bay Area by an equally-weighted ranking of seven key

²In some cases, jurisdictions had permitting processes that tracked solar permits separately from other electrical permit types.

³In the permitting data, we find very few cases of solar thermal adoption. Therefore, we focus almost exclusively on rooftop solar PV adoption except in Section 4 where we evaluate the specific compliance pathways under San Francisco’s citywide mandate.

⁴See <https://www.californiadgstats.ca.gov/> (last accessed, May 9, 2025).

Table 1: Residential New Construction and Solar Adoption, 2014-2016

	2014-2016	2017-2019
San Francisco		
Number of permits	509	116
Number of permits with solar	132	68
Share of permits with solar	25.93%	58.62%
Other Bay Area Cities		
Number of permits	3,567	2,610
Number of permits with solar	905	766
Share of permits with solar	25.37%	29.35%

Notes: (1) This table is restricted to completed residential single-family and multifamily new construction buildings with permits first filed between 2014 and 2019. (2) Dates are based on the year a project permit was first filed. (3) Building permits data were not available for all municipalities in all years.

criteria, including population size, population density, median household income, homeownership rate, solar irradiation, housing stock growth, and population growth. After processing data and removing municipalities with incomplete solar and building permits data, we generated a donor pool with 19 municipalities in the San Francisco Bay Area that most closely resembled San Francisco across these criteria.⁵ Among these municipalities, only San Francisco adopted a rooftop solar mandate prior to the statewide mandate beginning in 2020. Together, these 20 municipalities have a population of over 3.4 million, exceeding the population of some U.S. states and even some small countries. Our building permits data are summarized in Table 1. Our total sample includes 6,802 building permits for residential new construction buildings filed between 2014 and 2019 and subsequently completed. Prior to the implementation of San Francisco’s municipal solar mandate in 2017, solar adoption was only about 25 percent on residential new construction across all municipalities.

⁵Our full list of 20 San Francisco Bay Area municipalities includes San Francisco, Alameda, Berkeley, Burlingame, Cupertino, Daly City, Danville, Emeryville, Fairfield, Fremont, Oakland, Redwood City, Richmond, San Bruno, San Carlos, San Jose, Santa Clara, Santa Cruz, South San Francisco, and Sunnyvale. Note that our donor pool does not include Sebastopol, which implemented a solar mandate of its own prior to 2017. Our donor pool also does not include San Mateo and East Palo Alto for lack of data.

3.3 Engineering model data and parameterization

NREL’s System Advisor Model (SAM) as well as other public and proprietary engineering models are widely available to assist homeowners and firms in their solar installation decisions, and solar salespeople and installers present clients with engineering estimates to inform their solar adoption decision and choose the design parameters. Versions of SAM have been available to the public since 2008 as a user-friendly desktop application and would have been available to the builders during our study period. In our analysis, builder’s solve the problem of maximizing their expected private payoffs, including the solar adoption decision. Accordingly, the first step of our analysis is to parameterize SAM for our dataset of observed building projects using data that would have been available to these builders at the time of their solar adoption decision during their project planning and permitting process. That is, we use our data to parameterize SAM to estimate plausible expectations about the solar payoffs (net present values) for each project at the date the project was permitted. For each building project, we approximate the relevant decision time period as the year the builder first filed for their building permit. We provide additional details of SAM and our parameterization in Appendix B.1, including a formal statement of the builder’s problem in Appendix Equation B.1.

To properly parameterize the engineering model and more fully explore the variation in project payoffs, our analysis supplements the dataset of actual building projects with a rich array of additional data to account for the observed heterogeneity in project characteristics. In particular, we account for variation in climate at a fine scale (i.e., 4km by 4km grid cells); differences in rooftop characteristics and shading; electricity consumption patterns; electricity tariffs; changing state and federal tax incentives; and installation and financing costs. Following from Equation B.1 as introduced in Section 3.1, expected private payoff is the relevant statistic in builder’s solar adoption decision. Table 2 emphasizes the inputs that we are able to parameterize from our data as well as important default parameters to which our payoff estimates are particularly sensitive. In the remainder of this section, we summarize Table 2 and the data sources for the SAM parameterization with Appendix B.1 providing additional data details and citations.

Our installation cost data is from California’s interconnected project sites data, which we use to construct a time series of median cost-per-watt-DC by California city and year. As shown in

Table 2: Engineering Model Summary and Data Sources

Parameter	Data source and description
Building permits	Sample of California building permits described in Section 3.2
Installation costs	Estimated from California interconnection data. Primary parameterizations use city-year medians. Sensitivities include 95 th percentile costs and adjustments accounting for evidence that new construction installations are less expensive than retrofits
Financing costs	Annual time series of SAM defaults. Interest rates vary from 4 to 7.5%
Maintenance costs	SAM defaults
Electricity tariffs and net metering policies	NREL OpenEI database, public time series of PG&E tariffs, and CPUC information on net metering policies. Primary parameterization holds tariffs and policies constant over each project’s analysis period. Sensitivities include alternative tariff and net metering policy expectations
Compensation for excess generation	Public time series of PG&E compensation rates. \$0.03/kWh in primary parameterization, higher in sensitivities
Electricity consumption	PG&E hourly interval data for a sample of new premises. Primary parameterizations use city-level median hourly consumption profiles. Sensitivity parameterizations use address-matched consumption
Federal Investment Tax Credit (ITC)	Federal legislation. 30%
California State and local incentives	Sensitivities account for some additional incentives, such as the New Solar Home Partnership Program (NSHP)
Global horizontal irradiance (GHI)	Typical meteorological year (TMY) files from NSRDB database include GHI (kWh/m ² /day) and other weather variables at 4km×4km grid cell resolution
Shading	Estimated for each building’s roof using Google Sunroof satellite data and modeling
Analysis period	SAM default. 25 years
Solar owner discount rate	SAM default. 6.4%
<u>Decision variables in optimization</u>	
Solar adoption decision	Estimated. $solar_{ict} \in \{0, 1\}$
System size (capacity)	Estimated (kW)
<u>Engineering outputs</u>	
Electricity production	Estimated hourly
Electricity bills with and without solar	Estimated monthly
Net present value (NPV)	Discounted series of estimated costs and payoffs over analysis period

Notes: This table summarizes the data sources in our parameterization of NREL’s System Advisor Model (SAM). (1) The items in this table are only a subset of SAM’s numerous assumptions and available outputs. (2) Unless otherwise stated, we use SAM’s default assumptions. (3) We parameterize the model for each project i in city c where year t is the beginning of the building project (i.e., the year the builder first filed for a building permit). (4) Appendix B.1 provides additional data details and citations.

Appendix Figure B.1, solar costs have fallen substantially over our study period, so it is important to account for variation in builders’ cost expectations over time. Using city-level cost variation accounts for differences in material and labor costs. Section 3.4 and Appendix Section C present estimated payoffs with a variety of alternative cost assumptions, including 95th instead of 50th percentile costs from the interconnection data; 10 percent lower costs; and various alternative incentive levels. One consideration for the new construction mandates is that installation costs are less expensive on new construction than for retrofit projects on existing roofs, and therefore, targeting solar installations to new construction is more cost-effective. Accordingly, Appendix Section C presents results with 10 percent lower installation costs, which modestly increases estimated payoffs.⁶

Sunshine and other climate variables determine the availability of solar energy and panel performance. These data are available in Typical Meteorological Year (TMY) files from NREL’s National Solar Radiation Database (NSRDB), which include an entire year of hourly temperature, global horizontal irradiance (GHI), and other relevant climate variables for a representative year based on weather observations from 1998 to 2022. NSRDB is gridded to 4km by 4km cells, which we match to our geocoded building permits. To determine rooftop shading conditions, we extract a measure of annual “usable sunlight” hours from Google Project Sunroof for each building project in our data. Google Project Sunroof uses satellite data to identify rooftop area and shading conditions, which we use to construct a shading measure for each roof using its “usable sunlight” hours relative to the other roofs in our sample.

California’s net metering policy determines how solar owners are compensated for the electricity they produce and then use for their own consumption as well as any excess electricity they export to the grid. Since California introduced its first net metering policy, NEM 1.0, in 1996, there have been two major changes to the policy accompanied by substantial public debate and uncertainty. The most relevant policy change during our study period occurred in 2016 when California updated its net metering policy from NEM 1.0 to NEM 2.0, which required time-of-use electricity tariffs for solar customers. This policy became effective for customers of PG&E, the utility serving the

⁶Barbose et al. (2013) find evidence that solar installations are less expensive on new construction than retrofit buildings using California data from 2008 to 2012. Barbose et al. find a mean discount of \$0.75 per watt-DC, or approximately 10 percent. While not all solar mandates cost-effectiveness studies have adjusted their installation costs to reflect a new construction discount, one of the cost-effectiveness reports for San Francisco, Halberstadt (2014), did apply an adjustment based on Barbose et al. (2013). In our own data, we find mixed evidence for reduced installation costs associated with new construction, with discount estimates in the 5 to 22 percent range depending on the specification.

Bay Area, on December 15, 2016. Under NEM 1.0, solar customers were effectively compensated for their electricity production at the same retail rate during all hours of the day. Under NEM 2.0, solar customers were compensated at a higher retail rate during “peak” hours and a much lower retail rate during the rest of the day. In our primary parameterizations, builders expect the prevailing NEM policy at the time of their permit application to prevail for the life of their solar installation. Indeed, through the evolution of NEM, California has grandfathered solar systems into their original NEM policies. However, to account for the policy uncertainty around net metering, we have included model parameterizations with alternative policy expectations, which we describe in Section 3.4.⁷

Solar payoffs are largely determined by the interaction of a building’s electricity consumption pattern with the tariff structure and net metering policy. The relatively lower compensation for excess generation under California’s NEM policies described above generally disincentivizes builders from oversizing their solar systems. For the same reason, the largest electricity consumers can realize the greatest payoffs because they can offset their electricity consumption charges at a price closer to the retail electricity rate. Under NEM 2.0, where solar customers face tariffs that vary by time of day, the timing of electricity consumption is also an important determinant of the solar payoff. Historical Bay Area tariffs are available from the electric utility PG&E and archived in NREL’s Utility Rate Database (URDB). For electricity demand, we rely on a sample of restricted-access PG&E residential metering and billing data for new premises in Bay Area municipalities. This sample includes almost 63,000 addresses in 71 Bay Area municipalities with metering and billing data through 2018. The data were associated with premise IDs created in PG&E’s billing system after January 1, 2009, because new premises in PG&E’s billing data are typically associated with new construction. By using city-year sample means, we further increase our confidence that we are identifying electricity consumption associated with new and recently constructed buildings. For our primary parameterization, we use the PG&E sample city mean hourly electricity demand. Hourly consumption data is especially important for estimating payoffs under time-of-use tariffs and period-specific net excess generation exported back to the grid. Appendix Section C compares our payoff estimates from our primary parameterization to an alternative approach where we use building-level hourly interval electricity consumption data from our sample of new PG&E premises. However, our

⁷Appendix B.1 provides additional history and details of California’s evolving NEM policies.

model should reflect a builder’s expectations about the building’s future electricity consumption patterns. These expectations are not necessarily reflected in this alternative parameterization because a snapshot of observed building-level electricity consumption patterns are not necessarily predictive of *future* or even *expected* electricity consumption patterns. We also parameterize our model to 95th percentile load profiles in the PG&E data (at the granularity of city-year) to further demonstrate the importance of expectations about electricity demand.

In addition to California’s net metering policy, there were several other incentive programs available to solar projects during our study period. The 30 percent Federal Investment Tax Credit (ITC) has been available since 2006, so we include this incentive in all of our parameterizations. The well-known California Solar Initiative (CSI) expired for most Bay Area residential customers in 2013, so the incentives were not available for most projects in our sample. Some projects in our sample were qualified for the New Solar Homes Partnership Program (NSHP), which accepted applications through April 1, 2018. NSHP incentives varied from \$0.25 to \$3.50 per watt-DC depending on application timing and project characteristics. Appendix Section C presents alternative parameterizations where we account for the effect of the NSHP program on payoffs to complement our primary parameterization results in Section 3.4. Appendix B.1 provides additional details on solar incentive programs and our parameterization assumptions.

For other parameter values, we use SAM defaults, which NREL updates as frequently as once or twice per year. Updates to SAM defaults are based on market data, user feedback, and NREL staff expert judgement. These default parameters include panel lifetime; operation and maintenance costs; financing costs; and the private discount rate of the decision-maker (i.e., the builder or homeowner). SAM also includes dozens of parameters related to solar panel and inverter technical characteristics and performance.

The default behavior of SAM is to estimate the net present value of a particular project of a particular size. For example, the default system capacity in SAM has ranged from 4 to 8 kW. However, in our SAM parameterizations, we seek to represent the behavior of a payoff-maximizing economic agent. Therefore, we apply an optimization to identify the system size that maximizes payoff conditional on each project’s characteristics following Appendix Equation B.2 as discussed in Section 3.1 and Appendix B.2.

3.4 Evidence on private payoffs

We now present the results of parameterizing the engineering model to estimate the heterogeneity in the private payoffs of solar adoption on new construction buildings. As described in Section 3.3 above, we present distributions of payoffs for our primary parameterization as well as a variety of alternatives. These alternative parameterizations help us explore the robustness of our results, the sensitivity of the engineering model, and the most important determinants of solar payoffs. Lastly, we compare the estimated solar payoffs to the total overall costs of the new construction projects in our data as a possible measure of the relative importance of these solar payoffs in the context of an expensive and complex new construction building project. In this section, we present only payoffs of projects that filed for permits prior to 2017 before the applicability of any solar mandates. We start by focusing on absolute net present values. We then turn to relative benefits with respect to project costs.

From our primary parameterization of the engineering model (SAM) described in Section 3.3 above, in Figures 2a and 2b we estimate that payoffs are positive for every building project in our data. These payoffs range from \$2,700 to \$17,200 for single family homes and exceed \$100,000 for very large multifamily buildings. The multimodality in these distributions is generally driven by the higher electricity consumption patterns and larger homes as well as greater solar irradiance in some sample cities, such as Cupertino, Danville, and Fairfield at the higher end of the payoff distribution.

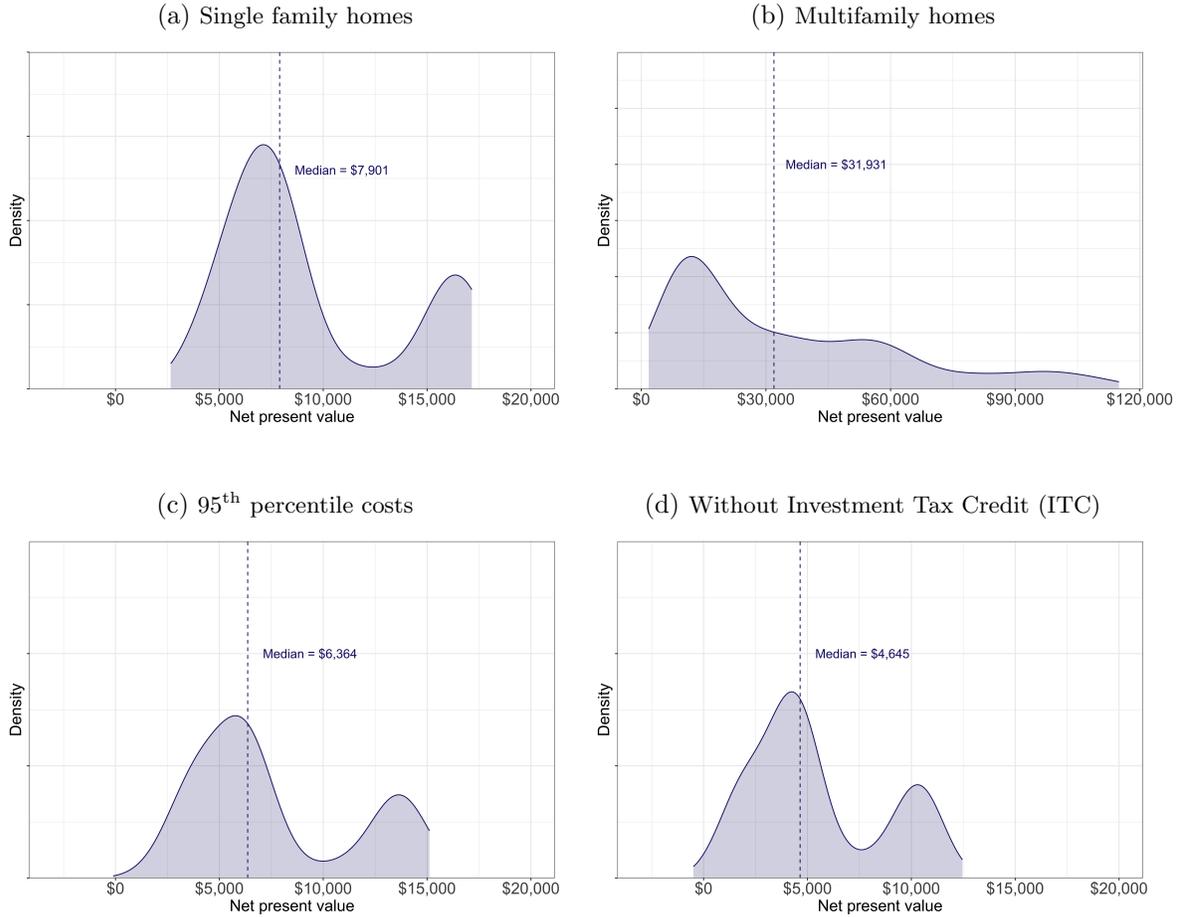
Despite the ubiquity of positive payoffs in our primary parameterization, we do identify negative payoffs for a minority of projects under alternative parameterizations, including higher installation costs, higher financing costs, larger solar systems, lower electricity consumption, and less favorable incentives and tariff structures. For example, the installation costs in the parameterization of Figure 2c are the 95th percentile costs (by city and year) from California’s interconnection data instead of the city-mean used in the primary parameterization. Even in our parameterization yielding the least valuable payoffs, still less than 1 percent of the sample has negative expected payoffs. That said, under these alternative parameterizations, such as a less valuable net metering policy, up to 74 percent of the sample have expected payoffs less than \$5,000, without even accounting for any hassle or other unobserved costs.

As discussed in Section 3.3, there were multiple changes to California’s net metering policy and substantial policy uncertainty during our study period, including uncertainty about the Federal Investment Tax Credit (ITC). Therefore, Figure 2d presents payoffs without the ITC, which substantially reduces solar payoffs. In 2015, Congress ultimately extended the ITC through 2019, still creating potential uncertainty for projects with longer construction timelines. To account for potential changes in California’s net metering policy under consideration prior to the implementation of NEM 2.0 in 2016, Appendix Figure C.3e is parameterized assuming that net metering is replaced with a less favorable net billing policy. Under this counterfactual net billing policy, consumer’s net excess generation is accounted for by month and time-of-day with an \$0.03/kWh compensation rate based on an approximation of the wholesale market. Recall that California’s NEM 1.0 policy already had this lower compensation rate for excess generation on an *annual* basis, discouraging investment in oversized systems. Therefore, while the alternative net billing policy in Figure C.3e reduces payoffs and incentivizes even smaller systems as compared to net metering, expected payoffs remain positive. Appendix Figure C.3d considers expectations of a more favorable net metering policy with compensation for excess generation closer to retail rates, which modestly increases payoffs.

Appendix Figure C.4 compares our payoff estimates from our primary parameterization to an alternative approach where we use building-level hourly interval electricity consumption data from our sample of new PG&E premises. As discussed in Appendix C, the results under this alternative generally confirm this section’s findings in that estimated payoffs are similar on average and continue to be positive for almost all projects. However, Appendix Figure C.4 does suggest that there may be even more heterogeneity in payoffs across buildings than suggested in Figure 2.

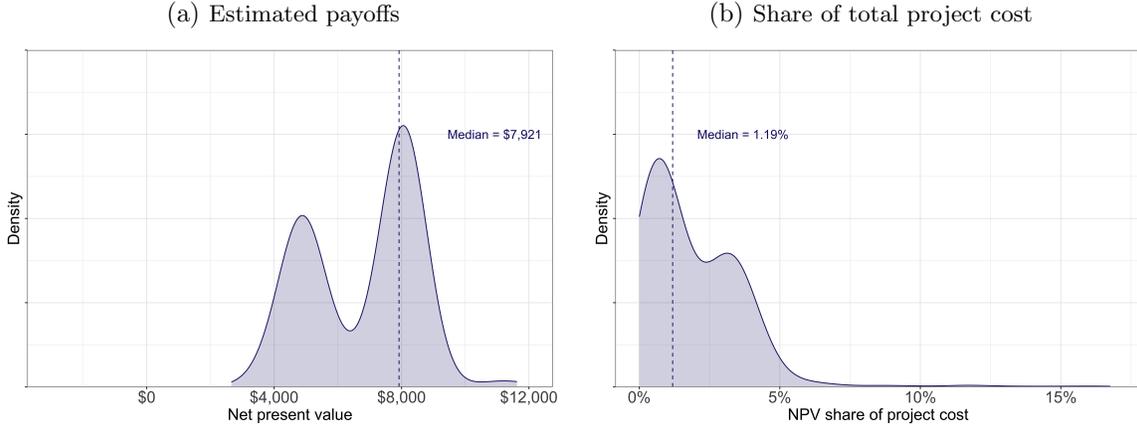
Figure 3 illustrates the possibility that solar payoffs may have low relative importance for many builders. Specifically, Figure 3a presents estimated payoffs under our primary specification for a subset of building permits for which permits data included estimated *total* new construction costs. Figure 3b then presents those payoffs as *shares* of the total project costs. The median residential building cost is \$500,000, so for the majority of projects, median payoffs are less than 2 percent of project costs. For almost all projects, solar payoffs are less than 5 percent of their project costs. Among the other demands of project management, builders may ignore the foregone solar payoffs, which are small even without accounting for administrative and hassle costs.

Figure 2: Estimated Distributions of Solar Payoffs



Notes: (1) Figures present medians and smoothed kernel density estimates of the distribution of estimated solar payoffs (net present values) from our parameterization of SAM across our sample of San Francisco Bay Area new construction building permits. (2) All distributions are for building permits issued between 2014 and 2016 (inclusive), prior to the implementation of solar mandates. (3) Panel 2a presents payoffs estimated under our primary parameterization, assuming all payoffs accrue to a single unit or building owner (i.e., single family homes). (4) Panel 2b is restricted to multifamily buildings where solar may offset electricity consumption at multiple units, though system size is bounded by the roof area observed from Google Sunroof. A small number of very large projects (fewer than 100) are omitted by the truncation of the x-axis. (5) Panel 2c is based on 95th percentile installation costs (by city and year) instead of 50th percentile costs in the primary parameterization. (6) Panel 2d assumes the expiration of the Federal Investment Tax Credit. (7) Axes scales differ. All distributions are truncated to show the actual estimates of minima and maxima.

Figure 3: Payoff Importance: Estimated Distributions of Solar Payoffs and Overall Project Costs



Notes: (1) In contrast to Figure 2, this figure is restricted to new residential buildings for which project cost data is available. (2) Panel 3a presents the estimated solar payoffs (net present values). (3) Panel 3b presents the solar payoffs as a share of total cost for the entire building project. (4) Both panels present medians and smoothed kernel density estimates of the distribution of estimated solar payoffs from our parameterization of SAM across our sample of San Francisco Bay Area new construction building permits. (5) All distributions are for building permits issued between 2014 and 2016 (inclusive), prior to the implementation of solar mandates.

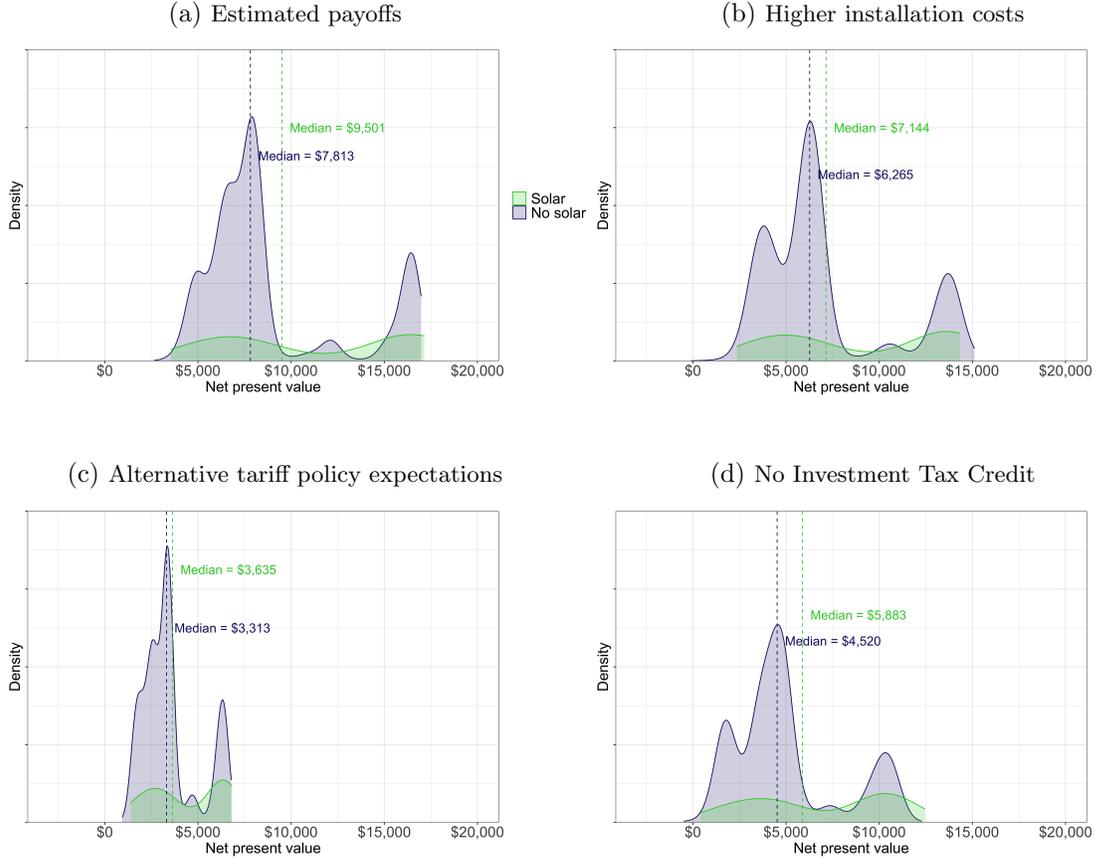
To summarize, we gather the following set of findings from Figures 2 and 3. First, we observe that there is substantial heterogeneity in payoffs across project characteristics, which underscores the value of project-specific data as opposed to the common practice in the engineering literature used for ex-ante policy simulations. Rather than detailed project-level data, these cost-effectiveness studies relied on a small number of representative projects with stylized characteristics. Second, we note that estimated payoffs are indeed positive for the great majority of the buildings in our study period, and this finding is robust to a battery of alternative parameterizations. Taken at face value, this finding supports the conclusions of the cost-effectiveness reports. Third, we also observe that the payoffs are modest, especially relative to the overall costs of new construction buildings. Even a relatively small amount of unobserved costs could exceed the estimated payoffs. For example, an extra few thousand dollars of administrative or architectural costs associated with the solar installation would exceed the payoffs for many projects.

3.5 Evidence of a “solar gap”

Next, we apply our estimated payoffs to assess the question of whether a solar gap exists. Specifically, we test whether expected payoffs determine builder solar adoption decisions following Equations B.1 and B.2 and the discussion in Section 3.1. The failure to adopt solar in the presence of positive estimated payoffs may be rationalized by additional characteristics unobserved in our data but considered by the builder, such as some aesthetic rooftop features that preclude solar. However, even in the circumstance of unobserved project characteristics or additional costs, we still expect a relationship, everything else equal, between estimates of net present values and observed solar adoption decisions.

Figure 4 presents the distributions of expected solar payoffs among new building projects adopting solar as compared to buildings that do not. As in Section 3.4, we continue to only consider years 2014 to 2016, before San Francisco implemented its mandates (and before California implemented the statewide policy). In Figure 4, we observe no apparent relationship between estimated payoff and the solar adoption decision, and building projects fail to adopt solar across the entire distribution of expected payoffs. Figure 4 suggests no evidence that the level of payoffs are a determinant of adoption. One exception is in the tails of the distribution where we observe that the very lowest NPV projects forego solar. That is, for the very lowest value project with the lowest expected payoffs, builders are least likely to adopt solar, especially after accounting for non-monetary costs. However, the evidence that many builders are foregoing profitable solar suggests a true “solar gap.” These findings are robust to a variety of alternative parameterizations, including higher installations costs (Figure 4b); expectations of a less valuable net metering policy (Figure 4c); and expectations of an expiring ITC (Figure 4d) as described in Section 3.4.

Figure 4: Estimated Distributions of Solar Payoffs and Observed Solar Adoption



Notes: (1) Figures present medians and smoothed kernel density estimates of the distribution of estimated solar payoffs (net present values) from our parameterization of SAM across our sample of San Francisco Bay Area new construction building permits. (2) The green-shaded distributions are for projects we observed adopting solar in the building permits data. Blue-shaded distributions are estimated payoffs for non-adopters. (3) All distributions are for new single family buildings with permits issued between 2014 and 2016 (inclusive), prior to the implementation of solar mandates. (4) Panel 4b is based on 95th percentile installation costs (by city and year) instead of 50th percentile costs in the primary parameterization. (5) Panel 4c assumes that net metering is replaced by a less valuable net billing policy. (6) Panel 4d assumes the expiration of the Federal Investment Tax Credit. (7) All distributions are truncated to show the actual estimates of minima and maxima.

Table 3 presents the results of estimating the linear probability model of the relationship between observed solar adoption decisions and the engineering model estimates of solar payoffs. The estimating equations are described in more detail in Appendix C. Confirming the visual evidence of Figure 4, we fail to identify a strong positive relationship between expected payoffs and observed adoption. To avoid the effect of solar mandates, these estimations only include the pre-policy years,

2014 to 2016. Across all models, the relationship between NPV and solar adoption is relatively modest — a \$1,000 change in NPV is associated with a smaller than 6 percent change in probability of solar adoption. NPV explains less than 8% of the variation observed in the data. Recall that in our sample we have almost ubiquitously positive expected payoffs, yet only 25 percent adoption. The inclusion of city and year fixed effects in models (2) and (3) substantially increases the standard errors, as these fixed effects also remove much of the variation in our engineering model parameterization. Principally, though, these estimations tend to lend support to the idea that there can be builders who forego profitable solar adoption, in line with the notion of a “solar gap.” While our results are directionally similar to the case of retrofits where adoption tends to follow expected payoffs (e.g., Tibebu et al. 2021), the solar adoption decision for a builder may differ substantially from a retrofit. As shown in Figure 3, the costs of building projects vastly exceed the cost of solar PV in our data, suggesting a different decision-making process.

Table 3: Linear Probability Models of Solar Adoption

	(1)	(2)	(3)
Net present value (NPV) (\$1,000s)	0.028***	0.057**	0.055*
	(0.002)	(0.021)	(0.023)
Income (\$1,000s)			0.002***
			(0.000)
City fixed effects	No	Yes	Yes
Year fixed effects	No	Yes	Yes
Num.Obs.	3632	3632	3553
R2	0.078	0.159	0.170
R2 Adj.	0.078	0.154	0.165

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: (1) Model 1 estimates the linear probability model (Equation C.3) of solar adoption regressed on estimated solar payoff (net present value) from the engineering model (SAM). (2) Observations are residential new construction building permits from our sample of Bay Area municipalities during the pre-policy years (2014 to 2016). (3) Model 2 adds city and year fixed effects. (4) Model 3 includes median income at the granularity of Census block group. (5) All models exclude buildings we identify as large multifamily buildings with more than four units due to their substantially higher potential payoffs.

In sum, as discussed in Section 3.4 above, we conclude that in most cases in our context, solar

payoffs are expected to be positive, supporting one of the rationales for solar mandates. However, the payoffs are also relatively small, suggesting that this rationale for mandates is weaker once considering payoffs relative to the overall costs and complexity of a building project as well as the associated administrative and hassle costs. At the same time, we do find evidence supporting the idea of a “solar gap,” whereby builders leave money on the table as conjectured in the cost-effectiveness reports. In fact, we find only modest evidence that larger estimated payoffs drive new construction solar adoption. However, the extent of the “solar gap” is relatively limited, bounded above by the most optimistic calculations of net present values provided in this section and Appendix C.

4 Policy evaluation

In this section, we describe the application of synthetic control methods to evaluate the San Francisco Better Roofs policy and the California solar mandate. We apply synthetic control methods consistently to the evaluation of both policies, including using the same sets of predictor variables.

4.1 Municipal mandate in San Francisco

We begin with the 2017 San Francisco solar mandate for newly constructed buildings. The outcome variable of interest is the share of new buildings with solar photovoltaics. The synthetic control method, introduced by Abadie and Gardeazabal (2003) and further developed in Abadie et al. (2010, 2015), provides a systematic way to construct a comparison unit using weighted combinations of control units. The synthetic control is constructed to approximate the characteristics of the treated unit before the intervention, providing a data-driven approach to counterfactual analysis.

The synthetic control is constructed by solving:

$$\min_W \sum_{k=1}^K v_k (X_{1k} - \sum_{j=2}^{J+1} w_j X_{jk})^2$$

subject to $w_j \geq 0$ and $\sum_{j=2}^{J+1} w_j = 1,$

where X_{1k} represents the value of predictor k for San Francisco, X_{jk} is the value for control unit j , v_k is the weight assigned to predictor k , and w_j is the weight assigned to control unit j . In this setting, the control unit geographies are municipalities.

As described in Section 3.2, recall that we constructed an initial donor pool by sorting municipalities in the San Francisco Bay Area by an equally-weighted ranking of seven key criteria, including population size, population density, median household income, homeownership rate, solar irradiation, housing stock growth, and population growth. After processing data and removing municipalities with incomplete solar and building permits data, we generated a donor pool with 19 municipalities in the San Francisco Bay Area that most closely resembled San Francisco across these criteria. Section 3.2 provides additional details on the building permits data. The donor pool municipalities are listed in Table 4. Prior to San Francisco’s policy implementation in 2017, most municipalities have relatively low levels of solar adoption on new residential construction, with a donor pool average of 15 percent. On a permit-weighted basis, however, the share is 25 percent as shown in Section 3.2. The cities with the lowest solar adoption shares, such as Berkeley and Emeryville, generally have the fewest new residential building permits during this period. Fairfield (51 percent), Santa Clara (36 percent), and Santa Cruz (35 percent) have the highest adoption shares.

To perform the synthetic control procedure with the San Francisco and donor pool building permits data, we select a number of predictors related to residential demographics as well as the markets for new construction and rooftop solar. Our data spans January 2014 to December 2019, with the outcome measure aggregated at the quarterly level. Our predictors are primarily from the U.S. Census American Community Survey (ACS) 5-year estimates, including owner-occupied housing share, electric heating housing share, average household size, total population, growth rate in the stock of housing units, and educational attainment (share of adults with bachelors degree or higher). For a measure of each municipality’s solar production potential, we use global horizontal irradiance (GHI) from NREL’s National Solar Radiation Database (NSRDB). In particular, we construct a population-weighted mean GHI from the centroid of each municipality.

Table 4: San Francisco Mandate: Donor Pool Characteristics

Municipality	Population	New construction solar share (%)
San Francisco	840,039	27.5
Alameda	76,208	3.9
Berkeley	74,023	0.0
Burlingame	29,892	20.8
Cupertino	61,229	14.2
Daly City	97,322	8.9
Danville	43,353	17.8
Emeryville	47,684	0.0
Fairfield	109,468	50.9
Fremont	162,900	14.6
Oakland	407,484	29.3
Redwood City	76,463	0.0
Richmond	102,049	5.0
San Bruno	63,502	4.6
San Carlos	36,846	14.1
San Jose	998,848	13.4
Santa Clara	121,208	35.6
Santa Cruz	62,702	34.5
South San Francisco	66,764	0.0
Sunnyvale	147,898	17.0
Donor Pool Average	146,623	15.0

Notes: This table lists our donor pool cities, their populations, and their new construction solar adoption shares. The shares of solar adoption on new construction are for the pre-policy period (i.e., before 2017).

Drawing from the donor pool in Table 4 and the predictors described above, we now discuss our construction of synthetic San Francisco. Table 5 shows that the synthetic control successfully replicates San Francisco’s pre-policy characteristics across important predictors. Specifically, the synthetic control matches San Francisco’s population fairly well (524,094 for the synthetic control vs. 839,108 for San Francisco), and tracks the housing unit growth rate (0% vs 0.1%) and electric heating housing share (31% vs. 30%) closely. As can be expected with synthetic control, when San Francisco tends to be an outlier, some imbalances remain, particularly for owner-occupied housing units (43% vs. 36%) and education (39% of adults with a bachelor degree vs. 54%).

Table 5: San Francisco Mandate: Balance Table for Predictors

Predictor	San Francisco	Donor Pool Avg.	Synthetic Control
Total population	839,108	181,294	524,094
Solar irradiance (GHI)	4.72	4.98	4.97
Housing unit growth rate	0.01	0.01	0.00
Electric heat housing units (share)	0.30	0.29	0.31
Owner occupied housing units (share)	0.36	0.50	0.43
Average household size	2.32	2.68	2.65
Bachelor degree (share)	0.54	0.49	0.39

Notes: For each of seven predictors used in our synthetic control, this table compares San Francisco characteristics to the donor pool averages as well as the synthetic control corresponding to the predictor and geographic weights in Table 6.

Table 6 presents the composition of synthetic San Francisco, including the predictor and geography weights. Two primary municipalities are used to construct the counterfactual. The synthetic control is composed of San José (19.8%) and Oakland (80.2%), the Bay Area cities with the highest and third-highest population, respectively. The most important predictors are total population (28.2%), share of owner occupied housing units (20.1%), and share of housing units with electric heat (16.5%).

Complementing the comparison of actual and synthetic San Francisco characteristics in Table 5, Figures 5 and 6 further illustrate the pre-period fit of our synthetic control. In Figure 5, the observed solar adoption share in San Francisco (solid green line) closely tracks synthetic San Francisco (dashed blue line) in the quarters preceding implementation of San Francisco’s Better Roofs Ordinance. Figure 6 then plots the differences between the two series, which are well-centered

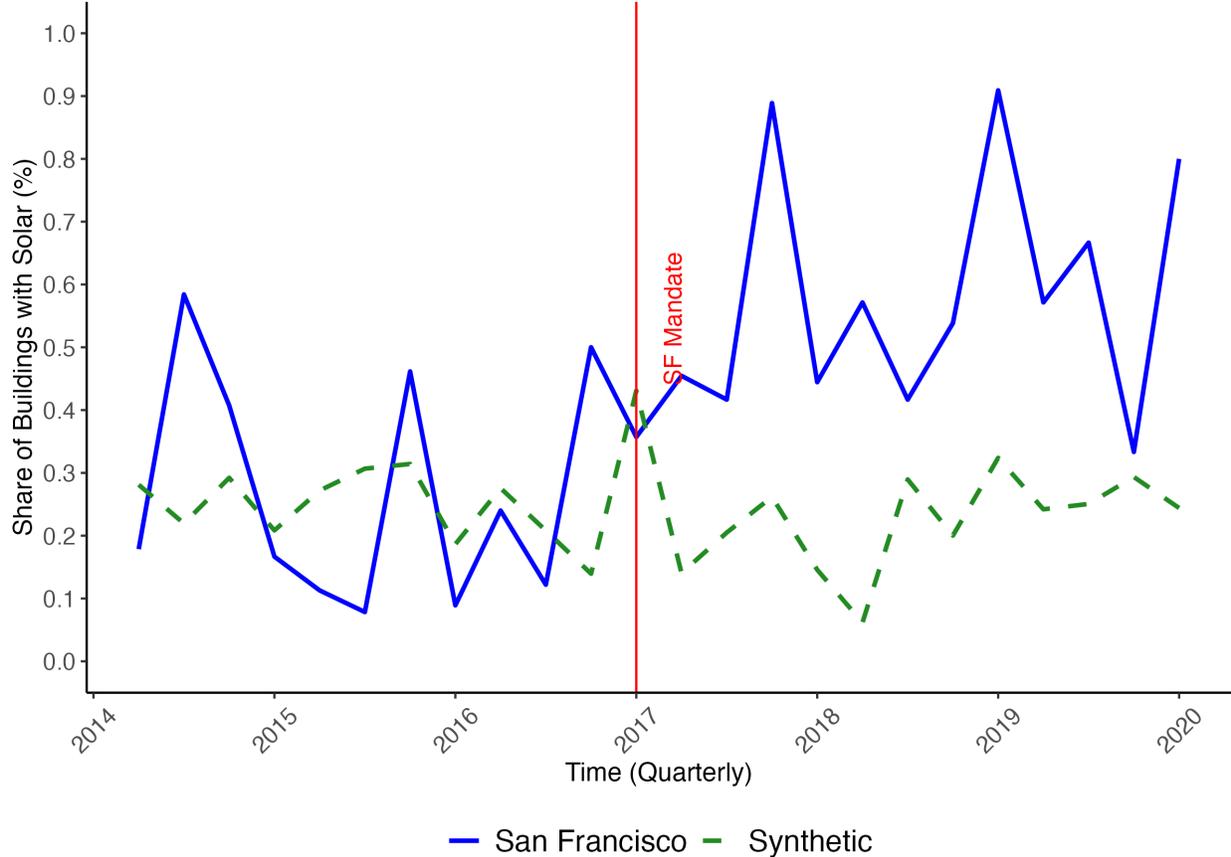
Table 6: San Francisco Mandate: Composition of the Synthetic Control and Estimated Effect

Estimate (Δ solar share on new construction)	0.33
<i>Standard Error</i>	(0.147)
<u>Predictor weights:</u>	
Total population	0.282
Owner occupied housing units (share)	0.201
Electric heat housing units (share)	0.165
Solar irradiance (GHI)	0.135
Bachelor degree (share)	0.112
Average household size	0.071
Housing unit growth rate	0.033
<u>Geographic weights:</u>	
Oakland	0.802
San Jose	0.198
Temporal scale	Quarterly

Notes: (1) This table presents the predictor and geographic weights used to construct the synthetic control to evaluate the San Francisco mandate. (2) The table also presents the estimated effect of the San Francisco mandate. The estimated effect is the change in solar share on new construction after the mandate in the observed data compared to the synthetic control. This estimated effect is illustrated graphically in Figures 5 and 6. (3) The table omits municipalities with zero or close-to-zero weights.

around zero in the pre-period. The relatively small and stable gap during the pre-treatment period supports our synthetic control construction, suggesting it effectively captures San Francisco’s solar adoption patterns before the mandate.

Figure 5: San Francisco Mandate: Actual and Synthetic San Francisco Solar Shares on New Construction

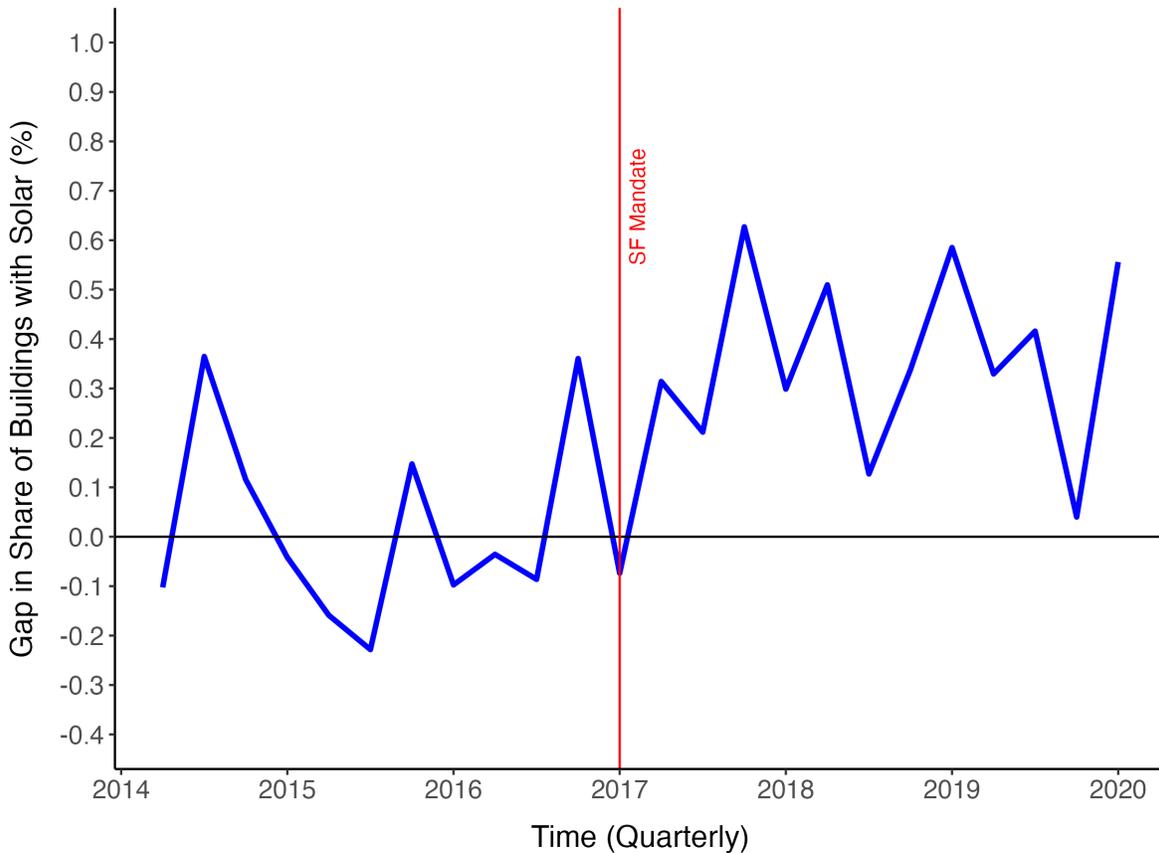


Notes: (1) This figure plots the observed time series of San Francisco solar shares on new construction as well as the synthetic control corresponding to the predictor and geographic weights in Table 6. (2) The time on the x-axis reflects building permit issuance, which is the relevant date for the policy.

Figures 5 and 6 show that San Francisco’s actual solar adoption shares increase markedly after the 2020 mandate implementation as compared to the synthetic counterpart, implying a clear and sustained divergence and positive gap. Table 6 presents the value of this estimated policy effect—an increase in the share of residential new construction solar adoption by 33 percentage points (120%),

with a standard error of 0.147. While post-treatment adoption at 55% is still far short of 100%, this outcome is consistent with the flexible policy design of the Better Roofs Ordinance, which allowed alternative compliance and exemptions as described in Section 2.2 and Appendix A.1. We return to the question of compliance shortly.

Figure 6: San Francisco Mandate: Gap Between Actual and Synthetic San Francisco Solar Shares on New Construction

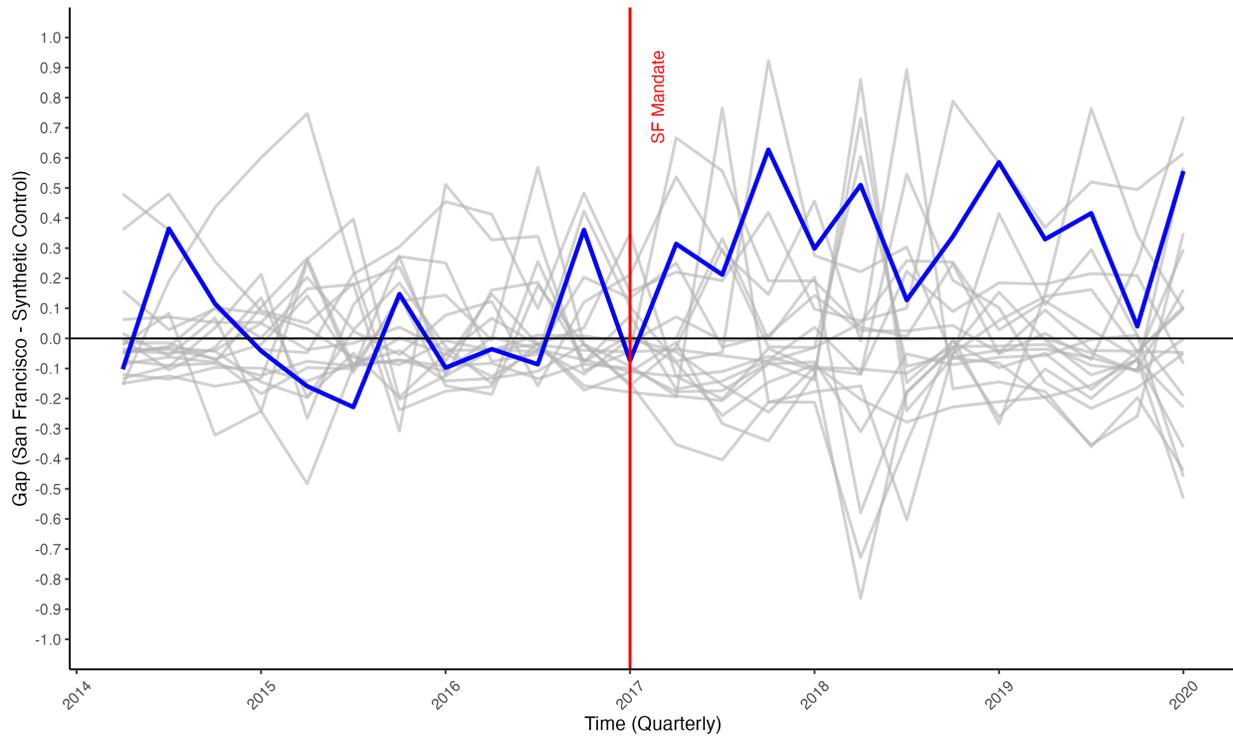


Notes: This figure plots the difference (or “gap”) between the observed time series of San Francisco solar shares on new construction and the synthetic control. These results correspond to the series in Figure 5 and the estimated effect in Table 6.

Abadie et al. (2010) recommend placebo tests to evaluate the robustness of synthetic control results. Figure 7 implements this placebo test by applying the synthetic control method to each control municipality in our donor pool. The blue line represents San Francisco’s gap, while the gray lines show the placebo gaps for control municipalities. The magnitude and consistency of

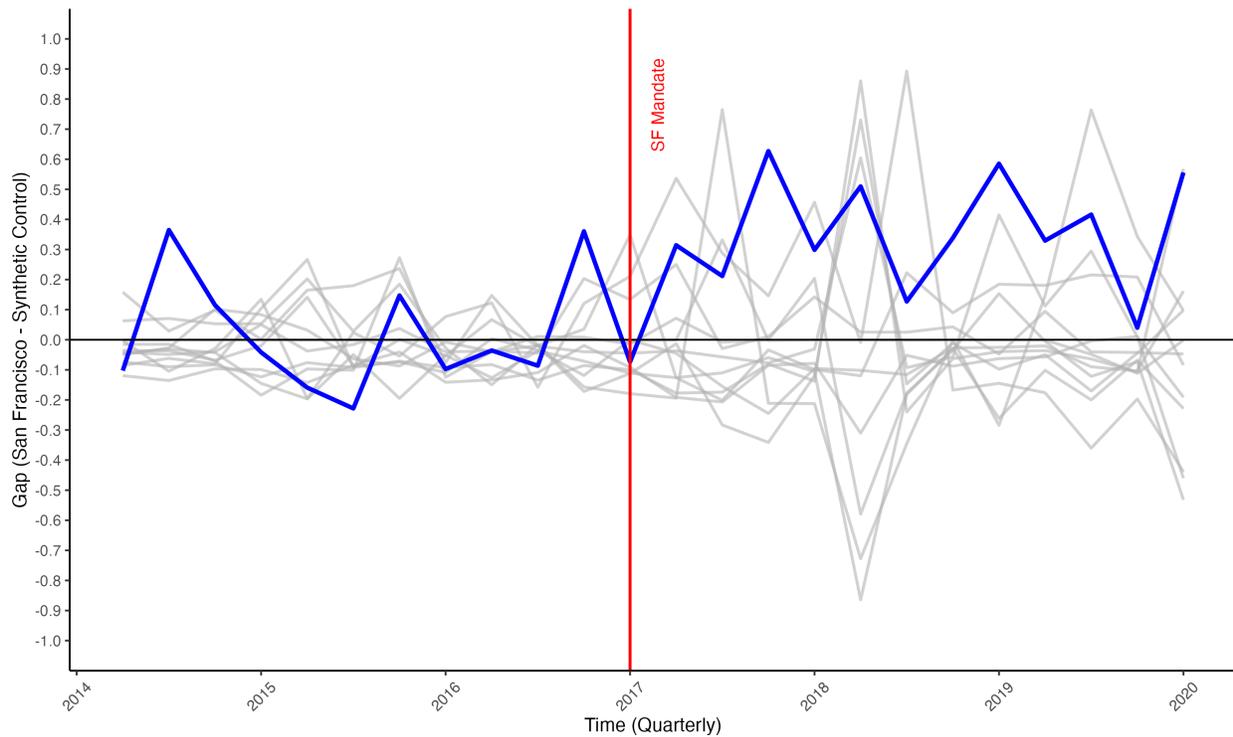
San Francisco's positive treatment effect stands out relative to the distribution of placebo effects, underlining the conclusion that the Better Roofs policy effectively increased the adoption of solar photovoltaics on new construction. Figure 8 shows the same placebo test for a subset of the donor pool, which has a pre-treatment mean squared prediction error no larger than San Francisco's.

Figure 7: San Francisco Mandate: New Construction Solar Share Gap and Placebo Gaps Using All Donor Pool Municipalities



Notes: This figure compares the gap produced from our synthetic control (also plotted in Figure 6) to a series of placebos where the synthetic control methodology is applied to each of the municipalities in the donor pool.

Figure 8: San Francisco Mandate: New Construction Solar Share Gap and Placebo Gaps Using Selected Donor Pool Municipalities



Notes: This figure compares the gap produced from our synthetic control (also plotted in Figure 6) to a series of placebos where the synthetic control methodology is applied to each of the municipalities. We discard control municipalities with a pre-treatment mean squared prediction error (MSPE) less than half of San Francisco's.

Complementing our main result using synthetic control, we implement a synthetic difference-in-difference (SDID) estimator (Arkhangelsky et al. 2021) to evaluate the effect of San Francisco's solar mandate on the share of new buildings with solar permits. Table 7 presents this SDID estimate. Table 7 shows that SDID estimates an impact of 25.4 percentage points (92%), slightly lower but statistically indistinguishable from the estimate of 33 percentage points from the standard synthetic control method, given a standard error for SDID of 0.115. The SDID estimate confirms that San Francisco's mandate substantially increased solar adoption on new construction. Notably, the geographic weights in the SDID approach are more evenly distributed across donor municipalities compared to the synthetic control method, with the largest weights assigned to Emeryville (16.6%), Danville (8.6%), and Burlingame (7.4%). The predictor weights also differ, with the population with

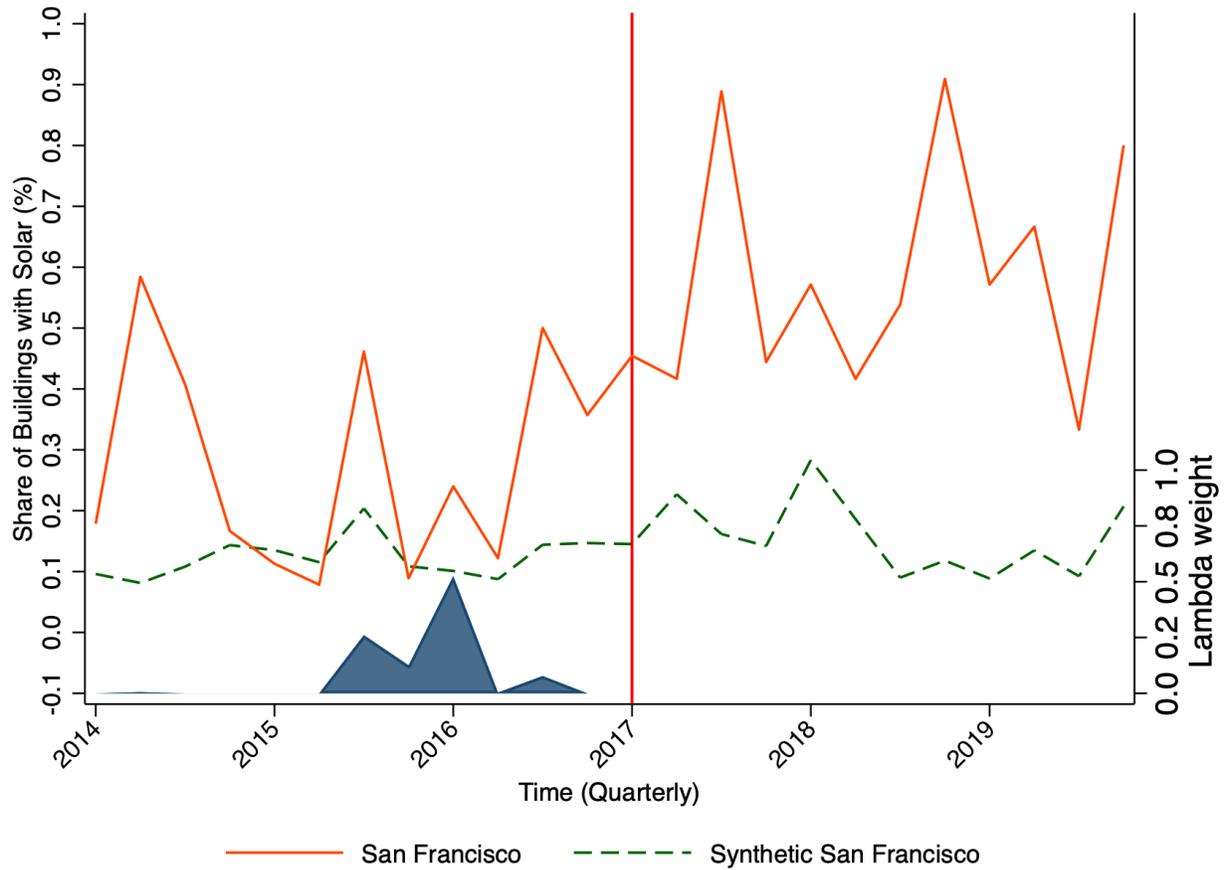
a bachelor degree (14.1%) and the share of housing units with electric heat (12.3%) receiving the highest weights, reflecting the SDID method's optimization that differs from standard synthetic control. Figure 9 illustrates our SDID to complement the results in Table 7. In Figure 9, we observe a clear divergence post-treatment, with substantially higher adoption of solar among new construction in San Francisco compared to the counterfactual.

Table 7: San Francisco Mandate: Composition of the Synthetic Difference-in-Differences Control and Estimated Effect

Predictor / Municipality	Weight
Estimate (Δ solar share on new construction)	0.254
<i>Standard Error</i>	<i>0.115</i>
<hr/>	
<u>Predictor weights:</u>	
Bachelor degree share	0.141
Electric heat housing units (share)	0.123
Owner occupied housing units (share)	0.038
Housing unit growth rate	0.037
Total population	0.017
Solar irradiance (GHI)	0.002
<hr/>	
<u>Geographic weights:</u>	
Emeryville	0.166
Danville	0.086
Burlingame	0.074
Daly City	0.067
Fairfield	0.067
Richmond	0.063
Redwood City	0.062
Oakland	0.061
San Bruno	0.060
San Carlos	0.053
Berkeley	0.052
Alameda	0.046
San Jose	0.037
Cupertino	0.024
Fremont	0.014
Santa Clara	0.001
<hr/>	
Temporal scale	Quarterly

Notes: (1) This table presents the predictor and geographic weights used to construct the synthetic difference-in-differences estimator to evaluate the San Francisco mandate. (2) The table also presents the estimated effect of the San Francisco mandate. The estimated effect is the change in solar share on new construction after the mandate in the observed data compared to the synthetic control. This estimated effect is illustrated graphically in Figures 9.

Figure 9: San Francisco Mandate: Actual and Synthetic Difference-in-Differences Solar Shares on New Construction

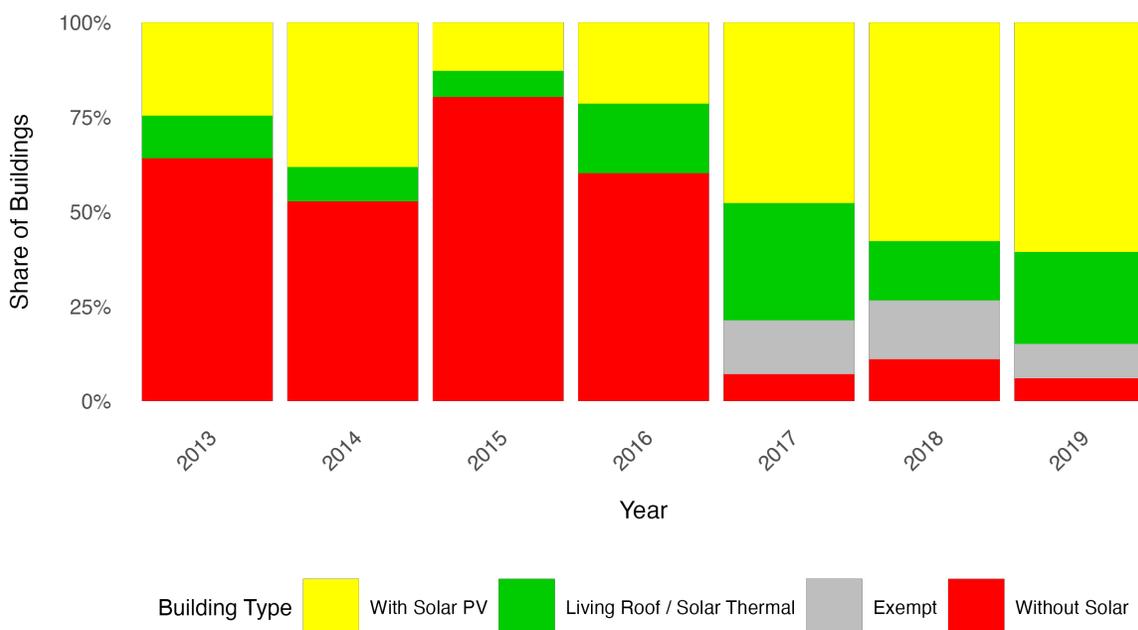


Notes: (1) This figure shows San Francisco’s actual solar adoption share (solid orange line) compared to the synthetic difference-in-differences control (dashed green line), with corresponding geographic and predictor weights in Table 7. (2) Shaded triangular regions represent lambda time weights assigned by the SDID procedure to improve pre-treatment fit. (3) The red vertical line marks policy implementation in 2017. (4) Standard errors calculated using 1,000 placebo replications.

The data that we have for San Francisco allow us to measure compliance with the citywide mandate. San Francisco’s Better Roofs Ordinance established three compliance pathways: solar photovoltaic installations, solar thermal systems, or engineered living roofs for stormwater management. This flexible design permits examination of both overall compliance rates within the scope of the mandate and builders’ revealed preferences among technologies following mandate implementation. Exemptions are also granted on a case-by-case basis. Hence, any residual non-compliance may be explained by such exemptions, narrowing the scope of the mandate.

The ordinance applies to residential buildings of ten occupied floors or fewer without minimum size thresholds, and non-residential buildings exceeding 2,000 square feet of gross floor area. Mixed-use buildings, common in San Francisco’s dense urban environment, may qualify for exemptions depending on their primary classification. Additional exemptions are made on a case-by-case basis for buildings with operational constraints.

Figure 10: San Francisco Better Roofs Ordinance Compliance by Compliance Method, 2013-2019



Notes: This figure presents the distribution of compliance outcomes for new construction projects in San Francisco from 2013 through 2019. Solar thermal installations are included with living roofs due to their limited adoption and frequent combination with living roofs. Data on new building construction and solar photovoltaics are from building permits submitted to the City of San Francisco. Supplementary public data (architect/contractor project descriptions, zoning classification, and satellite images) are used to identify buildings with living roofs and buildings that qualify for exemptions.

Figure 10 presents compliance outcomes for new construction projects from 2013 through 2019. The ordinance applies based on initial application dates, regardless of when the building project is completed and projects applying before January 1, 2017 remain exempt. Prior to 2017, voluntary solar and living roof adoption remained below 50 percent. Post-mandate adoption increased substantially, reaching approximately 80-85 percent by 2018-2019. The remaining buildings fall

into two categories. Some qualify for explicit exemptions under the ordinance. Only a small share cannot be verified to have compliant green technologies or an exemption.

These compliance patterns contextualize our synthetic control estimates as provided in this section. The estimated increase in solar adoption is congruent with the ordinances' flexible pathways for compliance, with approximately 58% of new construction featuring solar photovoltaics between 2017 and 2019. Compared to the pre-period average of approximately 25%, this represents more than a twofold increase in the share of new construction with solar photovoltaics in San Francisco compared to the pre-period.

California's statewide mandate implementation in January 2020 significantly altered San Francisco's requirements. The city deferred to state standards for residential buildings of three stories or fewer while maintaining local requirements for taller residential and non-residential projects. California's 2022 energy code expansion, based on Solar Access Roof Area calculations rather than occupied floors, prompted San Francisco to largely withdraw local requirements, retaining only limited provisions for living roof alternatives where they reduce required photovoltaic area. We now turn to the evaluation of the statewide mandate.

4.2 Statewide mandate in California

In this section, we evaluate California's statewide rooftop solar mandate, using synthetic control methods as in our evaluation of the San Francisco mandate in Section 4.1. We begin with constructing a synthetic control for California from a donor pool of the other states, which have not adopted such mandates. At the state level, we do not directly observe the share of solar installations on new construction buildings. However, we can construct a proxy measure, which is residential solar PV capacity additions (kW) per residential new construction building permit issued. In the construction of this measure of solar adoption, the numerator includes total additions to residential solar capacity, not only those associated with new construction. Therefore, in this section, we are interested in assessing changes in the ratio between new solar capacity additions and new construction.

We observe the monthly count of new residential buildings permits issued by state from the U.S. Census Building Permits Survey (BPS), and we observe monthly new residential solar capacity changes from the U.S. Energy Information Administration EIA-861M Monthly Electric Power In-

dustry Report. Our solar adoption measure is then the amount of added residential solar capacity per newly issued residential building permit (i.e., kW per building permit). Of course, there are lags between the issuance of a new building permit, the construction of a new residence, and the interconnection of a solar system. Our review of San Francisco Bay Area building permits identified a lag of generally 1 to 2 years between permit issuance and solar adoption, with some projects taking much longer. Therefore, more precisely, our solar adoption measure d_q is the amount of added solar capacity kW_q in quarter q per newly issued building permit b_q , 2 to 6 quarters prior. That is, our outcome variable is $d_q = \frac{kW_q}{(\sum_{t=q-6}^{t=q-2} b_t)/4}$. Since there are many more building permits at the state level than in our analysis of the San Francisco citywide mandate, we also have the option of monthly specifications where our outcome variable d_m in month m is $d_m = \frac{kW_m}{(\sum_{t=m-18}^{t=m-6} b_t)/12}$.

Similar to our evaluation of the San Francisco mandate, for the construction of our synthetic control for California, we select a number of predictors related to residential demographics as well as the markets for new residential construction and rooftop solar, given the outcome variable. Our predictors are primarily from the U.S. Census American Community Survey (ACS) 5-year estimates, including owner-occupied housing share, electric heating housing share, average household size, total population, growth rate in the stock of housing units, and educational attainment (share of adults with bachelors degree or higher). For a measure of each state’s solar production potential, we use global horizontal irradiance (GHI) from NREL’s National Solar Radiation Database (NSRDB), as previously described. In particular, we construct a population-weighted mean GHI from the centroid of the state’s three most populous counties to help ensure a relevant irradiance measure for each state as opposed to an alternative measure that might be biased by meteorological conditions in rural or undeveloped regions.

After processing data and removing states with incomplete solar and building permits data, we select a donor pool with 17 jurisdictions, which we list in Table 8. We restrict our donor pool to coastal states to ensure comparability with California on dimensions that affect both solar adoption and housing construction patterns. California’s climate creates structural differences in electricity generation and consumption that confound comparisons with inland states (Levinson 2014). Additionally, geography may constrain developable land and affect housing supply elasticity (Saiz 2010), which may influence builder investment decisions. By limiting our donor pool to coastal states, we construct a synthetic control that better reflects California’s climate and housing and

electricity market conditions.

Our data span January 2017 to July 2022. Table 8 shows that our average pre-policy solar adoption measure for California is 19.4 kW per new residential building permit, substantially higher than most donor pool states. As discussed above, recall that the numerator in our solar adoption measure is total residential solar capacity additions, so this value is not to be interpreted as the average size of a residential new construction solar installation. The magnitude of the measure reflects California's leading position in residential solar adoption prior to the mandate. Other jurisdictions with favorable conditions for solar energy and state-level policies supporting its adoption have comparable values, including Hawaii and New York.

Drawing from the donor pool in Table 8 and the predictors described above, we now discuss our construction of synthetic California. Table 9 shows that the synthetic control successfully replicates California's pre-policy characteristics across important predictors.

Specifically, the synthetic control closely matches California's share of owner-occupied housing units (57.5% vs. 55.6%), housing unit growth rate (both 0.6%) and the share of bachelor degree holders (35.5% vs. 33%). However, some imbalances remain, particularly for total population (17,167,844 vs. 39,138,368) and solar irradiance (4.27 vs. 5.52 GHI).

Table 8: California Mandate: Donor Pool Characteristics

State	Population	Avg. Solar PV kW Per New Building Permit
California	39,262,097	19.4
Connecticut	3,589,709	26.9
Delaware	965,614	1.4
Florida	20,994,906	3.0
Hawaii	1,431,657	12.6
Louisiana	4,658,984	0.6
Maine	1,343,880	2.3
Maryland	6,061,040	5.1
Massachusetts	6,886,521	15.9
New Hampshire	1,355,104	5.6
New Jersey	9,014,836	8.5
New York	19,768,829	10.6
North Carolina	10,282,754	0.9
Oregon	4,141,628	3.2
South Carolina	5,030,558	1.0
Texas	28,384,505	2.5
Virginia	8,491,756	3.0
Washington	7,447,798	2.0
Donor Pool Average	8,226,475	6.2

Notes: This table lists our donor pool states, their populations, and the mean of our solar adoption measure. As detailed in Section 4.2, our state-level solar adoption measure is residential solar capacity additions (kW) per residential new construction building permit issued. The solar adoption estimates are for the pre-policy period between 2017 and 2019.

Table 10 presents the composition of California’s synthetic control including predictor and geography weights for both a monthly and a quarterly specification. In both the quarterly and the monthly specifications, New York receives the highest weight, followed by Hawaii. For the quarterly specification, the most important predictor is owner-occupied housing share (81% weight), followed by adults with bachelor’s degree (14%), and household size (2%). In the monthly specification, owner-occupied housing share (60% weight), followed by the housing unit growth rate (35%). In both specifications, electric heating, population level, and household size take low weights (all less than 3%).

Complementing the comparison of characteristics between actual and synthetic California in Table 9, Figures 11 and 12 further illustrate the pre-period fit of our synthetic control. In Figure 11, the observed solar adoption measure for California (solid green line) closely tracks synthetic California (dashed blue line) in the quarters preceding the statewide solar mandate. Figure 12 then plots the differences between the two series, which are quite well-centered around zero in the pre-period. The relatively small gap during the pre-treatment period supports our synthetic control construction, suggesting it effectively captures California’s solar adoption patterns before the mandate. As discussed above, recall that our solar adoption measure on the y-axis includes *total* residential solar adoption in the numerator, so this measure cannot be interpreted as the *size* of solar systems on residential new construction.

Figures 11 and 12 show that California’s solar adoption measure increases markedly after the 2020 mandate implementation as compared to the synthetic trend with a clear and sustained positive divergence. Table 10 presents the values of the estimated effect, which are quite consistent between the monthly and quarterly specifications, with the two specifications estimating a policy effect of 5.83 (30%) and 6.67 kW (34%) per new construction building permit, respectively. These values are close to a typical system for a single family home.

Table 9: California Mandate: Balance Table for Predictors

Predictor	California	Donor Pool	Synthetic California
Total population	39,138,368	8,119,713	17,167,844
Solar irradiance (GHI)	5.52	4.35	4.27
Housing unit growth rate	0.6	0.9	0.6
Electric heat housing units (share)	26.5	40.3	13.9
Owner occupied housing units (share)	55.6	67.1	57.5
Average household size	2.96	2.61	2.66
Bachelor degree holders (share)	33.3	33.8	35.5

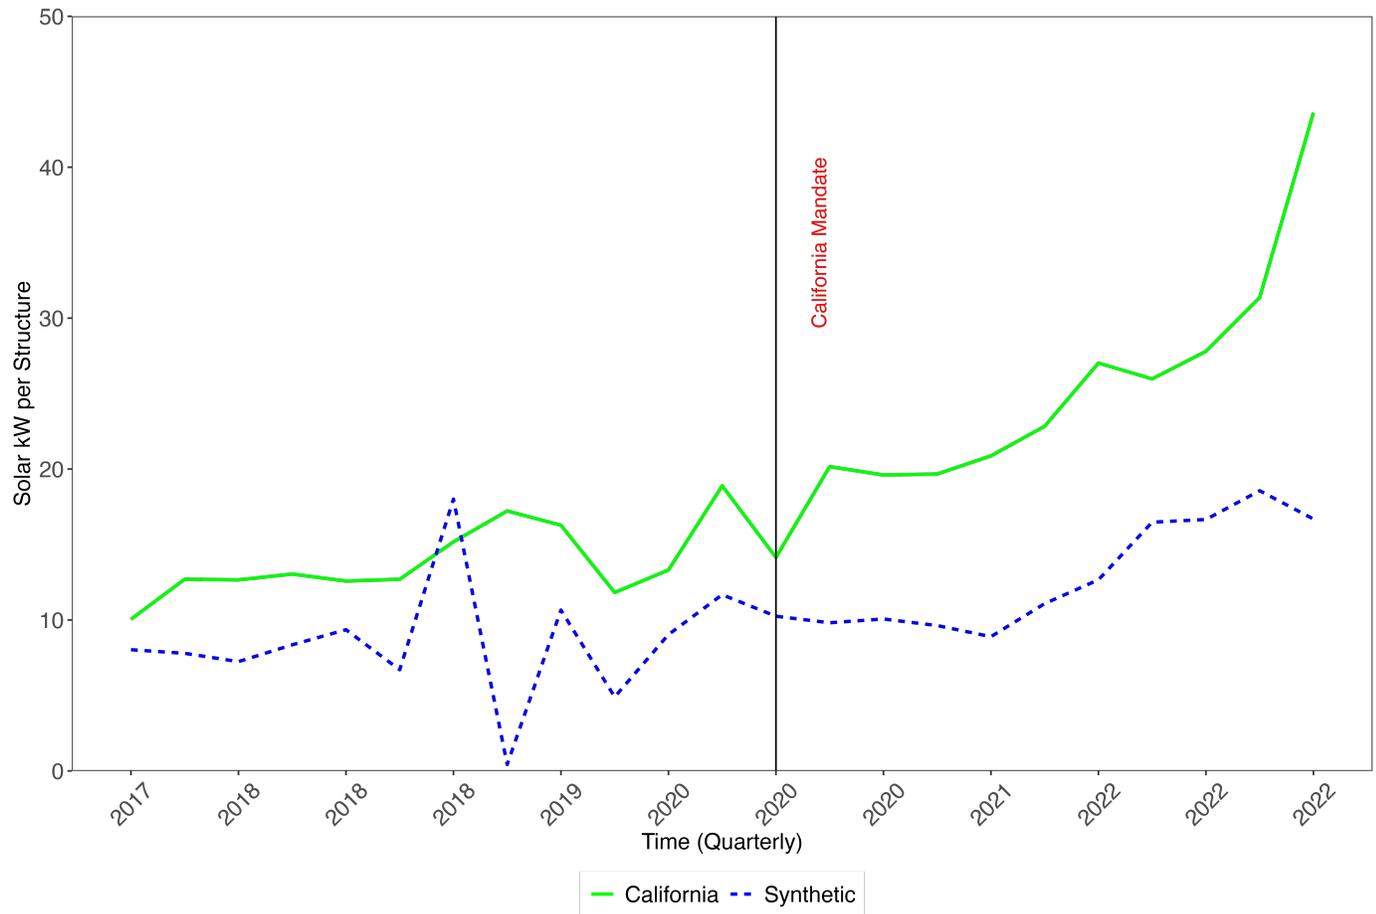
Notes: For each of seven characteristics considered in constructing our synthetic control, this table compares California’s characteristics to the donor pool averages and our synthetic control, as well as the synthetic control corresponding to the predictor and geographic weights from model (2) of Table 10, estimated with quarterly data.

Table 10: California Mandate: Composition of the Synthetic Control and Estimated Effect

	(1)	(2)
Estimate (kW per structure)	5.83 (2.21)	6.67 (2.37)
<u>Predictor weights:</u>		
Owner occupied housing units (share)	0.812	0.606
Adults with bachelor’s degree (share)	0.139	0.014
Household size (average)	0.024	0.000
Housing unit growth rate	0.022	0.349
Electric heat housing units (share)	0.002	0.001
Total population	0.000	0.000
Solar irradiance (GHI)	0.000	0.030
<u>State weights:</u>		
New York	0.969	0.863
Hawaii	0.031	0.137
<u>Temporal scale:</u>		
	Monthly	Quarterly

Notes: (1) Section 4.2 describes the construction of our solar adoption measure, which is residential solar capacity additions (kW) per new construction building permit issued. (2) We estimate models (1) and (2) with monthly and quarterly data, respectively. (3) The table omits states with zero or close-to-zero weights.

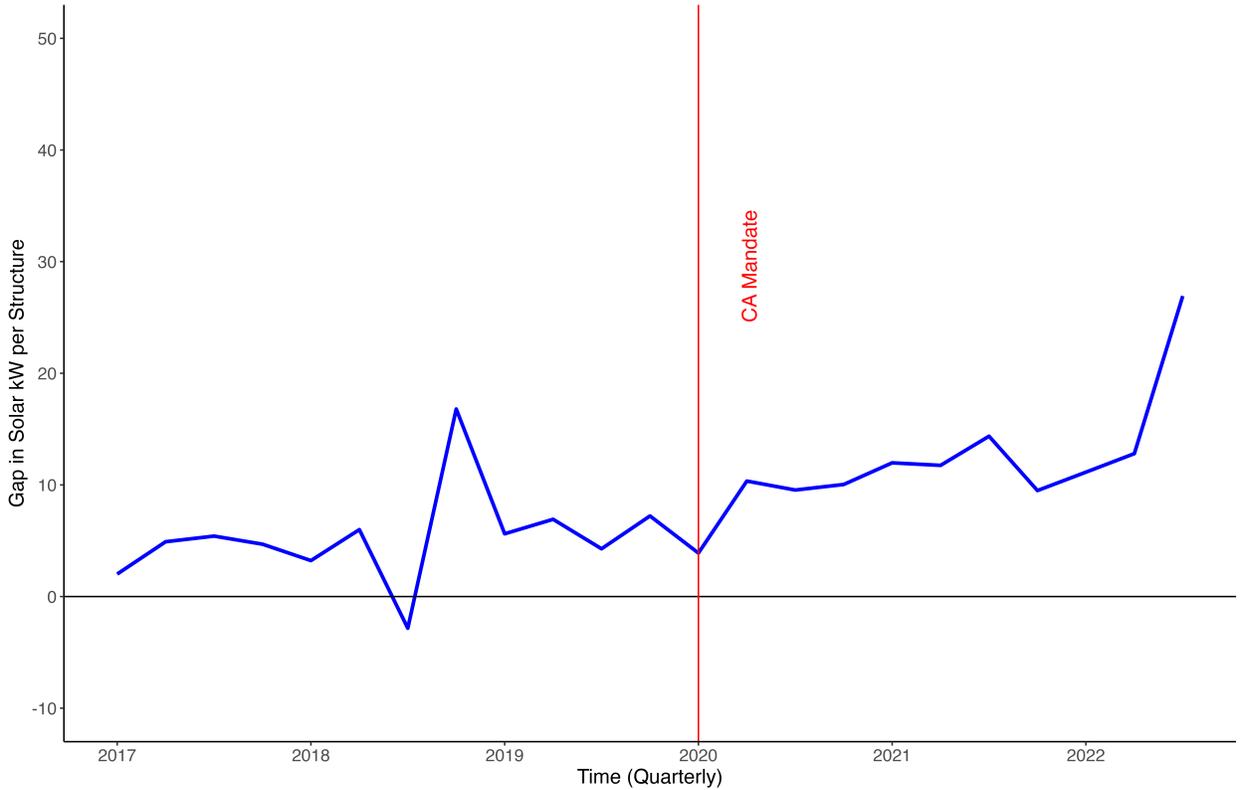
Figure 11: California Mandate: Actual and Synthetic California Solar Shares on New Construction



Notes: (1) Section 4.2 describes the construction of our solar adoption measure, which is total residential solar capacity additions (kW) per new construction building permit issued. (2) Table 10 presents the corresponding synthetic control predictor and geography weights. This figure corresponds to model (2) estimated with quarterly data. (3) The time on the x-axis reflects building permit issuance, not the date of solar adoption because date of building permit issuance is the relevant date for the policy.

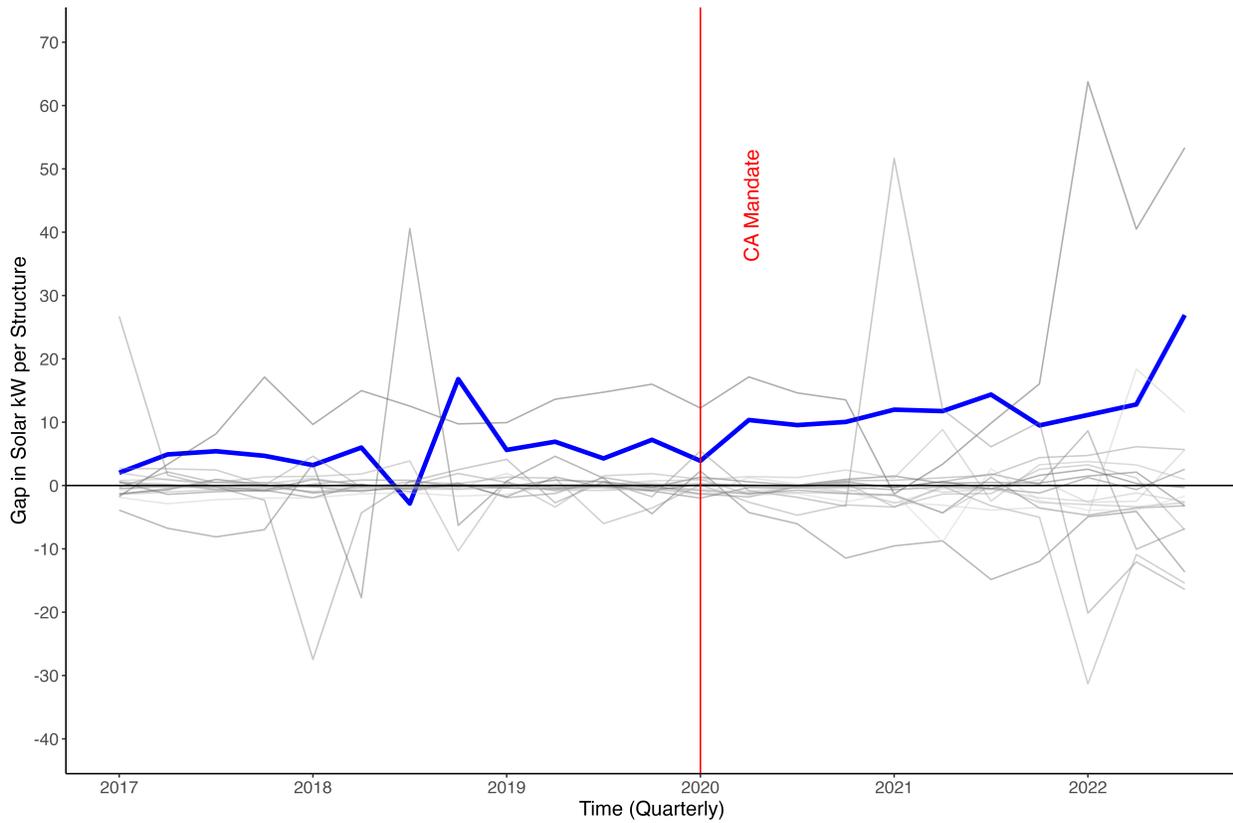
Figure 13 once again implements a placebo test by applying the synthetic control method to each control state in our donor pool. The blue line represents California’s gap, while the gray lines show the placebo gaps for control states. The magnitude and consistency of California’s positive treatment effect stands out relative to the distribution of placebo effects, underlining the conclusion that the state solar mandate effectively increased the adoption of solar photovoltaics on new construction. Figure 14 shows the same placebo test for a subset of the donor pool, which has a pre-treatment mean squared prediction error no larger than California’s.

Figure 12: California Mandate: Gap Between Actual and Synthetic California Solar Shares on New Construction



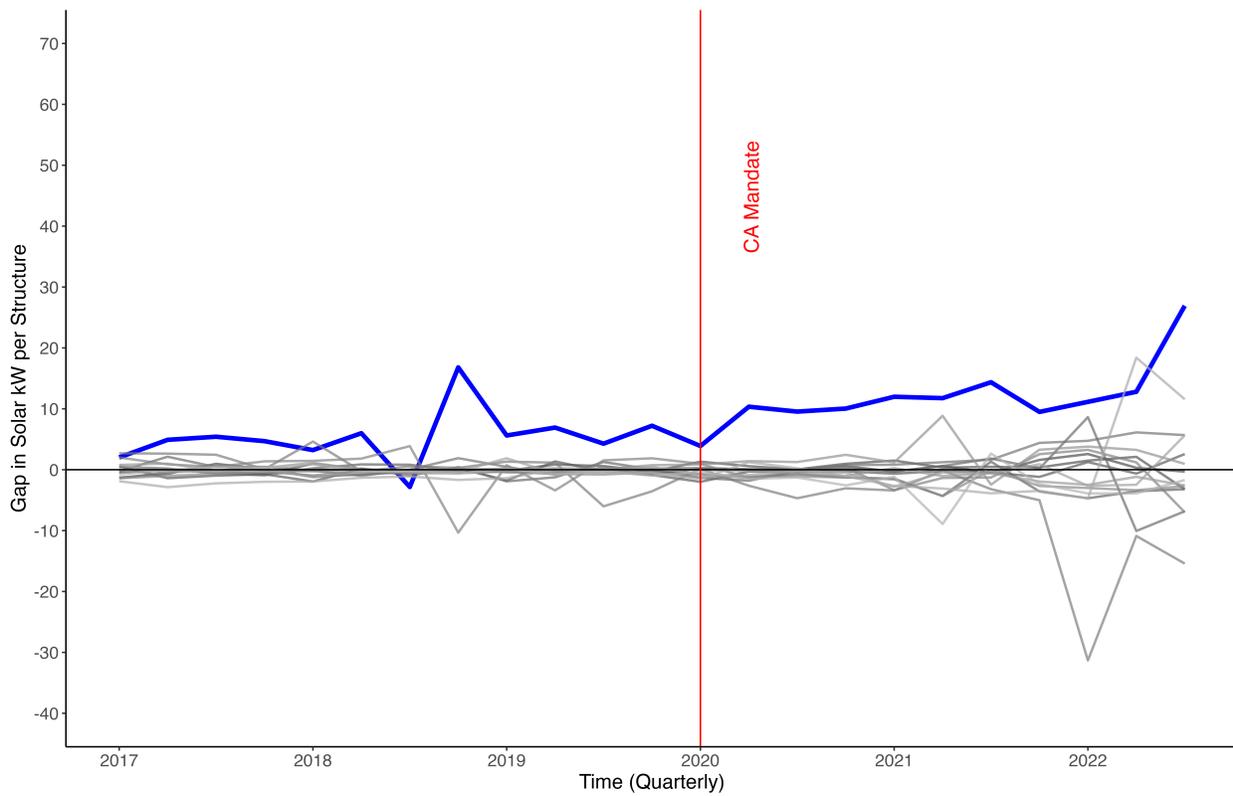
Notes: This figure plots the difference (or “gap”) between the observed time series of our solar adoption measure for California and the synthetic control. These results correspond to the series in Figure 11 and model (2) with quarterly data in Table 10.

Figure 13: California Mandate: New Construction Solar Share Gap and Placebo Gaps Using All Donor Pool Municipalities



Notes: This figure compares the gap produced from our synthetic control (also plotted in Figure 13) to a series of placebos where the synthetic control methodology is applied to each of the municipalities in the donor pool. This figure shows all 17 donor states.

Figure 14: California Mandate: New Construction Solar Share Gap and Placebo Gaps Using Selected Donor Pool Municipalities



Notes: This figure compares the gap produced from our synthetic control (also plotted in Figure 13) to a series of placebos where the synthetic control methodology is applied to each of the municipalities in the donor pool. We discard control states with a pre-treatment mean squared prediction error (MSPE) larger than California's.

As in the Section 4.1 evaluation of the San Francisco mandate, we once again implement a synthetic difference-in-differences (SDID) estimator (Arkhangelsky et al. 2021) to complement our standard synthetic control estimates of the causal effect of California’s solar mandate.

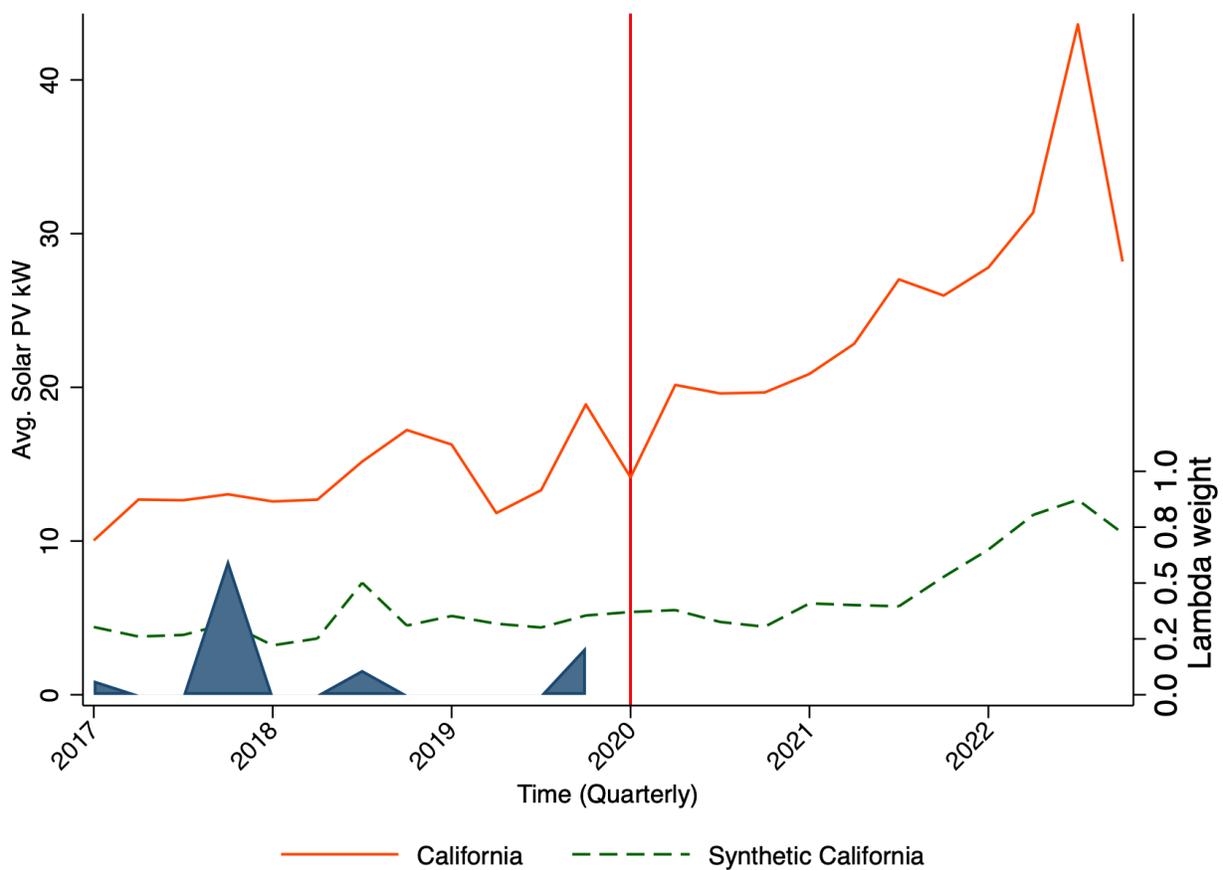
Figure 15 illustrates the SDID results, showing California’s outcome trajectory (solid orange line) compared to the reweighted SDID synthetic control (dashed green line). Table 11 presents the SDID estimate of the policy effect—an additional 8.36 kW (43%) of solar per new residential building permit. This estimate is higher than our standard synthetic control estimates above, but generally confirms the finding of a substantial positive policy effect.

Table 11: California Mandate: Composition of the Synthetic Difference-in-Differences Control and Estimated Effect

	(1)
Estimate (kW per structure)	8.357
<i>Standard Error</i>	<i>4.286</i>
<u>Predictor weights:</u>	
Owner occupied housing units (share)	0.670
Bachelor degree share	0.560
Housing unit growth rate	0.174
Electric heat housing units (share)	0.069
Total population	0.001
Solar irradiance (GHI)	0.000
<u>State weights:</u>	
Connecticut	0.072
New Hampshire	0.068
Texas	0.066
Florida	0.065
Virginia	0.063
North Carolina	0.063
Oregon	0.062
Washington	0.061
Maryland	0.061
Maine	0.061
Delaware	0.059
South Carolina	0.057
New Jersey	0.056
New York	0.055
Louisiana	0.053
Hawaii	0.049
Massachusetts	0.027
Temporal scale	Quarterly

Notes: (1) This table presents the predictor and geographic weights used to construct the synthetic control to evaluate the California mandate. (2) The table also presents the estimated effect of the California mandate. The estimated effect is the change in solar photovoltaic capacity (in kW) per structure after the mandate in the observed data compared to the synthetic control. Standard errors calculated using 1,000 placebo replications. This estimated effect is illustrated graphically in Figure 15.

Figure 15: California Mandate: Actual and Synthetic Difference-in-Differences Solar Shares on New Construction



Notes: (1) This figure shows California’s actual solar adoption share (solid orange line) compared to the synthetic difference-in-differences control (dashed green line), with corresponding geographic and predictor weights in Table 7. (2) The red vertical line marks policy implementation in 2020. (3) Shaded triangular regions represent lambda time weights assigned by the SDID procedure to improve pre-treatment fit. (4) Standard errors calculated using 1,000 placebo replications.

5 Conclusions

This paper examines the economics of solar mandates, which impose the use of solar panels on new constructions, and the case for their use. These mandates gained prominence in California, supported by engineering studies suggesting that solar panels are cost-effective and profitable on all new construction. These mandates have then expanded quite rapidly in Europe and elsewhere in the United States.

In this paper, we tackle three research questions. First, we examine whether solar is ubiquitous in the absence of a mandate, given its presumed cost-effectiveness. We find that it is not. Second, we leverage many data sources and use an engineering model to analyze the private payoffs of solar adoption for residential new construction and whether solar is consistently cost-effective when more fully accounting for building characteristics and the variety of building types in the setting where solar mandates were first introduced. We find much greater heterogeneity in project cost-effectiveness in parameterizing the model across our dataset of observed building projects and their detailed characteristics. Although estimated payoffs are generally positive, they tend to be small compared to overall new construction costs. We further compare payoff estimates with actual solar adoption and find little evidence of any strong relationship between solar adoption and the level of estimated payoffs, suggesting that factors other than solar payoffs dominate the adoption decision. This evidence does support the concept of a “solar gap,” whereby many builders do not consider profitable opportunities to install solar, thus supporting one of the rationalizations for mandates.

Third, we engage in a policy evaluation exercise and provide evidence that the 2017 municipal solar mandate in the frontrunner city of San Francisco as well as the 2020 statewide mandate in California have substantially increased the adoption of solar among new construction. Our data also allow us to examine compliance for the citywide mandate in San Francisco. Adoption remains below 100 percent, due to a system of exemptions.

Finally, we combine three main takeaways from our study to derive implications for the mandates. In particular, these are the evidence that solar mandates increase the adoption of solar in new construction, the evidence that solar adoption provides positive private payoffs, and the evidence that a substantial share of such positive private payoffs are paid by transfers, including the ITC and California’s net metering. As we show, changes to the extent of these transfers sub-

stantially affect private payoffs for adopters. Hence, it is fair to conclude that if solar mandates generate positive private payoffs, transfers are a large share.

Solar mandates may still generate external benefits, but many of these benefits can also be generated using utility-scale solar. Comparing the most recent data on the cost of distributed solar with the cost of utility-scale solar in NREL's 2024 Annual Technology Baseline (ATB), nationally representative residential solar capital costs are \$2,682/kW, while utility-scale solar capital costs are substantially lower at \$1,611/kW.⁸ Distributed solar may also provide benefits for grid management that utility-scale solar does not provide, but the literature has generally estimated these benefits to be small relative to the 66 percent higher capital costs (Cohen and Callaway 2016; Cohen et al. 2016). In sum, while we find that solar mandates are effective in increasing the adoption of solar and that by doing so they tend to generate positive payoffs for solar adopters, we find little support for their use to increase the adoption of solar energy relative to larger utility-scale projects. In conclusion, we also note that the implications of promoting the adoption of solar energy in terms of greenhouse gas reductions in a context wherein electricity generation is covered by a cap-and-trade system, as in California, depend on the exogeneity of the cap.

⁸See <https://atb.nrel.gov/electricity/2024b/about> (last accessed, May 9, 2025).

References

- Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program. *Journal of the American Statistical Association* 105(490), 493–505.
- Abadie, A., A. Diamond, and J. Hainmueller (2015). Comparative politics and the synthetic control method. *American Journal of Political Science* 59(2).
- Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review* 93(1), 113–132.
- Allcott, H. and M. Greenstone (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives* 26(1).
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review* 111(12).
- Aroonruengsawat, A., M. Auffhammer, and a. A. H. Sanstad (2012). The impact of state level building codes on residential electricity consumption. *The Energy Journal* 33.
- Barbose, G., N. Darghouth, B. Hoen, and R. Wiser (2024). Tracking the Sun: Pricing and Design Trends for Distributed Photovoltaic Systems in the United States, 2024 Edition. Technical report, Lawrence Berkeley National Laboratory, Berkeley, CA.
- Barbose, G., N. Dargouth, S. Weaver, and R. Wiser (2013). Tracking the Sun VI: An Historical Summary of the Installed Price of Photovoltaics in the United States from 1998 to 2012. Technical report, Lawrence Berkeley National Laboratory.
- Borenstein, S. (2012). The private and public economics of renewable electricity generation. *Journal of Economic Perspectives* 26(1).
- Brockman, K. and B. Hooper (2016). San Francisco Better Roofs Policy – Phase 1. Staff report, San Francisco Department of Environment.
- Bruegge, C., T. Deryugina, and E. Myers (2019). The distributional effects of building energy codes. *Journal of the Association of Environmental and Resource Economists* 6(S1), S95–S127.
- California Energy Commission (2016). Local Ordinance Applications - 2016 Standards. Docket 16-BSTD-07, California Energy Commission, Sacramento, CA.
- California Energy Commission (2019). Building Energy Efficiency Standards for Residential and Nonresidential Buildings: For the 2019 Building Energy Efficiency Standards Title 24, Part 6, and Associated Administrative Regulations in Part 1. Commission Final Report CEC-400-2018-020-CMF, California Energy Commission, Sacramento, CA.
- California Energy Commission (2021). New Solar Homes Partnership Program Quarterly Progress Report 2021 First Quarter. Technical report, California Energy Commission, Sacramento, CA.
- California Energy Commission (2022). 2022 Building Energy Efficiency Standards for Residential and Nonresidential Buildings: For the 2022 Building Energy Efficiency Standards Title 24, Part 6, and Associated Administrative Regulations in Part 1. Commission Final Report CEC-400-2022-010-CMF, California Energy Commission, Sacramento, CA.

- California Solar Initiative (2015). Customer-Sited Renewable Energy Generation. Technical report, California Solar Initiative (CSI) via Internet Archive.
- California Statewide Codes & Standards Program (2022). 2019 Code Cycle – Locally Adopted Energy Ordinances. Technical report, Pacific Gas and Electric Company and San Diego Gas & Electric Company and Southern California Edison Company, California.
- Carattini, S., B. Figge, A. Gordan, and A. Löschel (2024). Municipal building codes and the adoption of solar photovoltaics. *Journal of Environmental Economics and Management* 124.
- City and County of San Francisco (2016a). Ordinance No. 221-16: Planning, Green Building Codes – Better Roof Requirements, Including Living Roofs. Municipal ordinance, City and County of San Francisco, San Francisco, CA.
- City and County of San Francisco (2016b). Zoning Administrator Bulletin No. 11: Better Roofs Ordinance. Zoning Administrator Bulletin 11, City and County of San Francisco, San Francisco, CA.
- Cohen, M. and D. Callaway (2016). Effects of distributed pv generation on california’s distribution system, part 1: Engineering simulations. *Solar Energy* 128, 126–138.
- Cohen, M., P. Kauzmann, and D. Callaway (2016). Effects of distributed pv generation on california’s distribution system, part 2: Economic analysis. *Solar Energy* 128, 139–152. Special issue: Progress in Solar Energy.
- Congressional Research Service (2021). The Energy Credit or Energy Investment Tax Credit (ITC). Technical report, Congressional Research Service (CRS).
- CPUC (2025a). Customer-Sited Renewable Energy Generation. Technical report, California Public Utilities Commission (CPUC).
- CPUC (2025b). Net Energy Metering and Net Billing. Technical report, California Public Utilities Commission (CPUC).
- Crago, C. L. and I. 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.
- Darghouth, N. R., G. Barbose, and R. Wiser (2011). The impact of rate design and net metering on the bill savings from distributed PV for residential customers in California. *Energy Policy* 39(9), 5243–5253.
- Davis, L. W. (2025). What matters for electrification? Evidence from 70 years of U.S. home heating choices. *The Review of Economics and Statistics* 107(3), 668–684.
- Davis, L. W., A. Fuchs, and P. Gertler (2014). Cash for coolers: Evaluating a large-scale appliance replacement program in Mexico. *American Economic Journal: Economic Policy* 6(4).
- Davis Energy Group Inc. and Misti Bruceri & Associates LLC (2016). Local PV Ordinance Cost Effectiveness Study. Technical report, for Pacific Gas and Electric Company.
- Energy and Environmental Economics Inc. (2017). Rooftop Solar PV System Report. Technical report, for the California Energy Commission.

- Fowle, M., M. Greenstone, and C. Wolfram (2018). Do energy efficiency investments deliver? Evidence from the Weatherization Assistance Program. *Quarterly Journal of Economics* 133(3), 1597–1644.
- Gerarden, T. D., R. G. Newell, and R. N. Stavins (2017). Assessing the energy-efficiency gap. *Journal of Economic Literature* 55(4).
- Gillingham, K. and K. Palmer (2014). Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Review of Environmental Economics and Policy* 8(1), 18–38.
- Glaeser, E. L. and J. Gyourko (2003). The impact of building restrictions on housing affordability. *Economic Policy Review*, 21–39.
- Gould, C. F., M. L. Bejarano, B. De La Cuesta, D. W. Jack, S. B. Schlesinger, A. Valarezo, and M. Burke (2023). Climate and health benefits of a transition from gas to electric cooking. *Proceedings of the National Academy of Sciences* 120(34), e2301061120.
- Grimes, A., N. Preval, C. Young, R. Arnold, T. Denne, P. Howden-Chapman, and L. Telfar-Barnard (2016). Does retrofitted insulation reduce household energy use? Theory and practice. *Energy Journal* 37(4).
- Halberstadt, A. (2014). Report on cost-effectiveness and other analyses for proposed solar ordinance. Technical report, for the City and County of San Francisco.
- Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *The Bell Journal of Economics* 10(1).
- Hsieh, C.-T. and E. Moretti (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics* 11(2), 1–39.
- Hughes, J. E. and M. 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.
- Jacobsen, G. D. and M. J. Kotchen (2013). Are building codes effective at saving energy? Evidence from residential billing data in Florida. *Review of Economics and Statistics* 95(1).
- Jaffe, A. B. and R. N. Stavins (1994). The energy-efficiency gap – What does it mean? *Energy Policy* 22(10).
- Kotchen, M. J. (2017). Longer-run evidence on whether building energy codes reduce residential energy consumption. *Journal of the Association of Environmental and Resource Economists* 4(1).
- Levinson, A. (2014). California energy efficiency: Lessons for the rest of the world, or not? *Journal of Economic Behavior and Organization* 107(PA).
- Levinson, A. (2016). How much energy do building energy codes save? Evidence from California houses. *American Economic Review* 106(10).
- Levinson, A. and S. Niemann (2004). Energy use by apartment tenants when landlords pay for utilities. *Resource and Energy Economics* 26(1).

- Metcalf, G. E. and K. A. Hassett (1999). Measuring the energy savings from home improvement investments: Evidence from monthly billing data. *Review of Economics and Statistics* 81(3).
- Myers, E. (2019). Are home buyers inattentive? Evidence from capitalization of energy costs. *American Economic Journal: Microeconomics* 11(2).
- Novan, K., A. Smith, and T. Zhou (2022). Residential building codes do save energy: Evidence from hourly smart-meter data. *Review of Economics and Statistics* 104(3), 483–500.
- Pollakowski, H. O. and S. M. Wachter (1990). The effects of land-use constraints on housing prices. *Land Economics* 66(3), 315–324.
- Rai, V. and B. Sigrin (2013). Diffusion of environmentally-friendly energy technologies: buy versus lease differences in residential PV markets. *Environmental Research Letters* 8(1), 014022.
- Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics* 125(3), 1253–1296.
- Tibebu, T. B., E. Hittinger, Q. Miao, and E. Williams (2021). What is the optimal subsidy for residential solar? *Energy Policy* 155, 112326.
- Tsvetanov, T. and K. Segerson (2013). Re-evaluating the role of energy efficiency standards: A behavioral economics approach. *Journal of Environmental Economics and Management* 66(2).

A Supplemental solar mandate policy details

A.1 Supplemental details on the San Francisco’s Better Roofs Ordinance

The 2017 Better Roofs Ordinance (“Better Roofs,” Ordinance No. 221-16) was an amendment to San Francisco’s Planning Code and the Green Building Code. Better Roofs was intended to contribute to the city’s climate objectives (Ordinance No. 81-08), which include a greenhouse gas-free electric system by 2030 and emission reductions of 40% below 1990 levels by 2025 and 80% below 1990 levels by 2050 (Brockman and Hooper 2016).⁹

The ordinance made San Francisco the first major city in the United States to require solar installations on new buildings, though more than two dozen smaller California municipalities adopted solar photovoltaic requirements between 2017 and 2020 (California Energy Commission 2016; California Statewide Codes & Standards Program 2022). San Francisco also joined other major U.S. cities including Chicago, Portland, and Washington, D.C., that require living roofs on certain new buildings, though San Francisco’s approach offers more flexibility in compliance pathways.¹⁰

The Better Roofs policy built on existing city regulations. Since 2010, the San Francisco Green Building Code had included “a modest renewable energy requirement for new commercial buildings larger than 25,000 square feet,” but this could be met through purchasing Renewable Energy Certificates (RECs) or improving energy efficiency rather than on-site generation (Brockman and Hooper 2016). The Better Roofs mandate strengthened these requirements by mandating on-site renewable generation.

San Francisco’s Better Roofs Ordinance also extended California’s statewide Title 24 requirement to build “solar ready,” to which we return in Section A.2. San Francisco, and other local jurisdictions, were authorized under Title 24 to enact more stringent standards when reasonably necessary because of local conditions related to climate, geology, or topography. In 2013, Cali-

⁹The San Francisco Environment Department quantified the potential emissions impact as follows: “The aggregate impact of installing photovoltaics to minimally comply with the proposed Better Roof ordinance on all 200 major new construction projects in San Francisco Planning Department’s project pipeline as of third quarter 2014 would be to avoid over 26,000 metric tons of carbon dioxide emissions per year” (Brockman and Hooper 2016).

¹⁰According to the San Francisco Department of Environment, “considerable research by the Planning Department, Environment, SFPUC, and the 2013 Greener and Better Roofs Roadmap report from SPUR’s Green Roof Task Force recognizes that – in addition to being a prime location for renewable energy resources – roofs can host ‘green’ or ‘living roofs’ with many additional benefits such as reducing stormwater entering the sewer, enhancing biodiversity and habitat, sequestering carbon, capturing pollution, and connecting citizens with nature” (Brockman and Hooper 2016).

Appendix

California's Title 24 building efficiency standards established that 15% of the roof area in most new buildings must be reserved for future installation of solar energy systems. San Francisco's policy strengthened this state policy by mandating that the solar-ready area be activated with renewable energy systems during initial construction (City and County of San Francisco 2016a). Established through Planning Code Section 149, the minimum Better Roof area requirement mirrors the Solar Ready Area specifications in California Energy Standards. Specifically, the ordinance required that solar be applied to (1) 15% of total roof area for multifamily and non-residential buildings; or (2) 250 square feet of the roof area of single-family residential buildings (Brockman and Hooper 2016). The ordinance applied to site or building permit applications for new residential buildings, commercial or municipal buildings submitted on or after January 1, 2017 (City and County of San Francisco 2016b). Notable exceptions include (1) buildings over ten stories,¹¹ (2) non-residential buildings with gross floor area less than 2,000 square feet, and (3) data centers and laboratory buildings.

In addition to solar photovoltaic (PV) systems, San Francisco's Better Roofs ordinance also offered two alternative compliance pathways, providing developers flexibility in meeting the city's sustainability goals. These two alternative pathways were either (1) a solar thermal system; or (2) a living roof, which incorporates vegetation to improve insulation, reduce urban heat, manage stormwater runoff, and enhance biodiversity (City and County of San Francisco 2016b).¹² Project developers could also combine these pathways so far as they demonstrated productive use of rooftop space to deliver environmental and in principle also economic benefits (City and County of San Francisco 2016a).

¹¹Specifically, the ordinance applied to buildings of ten or fewer occupied floors (City and County of San Francisco 2016b).

¹²For each compliance pathway, San Francisco provided rationale for the expected economic and environmental benefits as well as specific compliance criteria. With respect to solar photovoltaic (PV) systems, the city cited the value of on-site renewable electricity in reducing reliance on fossil fuels and supporting the city's clean energy transition (Brockman and Hooper 2016). For water-heating solar thermal systems, the city discussed energy consumption related to water heating (Brockman and Hooper 2016). Living roofs were required to meet criteria for stormwater management. Living roofs were also required to achieve a compliance ratio of 2:1, meaning two square feet of living roof for each required square foot of solar (City and County of San Francisco 2016b). Solar PV systems were required to achieve a minimum performance standard of 10 watts-DC rated nameplate capacity per square foot of dedicated area (Brockman and Hooper 2016). Solar thermal systems were required to achieve a minimum performance standard 100 kBtu annual generation per square foot.

A.2 Supplemental details on the California’s solar mandate

In California, solar photovoltaics are regulated by the Title 24 Building Energy Efficiency Standards (Energy Code), which were first enacted in 1978.¹³ The standards operate under a conceptual framework divided into three basic sets: mandatory requirements applicable to all buildings, performance standards that vary by climate zone and building type, and prescriptive packages that provide alternative compliance pathways (California Energy Commission 2019). California’s building energy efficiency requirements are updated every three years. With the 2019 Energy Code, California implemented a solar mandate, which required all newly constructed low-rise residential buildings (defined as residential buildings up to three stories) to install solar PV systems (California Energy Commission 2019). The 2019 Energy Code became effective on January 1, 2020. The PV systems installed under the mandate need to meet performance and electricity output requirements. The requirements are based on the floor area of the building, the climate zone, and the goal to provide the full annual energy usage of the home. For residential buildings, the annual electrical output of the PV system shall be no less than the smaller of a PV system size determined using an equation specified in the code, Equation 150.1-C, or the maximum PV system size that can be installed on the building’s “Solar Access Roof Area” (California Energy Commission 2019).¹⁴ The 2019 standards incorporated exceptions to the solar PV installation requirement. In particular, the following exceptions apply: buildings with SARA less than 80 contiguous square feet; buildings where minimum required system size is less than 1.8 kW; buildings with significant shading (less than 70% annual solar access); buildings installing solar water heating that provides at least half the hot water heating needs; buildings with 3 or more habitable stories and less or equal than 2000 square feet meeting reduced solar zone requirements; buildings in Wildland-Urban Interface Fire Areas meeting specified conditions; and buildings meeting certain demand response and energy efficiency measures. PV system sizes may be reduced by 25% if paired with battery storage (of minimum 7.5 kWh usable capacity). Buildings may substitute approved community shared solar

¹³California’s Title 24 Building Energy Efficiency Standards trace their origins to the Warren-Alquist State Energy Resources Conservation and Development Act (Public Resources Code §§ 25000-25968), enacted by the California Legislature in 1974. By 1988, Title 24 applied to all occupancies in the state of California.

¹⁴Equation 150.1-C is given by $kW_{PV} = (CFA \times A)/1000 + (NDwell \times B)$, where: kW_{PV} is the size of the PV system, CFA stands for conditioned floor area, NDwell is the number of dwelling units, and A and B are adjustment factors based on the climate zone, which are provided in a table. Solar Access Roof Area (SARA) is the area of the roof that can accommodate a solar PV system, taking into account factors like the roof’s size, shape, and shading from nearby trees or structures.

Appendix

systems for on-site installation.

The 2019 Energy Code also established “solar ready” requirements for newly constructed buildings that do not have solar PV systems installed. “Solar ready” requirements prepare buildings for future solar installation by mandating designated solar zones on roofs, structural design considerations for future solar loads, electrical pathways and reserved panel space for future solar connections, and appropriate construction documentation. These requirements apply to single family residences in subdivisions with ten or more homes, low-rise multifamily buildings, high-rise multifamily buildings with ten habitable stories or fewer, hotel or motel occupancies with ten habitable stories or fewer, and nonresidential buildings with three habitable stories or fewer (except healthcare facilities).

The 2022 Energy Code, effective January 1, 2023, further extended California’s solar mandate to all residential buildings and most nonresidential buildings (California Energy Commission 2022). The 2022 standards also include a solar PV storage mandate. That is, all new buildings that are required to have a photovoltaic system also need a battery storage system.

B Supplemental details about data and engineering model

B.1 Supplemental details about data

Table B.1 lists our data and links to online sources followed by more detailed descriptions of each dataset to supplement the descriptions of our data in Sections 3.2, and 3.3. Section B.2 provides additional description of NREL’s System Advisor Model and our parameterization methodology.

Table B.1: Data Sources

Organization	Dataset	Data source information
California municipalities and counties	Building, electrical, and solar permits data	Individual municipal and county building departments
California Public Utility Commission (CPUC)	Interconnected Project Sites Data Set	https://www.californiadgstats.ca.gov/downloads/
Google	Project Sunroof	https://sunroof.withgoogle.com/
National Renewable Energy Laboratory (NREL)	National Solar Radiation Database (NSRDB) Typical Meteorological Year (TMY) Data	https://nsrdb.nrel.gov/datasets/tmy
NREL	Utility Rate Database (URDB)	https://openei.org/wiki/Utility_Rate_Database
PG&E	Energy Data Request Program (EDRP)	https://pge-energydatarequest.com/
PG&E	Net Surplus Compensation Rates for Energy	https://www.pge.com/assets/pge/docs/clean-energy/solar/AB920-RateTable.pdf
PG&E	Tariffs: Electric Rate Schedules	https://www.pge.com/tariffs/en.html

Notes: See Appendix Section B.1 text for detailed descriptions of each dataset. Links last accessed, May 9, 2025.

Building, electrical, and solar permits data: We construct a novel address-level dataset of new construction building projects in California municipalities, which we collect from the permitting offices of individual municipalities and counties. In addition to new construction permits, we also collect solar and electrical permits. In some cases, jurisdictions had permitting processes that tracked solar permits separately from other electrical permit types. In matching the solar electri-

Appendix

cal permits to building permits, we are able to identify the new construction projects that adopt rooftop solar. Specifically, we match the new construction building permit addresses to solar permits in the corresponding jurisdiction’s electrical permits. We are then able to validate these solar adoption statistics from individual jurisdictions’ electrical and solar permitting databases against California’s solar interconnection data. The building permits data from individual municipalities and counties vary in the level of data availability and details of each project’s characteristics. However, the general structure of the data tracks each project’s status from initial filing date through various approvals and inspections to final completion. We restrict the data to only completed projects, for which we are able to observe the final status of a building’s solar adoption decision. Our list of 20 San Francisco Bay Area municipalities includes San Francisco, Alameda, Berkeley, Burlingame, Cupertino, Daly City, Danville, Emeryville, Fairfield, Fremont, Oakland, Redwood City, Richmond, San Bruno, San Carlos, San Jose, Santa Clara, Santa Cruz, South San Francisco, and Sunnyvale.

California Interconnected Project Sites Data Set: These data include all grid-interconnected rooftop solar systems within the PG&E service territory since 1982. Primarily, we use these data to obtain a time series of solar installation costs, including median and 95th percentile cost-per-watt-DC by California city and year. We also use these data to validate the quality of our city-level electrical and solar permits data by comparing each city’s overall number of solar installations by year between the two datasets.

NREL National Solar Radiation Database: Sunshine and other climate variables are essential SAM inputs to determine the availability of solar energy and panel performance. These data are available in Typical Meteorological Year (TMY) files from NREL’s National Solar Radiation Database (NSRDB), which include an entire year of hourly temperature, global horizontal irradiance (GHI), and other relevant climate variables for a representative year based on weather observations from 1998 to 2022. NSRDB TMY data is gridded to 4km by 4km cells, which we match to our geocoded building permits.

Google Project Sunroof: To determine rooftop shading conditions, we extract a measure of annual “usable sunlight” hours from Google Project Sunroof for each building project in our data. Google Project Sunroof uses satellite data to identify rooftop area and shading conditions. We estimate a shading measure for each roof using its “usable sunlight” hours relative to the other roofs in our sample.

NREL Utility Rate Database (URDB): The Utility Rate Database (URDB) provides PG&E electricity tariff data for the years of our sample period.

PG&E public tariffs and solar compensation rates: Under the various iterations of California’s net metering policies, solar customers are compensated for any excess generation back to the grid (on an annual basis) at a wholesale rate of around \$0.03/kWh. PG&E tariffs and solar compensation rates are both publicly available on PG&E’s website. Historic tariffs and compensation rates are available on archived versions of these webpages from the Internet Archive.¹⁵

PG&E restricted-access data: In 2018, the electricity utility PG&E provided a restricted-access sample of billing and electricity consumption interval metering data for addresses from 73 Bay Area municipalities that were associated with premise IDs created in PG&E’s billing system after January 1, 2009. New premises in PG&E’s billing data are typically associated with new construction. The final sample included almost 63,000 addresses in 71 Bay Area municipalities with metering and billing data through December 31, 2017. We use these data for our hourly consumption profiles in the parameterization of the engineering model and estimation of project-level solar payoffs. Since PG&E can create new premise IDs that are not associated with new construction, we create city-mean, city-median, and 95th hourly electricity consumption profiles to increase our confidence that

¹⁵See <https://web.archive.org/> (last accessed, May 9, 2025).

Appendix

we are identifying electricity consumption associated with new and recently constructed buildings. Also, Appendix Section C compares our payoff estimates from our primary parameterization to an alternative approach where we use building-level hourly interval electricity consumption data from our sample of new PG&E premises. However, the preferred specification should reflect a builder's expectations about the building's future electricity consumption patterns, and a snapshot of observed building-level electricity consumption patterns are not necessarily indicative of *future* or even *expected* electricity consumption patterns. Therefore, the use of observed building-level consumption data is not necessarily strictly preferred in parameterization of the engineering model.

California net metering policies: California's net metering policy determines how solar owners are compensated for the electricity they produce and then use for their own consumption as well as any excess electricity they export to the grid. Since California introduced its first net metering policy, NEM 1.0, in 1996, there have been two major changes to the policy accompanied by substantial public debate and uncertainty. Especially notable in our study period, in 2016, California updated its net metering policy from NEM 1.0 to NEM 2.0, which required time-of-use electricity tariffs for solar customers. This policy became effective for customers of PG&E, the utility serving the Bay Area, on December 15, 2016. Under NEM 1.0, solar customers were effectively compensated for their electricity production at the same retail rate during all hours of the day. Under NEM 2.0, solar customers were compensated at a higher retail rate during "peak" hours and a much lower retail rate during the rest of the day. Under both policies, solar customers were compensated for any excess generation back to the grid (on an annual basis) at a wholesale rate of around \$0.03/kWh as discussed above (PG&E compensation rates). A NEM 3.0 policy went into effect in April 2023 after the end of our study period.¹⁶

Federal and California incentive programs: In addition to the net metering policy, there were several other incentive programs available to solar projects during our study period. The 30 percent Federal Investment Tax Credit (ITC) has been available since 2006, so we include this incentive in

¹⁶See CPUC (2025a) and CPUC (2025b) for more history and details of California's NEM policies.

Appendix

all of our parameterizations. The ITC was reduced to 26 percent in 2020 and 2021 before returning to the 30 percent level with the Inflation Reduction Act in 2022 (Congressional Research Service 2021 provides a history of the ITC). The well-known California Solar Initiative (CSI) expired for most Bay Area residential customers in 2013, so the incentives were not available for most projects in our sample. Initial CSI incentives were up to \$2.50 per watt-AC but declined over time. Archived tracking of the CSI phase-out schedule for PG&E's service territory is available via the Internet Archive.¹⁷ Some projects in our sample likely qualified for the New Solar Homes Partnership Program (NSHP), which accepted applications through April 1, 2018.¹⁸ NSHP incentives varied from \$0.25 to \$3.50 per watt-DC depending on application timing and project characteristics. Appendix Section C presents alternative parameterizations where we account for the effect of the NSHP program on payoffs to complement our primary parameterization results in Section 3.4.

¹⁷See California Solar Initiative (2015).

¹⁸See California Energy Commission (2021).

B.2 Supplemental details on engineering model parameterization

For the engineering estimates of private payoffs from a solar installation, we apply the System Advisor Model (SAM) from the U.S. National Renewable Energy Laboratory (NREL).¹⁹ SAM provides a rich representation of a rooftop solar project including geography-specific variation in climate and solar irradiance; installation materials, labor, and financing costs; taxes and subsidies; future electricity consumption and tariff structure; technological specifications of the solar module and inverter; and shading and other system losses. In fact, SAM accommodates dozens of possible user inputs and system configurations. In Sections 3.2 and 3.3 as well as Appendix Section B.1 above, we detail the specific data sources we use to parameterize the model. To estimate the private payoff of the solar installation relative to a counterfactual without the solar panels, the SAM outputs include total costs over the system life, electricity production, electricity bill savings, compensation for excess electricity generation, taxes, and subsidies. Principally, SAM estimates the net present value (NPV) of the installation. For other parameter values, we use SAM defaults, which NREL updates as frequently as once or twice per year. Updates to SAM defaults are based on market data, user feedback, and NREL staff expert judgement. These default parameters include panel lifetime; operation and maintenance costs; financing costs; and the private discount rate of the decision-maker (i.e., the builder or homeowner). SAM also includes dozens of parameters related to solar panel and inverter technical characteristics and performance.

Abstracting away from the exact specificities of the SAM model, Equation B.1 summarizes SAM as a generalized engineering model M that estimates the value of solar V_τ installed at time τ

$$V_\tau(s, X_{i,\tau}) = \mathbb{E}_\tau \left[\sum_{t=\tau}^{\tau+T} \beta^{(t-\tau)} M(s, X_i, \xi_\tau, \zeta_\tau) \right] \quad (\text{B.1})$$

where β is a discount factor and T is the useful life of the system. s is the size of the solar system (e.g., the capacity in kW), and $X_{i,\tau}$ is a matrix of characteristics of building project i initiated at time τ . $X_{i,\tau}$ might include the climate and solar irradiance at building's location, the building's expected electricity consumption, and roof characteristics including shading. $X_{i,\tau}$ may also include

¹⁹See <https://sam.nrel.gov/> (last accessed, May 9, 2025). Specifically, we use the Python implementation of the SAM model, NREL-PySAM version 5.1.0, which allows us to batch a large number of unique project characteristics and parameter assumptions. A description of the Python implementation is provided at <https://sam.nrel.gov/software-development-kit-sdk/pysam.html> (last accessed, May 9, 2025). NREL's well-known PVWatts model is nested within SAM, though SAM allows more flexible user inputs.

Appendix

the specific electricity tariff and government incentives for a building project initiated at time τ . ξ_τ and ζ_τ are vectors of the default assumptions and parameters of model M , which may also vary with project initiation date τ . ξ_τ are the model's default financial parameters, such as installation tax treatment and loan rate. ζ_τ are the model's default engineering assumptions, such as the panel and inverter efficiency of converting solar energy into electricity. In more detailed models (or model parameterizations), more parameters may be determined from $X_{i,\tau}$, so fewer parameters appear in ξ_τ and ζ_τ . Equation (B.1) applies the expectation operator \mathbb{E}_τ at time τ because the homeowner or builder forms their expectations about the system's future performance and payoff in the same period as their initial investment decision.

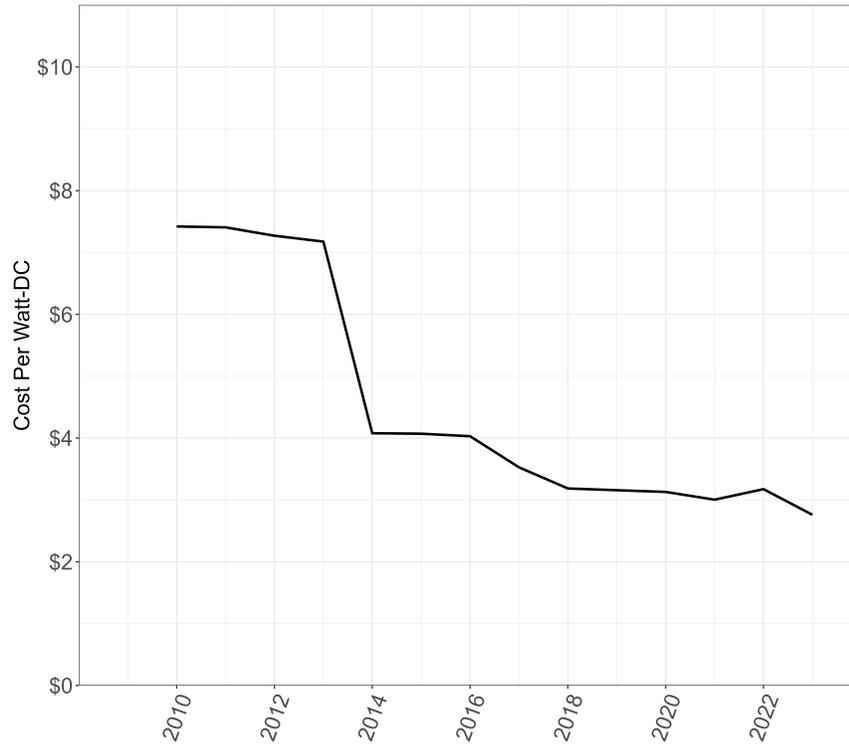
Now suppose that builder or homeowner uses model M in time τ to choose the size of their solar system s to maximize their expected payoff according to

$$\max_{s \geq 0} V(s, X_{i,\tau}). \quad (\text{B.2})$$

To solve Equation (B.2) with SAM, we initially test a variety of optimization routines, including a Bayesian optimization. However, we infer that there is a single global maximum for a given set of project characteristics $X_{i,\tau}$. Therefore, for computational efficiency, we ultimately maximize over a constrained integer grid search. As mentioned, we describe the specific details of our data and SAM parameterization in Sections 3.2 and 3.3. In particular, we use SAM to estimate expected private project payoffs (net present values) NPV_{ict} for each building project i initiated at time t in city c . We use the notation NPV_{ict} to differentiate our SAM estimates from the more general payoff notation $V(s^*, X_{i,\tau})$ in Equations (B.1) and (B.2). Having estimated the distribution of NPV_{ict} across all building projects i , we are able to illustrate the heterogeneity in project payoffs. We also compare these distributions to payoffs estimated by previous cost-effectiveness reports, which generally only considered a few stylized sets of building characteristics. The comparison to prior engineering studies as well as our sensitivities with alternative SAM parameterizations allow us to assess the importance of using such granular project-level data and accounting for differences in project characteristics $X_{i,\tau}$. We also compare our results to the finding of earlier cost-effectiveness studies that California solar projects have strictly positive payoffs across project types.

Appendix

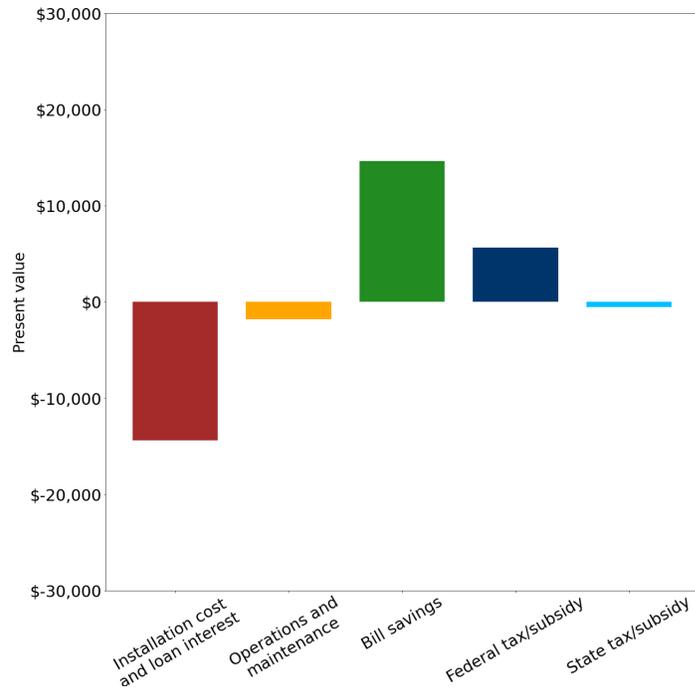
Figure B.1: Residential Solar Installation Costs (2023\$)



Notes: Residential solar panel costs in this figure are from historical versions of SAM. Since 2010, NREL has typically released new SAM versions once or twice per year. Installation costs include panels, inverter, other materials, labor, interconnection costs, and sales tax. The parameterizations in Section 3 use city-year costs from California Solar Interconnection data.

Appendix

Figure B.2: Engineering Model Net Present Value Components for an Illustrative Example



Notes: Components of SAM's net present value (NPV) calculation are for an illustrative building project from Section 3.

C Supplemental evidence on private payoffs and a “solar gap”

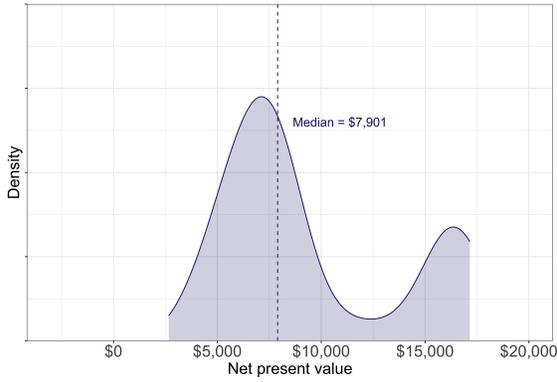
C.1 Estimated distributions of solar payoffs, alternative parameterizations

Figure C.3 presents estimated distributions of solar payoffs under alternative engineering model parameterizations to supplement Figure 2 in Section 3.4.

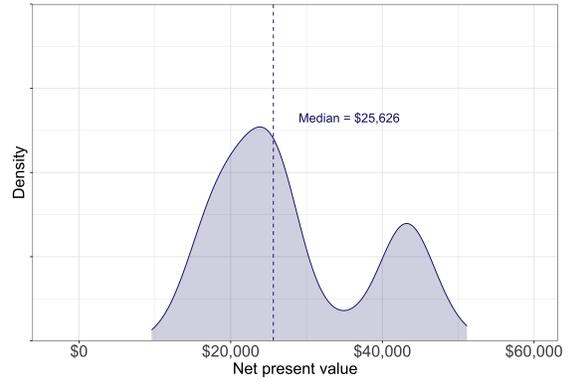
Appendix

Figure C.3: Estimated Distributions of Solar Payoffs, Alternative Parameterizations

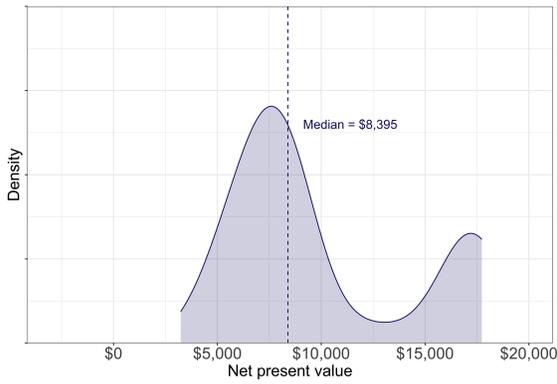
(a) Primary parameterization



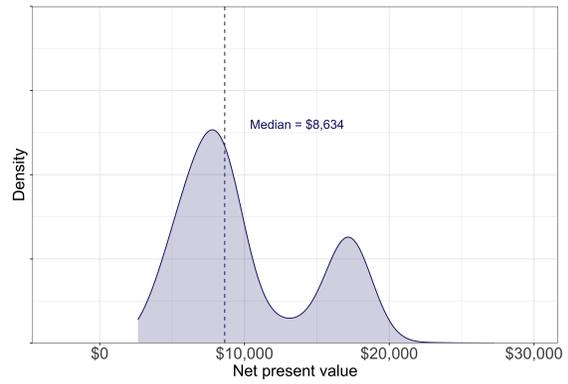
(b) Higher electricity consumption



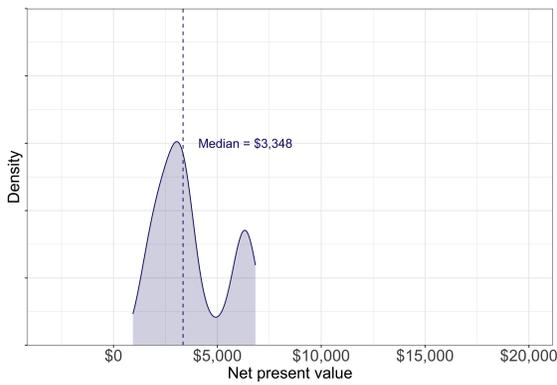
(c) 10 percent lower costs



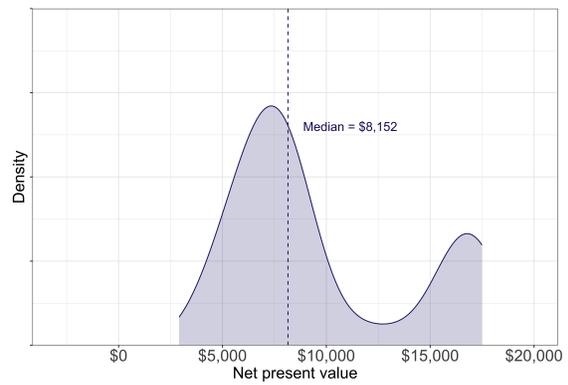
(d) Higher compensation for excess generation



(e) Alternative net metering policy



(f) \$0.25-per-watt-DC incentive



Notes: Continued on next page.

Appendix

Figure C.3 Notes: (1) Figures present medians and smoothed kernel density estimates of the distribution of estimated solar payoffs (net present values) from our parameterization of SAM across our sample of San Francisco Bay Area new construction building permits. These results supplement Figure 2 in Section 3.4. All distributions are for building permits for completed new residential buildings with building permits first filed between 2014 and 2016 (inclusive), prior to the implementation of solar mandates. (2) Panel C.3a presents the primary parameterization from Figure 2. (3) Panel C.3b presents substantially higher payoffs with higher electricity consumption expectations. In particular, these payoffs are estimated with 95th percentile load profiles at the city-level in our PG&E sample. (4) Panel C.3c presents payoffs with 10 percent lower installation costs, consistent with the findings that new construction installations are less expensive than retrofits as discussed in Section 3.3. (5) Panel C.3d presents payoffs with higher compensation rate for excess generation, consistent with possible expectations of a more valuable net metering policy. (6) Panel C.3e assumes that net metering is replaced by a less valuable net billing policy. (7) Panel C.3f represents an additional \$0.25-per-watt-DC incentive, the lowest rate available in California's New Solar Home Partnership Program (NSHP).

C.2 Solar payoffs and adoption decisions

Table 3 in Section 3 presents estimates of the relationship between solar adoption decisions and our engineering estimates of solar payoffs. Specifically, in Table 3, we estimate the linear probability model of solar adoption on net present values npv_{ict}

$$solar_{ict} = \alpha_0 + \alpha_1 npv_{ict} + \mathbf{x}'_{ict} \boldsymbol{\beta} + \epsilon_{ict} \quad (\text{C.3})$$

where $solar_{ict}$ is the observed solar adoption decision for project i in city c and year t . \mathbf{x}_{ict} is a vector of covariate, and ϵ_{ict} is an error term. To avoid the effect of solar mandates, we only consider years 2014 to 2016, before San Francisco implemented its mandates (and before California implemented the statewide policy).

C.3 Engineering model estimates with alternative electricity consumption (load profiles)

Figure C.4 explores the importance of a building’s electricity consumption pattern in determining the payoff of a solar investment. Payoffs are largely determined by the interaction of a building’s electricity consumption pattern with the tariff structure and net metering policy. Under California’s NEM 1.0 and NEM 2.0 policies, customers are compensated for any excess generation during an annual true-up at a much lower rate that approximates the wholesale market (historically around \$0.03 per kWh). Therefore, while customers are disincentivized from oversizing their systems, the largest electricity consumers can realize the greatest payoffs. Under NEM 2.0, where solar customers face tariffs that vary by time of day, the timing of electricity consumption is also an important determinant of the solar payoff.

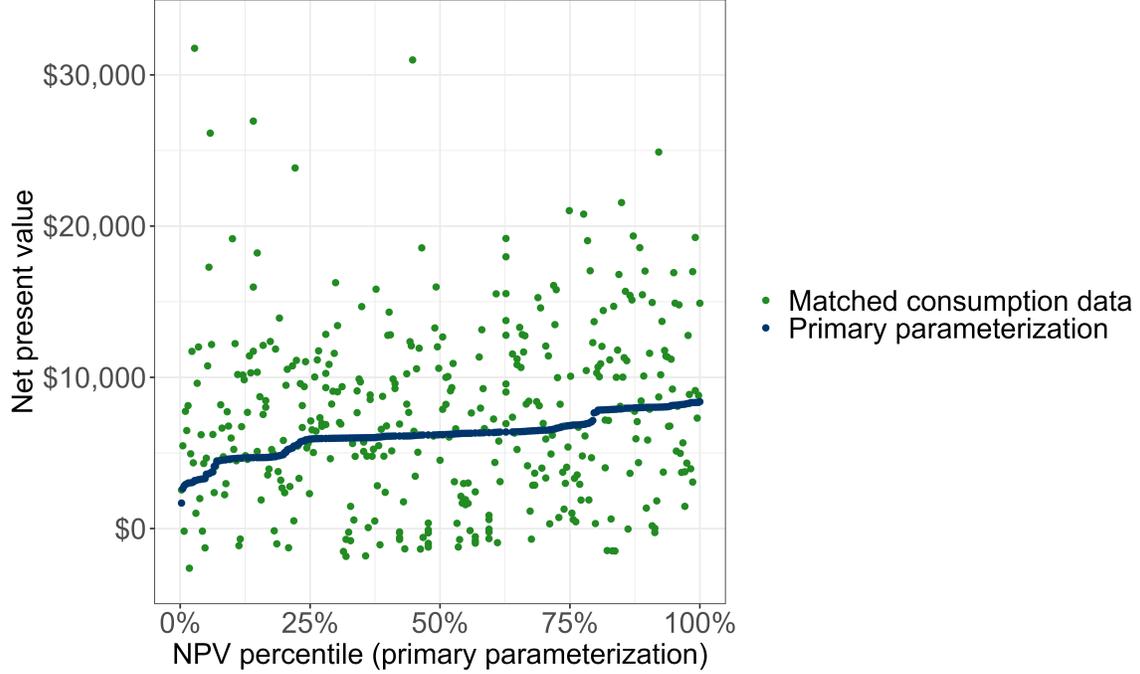
Figure C.4 compares our payoff estimates from our primary parameterization to an alternative approach where we use building-level hourly interval electricity consumption data from our sample of new PG&E premises. In contrast, recall that in our primary parameterization, we use city-mean hourly interval electricity consumption from the PG&E sample (described in Section 3.3). For our full sample of Bay Area new construction projects from municipal building permits, we only have corresponding electricity data available for a subset. Therefore, Figure C.4 only includes the single-family, non-solar new construction buildings for which we have electricity consumption data (398 buildings). We exclude solar buildings because we only observe *net* consumption in the PG&E data, and we do not directly observe solar production. Figure C.4 plots the payoffs estimated with our primary parameterization (the blue series) arranged from least to greatest estimated payoff. The green points are the corresponding payoff estimate for each respective project in the blue series using instead the observed electricity consumption interval data.

Our first observation from Figure C.4 is that the results generally confirm the findings of Sections 3.4 and 3.5. As found in the Section 3.4 primary parameterization results and numerous alternatives, Figure C.4 estimated payoffs continue to be positive for almost all projects. Confirming our findings of Section 3.5 and our discussion of the “solar gap,” Figure C.4 shows builders foregoing solar across the entire support of estimated payoffs (since Figure C.4 shows payoffs estimates for non-solar buildings only). However, Figure C.4 does suggest that there may be even more heterogeneity

Appendix

in payoffs across buildings than indicated in Section 3. That said, estimated payoffs using the building-level consumption data are on-average similar to the more aggregate consumption data of the primary parameterization (i.e., the median payoffs are \$6,831 and \$6,211 using the building-level electricity consumption data and primary parameterizations, respectively). Also, note that the observed building-level electricity consumption patterns are not necessarily indicative of *future* or even *expected* electricity consumption patterns. Therefore, the use of observed building-level consumption data is not necessarily strictly preferred in parameterization of the engineering model. For example, 12% of projects in Figure C.4 have negative estimated payoffs, but even this relatively small share is likely overstated because many of these observations with the lowest payoffs are missing electricity consumption data in some intervals or have especially low usage from not-yet-occupied new buildings.

Figure C.4: Estimated Distributions of Solar Payoffs, Project-Matched Electricity Consumption



Notes: For the subset of new construction projects where we observe electricity consumption interval data, we compare our payoff estimates using the building-level consumption data as compared to our primary parameterization using city-level mean consumption data from our sample of new premises. (1) The blue series presents the net present values estimated with our primary parameterization (i.e., Figure 2a) arranged from least to greatest estimated payoff. The green points are the corresponding payoff estimate for each respective project in the blue series using instead the observed electricity consumption interval data. (2) This figure is restricted to only single family homes without solar. We exclude solar buildings because we only observe net consumption in the PG&E data, and we do not directly observe solar production. (3) We exclude one project with exceptionally high electricity consumption and estimated payoffs.