

Equity in Transition: Analyzing the Distributional Impacts of Clean Energy Technology Subsidies

Yağmur Menzilioğlu*

October 20, 2025[†]

[Click here for the latest version](#)

Abstract

I study the distributional and welfare effects of clean energy subsidies in the context of US residential rooftop solar panel adoption. Although these subsidies are often criticized as regressive because wealthier households claim most benefits, I show that this conclusion changes once dynamic externalities from learning-by-doing and pollution abatement are incorporated. Using installation-level data on residential solar system installations, I provide new evidence of learning spillovers and estimate learning elasticities to discipline a heterogeneous-agent general-equilibrium model with incomplete markets, irreversible adoption, endogenous cost declines, and unequal pollution damage exposures. Calibrated to US data, the model quantifies how alternative subsidy designs and financing schemes affect adoption, inequality, and welfare. Uniform refundable subsidies financed by a flat labor income tax raise aggregate welfare and accelerate adoption, whereas progressive financing or nonrefundable credits reduce support among lower-wealth households. When pollution damages are incorporated, the same subsidy becomes universally welfare-improving and strongly progressive. Accounting for dynamic spillovers and unequal pollution exposure overturns the view that residential solar subsidies are inherently regressive.

Keywords: Clean energy transition, Fiscal policy design, Heterogeneous-agent model, Residential solar installations, Learning spillovers, Inequality, Pollution exposure.

JEL Classifications: Q48, Q52, E21, E62, H23.

*Department of Economics, Georgetown University, email: ym406@georgetown.edu.

[†]I am deeply grateful for the guidance and support of my advisor, Toshihiko Mukoyama, and mentors Gaston Navarro and William Peterman. I also thank Arik Levinson, Mark Huggett, Martin Bodenstein, Adele Morris, Jeffrey Wooldridge, Joakim Weill, Harun Alp, Cristina Fuentes-Albero, Dan Cao, Adrien Bilal, Sadhika Bagga, Maarten De Ridder, Simone Lenzue, Cesar Santos, Pascual Restrepo, Stephen Terry, Vincenzo Quadrini, and Devika Chirimar for their valuable feedback and suggestions. All errors are my own.

1 Introduction

One concern regarding low-carbon (“clean”) technology subsidies is that high-income households claim them at a disproportionately high rate compared to lower income households due to high initial costs. For example, on average, the top income quintile in the United States (US) received more than 50% of investment tax credits for residential energy efficiency improvements since introducing these credits, according the Internal Revenue Service’s (IRS) Statistics of Income (SOI) data. If these subsidies that mainly go to the richer households are funded by general taxpayer revenues, then they are pecuniary transfers from poor to rich households. Such transfers are perceived as regressive in monetary terms, and as such, appear inequitable and raise discontent among the populace.

At the same time, if we consider both the environmental and learning-by-doing external economy effects alongside the general equilibrium effects, these policies that subsidize richer households may have greater welfare benefits for poorer households than for richer ones. For example, the health co-benefits of reducing local air pollutants associated with decreasing greenhouse gas emissions are larger for poorer households who are more commonly exposed to air pollution. Since wealthy households are responsible for the majority of residential emissions, speeding up their transition to low-carbon technologies will yield the most environmental benefits. Moreover, subsidizing early adopters who are likely to be wealthy could reduce the high adoption costs at an accelerated rate, making adoption feasible for later adopters who are likely to be less able to afford high adoption costs.

Motivated by this regressivity concern, I ask: How is low-carbon technology adoption related to household income in the US? Focusing on residential rooftop solar panel adoption, I show a positive correlation between income and the aforementioned adoption rates. If richer households receive more of these subsidies, then what are the distributional welfare implications of existing residential clean technology subsidies in the US over the joint income and wealth distribution? By calibrating a model to reflect the US economy and policy in the early days of the energy transition, I can quantify the equity and efficiency of the recent energy transition policies that promote adoption of residential rooftop solar panels.

Finally, I ask: what is the best policy mix to achieve different policy objectives, such as emissions reduction, welfare maximization, and transition speed, together with majority support. I use my calibrated model to evaluate the optimal policy mix under these different objectives and their majority support. First, I evaluate the overall welfare effects and political feasibility of different financing strategies for a uniform subsidy payment schedule through policies including the following: a uniform labor income tax, a progressive labor income tax, a corporate income tax, an emissions tax, and finally, government debt. In a second set of

experiments, I evaluate the effectiveness of different subsidy programs’ ability to meet the outlined policy objectives with an example being an income-capped subsidy program. These comparative policy experiments can inform policy design and implementation of residential energy efficiency improvement tax credits in the US and other countries.

Versions of the questions posed here have been addressed in literature. Most of the previous research falls into two categories: (i) empirical analyses of the distributional effects of environmental policies, and (ii) the macroeconomic studies on the unequal economic consequences of climate change policies.

Borenstein and Davis (2024) have well documented the adoption patterns for various low-carbon energy technologies and tax credit receipts for these technologies across income groups in the US. However, their study neither goes beyond descriptive facts nor quantifies dynamic general equilibrium effects, such as the indirect effects on relative prices, or the non-pecuniary effects such as the health co-benefits of reducing local air pollutants. Vona (2023) summarize the multidimensional impacts that climate policies have on well-being which have been documented by various researchers and provides an overview of the empirical literature. The consensus of this empirical literature is that subsidies to clean and energy-efficient equipment exacerbate the regressivity of climate policies in the presence of financial constraints by benefiting the wealthier households at the expense of less wealthy ones. In a static modeling framework without uncertainty, Levinson (2019) show that taxing energy use would be both more cost-effective and less regressive than subsidizing energy-efficient appliances or taxing inefficient appliances. These static analyses ignore the demand response that is the intended result of these policies.

I will contribute to these empirical analyses of who benefits from clean technology subsidies by quantifying the overall distributional effects of these policies while taking into account multidimensional effects using a model calibrated to reflect the US economy. Specifically, I incorporate two additional indirect distributional effects of clean technology subsidies into a general equilibrium model: (i) the learning-by-doing spillover effects of early adopters, and (ii) the health co-benefits of reducing local air pollutants associated with reducing greenhouse gas emissions. The first dimension that I incorporate in my analysis is the learning-by-doing spillover effects from early adopters, who are likely to be wealthier. Gao, Rai, and Nemet (2022) is one of few empirical economic analyses which find that learning through economies of scale in production leads to cost reductions, not only in hardware costs, but also in non-hardware costs in US residential solar installations. Secondly, reducing local air pollutant emissions has well-known positive effects in terms of health co-benefits such as reducing local air pollutants and improving health outcomes. Banzhaf, Ma, and Timmins (2019) summarize the large body of economic literature that documents a strong positive correlation between

ambient air pollution, poverty, and race – the so-called environmental justice gap.

There are several quantitative macroeconomic studies that analyze the distributional effects of climate policies in a general equilibrium framework. The majority of these studies focus on the distributional effects of carbon pricing instead of low-carbon technology subsidies as my paper does. Känzig (2023), Benmir and Roman (2022), Fried, Novan, and Peterman (2024), Fried, Novan, and Peterman (2018), and Belfiori, Carroll, and Hur (2024) are some recent studies that analyze the distributional effects of carbon pricing.

Känzig (2023) find that a carbon tax is regressive using institutional features of the European emissions trading system and high-frequency data. Benmir and Roman (2022) study the economic consequences of carbon pricing that achieves the net-zero emissions target in the US by 2050 and find that the policy induces large redistribution of income and wealth from poor to rich households. Fried, Novan, and Peterman (2024) study the welfare and inequality implications of different ways to return carbon tax revenue back to households and find that the welfare-maximizing rebate uses two-thirds of carbon tax revenue to reduce the distortianary capital income tax while using the remaining one-third to increase the progressivity of the labor income tax. In an earlier work Fried, Novan, and Peterman (2018) study how different approaches for recycling carbon tax revenue affect the welfare of current and future generations. Their work highlights the importance of not only long-run outcomes, but also the transitional welfare effects of how carbon tax revenues are recycled. Using household expenditure and emissions data, Belfiori, Carroll, and Hur (2024) document that low-income households have higher emissions per dollar spent than high-income households, making a carbon tax regressive. Overall, the literature shows that carbon pricing is regressive, but the regressivity can be mitigated by recycling the carbon tax revenue in a progressive way.

Kuhn and Schlattmann (2024) is the closest work to this project in that they develop a quantitative life-cycle model with heterogeneous adoption rates of carbon-neutral commitment goods by income to quantify the reduction-redistribution trade-off of different policy mixes. Lanteri and Rampini (2025) is another similar macroeconomic study of adoption of clean technologies, but instead by heterogeneous firms in a dynamic general equilibrium model of firm dynamics and clean technology adoption with financial constraints. They find that financially constrained firms optimally invest in dirty new technologies as well as in older technologies, resulting in a positive relation between firm size and energy efficiency. They highlight that their proposed model could be a laboratory for studying the distributional effects of environmental policy across firms, however they leave these exercises for future work. The computational methods I propose in this paper could easily be extended to study heterogeneous firms rather than households.

This paper makes three contributions to understanding the distributional effects of clean energy subsidies in the US residential rooftop solar panel installation market. First, I provide new empirical evidence on localized learning-by-doing in residential solar installations. Using installation-level data merged with state and utility policy shocks, I find that each doubling of cumulative installed capacity reduces system costs by about 7%. An instrumental-variable strategy exploiting the timing of new policy introductions identifies learning effects that are stronger when adoption is policy-driven rather than market-driven, confirming that subsidies accelerate cost declines through learning spillovers.

Second, I develop a heterogeneous-agent dynamic general equilibrium model with incomplete markets, irreversible technology adoption, adoption-driven cost reductions, and unequal pollution damages. This framework captures both the private adoption incentives and the general equilibrium feedbacks that shape how residential clean energy subsidies affect adoption patterns, inequality, and welfare.

Third, I calibrate the model to reflect the US economy and policy environment in the early days of the energy transition, using detailed data on household demographics, income and wealth distributions, residential solar adoption patterns, and policy parameters. The calibrated model delivers several novel quantitative findings: Uniform refundable subsidies financed by a flat labor-income tax raise aggregate welfare and speed up adoption, with 94% of households benefiting. Financing subsidies through a progressive tax reduces overall welfare gains and disproportionately lowers gains for low-wealth households by depressing short-run wages and transfers. Nonrefundable tax credit, which mirrors the structure of the U.S. federal residential solar credit, further excludes low-income households, but does not slow down diffusion. Income-capped subsidies, while intended to improve fairness, slow down adoption, reduce learning spillovers and generate aggregate welfare losses that especially hurt the middle-wealth households.

When pollution damages are included, the nonrefundable uniform subsidy becomes universally welfare-improving and strongly progressive, as cleaner air disproportionately benefits poorer households. Together, these results show that the perceived regressivity of residential solar subsidies reflects a partial-equilibrium perspective. Once dynamic cost declines and pollution externalities are accounted for, the equity-efficiency trade-off in clean energy policy becomes much weaker.

The remainder of the paper is structured as follows. In Section 2, I summarize the data that motivates the research questions and provide background for the model. In Section 3, I outline the structural model that I use to answer my research questions. In Section 4, I describe the complete characterization of the model used for quantitative analysis, its calibration and fit to the data, and present the baseline model simulations. In Section 5, I

Table 1: Descriptive statistics for the income quintiles in the US

Income Percentile	Bottom 20%	20%-40%	40%-60%	60%-80%	Top 20%	Top 5%
Share of aggregate income	3.17	8.42	14.37	22.83	51.21	22.81
Share of residential energy consumption	12.59	18.16	13.07	19.75	24.42	12.01
Share of rooftop solar adopters	0.53	3.42	9.26	16.87	43.75	26.19
Share of residential clean energy credits	0.48	4.11	4.08	21.75	48.99	20.59
Mortality damages per capita (2020 dollars pp)	4,811	3,910	3,103	2,769	2,354	NA

Notes: Reported shares and rates are in percentages, except for the mortality damage per capita values, which is in 2020 US dollars per person (pp). NA indicates not available.

conduct policy experiments to evaluate the effectiveness of different policy mixes in achieving the outlined policy objectives and majority support. Finally in Section 6, I conclude and discuss the implications of the results for policy design and implementation.

2 Data and Empirical Motivation

To motivate my research questions, it is essential to understand the distributions of residential energy consumption, adopters of on-site solar power generation, the receipts of residential energy credits, ambient air pollution exposure, and energy expenditure shares across income quintiles in the US. I summarize these distributions in Table 1 using various cross-sectional data sources from the US in 2015. I focus on 2015 because it is the earliest year for which data on all these variables is available.

First, using data from the US Census Bureau’s (2023) 2015 American Community Survey (ACS) 5-Year Estimates, I construct income quintiles based on the upper income limits of quintiles summarized in Table B19080. For 2015, these limits are \$17,929, \$35,583, \$62,600, \$108,429, and the lower limit for top 5% is \$146,778. I report the share of aggregate income for each income quintile from the US Census Bureau’s (2023) 2015 ACS Table B19082 in the first row of Table 1. The top income quintile accounts for more than half of the aggregate income in the US.

Second, I use the US Energy Information Agency’s (2023) 2015 Residential Energy Consumption Survey (RECS) data to calculate the share of aggregate residential energy consumption for each income quintile. Annual household income is reported as a categorical variable in the RECS data, and I group households according to the income quintile’s upper limits as closely as possible. Thus, the upper income limits for the quintiles I report from the RECS data are \$20,000, \$40,000, \$60,000, \$100,000, and the lower limit for the top 5% is \$140,000. I use the household weights provided in the RECS data for all of my calculations. I calculate the share of aggregate residential energy consumption for each income quintile

through the following process: aggregating total energy consumption in British Thermal Units (BTUs) for each group; calculating the national total, and calculating the share of each group's total as a proportion of the national total. The results reported in the second row show that the top income quintile accounts for almost one quarter of the aggregate residential energy consumption in the US.

Third, I summarize some facts about rooftop solar panel deployment rates and the receipts of residential energy credits that cover installment of residential solar panels, across income quintiles. I use the 2015 RECS data to calculate the rooftop solar panel deployment rates for each income quintile. I calculate the deployment rate for each income quintile by dividing the number of households that generate power on-site using solar in that quintile by the total number of households that generate power on-site using solar across all income groups. The results reported in the fourth row show that the share of rooftop solar adopters increases with income, with the top income quintile having almost half of the adopters.

The US has several government-sponsored incentive programs that reduce the cost for people and businesses to use alternative energy sources. Eligible taxpayers meeting the criteria get the credit amount deducted from their total tax liability. There are two types of energy investment tax credits (ITCs) available to homeowners: the Energy Efficient Home Improvement Credit (EEHIC) and the Residential Clean Energy Credit (RCEC). Taxpayers can only use one or the other of these two residential energy tax credits in any one tax year.

The EEHIC offers a 30% ITC to cover some of the cost of eligible home improvements, such as alterations to exterior doors, windows, and electric or natural gas heat pumps. The RCEC, formerly known as the Residential Energy Credit, is an ITC that was extended and renamed under the 2022 Inflation Reduction Act. The credit covers 30% of the cost of installing solar panels, solar water heaters, geothermal heat pumps, small wind turbines, geothermal heat pumps, fuel cells and battery storage technologies of at least 3 kilowatts (kW) per hour (kWh). The credit is available for both existing homes and new construction, but the home must be the taxpayer's primary residence. The RCECs were enacted as part of the Energy Policy Act of 2005 and were extended and expanded several times since then. The final extension and expansion of the RCECs was under the Inflation Reduction Act, increasing the credit back to 30%, where it was set to remain until 2032 before dropping to 26% in 2033 and then 22% in its final year, 2034. However, the most recent 2025 One Big Beautiful Bill Act law eliminates RCECs completely after December 31, 2025. Importantly, the RCECs have been non-refundable, meaning that the credit cannot exceed the taxpayer's tax liability. The non-refundability is a key feature of the RCECs that makes them regressive, as only those with a tax liability, who tend to be wealthier, can benefit from the credit.

I use the IRS's (2023) 2015 Statistics of Income (SOI) data to calculate the share of

total value dispensed under each RCEC attributed to each income group. Unfortunately, the publicly available SOI data does only provide income categories and the ranges for these categories do not align perfectly with the income quintiles I use. Therefore, I group the income categories in the SOI data to match the income quintiles as closely as possible. Specifically, the upper income limits for the income categories I report from the SOI data are \$20,000, \$40,000, \$50,000, \$100,000, and the lower limit for top 5% is \$200,000. I calculate the share of the amount of RCECs claimed by each income group by dividing the total amount of residential energy credits issued to each group by the total amount of residential energy credits issued to all income groups. The results in fourth row show that the share of RCECs claimed by the income group increases with income, with the top income quintile receiving almost half of all residential energy credits issued in the US.

These four statistics describe that higher income households consume more energy and thus contribute more to pollution, adopt rooftop solar panels more, and receive more residential energy credits than lower income households. Thus, having richer households adopt clean technologies faster could yield the most emissions reductions. However, as I argued in the introduction, the health co-benefits of reducing local air pollutants is greatest for the poorer households who are more commonly exposed to air pollution. In order to provide evidence for this claim, I use Dennin et al.’s (2024) latest Air Pollution Emission Experiments and Policy Analysis model, the AP4 model, to calculate the county-level mortality damages per capita for 2017. I match these mortality damages with county-level median income data from the US Census Bureau’s (2022) 2017 ACS 5-Year Estimates. The results reported in the sixth row show that mortality damages per capita are inversely related to income, with the bottom income quintile experiencing the highest damages. Thus, reducing emissions by a unit would prevent greater damages for the poorer households. These are some of the distributional facts that motivate my research questions and which I will use for the calibration of my model.

2.1 Benefits of Residential Rooftop Solar Panel System Deployment

Deploying solar panels for on-site power generation has both private and social benefits. The private benefits include reduced electricity bills, increased property values, and reduced exposure to electricity price volatility. The social benefits include reduced emissions of greenhouse gases and local air pollutants, reduced strain on the electricity grid, and increased energy security. In this section, I describe these potential benefits of adopting solar panels for on-site power generation, provide descriptive evidence on the private benefits, and describe the potential social benefits that are not internalized by households.

Table 2: Average 2015 energy expenditure shares of income groups in the US

Income Percentile	Bottom 20%	20%-40%	40%-60%	60%-80%	Top 20%	Top 5%
Share of energy expenditure in total expenditure	8.5	6.4	5.0	4.2	3.3	2.2

Note: Reported shares are in percentages.

2.1.1 Private Benefits

First, I describe the private benefits of adopting solar panels for on-site power generation. The most immediate private benefit is a reduction in household electricity expenditures. To illustrate the magnitude of this accounting effect, I use the 2020 RECS. This dataset includes information on annual household electricity expenditures and consumption, as well as an indicator for whether the household has on-site solar generation.

Unsurprisingly, households with on-site solar report substantially lower grid electricity expenditures. For example, a simple regression of electricity expenditures on a solar indicator (presented in Appendix Table 12) shows that solar households spend roughly \$700 less annually on electricity. The average annual electricity expenditure in the RECS data is about \$1,400, so this represents a substantial reduction. However, this estimate should not be interpreted as a causal effect: it simply reflects the accounting identity that households generating their own electricity purchase less from the grid. Because solar adoption is endogenous to household characteristics and policy incentives, these regressions are not identified. I therefore report them only to illustrate magnitudes and relegate the tables to the Appendix A.

The key point for calibration is that these private pecuniary savings are meaningful in household budgets. Households with lower income and wealth levels have greater marginal utility of consumption and devote a larger share of their expenditures to energy. As a result, an identical dollar reduction in electricity expenditure has larger welfare consequences for poorer households. To discipline these heterogeneous effects in the model, I use the 2015 Consumer Expenditure Survey (CES) to calculate the energy expenditure shares of U.S. households by income quintile. The results in Table 2 show that the bottom income quintile devotes 8.5% of total expenditure to energy, compared with only 3.3% for the top quintile. These differences in energy budget shares will be reflected in the calibration of household heterogeneity in the model.

2.1.2 Social Benefits

Next, I describe the social benefits of adopting solar panels for residential on-site power generation. I will focus on the environmental benefits of reducing emissions of local air pollutants. For example in the US, the residential sector accounted for 15% of end-use energy

consumption 2023 according to US Energy Information Administration (2024).

The residential sector is also a significant source of air pollutants regulated by the Environmental Protection Agency (EPA) using human health-based and environment-based criteria under the Clean Air Act. Particle pollution, also known as particulate matter (PM), is a mixture of solid particles, such as dust, dirt, and soot, and liquid droplets found in the air. Breathing in particle pollution can be harmful to human health, as it can cause heart attacks, trouble breathing, lung cancer, and problems with babies.

Smaller particles, called $PM_{2.5}$, pose the greatest health risks, because they can penetrate deep into the lungs and the bloodstream. In 2020, fuel combustion by the residential sector accounted for 10.9% and fuel combustion by the electric power generation sector accounted for 1.3% of $PM_{2.5}$ emissions in the US, according to my calculations using EPA’s 2020 National Emissions Inventory (NEI) data. Dennin et al. (2024) estimate that the marginal damage associated with an additional ton of $PM_{2.5}$ emissions in the US to be between \$73,200 and \$133,000 per ton in 2020 dollars. Thus, given the significant share of $PM_{2.5}$ emissions from the residential sector, reducing emissions from this sector could yield health benefits for local communities.


2.2 Cost of Residential Rooftop Solar Panel System Installations

The cost of installing solar panels for on-site power generation is a significant barrier to adoption for many households. The total cost includes the following costs: the solar panels themselves, the inverter, the mounting hardware, the wiring, the installation labor, and the permitting and inspection. The price of solar panel installations has been decreasing over time due to technological advancements and economies of scale, even before accounting for government incentives.

Figure 1 shows the median price per watt and the median system size of residential solar panels installations in the US from 2000 to 2022. The data is from National Renewable Energy Laboratory’s (2023) 2022 Tracking the Sun report data, which covers all non-utility scale solar panel installations in the US. The figure shows that the median price of residential solar panels in the US has declined by almost 65% from 2000 to 2022. The median price per watt declined by almost 65%, while the median system size increased by nearly 75%. This joint trend highlights that, although unit costs fell, households increasingly adopted larger systems, so the decline in total installation costs was slower.

Using these two series on the price and capacity of residential solar installations, I calculate two median total costs of residential solar installation measures in the US from 2000 to 2022. The first measure is the median gross total cost, which is the product of the median total

Figure 1: Median price and size of residential solar panel system installations in the US per quarter, 2000-2022




Note: The shaded area represents the 25th and 75th percentiles of the distribution of prices of residential solar panel installations.

installation price per watt and the median system size. The second measure is the median net total cost, which accounts for the state- and utility-level incentives and rebates deducted from the gross total cost for residential solar installations. Importantly, the net total cost measure does not account for the federal investment tax credit (ITC) for residential solar installations, which is 30% of the gross total cost in 2022.

To assess affordability, Figure 2 reports the ratio of these cost measures to median household income in the US, using data from the US Census Bureau's (2022) ACS 5-Year Estimates (Table S1901). The figure shows that the median gross cost of a residential solar installation fell from about 65% of median household income in 2000 to around 36% in 2022. The gap between gross and net costs narrowed over time and eventually disappeared, reflecting the expiration of many state- and utility-level support programs during this period.

The decline in residential solar installation costs is widely attributed to learning effects. As more systems were produced and installed, both manufacturing and installation processes became more efficient, resulting in lower prices over time.

Figure 2: Ratio of median gross and median net system prices of residential solar PV systems to median annual household income



2.3 Did Rooftop Solar Panel Installations Experience Learning Effects?

The concept of learning effects suggests, as the cumulative experience with a technology increases, the costs associated with that technology tend to decrease. The potential learning effects from the increased adoption of technologies with positive externalities, such as rooftop solar panel installations, can lead to lower costs for future adopters. Thus, subsidies to early adopters of such technologies can have positive spillover effects on the costs of future adopters and has been used to justify public funding for subsidizing early adopters. In the context of rooftop solar panel installations, this could manifest in several ways, including improved manufacturing processes, better installation techniques, and more efficient supply chains.

There are several empirical studies that have found evidence of learning effects in rooftop solar panel installations. Nemet, O'Shaughnessy, et al. (2016) find that more experienced installers consistently quote lower prices than novice installers, with all else being equal. For instance, low-priced solar panel systems – the cheapest 10% of installs – are disproportionately completed by installers with extensive prior installations, highlighting how accrued experience translates into cost efficiency. O'Shaughnessy (2018) observed, in more concentrated markets, average installation costs tend to be lower, presumably because high-volume installers climb further down the learning curve and achieve economies of scale. However, if a market becomes

too concentrated, competitive pressures may diminish, leading to lower prices in moderately concentrated markets that balance scale benefits with competition. Nemet, Lu, et al. (2020) find knowledge spillovers between firms within a county to be a significant and substantial factor in reducing the costs of solar PV installations using data from 2008 to 2014. They find that these spillovers reduce costs only for firms over a certain size threshold and geographic spillovers within a firm across counties are also significant but smaller than the local between-firm effects.

Overall, literature strongly links cumulative installation experience with cost reductions: every doubling of US residential PV capacity has historically yielded a predictable price decline (a “learning rate”), though estimates of the learning rate for soft costs vary from study to study. Solar Energy Industries Association (2025) reports that the typical residential solar PV installation prices almost halved since early 2010s. Bollinger and Gillingham (2023) estimate that localized learning-by-doing accounted for a modest but measurable decline in installation prices – on the order of \$0.12 per watt reduction in soft costs from each doubling of installer experience in California’s early rooftop solar market data. They also found only limited spillovers between firms, meaning much of the learning was internal to each installer’s operations. On the other hand, Gao, Rai, and Nemet (2022) find that traditional learning-by-doing significantly reduced soft installations costs, although its effect is partly masked by other learning mechanisms such as accounting for installers’ learning-by-searching (e.g., innovation and R&D) and learning-by-interacting (e.g., knowledge spillovers via supplier networks). Their findings suggest that the traditional learning-by-doing effect may not be as dominant as previous literature suggests.

I test the hypothesis that the cost of residential solar panel installations has decreased over time due to learning effects by regressing the median price of residential solar panel installations on the cumulative installed capacity of residential solar panels in the US and quantify the learning-by-doing effect. There are a variety of statistical models for learning effects all based on the power law of learning, which states that the cost of a technology decreases by a constant percentage with each doubling of cumulative production. I implement the learning-by-doing effect as an exponential decay function, which is a common approach in existing literature. The underlying model of learning I assume is as follows:

$$p_t = p_0 \cdot \mathcal{I}_t^{-\xi} \cdot \exp(-\lambda t), \quad (1)$$

where p_t is the net installation price of residential solar panels per watt installed capacity at date t once rebates and incentives are taken into account, p_0 is the initial installation price of residential solar panels per watt installed capacity, $\mathcal{I}_t$ is the cumulative installed capacity of

residential solar panels before date t – could either be a count or a size measure – ξ is the learning-by-doing parameter that captures the rate of cost reduction with each doubling of cumulative installed capacity, and λ is the rate of exogenous time decay that captures the secular declines in PV costs due to global supply chain improvements, technological change, and other factors unrelated to local learning-by-doing. The learning rate implied by this model is given by $1 - 2^{-\xi}$, which represents the percentage reduction in cost associated with each doubling of cumulative installed capacity.

To allow for learning at multiple margins, I include cumulative experience at both the state and firm levels. Let $s(i)$ denote the locality of installation i and $f(i)$ the responsible installer firm. I estimate this model by regressing the natural logarithm of the installation price of residential solar panels on the natural logarithm of the cumulative installed capacity of residential solar panels.

To accommodate learning at multiple margins, I specify the following regression model:

$$\log p_i = \alpha_{j(i)} - \xi^{\text{state}} \log(\mathcal{I}_{s(i), t(i)-12}^{\text{state}}) - \xi^{\text{firm}} \log(\mathcal{I}_{f(i), t(i)-12}^{\text{firm}}) - \lambda t(i) + \gamma g_{s(i), t(i)} + X_i' \theta + \varepsilon_i, \quad (2)$$

where, p_i is the net installation price per watt of installation i , $\alpha_{j(i)}$ are fixed effects for the relevant entity j (state, county, or firm) depending on the specification, $\mathcal{I}_{s(i), t-12}^{\text{state}}$ is the cumulative installed residential capacity in state $s(i)$ lagged 12 months, $\mathcal{I}_{f(i), t-12}^{\text{firm}}$ is cumulative installed capacity by firm $f(i)$ lagged 12 months, ξ^{state} and ξ^{firm} are the corresponding learning elasticities, $t(i)$ is the installation month, λ is the exogenous time-decay parameter, $g_{s(i), t(i)}$ measures contemporaneous incentive generosity in the state of installation, X_i is a vector of installation-level controls (system size, hardware, financing), and ε_i is the error term.

In addition to the Ordinary Least Squares (OLS) estimates of equation 2, I consider Instrumental Variable (IV) approaches to address potential endogeneity concerns. Cumulative installations are likely endogenous due to simultaneity bias (e.g., areas or firms with lower costs may attract more installations, making it look like learning when in fact it is selection bias). Moreover, policy support, local demand shocks, or installer entry and exit could drive both installations and prices. An IV approach could isolate the variation in installations that comes from the exogenous policy changes. Finally, an IV approach could provide external validity by leveraging variation across different policy contexts.

I instrument for local cumulative installed capacity using local-level policy changes as an instrument. To capture the exogenous variation in incentives, I construct a monthly binary shock variable that takes the value of 1 if a new residential solar policy starts in location j in month t , and 0 otherwise. To construct this policy shock variable, I use the North Carolina Clean Energy Technology Center’s (2025) Database of State Incentives for

Renewables & Efficiency (DSIRE) data, which provides detailed information on state and local policies and incentives for renewable energy and energy efficiency in the US. I provide a detailed description of the construction of this instrument and the policy generosity measure in Appendix [A.1.2](#).

The IV design follows a standard two-stage least squares (2SLS) approach. The core concern is that cumulative installed capacity at the state or county level may be endogenous: areas with falling prices for other reasons may also attract more installations, biasing OLS estimates downwards. I use lagged policy shocks as an instrument for local cumulative installed capacity, to capture the introduction of new state or county incentives (rebates, grants, tax credits, net metering measures), lagged by 12 months to ensure that they affect the stock of cumulative installations but not contemporaneous prices directly. The first stage regression links cumulative installed capacity (the endogenous regressor) to the lagged policy shock, controlling for firm and time fixed effects. A strong first stage relationship is crucial for the validity of the IV approach, as it ensures that the instrument is correlated with the endogenous regressor and satisfy the relevance requirement, and can effectively isolate exogenous variation in installations.

In the second stage, I estimate equation (2) using the predicted values of cumulative installed capacity from the first stage regression, together with additional controls and fixed effects. This approach isolates variation in learning-by-doing that is plausibly exogenous, coming from policy shocks rather than unobserved demand or cost shocks. The exclusion restriction assumes that lagged policy shocks affect installation prices only through their effect on cumulative installed capacity, not directly. This assumption would be reasonable if policies were targeted at adoption rather than subsidies to system prices conditional on installation. However, the majority of the policies do indeed provide direct financial incentives that affect net installation prices, potentially violating the exclusion restriction.

To ensure validity, I implement a refinement of the IV strategy. Specifically, I exploit policy timing shocks by using lagged indicators for the onset of a new policy as instruments for cumulative installations, while simultaneously controlling for the contemporaneous generosity of all active incentives (e.g., rebate amount per watt). This design ensures that any direct effect of subsidies on prices in period t is absorbed by controls, while variation in the timing of policy introductions (lagged L months) provides exogenous shocks to the cumulative stock. After this refinement, my instruments are indicators for whether a new policy incentive began L months earlier, with contemporaneous generosity always included as a control. Identification therefore relies on the assumption that, conditional on current subsidy generosity and fixed effects, the lagged timing of a policy’s onset influences current installation prices only by raising the cumulative stock of past installations (learning-by-doing), not by directly lowering

system costs at the time of observation.

Thus, the first stage regression is as follows:

$$\log(\mathcal{I}_{s(i),t-12}^{\text{state}}) = \mu_{s(i),t-12} + \pi Z_{s(i),t-12} + \rho g_{s(i),t-12} + \delta t + W_{s(i),t-12}'\theta + u_{s(i),t-12}, \quad (3)$$

where $\mu_{s(i),t-12}$ is a vector of entity fixed effects, $Z_{s(i),t-12}$ is the policy shock variable for entity $s(i)$ lagged by 12 months, $W_{s(i),t-12}$ is a vector of additional controls, and $u_{s(i),t-12}$ is the error term. The second stage regression is similar to equation 2, but uses the predicted values of cumulative installed capacity from the first stage regression, and is as follows:

$$\log p_i = \alpha_{j(i)} - \xi^{\text{state}} \log(\widehat{\mathcal{I}}_{s(i),t(i)-12}^{\text{state}}) - \xi^{\text{firm}} \log(\mathcal{I}_{f(i),t(i)-12}^{\text{firm}}) - \lambda t(i) + \gamma g_{s(i),t(i)} + X_i'\theta + \varepsilon_i, \quad (4)$$

where $\widehat{\mathcal{I}}^{\text{state}}$ are the predicted values from the first stage regression. The causal effect of learning is identified by ξ .

Equations (3) and (4) constitute a two-stage least squares (2SLS) design. In the first stage, lagged policy shocks serve as instruments for cumulative installed capacity, exploiting the timing of new incentive introductions as an exogenous source of variation in adoption. I include contemporaneous policy generosity, $g_{j,t}$, in both stages to absorb any direct price effects of currently active subsidies so that the identifying variation arises only from past policy onsets. I also control for a global time trend and a rich set of technological controls to account for background cost declines and other market dynamics. The second stage then regresses log net prices on the predicted cumulative installations from the first stage, controlling for the same covariates and fixed effects. Under the exclusion restriction, the lagged policy shocks affect current installation prices only through their effect on cumulative adoption, not directly, once contemporaneous generosity and time trends are held constant. The coefficient ξ is thus identified as the causal elasticity of installation prices with respect to cumulative installed capacity which I interpret as the learning-by-doing effect.

Table 3 presents results. Columns (1), (3), (5), and (7) report OLS estimates with different fixed effects (FE) specifications, while columns (2), (4), and (6) report IV estimates, alternating OLS-IV within each entity level. State FE appear in columns (1)-(4), county FE in (5)-(6), and firm FE in (7). The alternating structure highlights how coefficient magnitudes shift when moving from OLS to IV specifications. The first-stage F-statistics in the IV columns (2), (4), and (6) exceed 10, indicating strong instruments. The coefficients on lagged cumulative installations at both the state and firm levels are the primary parameters of interest, capturing learning effects.

The results in Table 3 show mixed but broadly supportive evidence for learning-by-doing in residential rooftop solar panel installations. All specifications yield negative and statistically

Table 3: Learning-by-Doing in Residential PV with Exogenous Unexplained Decay: OLS and IV with Fixed Effects

Independent variables	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV	(7) OLS
log Cumulative Installs (Firm, 12m lag)	−0.0073 (0.0003)	−0.0073 (0.0003)	−0.0097 (0.0003)	−0.0097 (0.0003)	−0.0073 (0.0003)	−0.0073 (0.0003)	−0.0110 (0.0006)
log Cumulative Installs (State, 12m lag)	0.0226 (0.0022)	0.1027 (0.0054)	0.0244 (0.0022)	0.0996 (0.0055)			−0.0033 (0.0010)
log Cumulative Installs (County, 12m lag)					0.0145 (0.0014)	0.1249 (0.0055)	
t	0.0029 (0.0001)	0.0029 (0.0000)	0.0030 (0.0001)	0.0030 (0.0000)	0.0031 (0.0000)	0.0031 (0.0000)	0.0042 (0.0000)
Policy Generosity $\times 10^{-4}$	−0.0034 (0.0002)	−0.0034 (0.0002)	−0.0034 (0.0002)	−0.0034 (0.0002)	−0.0091 (0.0005)	−0.0091 (0.0005)	−0.0021 (0.0002)
Has DC Optimizer	0.0303 (0.0022)	0.0303 (0.0022)			0.0347 (0.0022)	0.0347 (0.0022)	0.0105 (0.0022)
Ground Mounted	0.0229 (0.0053)	0.0229 (0.0069)			0.0308 (0.0054)	0.0308 (0.0069)	0.0673 (0.0047)
Has Microinverter	0.0098 (0.0021)	0.0098 (0.0024)			0.0109 (0.0022)	0.0109 (0.0024)	0.0635 (0.0023)
Inverter Loading Ratio	0.0418 (0.0038)	0.0418 (0.0040)			0.0425 (0.0038)	0.0425 (0.0040)	−0.0346 (0.0036)
log Size	−0.1040 (0.0016)	−0.1040 (0.0019)			−0.0897 (0.0016)	−0.0897 (0.0020)	−0.0952 (0.0015)
Has Tracking Bin	−0.3999 (0.0089)	−0.3999 (0.0298)			−0.4023 (0.0089)	−0.4023 (0.0298)	−0.0833 (0.0090)
First stage							
Policy Shock (State, 12m lag)		−0.7520 (0.0070)		−0.7518 (0.0069)			
Policy Shock (County, 12m lag)						−1.1759 (0.0141)	
F - Statistic		11,705.12		11,701.16		6,995.60	
Number of observations	874,991	874,991	874,991	874,991	874,991	874,991	874,991
R^2	0.060	0.013	0.053	0.013	0.058	0.008	0.054
Location FE	state	state	state	state	county	county	firm

Notes: Robust standard errors in parentheses. Coefficients with robust standard errors in parentheses. Columns are numbered with OLS and IV alternating. First-stage coefficients appear only under IV columns. “log Cumulative Installs (Firm/State/County, 12m lag)” correspond to $\log(\mathcal{I}_{f(i),t-12}^{\text{firm}})$, $\log(\mathcal{I}_{s(i),t-12}^{\text{state}})$, and $\log(\mathcal{I}_{c(i),t-12}^{\text{county}})$. “Policy Shock (State/County, 12m lag)” equals 1 if a new residential PV incentive began in that entity 12 months earlier.

significant coefficients on lagged cumulative firm capacity, suggesting that more experienced firms charge higher prices consistent with market power rather than learning. In contrast, the state- and county-level cumulative installation coefficients are consistently positive and significant when location fixed effects are included in columns (1)-(6), suggesting that local market experience reduces prices. However, the magnitude and significance of these coefficients vary across specifications, ranging from 0.0145 to 0.1249, corresponding to learning rate estimates ranging from 1 to 8% cost reduction per doubling of cumulative capacity.

The difference in relative magnitudes of state- and county-level learning effects across the

OLS and IV specifications is notable. The state-level learning elasticity is four times larger in the IV specifications in columns (2) and (4) than in the OLS specifications in columns (1) and (3). This suggests that OLS may understate state-level learning effects, potentially due to simultaneity and measurement error biases. Instrumenting for cumulative installations with lagged policy shocks appears to amplify the estimated learning effect, indicating that exogenous policy-driven adoption has a stronger impact on cost reductions than organic market growth. The difference between the OLS and IV estimates is even more pronounced at the county level, where the IV estimate in column (6) is nearly six times larger than the OLS estimate in column (5).

The estimates from the OLS specification with firm fixed effects in column (7) are puzzling: both coefficients on cumulative firm capacity and cumulative state capacity are negative and significant, suggesting that more firm or local experience leads to higher prices. This finding contradicts the learning-by-doing hypothesis, which would predict that increased experience leads to lower prices. This counterintuitive result may reflect unobserved heterogeneity among firms, such as market power or strategic pricing behavior, rather than true learning effects. Firms with more experience may have established reputations or customer bases that allow them to charge higher prices, regardless of their actual cost structures and states with more installations may have higher demand and thus higher prices, confounding the learning effect.

My preferred specification is the IV model with state fixed effects in column (2). This specification balances the need to control for unobserved heterogeneity at the state level while still allowing for meaningful variation in cumulative installations. The IV approach helps address endogeneity concerns, and state-level fixed effects capture important policy and market differences across states. The estimated state-level learning elasticity of 0.1027 implies a learning rate of approximately 7% cost reduction per doubling of cumulative capacity.

The inclusion of an exogenous time decay term (λ) reveals a strong and highly significant negative trend across all models, with magnitudes of -0.0015 to -0.003 per month. This implies secular annual declines in PV costs of roughly 2-4% annually, independent of local learning effects. The policy generosity coefficient (γ) is consistently negative, significant and small, indicating that more generous current incentives are associated with lower prices, as expected. Additional controls behave as expected: larger systems are cheaper per watt, while technology upgrades such as optimizers and microinverters increase costs.

My regression results support the consensus view in the literature that rooftop solar panel installations have experienced learning effects, and state-level experience and markets matter more than purely local and firm learning. Moreover, exogenous time trends are a major driver of cost declines, reflecting global supply chain improvements and technological change. The contribution of these results to existing literature is to provide estimates for both OLS and

IV specifications, confirming that OLS may overstate firm-level learning and that state-level cumulative capacity passes IV robustness.

3 Model

Motivated by these empirical observations, I develop a heterogeneous agent dynamic stochastic general equilibrium (DSGE) model with incomplete markets (i.e., Aiyagari-Bewley-Huggett models, building on Bewley (1977), Huggett (1993), and Aiyagari (1994)), with the inclusion of costly clean energy adoption decisions and environmental externalities. The model has discrete time and infinite horizon. There is a continuum of households indexed by $i \in [0, 1]$. Consumers supply an exogenous amount of labor, receive wage income, accumulate assets in physical capital, and rent out capital to firms. The labor supply is stochastic and is governed by an idiosyncratic shock. Consumers can self-insure by accumulating assets subject to a borrowing constraint.

Households have preferences over consumption and ambient air pollution. Consumption requires energy use, and energy use is assumed to be an affine function of final good consumption. This assumption implies that the expenditure share of energy is larger for households with lower consumption levels, which is an empirical fact documented in Table 2. Consumers have two energy technology options: old fossil (dirty) energy and new renewable (clean) energy. Energy produced by the two technologies are perfect substitutes for each other. However, energy consumption through the old technology creates ambient air pollution – a flow variable that results in utility damages to consumers – a negative externality to society. Before the clean energy technology is available for household adoption, all households use the old energy technology and the economy is in the initial steady state.

Once the aforementioned clean energy technology is available, consumers can switch to a clean energy source by deploying technology. For example, a consumer can adopt rooftop solar panels to generate electricity with a one-time investment cost, and after having adopted, they will use the clean energy in perpetuity. The unit cost of clean energy is lower than the unit cost of dirty energy, but the initial investment cost is very high. Specifically, the clean energy technology investment is lumpy: it requires a large one-time and irreversible upfront cost, and once the investment is made, it cannot be recovered. As the scale of clean energy adoption increases, the investment cost decreases due to learning-by-doing spillover effects. After a sufficiently long time, all households will have adopted the clean energy technology and the economy will be in a new steady state: the terminal steady state.

Firms combine capital and labor to produce the final good (used for consumption and investment). The negative externality generated by dirty energy consumption is not internal-

ized by the firms or consumers, but leads to a fraction of output being lost due to damages. The loss of utility will be greater for poorer households due to decreasing marginal utility of consumption. Both firms and consumers act competitively and take prices as given. The government collects labor income taxes and uses the revenue to subsidize the investment cost of adopting clean energy technology and returns any excess tax revenues to consumers lump-sum.

3.1 Consumers

Each household is infinitely lived and has preferences over consumption and ambient air pollution. At time t , a household's individual state is described by a vector z_t defined as:

$$z_t = (a_t, \ell_t, s_t) \in \mathcal{Z},$$

where $a_t \in \mathcal{A} = [0, \infty)$ is the household's risk-free asset holding at the beginning of period t , $\ell_t \in \mathcal{L}$ is the idiosyncratic labor productivity endowment at time t , and $s_t \in \{0, 1\}$ is the household's utilization status of the clean energy technology at time t , where $s_t = 0$ indicates that the household is using the fuel combusting old energy technology and $s_t = 1$ is the new clean energy technology. Define the measurable space $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$, where:

$$B(\mathcal{Z}) = B(\mathcal{A}) \times P(\mathcal{L}) \times P(\{0, 1\}),$$

with $B(\mathcal{A})$ being the Borel σ -algebra on $\mathcal{A}$ and $P(\cdot)$ the power set. The cross-sectional distribution of households over the state space at time t is represented by a probability measure $\Phi_t \in \mathcal{M}$, where $\mathcal{M}$ is the set of all Borel probability measures on $(\mathcal{Z}, B(\mathcal{Z}))$. For any measurable set $B \in B(\mathcal{Z})$, $\Phi_t(B)$ is the fraction of households with states in B at time t . I will denote $\Phi_t(B)$ by Φ_t when there is no ambiguity. Aggregate objects are computed as integrals with respect to the invariant cross-sectional measure over states, following Huggett (1993).

Each household supplies one unit of time endowment inelastically to the labor market with labor productivity ℓ_t that follows a finite-state Markov chain with transition matrix $\pi(\ell'|\ell)$ and a unique invariant distribution $\Pi(\ell)$. Households derive utility from consumption and ambient air pollution according to:

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(c_t, X_t) \right],$$

where $\beta \in (0, 1)$ is the discount factor, c_t is consumption at time t , and X_t denotes the

aggregate ambient air pollution at time t . $U(\cdot, \cdot)$ is a strictly increasing and concave one-period utility function in its first argument, strictly decreasing and convex in its second argument, and $\mathbb{E}_0$ is the mathematical expectation conditioned on the consumer's time-0 information.

The household budget constraint is:

$$c_t + a_{t+1} + \bar{q}_t e_t(1 - s_t) + \underline{q}_t e_t s_t + p_t(1 - \tau_t)S_t = w_t(1 - \tau^\ell)\ell_t + (1 + r_t)a_t + T_t,$$

subject to the borrowing constraint:

$$a_{t+1} \geq \underline{a},$$

where $e(\cdot)$ is the energy consumption function that maps household's consumption to energy demand, w_t and r_t denote the wage and interest rate, respectively, $\bar{q}_t$ and $\underline{q}_t$ are the exogenous unit energy prices under dirty and clean energy technologies, respectively, with $\underline{q}_t < \bar{q}_t$ for all t , $s_t^i \in \{0, 1\}$ is the household's utilization status of the clean technology, where $s_t^i = 0$ indicates that the household is using the fuel combusting old energy technology and $s_t^i = 1$ is the new clean energy technology, $S_t^i \in \{0, 1\}$ is the irreversible binary technology adoption decision, p_t is the one-time clean energy technology adoption cost, τ_t is the uniform tax credit (subsidy) for the clean energy technology investment cost, τ^ℓ is the exogenous labor income tax rate, and T_t is the lump-sum transfer. The borrowing limit $\underline{a} \leq 0$ is exogenous and the same for all households.

The discrete technology adoption choice is the main extension of the model from the standard Aiyagari model. The adoption decision, S_t , is a binary choice that is made at the beginning of each period. The adoption decision is irreversible and households who adopt the clean technology in period t will use the clean technology in all future periods. This irreversible adoption decision can be expressed as the difference between the household's next period utilization status of the clean technology and the current period's utilization status of the clean technology:

$$S_t = s_{t+1} - s_t, \text{ with } s_{t+1} \geq s_t \text{ for all } t.$$

The assumption of irreversibility is a reasonable first step, but future work could consider extensions such as depreciation and replacement of equipment.

3.2 Producers

There is a continuum of competitive firms (with measure normalized to one) producing the consumption good using capital and labor. The production function is:

$$Y_t = F(K_t, L_t),$$

where Y_t is the final output, K_t and L_t are capital and labor demands, respectively, and $F(\cdot, \cdot)$ is a constant returns to scale production function with inputs K_t and L_t . The representative firm takes the factor prices r_t and w_t as given and maximizes its profits each period.

3.3 Government

The government collects tax revenue by taxing every household's labor earnings at rate τ_t^ℓ , uses a portion of its tax revenue to finance the subsidies outlayed for technology adoption investment at rate τ_t , and returns the remainder of its revenue as a uniform lump-sum transfer T_t to each household in each period t . Thus, the government budget constraint is:

$$\int_{\mathcal{Z}} \tau_t^\ell \ell_t d\Phi_t = \int_{\mathcal{Z}} (T_t + \tau_t p_t S_t) d\Phi_t, \quad \forall t. \quad (5)$$

3.4 Ambient Air Pollution

The ambient air pollution each period, X_t , is determined by the flow of energy consumption during the period with the following mapping:

$$X_t = \Omega \left(\int_{\mathcal{Z}} e(c_t)(1 - s_t) d\Phi_t \right), \quad (6)$$

where $\Omega(\cdot)$ is an increasing function that maps the total period t energy consumption of old energy technology users to ambient air pollution. In summary, the air pollution is a function of the total energy consumption of the population that has not adopted low-carbon energy technology, and dirty energy consumption only creates a negative externality within the same period of its generation.

3.5 Learning-by-Doing Spillover

The one-time investment cost of the clean energy technology, p_t , decreases as the cumulative adoption of the clean energy technology increases due to learning-by-doing spillover effects. As I introduced in subsection 2.3, the learning-by-doing spillover effects can be modeled using a power-law of learning function. Specifically, I assume that the adoption cost is a function of the cumulative adoption of the clean energy technology in the economy before period t , denoted by Z_t , defined as

$$Z_t = \int_{\mathcal{Z}} s_t d\Phi_t, \quad (7)$$

and an exogenous time decay component, as stated in equation (1). Thus, the adoption cost function is:

$$p_t = p_0 \cdot Z_t^{-\xi} \cdot \exp(-\lambda t), \quad (8)$$

where $p_0 > 0$ is the initial adoption cost when the technology is first introduced, $\xi > 0$ is the learning-by-doing parameter that captures the rate of cost reduction with each doubling of cumulative adoption, and $\lambda > 0$ is the rate of exogenous time decay that captures the cost declines unexplained by local learning-by-doing spillover effects, such as global supply chain improvements and technological change.

The learning-by-doing spillover effects create a positive externality to society, because as more households adopt the clean energy technology, the adoption cost decreases for all future adopters. Modeling the learning-by-doing spillover allows me to evaluate the progressivity of subsidies in a richer sense than just pecuniary transfers, as subsidies also accelerate the learning-by-doing process and thus, the cost declines for future adopters.

3.6 Feasibility

In equilibrium, the market clearing conditions for the capital and labor markets are:

$$\begin{aligned} K_t &= \int_{\mathcal{Z}} a_t d\Phi_t, \\ L_t &= \int_{\mathcal{Z}} \ell_t d\Phi_t \end{aligned}$$

where the left-hand side of each equation is the total demand for the factor and the right-hand side is the total supply of the factor. Denote the market-clearing quantities of aggregate capital and labor by K_t and L_t , respectively. The goods market clearing condition is:

$$\int_{\mathcal{Z}} [c_t + \bar{q}e(c_t)(1 - s_t) + \underline{q}e(c_t)s_t + p_t S_t] d\Phi_t = F(K_t, L_t) + (1 - \delta)K_t - K_{t+1},$$

where δ is the depreciation rate of aggregate capital stock.

3.7 Recursive Formulation

The model can be formulated recursively. A household's consumption, saving, and adoption decision is governed by its three individual state variables, z_t . In addition, there is one aggregate state variable, the distribution of households across the individual state variables, Φ_t . The ambient air pollution index is not a state variable because it is assumed to be a flow variable and a function of only the current period's aggregate emissions from dirty energy consumption,

and thus, is a function of Φ_t alone. The definition of the initial and terminal recursive and stationary competitive equilibria are provided in the Appendix [B.1](#).

3.7.1 Transitional Dynamics

The main focus of the analysis is the transition between the initial steady state with no clean energy technology available to the terminal steady state with every household utilizing the clean energy technology. During the transition, households face the adoption choice $S_t \in \{0, 1\}$, which is irreversible, such that:

$$s_{t+1} = s_t + S_t \geq s_t.$$

The Bellman equation of a household who has not adopted the clean technology and is in state $(a_t, \ell_t, 0)$ at the beginning of period t is given by:

$$\begin{aligned} V_t(a_t, \ell_t, 0; \Phi_t) = & \\ \max & \begin{cases} \max_{c_t \geq 0} U(c_t, X_t) + \beta \mathbb{E}_t \left\{ V_{t+1}[w_t(1 - \tau^\ell)\ell_t + (1 + r_t)a_t + T_t - c_t - \bar{q}_t e(c_t), \ell_{t+1}, 0; \Phi_{t+1}] | \ell_t \right\}, \\ \max_{c_t \geq 0} U(c_t, X_t) + \beta \mathbb{E}_t \left\{ V_{t+1}[w_t(1 - \tau^\ell)\ell_t + (1 + r_t)a_t + T_t - c_t - \bar{q}_t e(c_t) - p_t(1 - \tau_t), \ell_{t+1}, 1; \Phi_{t+1}] | \ell_t \right\} \end{cases} \\ \text{subject to } & \Phi_{t+1} = \Gamma_t(\Phi_t), \end{aligned} \tag{9}$$

where $\Gamma_t : \mathcal{M} \rightarrow \mathcal{M}$ is the aggregate law of motion in period t governing the distribution of households across the state variables' tomorrow as a function of the distribution today, and $\mathbb{E}_t$ is the expectation operator conditioned on the consumer's time t information. A household in state $s_t = 0$ at the beginning of period t will choose to adopt the clean energy technology, i.e., set $S_t = 1$ and be in state $s_{t+1} = 1$ at the beginning of period $t + 1$, if the value of adopting is greater than the value of not adopting, i.e., if the second term in the maximization operator is greater than the first term. In the period that follows the decision to adopt, the household is in state $(a_t, \ell_t, 1)$ at the beginning of period t and have the Bellman equation given by:

$$\begin{aligned} V_t(a_t, \ell_t, 1; \Phi_t) = & \max_{c_t \geq 0} U(c_t, X_t) + \beta \mathbb{E}_t \left\{ V_{t+1}[w_t(1 - \tau^\ell)\ell_t + (1 + r_t)a_t + T_t - c_t - \underline{q}_t e(c_t), \ell_{t+1}, 1; \Phi_{t+1}] | \ell_t \right\}, \\ \text{subject to } & \Phi_{t+1} = \Gamma_t(\Phi_t). \end{aligned} \tag{10}$$

Definition 1 *Given an initial distribution $\Phi_0 \in \mathcal{M}$, fiscal policies $\tau^\ell, \{\tau_t\}_{t=0}^\infty$, and energy prices $\{\bar{q}_t, \underline{q}_t\}_{t=0}^\infty$, a competitive equilibrium is a sequence of: household value and policy*

functions $\{V_t, c_t, a_{t+1}, S_t, s_{t+1}\}_{t=0}^\infty$, aggregate factor stocks, $\{K_t, L_t\}_{t=0}^\infty$, prices $\{w_t, r_t, p_t\}_{t=0}^\infty$, government transfers $\{T_t\}_{t=0}^\infty$, ambient air pollution levels $\{X_t\}_{t=0}^\infty$, adoption stocks $\{Z_t\}_{t=0}^\infty$, and distributions $\{\Phi_t\}_{t=0}^\infty \subseteq \mathcal{M}$, such that for all t :

1. **Household optimization.** The household's value function V_t solves the household Bellman equation given $(w_t, r_t, p_t, \bar{q}_t, \underline{q}_t, \tau^\ell, \tau_t, T_t, X_t, \Phi_t)$, with policy functions $(c_t, a_{t+1}, s_{t+1}, S_t)$ satisfying the budget and borrowing constraints, and the adoption irreversibility constraint.

2. **Factor prices.** Factor prices are determined by the marginal products of capital and labor:

$$\begin{aligned} r_t &= F_K(K_t, L_t), \\ w_t &= F_L(K_t, L_t). \end{aligned}$$

where F_j is the first order partial derivative of F with respect to its input j .

3. **Government budget constraint.** Government budget constraint holds with equality:

$$\int_{\mathcal{Z}} \tau^\ell \ell_t d\Phi_t = T_t + \int_{\mathcal{Z}} \tau p_t S_t(a_t, \ell_t, s_t) d\Phi_t.$$

4. **Ambient air pollution.** Given the distribution of households Φ_t and the policy function $c_t(a_t, \ell_t, s_t)$, the ambient air pollution level X_t satisfies:

$$X_t = \Omega \left(\int_{\mathcal{Z}} e(c_t(a_t, \ell_t, 0)) d\Phi_t \right).$$

5. **Cumulative adoption.** Given the distribution of households Φ_t , the cumulative adoption level Z_t satisfies:

$$Z_t = \int_{\mathcal{Z}} s_t d\Phi_t.$$

6. **Adoption cost.** Given the cumulative adoption level Z_t , adoption cost satisfies:

$$p_t = p_0 \cdot Z_t^{-\xi} \cdot \exp(-\lambda t).$$

7. **Market clearing.** The factor markets clear:

$$\begin{aligned} K_{t+1} &= \int_{\mathcal{Z}} a_{t+1}(a_t, \ell_t, s_t) d\Phi_t, \\ L_t &= \int_{\mathcal{Z}} \ell_t d\Phi_t, \end{aligned}$$

and the goods market clearing condition:

$$\begin{aligned} \int_{\mathcal{Z}} [c_t(a_t, \ell_t, s_t) + a_{t+1}(a_t, \ell_t, s_t) + \bar{q}e(c_t(a_t, \ell_t, s_t))(1 - s_t) + \underline{q}e(c_t(a_t, \ell_t, s_t))s_t - p_t S_t(a_t, \ell_t, s_t)] d\Phi_t \\ = F(K_t, L_t) + (1 - \delta)K_t - K_{t+1}. \end{aligned}$$

8. **Aggregate law of motion.** The aggregate law of motion Γ_t is induced by the transition probabilities and optimal policies $a_{t+1}(a, \ell, s)$, $S_t(a, \ell, s)$, and is explicitly stated in Appendix B.3.

The model's key innovations are the irreversible binary adoption decision with subsidies and the learning-by-doing spillover effects that endogenously lower adoption costs as cumulative adoption rises. Together, these features allow me to study the welfare implications of the clean energy transition and the progressivity of subsidies in a richer sense than pecuniary transfers alone: subsidies not only redistribute resources, but also accelerate learning-by-doing and thereby reduce costs for future adopters. Furthermore, my model allows for heterogeneous impacts of pollution on households with different income levels, which is another feature to study the progressivity of subsidies. The next section describes the calibration of the model for the quantitative analysis. Standard parameters are taken from existing literature, while the parameters associated with the novel adoption and learning-by-doing mechanisms are calibrated using microdata and reduced-form estimates.

4 Quantitative Analysis

Having laid out the structure of the model, I now turn to its quantitative implementation. To evaluate the distributional and welfare effects of clean energy subsidies, the model must be calibrated to match salient features of the US economy and the residential energy sector. The quantitative analysis proceeds in two steps. First, I describe the parameterization of functional forms and the calibration of the model parameters, distinguishing between those taken from the macroeconomics literature, those pinned down by empirical moments from household- and installation-level data, and those estimated in my own empirical analysis

(such as the learning-by-doing elasticity). Second, I outline the computational methods used to solve the model, in the initial and terminal steady states, and during the transition between them under alternative policy scenarios. I defer the discussion of the pollution preference block and its calibration to Section 6.2, where I revisit the baseline results with pollution preferences activated.

I make functional assumptions for the household's utility function, the final goods production function, the dirty and clean energy production functions, the pollution function, and the pollution damage function. I assume that the household's preferences are represented by a constant relative risk aversion (CRRA) utility function of the form:

$$u(c, X) = \frac{c^{1-\sigma} - 1}{1 - \sigma} - \nu \frac{\max\{0, X - \bar{X}\}}{(c/\bar{c})^\omega}, \quad (11)$$

where $\sigma > 0$ is the coefficient of relative risk aversion for consumption, $\nu > 0$ scales ambient pollution to utility units, $\bar{X}$ is the pollution threshold above which pollution starts to cause utility losses, $\omega > 0$ makes damages to amplify at lower consumption levels, and $\bar{c}$ is a reference consumption level used to normalize the pollution damage term. Even though the disutility of pollution is separable from the utility of consumption, the pollution damage term is nonseparable in consumption and pollution, as pollution damages are larger when consumption is lower. Importantly, marginal utility of consumption remains positive, $\frac{\partial u}{\partial c} > 0$, and diminishing, $\frac{\partial^2 u}{\partial c^2} < 0$, for all $c > 0$ and $X \geq 0$.

Energy consumption function is an affine function of goods consumption with a nonzero intercept term, with parameters η_0 and η_1 :

$$e(c) = \eta_0 + \eta_1 c,$$

where $\eta_0 > 0$ captures the baseline energy consumption that is independent of goods consumption level, such as energy used for grid connection, while η_1 captures the marginal energy consumption associated with additional final goods consumption. The goods production function is of Cobb-Douglas form:

$$F(K, L) = AK^\alpha L^{1-\alpha},$$

where α is the output share of capital, and A is total factor productivity.

The ambient air pollution function is assumed to be linear in aggregate dirty energy

consumption:

$$\begin{aligned} X &= \Omega \left(\int_{\mathcal{Z}} e(c_t)(1 - s_t) d\Phi \right) \\ &= \gamma \left(\int_{\mathcal{Z}} e(c_t)(1 - s_t) d\Phi \right), \end{aligned}$$

where γ is the pollution intensity parameter of dirty energy consumption.

4.1 Calibration

The parameters disciplining the model to match the US economy can be divided into three categories: (i) standard macroeconomic parameters; (ii) parameters pinned down to match moments from data; and, (iii) estimated parameters. One period in the model corresponds to one year, and the initial steady state corresponds to the US economy in 2000, before the widespread adoption of residential solar panels. I summarize the baseline calibration of the structural parameters in Table 4 and provide details below.

4.1.1 Calibration for the Baseline Economy

In the baseline quantitative analysis, I shut down the pollution preference block in utility. Household utility is $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$ and ambient pollution X does not affect utility in the baseline. This choice isolates the adoption, affordability, and general equilibrium channels of the residential solar transition.

Standard Macroeconomic Parameters

I begin by assigning values to the set of standard macroeconomic parameters that are commonly used in the heterogeneous agent macroeconomics literature. I follow Aiyagari (1994) and set the output share of capital, α , to 0.36, the discount factor, β , to 0.96, and the capital depreciation rate, δ , to 0.08. I set the total factor productivity of goods production, A , to 1. I assume that the idiosyncratic labor endowment process, ℓ_t , follows a persistent autoregressive process with a persistence parameter of ρ and a standard deviation of σ_ε :

$$\log(\ell_t) = \rho \log(\ell_{t-1}) + \sigma_\varepsilon \sqrt{1 - \rho^2} \varepsilon_t,$$

where $\varepsilon_t \sim \mathcal{N}(0, 1)$. I discretize this earnings process using the Tauchen method. I parameterize this AR(1) labor productivity process with persistence parameter $\rho = 0.9$ and innovation standard deviation $\sigma_\varepsilon = 0.04$ as in one of the parameter combinations considered by Aiyagari (1994). Preferences over consumption are represented by a constant relative risk aversion

Table 4: Calibration summary and data sources

Parameter	Description	Value	Source
Standard macroeconomic parameters			
α	Output share of capital	0.36	Aiyagari (1994)
β	Discount factor	0.96	Aiyagari (1994)
δ	Capital depreciation rate	0.08	Aiyagari (1994)
σ	CRRA parameter (goods consumption)	1	Aiyagari (1994)
ρ	Persistence of labor productivity process	0.9	Aiyagari (1994)
σ_ε	Std. dev. of labor productivity shocks	0.04	Aiyagari (1994)
$\underline{a}$	Borrowing limit	-0.5	PSID (2000)
A	Total factor productivity	1	normalization
Moment matching parameters			
$\bar{q}$	Unit energy price without solar panels	0.04	RECS (2023)
$\underline{q}/\bar{q}$	Unit energy price with solar panels	0.5	RECS (2023)
θ	Fraction of households that can adopt	0.001 $\rightarrow$ 0.15 (linear)	RECS (2023)
τ^ℓ	Uniform labor income tax rate	0.1953	IRS (2000)
τ	Subsidy rate for solar adoption cost	0.3	Lane (2025)
$\pi_{p/y}$	Initial investment cost to income ratio, 2000	0.7	NREL (2023), Census (2024)
Estimated parameters			
ξ	Learning-by-doing (LBD) elasticity	0.1027	NREL (2023), DSIRE (2025)
λ	Exogenous cost decay parameter	0.0213	NREL (2023), DSIRE (2025)
η_0	Energy consumption function constant	0.87	BLS (2024), RECS (2023)
η_1	Energy consumption function slope	0.74	BLS (2024), RECS (2023)
Pollution block (extension)			
$\bar{X}$	PM _{2.5} threshold	9	EPA (2024b)
$\bar{c}$	Median baseline consumption	1.4590	Model moment
ν	Pollution utility scale (MWTP match)	0.0471	Vodonos, Awad, and Schwartz (2018), EPA (2024c), CDC (2023), BLS (2024)
ω	Inequality lever	2.838781	Dennin et al. (2024), BLS (2024), Census (2024), Census (2025)
γ	PM _{2.5} intensity of dirty energy consumption	6.8163	EPA (2024a)

(CRRA) utility function with a coefficient of relative risk aversion, σ , equal to 1, corresponding to log-utility.

Agents face a non-state-contingent borrowing constraint $a' \geq a$. The theoretical benchmark is the natural borrowing limit (NBL), defined as the present value of the lowest realizable future labor income stream under the no default (solvency) condition; see, Aiyagari (1994).¹ Under my baseline parameters this implies $\underline{a}^{\text{NBL}} \approx -5$ in consumption units. Using the NBL as the operative constraint, however, leads to unstable dynamics and excessive borrowing on a finite grid in this environment (large mass at the constraint and slow convergence along transitions). Following standard practice in incomplete-markets models, I therefore adopt a tighter, empirically plausible limit of $\underline{a} = -0.5$, which corresponds to an average debt-to-income ratio of approximately 31% in the initial steady state, consistent with the

¹In a stationary environment with net interest rate r and a lower bound on labor income $\underline{\ell}$, the natural borrowing limit is given by $\underline{a}^{\text{NBL}} = -\sum_{s=1}^{\infty} \frac{w \underline{\ell}}{(1+r)^s}$, so that $\underline{a}^{\text{NBL}}$ is the present value of the minimum feasible earnings path.

average US household debt-to-income ratio of 30% in 2000 University of Michigan’s (2000) Panel Study of Income Dynamics (PSID).

Moment Matching Parameters

The second set of parameters is disciplined by empirical moments from household energy consumption, residential solar panel installation data, and macroeconomic aggregates. Energy prices are pinned down to match the average unit electricity prices reported in the RECS (2023). In particular, the unit electricity price without on-site solar power generation is normalized to $\bar{q} = 0.04$ (in 2020 US dollars per BTU), while households without on-site solar generation face an effective price that is 50% lower, $\underline{q} = 0.5\bar{q}$, reflecting the reduction in electricity consumption from the grid due to on-site solar generation as described in Appendix Table 12. I calibrate the fraction of households that can adopt the clean technology, θ , to 0.5% in the first transition period and linearly increase it to 15% by the terminal steady state, in order to match the observed slow diffusion path of residential solar panels in the US between 2000 and 2020 observed in the RECS (2023).

The uniform labor income tax rate τ^ℓ is set to 19.53%, which corresponds to the revenue-neutral flat tax rate that raises the same total tax revenue as the US federal progressive income tax schedule in 2000, when applied to the model’s steady-state income distribution. The progressive schedule is taken from the Internal Revenue Service’s (IRS) published 2000 tax brackets and rates.² Differences between empirical effective tax rates reflect differences in tax base definitions and sample coverage between the model and the data. For reference, the Congressional Budget Office reports an average effective federal labor income tax rate of 12% in 2000; the gap with the model’s 19.53% reflects differences in the tax base.³ For comparison, I also specify a progressive labor income tax schedule, denoted $\tau^\ell(y)$, that reproduces the US federal marginal tax brackets in 2000 described in Appendix Table 13: This progressive tax specification is used in policy experiments to evaluate the implications of progressivity in the labor income tax system for adoption incentives and welfare. The subsidy rate for residential solar panel adoption cost, τ , is set to 30%, reflecting the average federal investment tax credit (ITC) rate for residential solar installations in the US between 2006 and present as summarized in the Lane (2025). Following the 2025 One Big Beautiful Bill, the ITC rate is set to be eliminated at the end of 2025.

To calibrate the initial investment cost of solar panels, p_0 , I use the ratio of median gross system price of residential solar panel systems to median annual household income in 2000 and set p_0 such that the model reproduces this ratio internally. I obtain the median gross

²<https://www.irs.gov/pub/irs-prior/i1040tt--2000.pdf>

³<https://www.cbo.gov/sites/default/files/108th-congress-2003-2004/reports/08-29-2003AvergeTaxRates.pdf>

system price of residential solar panel system installations in 2000 from NREL (2023), and the median annual household income in 2000 from the Census (2024). The time series of this ratio between 2000 and 2022 is plotted in Figure 2. I choose initial ratio in 2000 to equal to the largest value in the sample period I denote this ratio by $\pi_{p/y}$, and choose it to equal to the largest value in the sample period, which is approximately 70 %, and set the initial investment cost of solar panels as:

$$p_0 = \pi_{p/y} \times \bar{y}_0,$$

where $\bar{y}_0$ is the median annual household income in the initial steady state of the model. Calibrating p_0 in this way ensures that the model matches how adoption decisions are often framed in practice, i.e., in terms of the affordability of the upfront investment cost relative to annual income. These parameters ensure that the model reproduces both the affordability and heterogeneity of residential solar panel adoption in the US.

Estimated Parameters


The final set of baseline model parameters is estimated using reduced-form evidence developed in section 2.3 and moments from household energy consumption data. The key parameter of interest is the learning-by-doing elasticity, ξ , which governs how cumulative adoption reduces subsequent installation costs. The elasticity is estimated using state-level panel regressions of residential solar installation prices net of subsidies on local cumulative installed capacity, drawing on data from NREL (2023) and DSIRE (2025) policy database. In addition, I calculate an exogenous time decay parameter, λ , to capture the declines in average costs unrelated to local learning spillovers, such as global technology improvements, economies of scale, and supply chain optimizations. Both ξ and λ are central to quantifying the dynamic effects of subsidies and adoption spillovers in the model.

My baseline learning elasticity ξ estimate comes from the state-level IV regression results reported in column 2 of Table 3, which aligns with the model’s national cost curve while mitigating policy and soft-cost endogeneity concerns. The implied elasticity of 0.1027 implies a learning rate $1 - 2^{-\xi} = 0.0687$, meaning that each doubling of cumulative installed capacity leads to an 6.9% reduction in installation costs. Table 3 also provides results from regression specifications without additional controls and firm-level estimates as robustness checks.

To discipline the residual cost trend, I estimate λ from the full time series rather than only two endpoints. Specifically, I use the transformed learning-curve equation:

$$\log(p_m) + \xi \log(Z_{m-12}) = a - \lambda_m m + \epsilon_m,$$

Figure 3: Fitted cost curve from estimated learning-by-doing parameters



where m is months since the initial sample period and Z_{m-12} is the one-year lagged cumulative installed capacity in month m . The capacity-weighted pre-incentive national monthly price series p_m is constructed using the installation-level data from NREL (2023) as follows:

$$p_m^{\text{gross}} = \frac{\sum_{i \in I_t} ppw_{i,m}^{\text{gross}} \cdot w_{i,m}}{\sum_{i \in I_m} w_{i,m}},$$

where I_m is the set of all residential solar panel installations in month m , $ppw_{i,m}^{\text{gross}}$ is the gross price per watt of installation i in month m before incentives, and $w_{i,m}$ is the size of installation i in watts. The prices are adjusted for inflation using the price deflator.

Together with the estimated elasticity ξ , I run an OLS regression of the left-hand side on a constant and a linear time trend to estimate λ . Specifically, I back out λ_m rearranging the estimated learning-curve equation to obtain:

$$\lambda_m = \frac{\log\left(\frac{p_{m_a}}{p_{m_b}}\right) - \xi \log\left(\frac{Z_{m_b-12}}{Z_{m_a-12}}\right)}{m_b - m_a},$$

where m_a and m_b are the first and last months in the sample period, respectively and p_{m_a} and p_{m_b} are the corresponding capacity-weighted average national gross prices of residential solar panel systems in those months. I convert the estimated monthly λ_m to an annual λ by multiplying by 12 and obtain $\lambda = 0.0213$, indicating an average annual cost increase of 2.5% unexplained by local learning-by-doing spillover effects. Figure 3 plots the fitted cost curve implied by the estimated learning-by-doing elasticity $\xi = 0.1027$ and exogenous time decay parameter $\lambda = 0.0213$ against the actual average gross prices of residential solar panel systems between 2000 and 2022.

The per household adoption cost follows a multiplicative index that combines learning-by-

doing and an exogenous time trend. Let $Z^{\text{pre}} > 0$ denote the pre-adoption experience shifter. I force the cumulative adoption stock $Z_t \geq 0$ to be weakly increasing over time, with $Z_0 = 0$ in the initial steady state. Adoption cost is left-continuous in the experience stock to avoid simultaneity with current installations.

Define the effective experience stock

$$\text{eff}_t \equiv Z^{\text{pre}} + Z_{t-1}, \quad \text{eff}_0 \equiv Z^{\text{pre}} + Z_0,$$

and the (unit-free) cost index

$$\text{idx}_t = \left(\frac{\text{eff}_t}{\text{eff}_0} \right)^{-\xi} \exp(-\lambda \tilde{t}), \quad \tilde{t} \equiv \min\{t, 99\}. \quad (12)$$

The cap $\tilde{t} = \min\{t, 99\}$ halts the pure time trend after 100 periods to prevent implausibly low long-run costs. Numerically, I bound idx_t away from zero by $\text{idx}_t \leftarrow \max\{\text{idx}_t, 10^{-12}\}$ to maintain positivity.

Given a baseline level p_0 , the adoption cost path is

$$p_t = p_0 \times \text{idx}_t, \quad (13)$$

so that p_t falls with accumulated experience (learning) and with the exogenous trend. Let δ^{target} denote the average annual decline in adoption costs over 2000-2020, measured from the data, and let Δ_Z be the average per period increment of the normalized cumulative adoption stock over the same window. For fixed (ξ, λ) , I pin down Z^{pre} so that the model's average log change matches δ^{target} . Under a small-step approximation with mean growth in effective experience, this yields

$$Z^{\text{pre}} = \frac{\Delta_Z}{\exp((\delta^{\text{target}} - \lambda)/\xi) - 1}, \quad (14)$$

and I set $Z^{\text{pre}} = \max\{Z^{\text{pre}}, 10^{-6}\}$ in implementation. In the baseline calibration I use $(\xi, \lambda) = (0.1027, 0.0213)$ and the data moments $(\delta^{\text{target}}, \Delta_Z) = (0.0242, 0.032)$ from 2000-2020 to compute Z^{pre} , which is then held fixed throughout the transition computations.

Finally, I parameterize the affine energy consumption function as follows:

$$e(c) = \eta_0 + \eta_1 c,$$

where $e(c)$ is a household's annual energy expenditure, c is the annual consumption expenditure, and η_0 and η_1 are parameters to be calibrated. I estimate, η_0 and η_1 , to fit the average energy expenditure shares by net worth quintiles from the 2000 PSID to be replicated in the

model's initial steady state.

Let $z \in \mathcal{Z}$ denote the household state (e.g., assets and idiosyncratic labor productivity), and let Φ be the associated invariant probability measure on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$. Let $c(z)$ be goods consumption and let $e(c(z); \eta)$ denote energy services as a function of consumption, parameterized by $\eta = (\eta_0, \eta_1)$. The energy price $\bar{q}$ is taken as exogenous and constant in the baseline.

To map model implications to income quintiles, define an income mapping $\iota : \mathcal{Z} \rightarrow \mathbb{R}_+$ and quantile cutoffs $\{\kappa_j\}_{j=0}^5$ such that

$$\Phi(\{z : \iota(z) \leq \kappa_j\}) = \frac{j}{5}, \quad j = 0, 1, \dots, 5,$$

with $\kappa_0 = -\infty$ and $\kappa_5 = \infty$. The quintile sets are then

$$Q_j := \{z \in \mathcal{Z} : \kappa_{j-1} < \iota(z) \leq \kappa_j\}, \quad j = 1, \dots, 5.$$

For each quintile j , the model-implied energy expenditure share is

$$\epsilon_{Q_j}(\eta) = \frac{\int_{Q_j} \bar{q} e(c(z); \eta) d\Phi(z)}{\int_{Q_j} [\bar{q} e(c(z); \eta) + c(z)] d\Phi(z)}. \quad (15)$$

Let $\widehat{\epsilon}_{Q_j}$ denote the empirical targets from the 2000 PSID. The calibration chooses η to minimize the weighted sum of squared deviations:

$$\eta^* \in \arg \min_{\eta \in \mathbb{R}^2} \sum_{j=1}^5 \omega_j \left(\epsilon_{Q_j}(\eta) - \widehat{\epsilon}_{Q_j} \right)^2, \quad (16)$$

where $\omega_j = 1$ by default (equal weighting); when available, I set $\omega_j = 1/\widehat{\sigma}_{Q_j}^2$ using the sampling variances from PSID (inverse-variance weighting).

On a finite grid $\{z_m\}_{m=1}^M$ with probabilities $\{\Phi_m\}_{m=1}^M$ (so that $\sum_m \Phi_m = 1$), (15) becomes

$$\epsilon_{Q_j}(\eta) = \frac{\sum_{m=1}^M \mathbf{1}\{z_m \in Q_j\} \Phi_m \bar{q} e(c(z_m); \eta)}{\sum_{m=1}^M \mathbf{1}\{z_m \in Q_j\} \Phi_m [\bar{q} e(c(z_m); \eta) + c(z_m)]}.$$

All objects are evaluated at the stationary distribution Φ used for calibration; along transitions one would replace Φ with the relevant Φ_t . The values of η that minimize (16) are $\eta_0 = 0.87$ and $\eta_1 = 0.74$.

4.1.2 Calibration for Pollution Damages

For transparency, I outline here how the pollution-damage parameters will be identified and report the external data sources. These parameters are *not* used in the baseline calibration or baseline policy experiments; they are activated only in Section 6.2.

Pollution Metric and Mapping

I measure X as population-weighted annual $\text{PM}_{2.5}$ ($\mu\text{g}/\text{m}^3$), and map dirty energy to ambient pollution by

$$X_t \equiv \gamma \cdot \left[\int_{\mathcal{Z}} e(c_t(a_t, \ell_t, s_t))(1 - s_t) d\Phi_t \right].$$

I set γ to match the observed baseline $\text{PM}_{2.5}$ in year 2000 using the baseline model's implied dirty-energy integral with pollution preferences shut down. The baseline $\text{PM}_{2.5}$, denoted by X_0 , is obtained from the EPA (2024a) and is $13.52 \mu\text{g}/\text{m}^3$ in 2000. I calibrate γ as:

$$\gamma = \frac{X_0}{\int_{\tilde{\mathcal{Z}}} e(c_0(a_0, \ell_0)) d\Phi_0},$$

where c_0 is the initial steady state consumption policy function, Φ_0 is the associated invariant distribution, and $\tilde{\mathcal{Z}}$ is the state space at the initial steady state, stated in Appendix B.1. The aggregate energy consumption at the initial steady state is 1.9835 in model consumption units, so $\gamma = 13.52/1.9835 = 6.8163$.

Damage Function in Utility

When activated, utility is

$$u(c, X) = \frac{c^{1-\sigma} - 1}{1 - \sigma} - \nu \frac{\max\{0, X - \bar{X}\}}{(c/\bar{c})^\omega},$$

where $\bar{X}$ is set to the health-based annual $\text{PM}_{2.5}$ standard, equal to $9 \mu\text{g}/\text{m}^3$, based on the EPA's (2024) National Ambient Air Quality Standards (NAAQS), and $\bar{c}$ is the median consumption in the baseline steady state.

Utility Scale of Pollution Concentration

I set the parameter ν to equalize the model generated marginal rate of substitution between consumption and pollution at the initial steady state to match an external estimate of marginal willingness to pay (MWTP) for a small reduction in $\text{PM}_{2.5}$ at the baseline pollution level. When $X > \bar{X}$, consumption dollars a representative household would give up

for a unit drop in pollution, or the MRS between consumption and pollution, $MRS_{X,c}$, is:

$$\begin{aligned} MRS_{X,c} &= \frac{\partial u(c, X)/\partial X}{\partial u(c, X)/\partial c}, \\ &= \nu \frac{(c/\bar{c})^\omega}{c^{-\sigma}}, \\ &= \nu \bar{c}^{-\omega} c^{\sigma+\omega}. \end{aligned}$$

Finally, I set ν to equalize the model-implied $MRS_{X,c}$ at the median household's consumption level, to the MWTP estimate:

$$\nu \bar{c}^{-\omega} (c_0^{\text{med}})^{\sigma+\omega} = MWTP_{X,c},$$

where c_0^{med} is the median consumption in the initial steady state, and $MWTP_{X,c}$ is the external MWTP estimate of a 1 $\mu\text{g}/\text{m}^3$ reduction in $\text{PM}_{2.5}$ at the baseline pollution level. Since baseline $\bar{c}$ equals c_0^{med} , this simplifies to:

$$\nu = \frac{MWTP_{X,c}}{(c_0^{\text{med}})^\sigma}.$$

I calculate the MWTP per 1 $\mu\text{g}/\text{m}^3$ reduction in $\text{PM}_{2.5}$ at the baseline pollution level using: (i) a long-term all cause mortality $\text{PM}_{2.5}$ concentration-response (C-R) of 7% per 10 $\mu\text{g}/\text{m}^3$ from Vodonos, Awad, and Schwartz (2018), (ii) a value of statistical life (VSL) equal to 10 million (2000 dollars) from the EPA (2024c), and (iii) the 2000 mortality rate of 845.2 deaths per 100,000 people from the CDC (2023). To calculate MWTP, I use the following formula:

$$MWTP_{X,c} = \text{VSL} \times \text{Baseline mortality rate} \times \text{C-R per } 1 \mu\text{g}/\text{m}^3,$$

which yields a MWTP of approximately \$591 (in 2000 dollars). Finally, I scale the MWTP value with the ratio between the annual consumption expenditures in 2000 dollars of the median income households from the BLS's (2024) CES, which is \$18,323, and the median income household's consumption expenditure at the initial model steady state, which is 1.4590, to obtain a MWTP value in model consumption units. The ratio of median income household's expenditure in the data to expenditure in the model is approximately $\frac{18,323}{1.4590} \approx 12,599$, so the final MWTP value in model consumption units is approximately $MWTP_{X,c} = 0.0471$. The implied ν is then $\nu = 0.0471/(1.4590)^1 = 0.0323$.

Inequality Lever

The inequality level ω amplifies pollution disutility when a household's consumption is

lower than the reference consumption level $\bar{c}$, and reduces it when consumption is higher. I set ω to match the ratio of pollution damages as a share of consumption for the bottom versus top income distribution terciles. I partition the population into three groups based on income terciles at the initial steady state without pollution damages, and compute the model-implied pollution damage mass for each group:

$$D_g(\omega) = \int_{T_g} \nu \frac{\max\{0, X_0 - \bar{X}\}}{[c_0(a, \ell)/\bar{c}]^\omega} d\Phi_0, \quad g = 1, 2, 3,$$

where T_g is the income tercile set g at the initial steady state, and c_0 and Φ_0 are the associated consumption function and invariant distribution, respectively. Similarly, I define the model-implied consumption mass for each group:

$$C_g = \int_{T_g} c_0(a, \ell) d\Phi_0, \quad g = 1, 2, 3.$$

Prior to calculating the pollution burden as a ratio of damages to consumption, I calculate damage and consumption shares by group as:

$$\psi_g^D(\omega) = \frac{D_g(\omega)}{\sum_{g=1}^3 D_g(\omega)}, \quad \psi_g^C = \frac{C_g}{\sum_{g=1}^3 C_g}, \quad g = 1, 2, 3.$$

I then compute the pollution burden ratio for the bottom versus top income terciles as:

$$\Upsilon(\omega) = \frac{\psi_1^D(\omega)/\psi_1^C}{\psi_3^D(\omega)/\psi_3^C}.$$

I use pollution and consumption shares in pollution burden calculation to abstract away from units and maintain comparability with alternative calibrations.

I set ω to match the pollution burden ratio that I calculate using estimates of US county-level mortality damages from PM_{2.5} pollution and consumption expenditure data across income terciles. To obtain the weighted average of mortality damages by income terciles, I combine county-level mortality damage estimates from Dennin et al. (2024) with county-level median income from Census (2024) and population data from Census (2025). I obtain weighted average of consumption expenditure by income terciles from the BLS's (2024) CES. I compute the empirical pollution burden ratio as follows:

$$\hat{\Upsilon} = \frac{\hat{\psi}_1^D/\hat{\psi}_1^C}{\hat{\psi}_3^D/\hat{\psi}_3^C},$$

where $\hat{\psi}_g^D$ is the share of total mortality damages borne by income tercile g from Dennin et al. (2024), and $\hat{\psi}_g^C$ is the share of total consumption expenditure by income tercile g from the 2014 CES. The empirical pollution burden ratio is approximately $\hat{\Upsilon} = 5.83$, indicating that the bottom income tercile bears significantly a larger pollution burden than the top income tercile. I then choose ω to solve the minimization problem:

$$\omega^* = \arg \min_{\omega \geq 0} (\Upsilon(\omega) - \hat{\Upsilon})^2.$$

The implied ω is $\omega = 2.838781$.

Having established the calibration of baseline and extended model parameters, the next step is to describe how the model is solved and simulated. The computational procedure involves characterizing the household decision problem under the calibrated environment, solving for the stationary equilibrium of the initial and terminal economies, and then tracing out the transition dynamics in response to policy interventions. In what follows, I outline the numerical methods used to solve the model, describe the construction of both the initial and long-run steady states, and detail how transitional paths are computed under baseline subsidy scenario.

4.2 Computation

The model is solved numerically in three stages: (i) the initial steady state without the clean technology, (ii) the terminal steady state with only the clean technology available, and (iii) the transition path connecting the two. Each stage ensures consistency between individual decisions, aggregate quantities, and market clearing conditions.

In the initial steady state, households do not have access to the clean technology, and the economy settles into a stationary equilibrium given exogenous energy prices and fiscal policies. In the terminal steady state, all households are equipped with the clean energy technology, and the economy again reaches a stationary equilibrium under the new energy price regime.

The transition path is computed under perfect foresight. Given an initial guess for the sequences of aggregates, the model is solved by backward induction on household value functions and forward iteration on the distribution of households. Paths of aggregate capital stock, lump-sum transfers, and the cumulative stock of adopters are updated iteratively until factor and goods markets clear at each point along the path. The equilibrium path thus describes the joint evolution of prices, adoption, and welfare as the economy transitions to the new steady state. Further computational details—including the recursive formulation, iteration schemes, and convergence criteria—are provided in Appendix C.2. I report the key steady-state moments and transition dynamics in Section 5.

The computation of the extended model with pollution preferences follows the same steps as above, with the addition of the pollution preference block in utility and the pollution mapping. As mentioned in section 3, incorporating pollution preferences does not alter computation significantly, as pollution is a deterministic function of aggregate dirty energy consumption. Thus, the number of state variables remains unchanged, and the household problem retains its recursive structure. The main difference is that the household value functions and policy functions now depend on the pollution level, which in turn depends on the aggregate dirty energy consumption. This adds an additional layer of general equilibrium feedback, as households' adoption decisions affect pollution, which affects utility, which in turn affects adoption incentives. The computational algorithm is adjusted to account for this feedback loop, ensuring that the pollution level is consistent with the aggregate dirty energy consumption at each point in time. To implement this extension, I add pollution level as a fourth variable to guess and update along the transition path, in addition to aggregate capital stock, lump-sum transfers, and cumulative adopters. In Section 6.2, I report the results of the extended model, re-solve the steady states and transition paths, and revisit policy experiments.

Since the objective of this quantitative analysis is to understand the distributional and welfare implications of the residential solar transition, I compute a range of household-level welfare measures along the transition path. These include the consumption equivalent variation (EV), which measures the percentage change in initial consumption that would make a household indifferent between the baseline and counterfactual scenarios, consumption compensating variation (CV), which measures the percentage change in final consumption that would make a household indifferent between the baseline and counterfactual scenarios, and the lifetime utility change, which measures the absolute change in lifetime utility from the baseline to counterfactual scenarios. These welfare measures are computed for each household along the transition path, allowing for a detailed analysis of how different households are affected by the transition and the associated policies. The computation of these welfare measures is described in detail in Appendix C.3.

The preceding section laid out the calibration strategy and computational methods necessary to map the model to the data. By aligning key parameters with observed macroeconomic aggregates, energy expenditure statistics, and installation cost dynamics, the calibration ensures that the model reflects both the technological and distributional features central to the adoption decision. The computational approach, in turn, allows for the characterization of households' dynamic optimization under incomplete markets, as well as the general equilibrium interactions that shape prices and policy incidence. With this foundation in place, the analysis is now equipped to assess the model's performance relative to the data and to

examine the distributional and welfare implications of policy interventions.

5 Quantitative Results

This section presents the quantitative results of the model. I begin by evaluating the fit of the initial stationary equilibrium, focusing on how well the model reproduces salient features of the joint income-wealth distribution observed in the data. This validation step establishes the credibility of the model as a tool for policy analysis. I then turn to the main policy experiment: the introduction of an adoption subsidy for clean energy technologies. The results highlight how the subsidy reshapes adoption incentives across the income and wealth distribution and trace the resulting welfare differences among heterogeneous households. Together, these findings provide a benchmark for understanding both the equity and efficiency consequences of adoption subsidies in the transition to clean energy even before considering the role of pollution externalities.

5.1 Model Fit

Before analyzing the welfare effects of adoption subsidies, I first evaluate how well the initial stationary equilibrium of the model matches key features of the wealth distribution in the 2000 cross-section of the PSID data. Table 5 compares selected variables across the net worth quintiles in the data and the initial stationary equilibrium of the model. The variables include the share of total income, expenditure, and wealth held by each net worth quintile as well as the ratio of energy expenditure to total expenditure.

The model captures the overall distribution of income, expenditure, and wealth reasonably well. The model slightly overpredicts the share of income and expenditure held by the lower net worth quintiles and underpredicts the share held by the highest net worth quintile, but the discrepancies are relatively small. The model also matches the wealth distribution well, although it underpredicts the wealth share of the highest net worth quintile. This discrepancy is a common challenge in heterogeneous-agent models that could be addressed by introducing additional features such as persistent discount factor heterogeneity across households. Finally, the model closely replicates the ratio of energy expenditure to total expenditure across the net worth distribution, indicating that it captures the consumption patterns of households well.

Table 5: Selected variables across the net worth quintiles from the data vs the initial model stationary equilibrium

NW Quintile	% Share of						% Ratio	
	Income		Expenditure		Wealth		Energy-Expend	
	Data	Model	Data	Model	Data	Model	Data	Model
Q1	10.9	13.3	12.6	13.8	-0.9	0.3	5.9	6.2
Q2	12.2	18.2	15.4	18.4	1.2	6.8	6.0	5.6
Q3	17.8	17.4	19.2	17.4	5.2	12.6	5.8	5.1
Q4	23.1	22.6	23.2	22.4	14.8	25.6	5.1	4.8
Q5	35.9	28.5	29.6	28.1	79.6	54.7	4.5	4.4

Notes: Data is from the 2000 PSID.

5.2 Baseline Results

In order to assess the welfare implications of the energy transition and the role of adoption subsidies during the transition, I compute the welfare changes for each household in the economy. I provide the details of the various welfare metrics used and their calculations in Appendix C.3. The baseline results focus on the case without pollution in the utility function to isolate the macroeconomic and adoption mechanisms that drive welfare changes. I discuss the role of general equilibrium effects and learning-by-doing in shaping the welfare outcomes. Finally, I explore how the introduction of an adoption subsidy affects welfare across different income and wealth groups.

Aggregate Dynamics

First, I describe the aggregate dynamics of the economy in response to the introduction of the clean energy technology.

Figure 4 presents the aggregate dynamics of the economy under the baseline uniform labor income tax policy with LBD in adoption cost, and with and without adoption subsidies. The panels show the paths of: (a) the cumulative stock of adopters, (b) technology cost to median income ratio, (c) capital-labor ratio, and (d) lump-sum transfers with subsidies (blue lines) and without subsidies (orange lines) when LBD is active.

In the long-run, the aggregate capital-labor ratio decreases by 12%, which is entirely driven by the reduction in energy use, in consumption units, due to improved energy efficiency. This decline in the capital-labor ratio corresponds to a 4% decline in the real wage and a 24% increase in the real interest rate. Whereas, the 50% decline in energy price leads to 4.9% increase in good consumption of the poorest households and 1.6% increase in good consumption of the richest households. In the long-run, although the aggregate capital, and thus output decreases, aggregate consumption increases because of the large decrease in

energy price.


Along the transition path, the paths of the aggregate variables are not monotonic and exhibit significant differences when subsidies are introduced. In the short-run (the first 5-10 years of transition path), without subsidies, the aggregate capital-labor ratio increases, driven by the saving motive of households to afford the adoption cost of the new technology. This increase in the capital-labor ratio leads to an increase in the real wage over the short-run, which benefits all households, but especially the low-wealth households who rely more on labor income. The real interest rate decreases in the short-run, which hurts savers, especially high-wealth households. As the adoption of the new technology increases over time, the aggregate capital-labor ratio starts to decrease, leading to a decrease in the real wage and an increase in the real interest rate. The immediate short-run effect on capital-labor ratio reverses when subsidies are introduced, as subsidies lower the adoption cost and allow households to adopt earlier without having to save as much.

The aggregate path of adoption has an S-shape, which is a common feature in the technology adoption literature. The adoption starts slowly, as only high-wealth households can afford to make the expensive adoption investment. As the adoption cost decreases due to LBD, more households can afford the adoption, leading to a rapid increase in adoption. Finally, as adoption saturates, the adoption rate slows down.

In the short run without subsidies, when adoption is low and the capital-labor ratio is high, the lump-sum transfers to households increase, which benefits all households, but especially low-wealth households who have higher marginal utility of consumption. As adoption increases and adoption cost declines, the capital-labor ratio decreases and the lump-sum transfers decrease, which hurts all households, but especially low-consumption households. However, when subsidies are introduced, the aggregate capital-labor ratio increases more and for longer in the short run (15 years) as early adopters spend less on adoption and can save up more. Although this larger increase in the aggregate capital-labor ratio results in a larger increase in wages, which benefits all more, and a fall in interest rate, which hurts the wealthy and benefits borrowers. The larger increase in the capital-labor ratio and thus wages increase tax revenues, but financing of the subsidies greatly lower the net of subsidy lump-sum transfers that the households receive. The effect of subsidies on lump-sum transfers is very large and hurts everyone, especially poorer households, because their marginal utility of consumption is higher. Overall, the transition path is characterized by non-monotonic dynamics in the aggregate variables and are affected by subsidies, which have important implications for the welfare of heterogeneous households.

Aggregate Welfare Effects


Figure 4: Aggregate dynamics under baseline policy with LBD, with and without adoption subsidies.



First, I analyze the overall welfare consequences of subsidizing the clean energy transition. The main welfare change metric I consider is the equivalent variation (EV). I describe the definition and calculation of EV in Appendix C.3. The decomposition of welfare changes by asset and income groups in Figure 5 reveals that on average almost every group benefits from adoption subsidies but welfare gains increase as income and wealth increase. Only the lowest income and wealth group experience a welfare loss. The welfare changes are most pronounced for middle income group.

Table 6 presents a detailed decomposition of the welfare changes induced by adoption subsidies, breaking down the contributions from direct subsidies, learning-by-doing (LBD)-induced cost reductions, price effects, and fiscal transfer effects. Before I discuss the welfare decomposition, I first highlight the total welfare effects and their distributional patterns

Figure 5: Welfare impact of subsidizing solar panel adoption cost across the income-wealth distribution



across the wealth terciles.

Panel A of Table 6 summarizes various welfare metrics for each welfare changing component of the subsidy. These welfare change metrics include the average expected lifetime utility change (Avg. ΔV), the average EV (Avg. EV), and the average compensating variation (Avg. CV). I provide the definitions of these welfare metrics in Appendix C.3. The last row of Panel A of Table 6 summarizes the total averages of these three welfare change measures across all households. On average, I find that the subsidy leads to a small welfare gain across all households, with an average of 1.29×10^{-2} utils in lifetime utility, an average CV of -0.05% , and an average EV of 0.05% .

I define a household as strictly benefiting, i.e., a winner from the subsidy, if its EV is strictly positive. The final column of Panel A of Table 6 reports the share of all households that benefit from each welfare change induced by the subsidy. Overall, I find that a strong majority of 93.9% of households experience net welfare gains from subsidizing the transition, while the remaining 6.1% experience losses or are indifferent relative to the baseline. However, the distribution of winners is unequal.

Panel B of Table 6 breaks down the share of winners by asset terciles. The bottom row of Panel B shows the within asset tercile winner shares. The results reveal unequal support. Although the majority of each wealth tercile wins from the subsidies, the strict support decreases with wealth. The within-tercile shares of strict winners are 84.8%, 98.1%, and

Table 6: Decomposition of welfare effects and distribution of winners by asset tercile

Panel A. Aggregate Welfare Decomposition				
Component	Avg. ΔV	Avg. CV (%)	Avg. EV (%)	Winner (%)
Direct subsidy	$+5.89 \times 10^{-2}$	-0.24	+0.24	100.0
LBD-induced cost change	$+2.43 \times 10^{-4}$	-0.00	+0.00	99.5
Price effect	-4.95×10^{-3}	+0.02	-0.02	39.3
Transfer effect	-4.13×10^{-2}	+0.17	-0.17	0.0
Total	$+1.29 \times 10^{-2}$	-0.05	+0.05	93.9

Panel B. Winner Shares by Asset Tercile (%)				
	Bottom	Middle	Top	Aggregate
Direct subsidy	100.0	100.0	100.0	100.0
LBD-induced cost change	100.0	100.0	98.3	99.5
Price effect	92.9	18.2	0.0	39.3
Transfer effect	0.0	0.0	0.0	0.0
Total	84.8	98.1	99.9	93.9

Notes: Winner share is the fraction of households with strictly positive consumption-equivalent variation (EV). Tercile population masses are [0.360, 0.323, 0.317].

99.9%, for bottom, middle, and top wealth terciles respectively. This pattern indicates that although the subsidies enhance aggregate welfare, low-wealth households are not unanimously benefiting and may be disproportionately burdened by the costs of financing the subsidy, while high-wealth households are strictly better off.

Welfare Decomposition

To understand the distributional mechanisms underlying these aggregate welfare effects, I decompose the total welfare change into four components: (1) the direct subsidy effect, which captures the immediate benefit to adopters from the subsidy; (2) the LBD-induced cost change effect, which reflects how the subsidy accelerates adoption and thereby reduces future technology costs for all households; (3) the price effect, which accounts for changes in equilibrium prices (wages, interest rates, energy prices) induced by the subsidy; and (4) the transfer effect, which captures changes in fiscal transfers due to altered government budget constraints. The first two components represent the direct benefits of the subsidy, while the latter two capture general equilibrium feedback effects. Table 6 presents the decomposition by these four components for the aggregate economy in Panel A and by asset terciles in Panel B.

The direct subsidy and LBD-induced cost change effects are unambiguously positive for all households, as they directly lower the cost of adoption, either directly through the subsidy or indirectly via accumulated learning that reduces future costs. On average, these two channels raise household welfare by 5.89×10^{-2} and 2.43×10^{-4} utils, respectively, corresponding to EV

gains of 0.24% and 0.004%. Note that the average welfare gain from the LBD-induced cost change is relatively small. Panel B shows that these benefits are uniformly distributed across the asset distribution, with all terciles weakly benefiting from these two channels.⁴

The price effect, which reflects changes in equilibrium wages and interest rates, introduces some heterogeneity in welfare outcomes. On average, the price effect yields a welfare loss of 4.95×10^{-3} utils, corresponding to an EV loss of 0.02%, and benefits strictly benefits 39.3% of households. However, Panel B reveals that the price effect disproportionately benefits households in the bottom asset tercile, with winner shares of 92.9%, 18.2%, and 0.0% for the bottom, middle, and top asset terciles, respectively. This pattern arises because wealthier households rely more on asset income and are hurt by the fall in interest rates. In the short run, Figure 4 showed that without subsidies, the capital-labor ratio increases, leading to an increase in wages and a decrease in interest rates, while the short run effect on capital-labor ratio is more amplified and persistent when subsidies are present. As a result, low-wealth households, who depend more on labor income, benefit from the price effect, while high-wealth households, who derive more income from assets, are more likely to gain.

Finally, the transfer effect is unambiguously negative for all households, as the subsidy decreases the lump-sum transfers that households receive, even though tax revenues increase as wages increase. On average, the transfer effect reduces household welfare by -4.13×10^{-2} utils, corresponding to an EV change of -0.17% . No households benefit from this channel, as indicated by the 0% winner share. The unequal impact of the transfer effect is seen when looking at the overall welfare change support across asset terciles in the final row of Panel B. Although low-asset households benefit more from the price effect, they are also more adversely affected by the transfer effect, as they rely more on lump-sum transfers for their consumption as implied by their lower overall subsidy support.

Aggregating across all four components, the total welfare effect of the subsidy is positive on average. These results highlight that while the learning mechanism generates some efficiency gains, the welfare change due to the direct subsidy effect is the main driver of the positive average welfare outcome compensating for the strong negative transfer effect. However, as highlighted earlier, the transfer effect disproportionately burdens low-wealth households, leading to lower support for the subsidy among this group.

In summary, the decomposition reveals that the subsidy generates direct benefits for the majority through reduced adoption costs and learning effects, but these benefits are subdued by the negative transfer effects for everyone and negative price effects for the asset-rich.

⁴Note that the top asset tercile has a 98.3% winner share from LBD-induced cost change. The winner shares describe the fraction of households who strictly benefits (i.e., have a strictly positive welfare change). The remaining 1.7% top asset tercile households are indifferent to the policy, meaning their welfare is unchanged.

The net result is although the subsidy has majority support, it overburdens low-wealth households, potentially leading to a regressive distribution of welfare outcomes, with low-wealth households are burdened more by the negative transfer effect. Importantly, these results rely on a simplification of the model that abstracts from the progressivity of labor income taxation, nonrefundability of subsidies, and pollution externalities. In the next section, I explore how changes in these assumptions affect the welfare outcomes of the subsidy.

6 Sensitivity Analysis

The baseline results discussed in the previous section evaluated the welfare changes from subsidizing the adoption of clean energy technologies under a simplified policy design with a uniform labor income tax financing mechanism and a permanent, refundable adoption subsidy. However, real-world policies are often more complex and multifaceted. The US income tax system is progressive, and the federal investment tax credit for residential solar has been nonrefundable and temporary. Moreover, the previous analysis abstracted from pollution damages, which are a key motivation for environmental policies. This section extends my quantitative analysis to address these limitations. First, section 6.1 explores alternative financing and subsidy designs to assess how different policy instruments affect the efficiency and equity of the energy transition. Second, section 6.2 extends the analysis by activating pollution damages in the utility function, allowing me to quantify the additional welfare effects arising from reduced pollution exposure.

6.1 Alternative Financing and Subsidy Designs

To move beyond the benchmark subsidy, I evaluate how alternative policy designs affect both the efficiency and equity of the energy transition. These policies are motivated by real-world policy designs debated or implemented in the US. I consider three alternative policies, each building on the previous one, to isolate the effects of specific design features: (i) a progressive labor income tax financing mechanism, (ii) a nonrefundable adoption subsidy under progressive financing, (iii) an income-capped nonrefundable subsidy. These policies vary along two dimensions: how the policy is financed and who is eligible for the subsidy. I will decompose welfare effects and winner shares for each policy as in section 5 and evaluate which designs best balance efficiency and equity.

Experiment 1: Progressive Labor Income Tax Financing

First, I replicate the baseline results under a progressive labor income tax. Table 7

Table 7: Decomposition of experiment 1 welfare changes and distribution of winners by asset tercile

Panel A. Aggregate Welfare Decomposition				
Component	Avg. ΔV	Avg. CV (%)	Avg. EV (%)	Winner (%)
Direct subsidy	$+5.7 \times 10^{-2}$	-0.23	+0.23	100.0
LBD-induced cost change	$+5.16 \times 10^{-4}$	-0.00	+0.00	99.8
Price effect	-3.9×10^{-3}	+0.02	-0.02	32.3
Transfer effect	-4.21×10^{-2}	+0.17	-0.17	0.0
Total	$+1.15 \times 10^{-2}$	-0.05	+0.05	83.3


Panel B. Winner Shares by Asset Tercile (%)				
	Bottom	Middle	Top	Aggregate
Direct subsidy	100.0	100.0	100.0	100.0
LBD-induced cost change	100.0	100.0	99.2	99.8
Price effect	87.9	8.4	0.0	32.3
Transfer effect	0.0	0.0	0.0	0.0
Total	55.8	94.4	100.0	83.3

Notes: Winner share is the fraction of households with positive consumption-equivalent variation (EV). Tercile population masses are [0.335, 0.343, 0.322].

presents the detailed welfare decomposition for this experiment, analogous to Table 6 for the benchmark. First, note that the initial and terminal states under the progressive tax schedule differ from those under uniform financing, as households face different effective tax rates that affect their savings and adoption decisions even in the absence of subsidies. Subsidizing adoption under a progressive financing mechanism increases aggregate welfare by less than it does under uniform financing in utils; however, the average EV increase is equal at 0.05% in both cases. However, the aggregate winner share falls to 83.3% from 93.9% under uniform financing. The within-tercile shares of winners show that the decline in aggregate winner shares comes from the low-wealth households.

I argue that this decline in support among low-wealth households stems from the smaller relative increase in short-run wages. To understand the changes in welfare effects under progressive financing, first, I compare the aggregate and across-asset-tercile adoption dynamics in Figure 6 between the benchmark (left) and this experiment (right). The path of stock of adopters shows that adoption is slower under progressive financing absent subsidies. Progressive financing increases the effective tax rates for higher-income households, who are wealthier and initial adopters, thus reducing their liquidity to finance adoption and slowing down early aggregate adoption. Progressive financing decreases the effective tax rates for lower-income households, who are less likely to be initial adopters, thus the progressivity of the tax system does not increase their adoption speed.


Figure 6: Aggregate and across asset tercile adopter stock, benchmark and experiment 1.



Composition of adopters across wealth terciles shows that subsidies increase adoption speed for all wealth groups under both financing mechanisms. However, under progressive financing, adoption speed increases relatively more for middle- and high-wealth households, given that they had a lower adoption speed absent subsidies compared to uniform financing. The speed of adoption in the aggregate and across wealth groups is very similar between the two financing mechanisms once subsidies are introduced, as seen in Figure 6.

Next, I investigate how these adoption dynamics affect aggregate savings and transfers in Figure 7. Under progressive financing, subsidizing adoption does not lead to as large a short-run increase in aggregate savings as it does under uniform financing, as seen in Figure 7 (a). As seen in Figure 6, the speed and composition of adoption are similar between the two financing mechanisms once subsidies are introduced; however, under the progressive system, the early adopters have larger effective tax rates, thus they save less after adopting the technology

Figure 7: Aggregate dynamics under progressive financing policy with LBD, with and without adoption subsidies, experiment 1.



compared to uniform financing. The more muted increase in aggregate savings leads to smaller increases in the capital-labor ratio and wages in the short run under progressive financing. The smaller rise in wages creates fewer benefits for low-wealth households, who rely more on labor income, reducing their support for the subsidy. While the adoption is similar between the two financing mechanisms once subsidies are introduced, the relative increase in wages is larger under uniform financing, which implies greater tax revenues returned lump-sum to households net of subsidy payments compared to progressive financing, as seen in Figure 7 (b). Comparing Figure 4 (d) and Figure 7 (b), the net-of-subsidy lump-sum transfers fall more under progressive financing, leading to larger negative transfer effects for all households.

Comparing the welfare decompositions in Tables 6 and 7, and the aggregate dynamics in Figures 4 and 7, I infer that the source of declined support among low-wealth households under progressive financing stems from the smaller relative increase in short-run wages and larger relative decrease in lump-sum transfers. The LBD-induced cost change effect still remains insignificant under progressive financing.

Experiment 2: Nonrefundable Subsidy under Progressive Financing

In a second experiment, I explore the implications of a nonrefundable technology adoption subsidy under progressive financing, which mirrors the structure of the US federal solar investment tax credits before their recent elimination. Under a nonrefundable subsidy, low-income households with labor income below the standard deduction are ineligible for the subsidy, as subsidies are deducted from tax liabilities, and eligible households must have

enough tax liability to receive the full benefit. Compared to the baseline refundable subsidy, this nonrefundable design intends to reduce the negative transfer effects, but sharply limits access for lower-income households and could slow down adoption and its LBD-induced cost declines, further increasing regressivity concerns.

Table 8 presents the welfare decomposition of this experiment. It shows that nonrefundable subsidies still improve aggregate welfare by a similar magnitude as the previous experiment in utils and EV, at 0.05% change in EV. Direct subsidy and LBD-induced cost change effects remain positive and strictly benefit every household, but the magnitude of the average direct subsidy benefit declines slightly, and benefits from LBD remain insignificant. The average price effect remains negative, but its magnitude decreases compared to the previous experiment, while the share of price effect beneficiaries declines to 23% from 32.3% in experiment 1. Finally, the magnitude of the negative transfer effect decreases slightly, but it remains non-positive for all households.

Although the aggregate welfare change remains positive, the aggregate share and distribution of winners shift under nonrefundable subsidies. The aggregate winner share falls to 78.1% from 83.3%. The within-tercile shares of winners show that the decline in aggregate winner shares comes entirely from the low-asset households, whose within-tercile winner share falls to 44.4% from 55.8% in experiment 1. This sharp decline in support among low-wealth households indicates that making subsidies nonrefundable under progressive financing significantly reduces the equity of the policy, as low-income households are excluded from subsidy eligibility.

Figure 8 compares the adoption dynamics under refundable (experiment 1) and nonrefundable (experiment 2) subsidy regimes. Panels (a)-(b) show that aggregate adoption evolves nearly identically across the two experiments, suggesting that refundability does not affect overall diffusion. The difference lies in the timing of adoption across wealth groups. Panels (c)-(d) show that under the nonrefundable subsidy, early adoption (within 5-10 years) slows among high-wealth households and slightly accelerates among middle- and low-wealth households. Because the subsidy can only be claimed when households owe taxes and is capped by their tax liability, its benefit becomes state-contingent: High-wealth households in temporarily low-income states postpone adoption, while liquidity-constrained households adopt earlier if they draw a high-income state that makes them eligible. Although all households eventually cycle through income states due to the persistent AR(1) income process, these timing frictions temporarily alter who adopts first.

The composition of adopters matters for the dynamics of economic variables, which matter for welfare outcomes. Figure 9, panel (a) shows that when subsidies are nonrefundable, aggregate savings do not increase as much in the short run compared to cases without

Table 8: Decomposition of experiment 2 welfare changes and distribution of winners by asset tercile

Panel A. Aggregate Welfare Decomposition				
Component	Avg. ΔV	Avg. CV (%)	Avg. EV (%)	Winner (%)
Direct subsidy	$+5.24 \times 10^{-2}$	-0.21	+0.21	100.0
LBD-induced cost change	$+4.83 \times 10^{-4}$	-0.00	+0.00	100.0
Price effect	-1.08×10^{-3}	+0.00	-0.00	23.0
Transfer effect	-4×10^{-2}	+0.16	-0.16	0.0
Total	$+1.18 \times 10^{-2}$	-0.05	+0.05	78.1


Panel B. Winner Shares by Asset Tercile (%)				
	Bottom	Middle	Top	Aggregate
Direct subsidy	100.0	100.0	100.0	100.0
LBD-induced cost change	100.0	100.0	100.0	100.0
Price effect	64.8	3.7	0.0	23.0
Transfer effect	0.0	0.0	0.0	0.0
Total	44.4	90.6	99.8	78.1

Notes: Winner share is the fraction of households with strictly positive consumption-equivalent variation (EV). Tercile population masses are [0.335, 0.343, 0.322].


subsidies and refundable subsidies. This muted increase in savings leads to smaller increases in capital-labor ratio and wages in the short run. The smaller rise in wages creates fewer benefits for low-wealth households, who rely more on labor income, reducing their support for the subsidy. The savings increase by less with subsidies under a nonrefundable design because low-wealth but high-income households increased their adoption in the short run, as they may lose their subsidy eligibility when they have enough savings buffer to smooth consumption. In the medium run, the aggregate capital-labor ratio falls more under nonrefundable subsidies than under refundable subsidies. This steeper decline occurs as adoption speed increases among high- and medium-wealth households and reduces their savings to finance adoption. Panel (b) shows that lump-sum transfers net of subsidy payments fall less when subsidies are nonrefundable compared to refundable subsidies, even though a smaller rise in wages leads to smaller overall tax revenues relative to experiment 1.

Thus, overall, making subsidies nonrefundable under progressive financing reduces aggregate welfare gains only slightly compared to refundable subsidies, but it reduces support among low-wealth households significantly. This policy still generates positive welfare gains for the majority of households, but it raises equity concerns as low-income households are excluded from subsidy eligibility, and wages don't rise as much in the short run to benefit them. These two negative effects outweigh the benefit of smaller negative transfer effects under nonrefundable subsidies for low-wealth households.


Figure 8: Aggregate and across asset tercile adopter stock, experiment 1 and experiment 2.




(a) Aggregate adoption, experiment 1



(b) Aggregate adoption, experiment 2



(c) Asset-tercile adoption, experiment 1




(d) Asset-tercile adoption, experiment 2

Experiment 3: Income-Capped Subsidy under Progressive Financing

Finally, I analyze income-capped subsidies that restrict eligibility to middle- and lower-income households. I explore income-capped subsidies that are refundable but only available to households with income below a certain threshold. Income-capped subsidies aim to target adoption support toward liquidity-constrained households, while avoiding fiscal costs from subsidizing wealthier households who may adopt regardless of subsidies. Linn (2022) discusses welfare and distributional implications of income-capped subsidies in the context of electric vehicle adoption. In a model disciplined by the US new vehicles market, he finds that income-based subsidies are more effective and more equitable than uniform subsidies in theory.

Such income-capped subsidies have been briefly implemented in the US for clean vehicle

Figure 9: Aggregate dynamics under progressive financing policy with LBD, with and without nonrefundable adoption subsidies, experiment 2.



purchases between 2023 and 2025 (see US Department of Energy [2024](#) and IRS ([2025](#))). The Inflation Reduction Act (IRA) established, starting in 2023, modified adjusted gross income (AGI) limits for clean vehicle credits, with different caps for new and used electric vehicles. DOE ([2024](#)) summarizes that to qualify for a credit, a filer’s modified AGI must be below a limit based on their tax filing status, using either their income from the year they purchased the vehicle or the year prior, whichever is lower. This design targets adoption support toward liquidity-constrained households, but it may slow aggregate adoption if wealthier households are the primary drivers of learning-by-doing spillovers. Allcott et al. ([2024](#)) argues that the IRA’s income caps were generous enough that the majority of new vehicle buyers still qualified for the credit. Since the passage of the One Big Beautiful Bill, clean vehicle credits were eliminated starting on September 30, 2025, regardless of income, according to IRS ([2025](#)).

Table 9 presents the welfare decomposition of this experiment. Among all four experiments with alternative policy designs, refundable income-capped subsidy is the only design that fails to increase aggregate welfare, resulting in a large aggregate welfare loss of -0.09% in EV. Overall, only 2.2% of households strictly benefit from this policy, while the majority of households experience welfare losses. The decomposition of welfare changes shows that direct subsidy and LBD-induced cost change effects remain positive, but the relative magnitude of direct subsidy benefits declines sharply compared to previous experiments, as all households greater than or equal to the income cap, which is set to equal the within-period median income, are ineligible for the subsidy. The average price effect changes sign to positive, indicating that general equilibrium price changes create larger benefits than costs on average. Finally, the magnitude of the negative transfer effect decreases slightly compared to the previous

Table 9: Decomposition of experiment 3 welfare changes and distribution of winners by asset tercile

Panel A. Aggregate Welfare Decomposition				
Component	Avg. ΔV	Avg. CV (%)	Avg. EV (%)	Winner (%)
Direct subsidy	$+8.77 \times 10^{-3}$	-0.04	+0.04	100.0
LBD-induced cost change	$+1.06 \times 10^{-5}$	-0.00	+0.00	100.0
Price effect	$+6.1 \times 10^{-3}$	-0.02	+0.02	77.0
Transfer effect	-3.67×10^{-2}	+0.15	-0.15	0.0
Total	-2.18×10^{-2}	+0.09	-0.09	2.2

Panel B. Winner Shares by Asset Tercile (%)				
	Bottom	Middle	Top	Aggregate
Direct subsidy	100.0	100.0	100.0	100.0
LBD-induced cost change	100.0	100.0	99.9	100.0
Price effect	33.6	97.9	100.0	77.0
Transfer effect	0.0	0.0	0.0	0.0
Total	0.0	0.0	6.9	2.2


Notes: Winner share is the fraction of households with strictly positive consumption-equivalent variation (EV). Tercile population masses are [0.335, 0.343, 0.322].

experiment, but it remains non-positive for all households.

Figure 10 compares the adoption dynamics induced by refundable subsidies with (experiment 3) and without (experiment 1) income caps. Panels (a)-(b) show that aggregate adoption slows significantly under income-capped subsidies compared to non-capped subsidies. Moreover, panel (b) indicates that income-capped subsidies do not accelerate adoption at all relative to the no-subsidy case. Panels (c)-(d) show that the income-capped subsidy design does not generate the short-run adoption acceleration among the medium- and high-wealth households that the non-capped subsidy does. Conversely, in the short run, the adoption speed among low-wealth households increases slightly under income-capped subsidies compared to non-capped subsidies. However, in the medium run, adoption speed among low-wealth households declines to its no-subsidy level, whereas adoption accelerates among high-wealth households.

Looking at the aggregate economic dynamics in Figure 11, reveals that general equilibrium effects create this unintended outcome of speeding up adoption among high-wealth and slowing down adoption among low- and medium-wealth households in the medium run. Panel (a) shows that aggregate savings in the short run do not change much with income-capped subsidies compared to no subsidies, as adoption speed and composition remain similar. However, in the medium run, aggregate savings fall significantly compared to no subsidies. Several factors are driving this significant decline in savings in the medium run.

Figure 10: Aggregate and across asset tercile adopter stock, experiment 1 and experiment 3.



First, as adoption speed remains the same, the subsidy does not induce additional LBD-induced cost declines, thus subsidy payments that are proportional to technology cost remain high. Thus, net of subsidy lump-sum transfers fall significantly more compared to non-capped subsidies, as seen in panel (b). Second, this larger adverse transfer effect reduces households' ability to save, further lowering aggregate savings. The larger decline in savings leads to a significant decrease in capital-labor ratio and wages in the medium run, which significantly reduces labor income for low- and medium-wealth households, further dampening their adoption speed. Moreover, the significant decline in the capital-labor ratio increases interest rates, which benefits high-wealth households who rely more on capital income, increasing their adoption speed. Thus, the price effect benefits middle- and high-wealth households on average, as they benefit from higher interest rates, while low-wealth households are harmed on average due to lower wages, as seen in Table 9

Figure 11: Aggregate dynamics under progressive financing policy with LBD, with and without income-capped adoption subsidies, experiment 3.





Figure 12: Welfare impact of subsidizing solar panel adoption cost across the income-wealth distribution, experiment 3



Finally, I visualize the welfare impact of income-capped subsidies across the income-wealth distribution in Figure 12. The decomposition of welfare changes across labor income and wealth groups highlights unintended distributional consequences of income-capped refundable subsidies. Virtually all households experience welfare losses from this policy, with the most significant losses concentrated among middle-income and middle-wealth households. These

households face significant adverse transfer effects and price effects, as they do not benefit from the subsidy but are adversely affected by general equilibrium price changes. On average, low-wealth households experience more minor welfare losses compared with the other groups, as they are more likely to qualify for the subsidy and benefit from direct subsidy effects, partially offsetting negative transfer and price effects. High-wealth households also experience more minor welfare losses, as they benefit from positive price effects due to higher interest rates, partially offsetting adverse transfer effects. Overall, this experiment demonstrates that income-capped subsidies under progressive financing can lead to significant and unequal welfare losses, highlighting the unintended distributional consequences of a policy designed to enhance equity.

6.2 Extension: Activating Pollution Damages and Inequality

The baseline and alternative policy analyses abstracted from pollution damages, which are a key motivation for environmental policies. These damages are particularly relevant for distributional assessment of policies, as pollution exposure and its health impacts are often concentrated among lower-income and marginalized communities. To incorporate pollution damages, I activate the pollution term in the utility function given by equation 11, which captures the disutility from pollution exposure. Importantly, disutility from pollution increases with a household’s relative consumption level, reflecting the pollution exposure inequality documented in Banzhaf, Ma, and Timmins (2019) and Sergi et al. (2020). This extension allows me to quantify the additional welfare effects arising from reduced pollution exposure due to clean technology adoption.

I build this extension on experiment 2, which features a nonrefundable subsidy under progressive financing. Table 10 presents the welfare decomposition of this extension. Accounting for abated pollution damages increases the aggregate welfare gain from subsidizing adoption significantly by an EV of 2.02%, compared to 0.05% without pollution damages in experiment 2. Moreover, every household strictly benefits from this policy when pollution damages are considered, compared to only 78.1% without pollution damages.

Figure 13 compares the adoption dynamics between experiment 2 without pollution damages (left) and this extension with pollution damages (right). This comparison shows that accounting for pollution damages slows aggregate adoption significantly in the no-subsidy case, as households face higher disutility from pollution exposure, reducing their effective consumption and savings. To understand how the introduction of pollution damages affects adoption dynamics, I consider the effect of the pollution term on the steady state equilibrium.


Table 10: Decomposition of welfare changes and distribution of winners by asset tercile when pollution damages are incorporated.

Panel A. Aggregate Welfare Decomposition				
Component	Avg. ΔV	Avg. CV (%)	Avg. EV (%)	Winner (%)
Direct subsidy	$+4.97 \times 10^{-2}$	-0.20	+0.20	100.0
LBD-induced cost change	$+5.85 \times 10^{-4}$	-0.00	+0.00	100.0
Price effect	-2.75×10^{-3}	+0.01	-0.01	56.7
Pollution effect	$+5.3 \times 10^{-1}$	-2.09	+2.15	100.0
Transfer effect	-7.92×10^{-2}	+0.32	-0.32	0.0
Total	$+4.99 \times 10^{-1}$	-1.97	+2.02	100.0

Panel B. Winner Shares by Asset Tercile (%)				
	Bottom	Middle	Top	Aggregate
Direct subsidy	100.0	100.0	100.0	100.0
LBD-induced cost change	100.0	100.0	100.0	100.0
Price effect	0.1	75.1	100.0	56.7
Pollution effect	100.0	100.0	100.0	100.0
Transfer effect	0.0	0.0	0.0	0.0
Total	100.0	100.0	100.0	100.0

Notes: Winner share is the fraction of households with strictly positive consumption-equivalent variation (EV). Tercile population masses are [0.355, 0.312, 0.333].


Figure 13: Aggregate adopter stock, experiment 2 and extension with pollution damages.



Differentiating the utility function given by equation (11) yields:

$$\begin{aligned}
 u_c(c, X) &= c^{-\sigma} + \nu \omega (X - \bar{X}) \bar{c}^\omega c^{-(\omega+1)}, \\
 u_{cc}(c, X) &= -\sigma c^{-(\sigma+1)} - \nu \omega (\omega + 1) (X - \bar{X}) \bar{c}^\omega c^{-(\omega+2)}, \\
 u_{ccc}(c, X) &= \sigma(\sigma + 1) c^{-(\sigma+2)} + \nu \omega (\omega + 1)(\omega + 2) (X - \bar{X}) \bar{c}^\omega c^{-(\omega+3)}.
 \end{aligned}$$

Figure 14: Aggregate dynamics under progressive financing policy with LBD, with and without nonrefundable adoption subsidies, for pollution extension of experiment 2.



For $X > \bar{X}$, both $|u_{cc}|$ and u_{ccc} increase, implying stronger curvature of marginal utility and hence higher risk aversion $A(c; X) = -u_{cc}/u_c$ and prudence $P(c; X) = -u_{ccc}/u_{cc}$. Pollution therefore, raises the marginal value of consumption and increases the incentive to hold liquid assets for self-insurance. In a steady state, this manifests as higher buffer-stock savings and a greater demand for liquidity, even if aggregate consumption falls. Panel (a) of Figure 14 shows that the initial steady state capital-labor ratio is higher when pollution damages are active, reflecting this increased demand for savings. The higher initial steady state capital-labor ratio implies that equilibrium wages are also higher initially. Thus the steady state level of lump-sum transfers is higher, as seen in panel (b) of Figure 14.

The rise in prudence has a direct implication for irreversible investments such as solar adoption. Let $V^N(a, \ell, 0)$ and $V^A(a, \ell, 0)$ denote the value functions for not adopting and adopting, respectively. The up-front cost p_t reduces liquid assets from a to $a - p_t$. The difference between these value functions is the expected next-period continuation value, where adopters switch to the clean-technology state:

$$\Delta(a, \ell; p_t) = V^A(a - p_t, \ell, 0) - V^N(a, \ell, 0).$$

A second-order expansion around a gives

$$V^A(a - p_t, \ell, 0) \approx V^A(a, \ell, 0) - V_a^A(a, \ell, 0)p_t + \frac{1}{2}V_{aa}^A(a, \ell, 0)p_t^2,$$

thus adopting entails an additional second-order welfare loss from giving up liquid wealth, $\frac{1}{2}|V_{aa}^A(a, \ell, 0)|p_t^2$. Because $V_{aa}^A < 0$, this term measures the curvature-induced penalty from

converting liquid assets into an illiquid investment. The curvature $|V_{aa}^A|$ is tied to $|u_{cc}|$ through the Euler condition, so when pollution damages are active and $\omega > 1$, the value function becomes more concave, and the shadow value of liquidity rises. Households therefore, face a more substantial option value of waiting: they postpone adoption until wealth is sufficiently high or technology costs have declined further.

Let $a^*(\ell; p_t)$ denote the adoption threshold such that $\Delta(a^*, \ell; p_t) = 0$. Because the adoption surplus falls with greater curvature, $\partial a^*(\ell; p_t)/\partial |V_{aa}^A| > 0$, pollution raises the wealth level required for adoption and slows the extensive-margin response in the early transition, as seen in Figure 13. This delay feeds back through learning-by-doing, flattening cost reductions and further slowing aggregate adoption.


The same higher curvature that discourages early adoption also alters aggregate saving behavior. In the pollution extension, households are more prudent in a steady state and ultimately accumulate more wealth, but during the transition, they expect pollution to decline as adoption expands. Because $\partial u_c/\partial X > 0$, an expected improvement $\mathbb{E}_t[X_{t+1}] < X_t$ lowers the expected future marginal utility on the right-hand side of the Euler equation,

$$u_c(c_t, X_t) = \beta(1 + r_{t+1}) \mathbb{E}_t[u_c(c_{t+1}, X_{t+1})].$$

To restore equality, households increase current consumption relative to c_{t+1} , temporarily reducing saving. As a result, aggregate capital falls at the start of the transition (Figure 14, panel (a)), before recovering as pollution stabilizes and the precautionary motive dominates. This pattern contrasts with the benchmark without pollution damages (Figure 9, panel (a)), where capital initially rises because agents expect higher future returns from adoption. Hence, pollution damages raise long-run prudence and steady-state saving, but in the short run, the anticipation of cleaner future conditions induces a front-loading of consumption that depresses capital and slows the overall diffusion of adoption.

The stronger curvature of preferences under pollution damages also amplifies the responsiveness of adoption to subsidies and the resulting welfare gains. Figure 13 shows that while overall adoption eventually converges in both environments, the relative acceleration of adoption due to the subsidy is much larger when pollution damages are active. With pollution damages unaccounted for, the subsidy primarily affects adoption through its financial channel—by lowering effective installation costs and speeding up learning-by-doing—yielding modest aggregate welfare gains. When pollution damages are included, however, the same subsidy additionally reduces future pollution exposure, generating a direct utility improvement and indirectly mitigating the precautionary motive that suppresses early saving. As a result, households adopt more rapidly, pollution declines sooner, and the welfare impact

Figure 15: Welfare impact of subsidizing solar panel adoption cost across the income-wealth distribution, extension with pollution



of the subsidy expands well beyond its pure learning-by-doing effects. Table 10 quantifies this amplification: the equivalent variation rises from 0.05% in experiment 2 to 2.02% in the pollution extension, and every household benefits once pollution damages are accounted for.

In general equilibrium, the subsidy alleviates the short-run contraction in capital observed in Figure 14(a) by accelerating the decline in X_t . Faster abatement raises effective lifetime wealth and allows precautionary saving to recover earlier, reinforcing the positive income effect of the subsidy. Hence, when pollution damages are active, the subsidy internalizes both the learning-by-doing and pollution externalities: It increases the speed of adoption, smooths the short-run adjustment of capital, and yields substantially larger welfare gains for all households. Because households do not internalize the social benefit from lower pollution, the stronger precautionary motive amplifies the wedge between private and social incentives to adopt, which further magnifies the aggregate welfare benefit of subsidizing clean technology adoption once pollution damages are accounted for.

The welfare gains from subsidizing adoption are also more progressive when pollution damages are accounted for. Figure 15 shows that the average welfare improvement declines monotonically with wealth, with the most significant gains accruing to lower- and middle-wealth households. This pattern reflects two reinforcing mechanisms. First, low-wealth households are more exposed to the disutility from pollution in the baseline, thus the reduction in pollution damages brought about by faster adoption generates a larger direct welfare benefit

for them. Second, pollution damages increase prudence and the marginal value of consumption more strongly for liquidity-constrained households. The subsidy’s income effect—through higher effective wealth and lower precautionary saving demand—is proportionally greater for these groups. As a result, the welfare impact of the subsidy is both larger in aggregate and more progressive when pollution damages are active. Unlike in the baseline without pollution, where gains were concentrated among higher-wealth adopters, every household type benefits once pollution exposure is internalized, and the relative improvement is most significant among those who were initially most vulnerable to pollution and liquidity constraints.

7 Conclusion

This paper re-examines the equity and efficiency of clean energy subsidies through the case of US residential rooftop solar panel installations. Combining new empirical evidence on learning-by-doing with a heterogeneous-agent general equilibrium model featuring incomplete markets, irreversible adoption, and unequal pollution damages, I quantify how different subsidy and financing designs shape adoption dynamics and welfare across the income-wealth distribution.

The analysis shows that static incidence measures overstate the regressivity of residential solar subsidies. When learning-by-doing spillovers are included, subsidies accelerate cost declines and generate broad welfare gains, even if the direct fiscal transfers favor wealthier households. Moreover, when pollution exposure and its unequal health impacts are incorporated, the welfare effects become universally positive and strongly progressive: all households gain, and the poorest benefit most from cleaner air and faster cost declines.

The results highlight that policy design—particularly the choice of financing mechanism and refundability—plays a decisive role. Progressive tax financing, though aimed at fairness, can inadvertently slow down early adoption and short-run wages, while uniform refundable subsidies promote faster diffusion and higher aggregate welfare. The most equitable transition therefore requires a portfolio approach: broad early subsidies to exploit learning externalities, complemented by targeted measures that alleviate liquidity constraints and internalize pollution damages.

More broadly, the paper provides a quantitative framework for evaluating the joint efficiency and equity effects of climate policy in economies with heterogeneous households, incomplete markets, and dynamic externalities. The evidence from US residential solar suggests that when policies internalize both learning and pollution externalities, a fast energy transition can also be an equitable one.

References

- Aiyagari, S. Rao (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving”. *The Quarterly Journal of Economics* 109.3, pp. 659–684.
- Allcott, Hunt, Reigner Kane, Maximilian S. Maydanchik, Joseph S. Shapiro, and Felix Tintelnot (2024). *The Effects of “Buy American”: Electric Vehicles and the Inflation Reduction Act*. Working Paper. Accessed: 2025-10-10.
- Banzhaf, Spencer, Lala Ma, and Christopher Timmins (2019). “Environmental Justice: The Economics of Race, Place, and Pollution”. *Journal of Economic Perspectives* 33.1, pp. 185–208.
- Barboase, Galen, Naïm Darghouth, Eric O-Shaughnessy, and Sydney Forrester (2023). *Tracking the Sun*. Tech. rep. Accessed: 2024-05-24.
- Belfiori, E., D. Carroll, and S. Hur (2024). *Unequal Climate Policy in an Unequal World*. Working Paper 427. Federal Reserve Bank of Dallas.
- Benmir, G. and J. Roman (2022). *The Distributional Costs of Net-Zero: A Heterogeneous Agent Perspective*. Working Paper.
- Bewley, Truman (1977). “The permanent income hypothesis: A theoretical formulation”. *Journal of Economic Theory* 16.2, pp. 252–292.
- Bollinger, B. and K. Gillingham (2023). *Learning-by-Doing in Solar Photovoltaic Installations*. Working Paper.
- Borenstein, Severin and Lucas W Davis (2024). *The Distributional Effects of U.S. Tax Credits for Heat Pumps, Solar Panels, and Electric Vehicles*. Working Paper 32688. National Bureau of Economic Research.
- Bureau of Labor Statistics (2024). *Consumer Expenditure Surveys, Public Use Microdata*. Tech. rep. Accessed: 2024-11-05.
- Dennin, Luke, Nicholas Muller, Medinat Akindele, and Peter Adams (2024). *AP4 and AP4T Integrated Assessment Model*. Technical Appendix. Accessed: 2025-01-10.
- Environmental Protection Agency (2024a). *Particulate Matter (PM) Trends*. <https://www.epa.gov/air-trends/particulate-matter-pm25-trends>. Accessed: 2025-10-05.
- (2024b). *PM Standards*. <https://www.epa.gov/pm-pollution/national-ambient-air-quality-standards-naaqs-pm>. Accessed: 2025-10-05.
- Fried, Stephanie, Kevin Novan, and William B. Peterman (2018). “The distributional effects of a carbon tax on current and future generations”. *Review of Economic Dynamics* 30, pp. 30–46.
- (2024). “Understanding the Inequality and Welfare Impacts of Carbon Tax Policies”. *Journal of the Association of Environmental and Resource Economists* 11.S1, S231–S260.

- Gao, Xue, Varun Rai, and Gregory F. Nemet (2022). “The roles of learning mechanisms in services: Evidence from US residential solar installations”. *Energy Policy* 167, p. 113003.
- Huggett, Mark (1993). “The risk-free rate in heterogeneous-agent incomplete-insurance economies”. *Journal of Economic Dynamics and Control* 17.5, pp. 953–969.
- Internal Revenue Service (2000). *Tax Tables and Tax Rate Schedules (Form 1040 and 1040A Instructions, 2000)*. <https://www.irs.gov/pub/irs-prior/i1040tt--2000.pdf>. Accessed: 2025-10-05.
- (2025). *Credits for new clean vehicles purchased in 2023 or after*. <https://www.irs.gov/credits-deductions/credits-for-new-clean-vehicles-purchased-in-2023-or-after>. Accessed: 2025-10-05.
- Känzig, Diego R. (2023). *The Unequal Economic Consequences of Carbon Pricing*. Working Paper.
- Kuhn, Moritz and Lennard Schlattmann (2024). *Distributional Consequences of Climate Policies*. Working Paper. Accessed: 2024-11-27.
- Lane, Catherine (2025). *Solar Tax Credits Over Time: The History of Home Solar’s Best Incentive*.
- Lanteri, Andrea and Adriano Rampini (2025). *Financing the Adoption of Clean Technology*. Working Paper.
- Levinson, Arik (2019). “Energy Efficiency Standards Are More Regressive Than Energy Taxes: Theory and Evidence”. *Journal of the Association of Environmental and Resource Economists* 6.S1, S7–S36.
- Linn, Joshua (2022). *Balancing Equity and Effectiveness for Electric Vehicle Subsidies*. Working Paper. Accessed: 2024-06-10.
- Nemet, Gregory F., Jiaqi Lu, Varun Rai, and Rohan Rao (2020). “Knowledge spillovers between PV installers can reduce the cost of installing solar PV”. *Energy Policy* 144, p. 111600.
- Nemet, Gregory F., Eric O’Shaughnessy, Ryan H Wiser, Naïm R Darghouth, Galen L Barbose, Kenneth Gillingham, and Varun Rai (2016). *What Factors Affect the Prices of Low-Priced US Solar PV Systems?* Tech. rep.
- North Carolina Clean Energy Technology Center (2025). *Database of State Incentives for Renewables and Efficiency (DSIRE)*. Tech. rep. North Carolina State Univ., Raleigh, NC (United States).
- O’Shaughnessy, Eric (2018). *The Effects of Market Concentration on Residential Solar PV Prices: Competition, Installer Scale, and Soft Costs*. Tech. rep.

- Sergi, Brian, Inês Azevedo, Steven J Davis, and Nicholas Z Muller (Oct. 2020). “Regional and county flows of particulate matter damage in the US”. *Environmental Research Letters* 15.10, p. 104073.
- Social Research Center (2000). *Panel Study of Income Dynamics (PSID): Main Interview, 2000*. Accessed: 2025-10-05.
- Solar Energy Industries Association (2025). *Solar Industry Research Data*. Tech. rep. Accessed: 2025-06-11.
- US Census Bureau (2022). *Income in the Past 12 Months, Table S1901*. Tech. rep. Accessed: 2024-06-06.
- (2023a). *Household Income Quintile Upper Limits, Table B19080, 2015*. Tech. rep. Accessed: 2025-01-06.
- (2023b). *Shares of Aggregate Household Income by Quintile, Table B19082, 2015*. Tech. rep. Accessed: 2025-01-06.
- (2024). *Real Median Household Income in the United States*. Tech. rep.
- (2025). *Population Estimates Program*. Accessed: 2025-10-06. US Census Bureau.
- US Centers for Disease Control and Prevention (2023). *Natality and Mortality Trends in the United States, 2000-2023*. <https://www.cdc.gov/nchs/data-visualization/birth-to-death-ratios/natality-mortality-trends.htm>. Accessed October 6, 2025.
- US Department of Energy (2024). *New and Used Clean Vehicle Tax Credits*. <https://www.energy.gov/energysaver/new-and-used-clean-vehicle-tax-credits>. Accessed: 2025-10-05.
- US Energy Information Administration (2023). *Residential Energy Consumption Survey Microdata*. Tech. rep. Accessed: 2024-06-11.
- (2024). *US energy consumption by source and sector, 2023, Monthly Energy Review*. Tech. rep. Accessed: 2025-02-11.
- US Environmental Protection Agency (2024c). *Guidelines for Preparing Economic Analyses, Appendix B*. Tech. rep. EPA-240-R-24-001.
- US Internal Revenue Service (2023). *Individual Statistical Tables by Size of Adjusted Gross Income*. Tech. rep. Accessed: 2024-06-06.
- Vodonos, Alina, Yara Abu Awad, and Joel Schwartz (2018). “The concentration-response between long-term PM_{2.5} exposure and mortality; A meta-regression approach”. *Environmental Research* 166, pp. 677–689.
- Vona, Francesco (2023). “Managing the Distributional Effects of Climate Policies: A Narrow Path to a Just Transition”. *Ecological Economics* 205, p. 107689.

Appendices

Appendix A Empirical Appendix

A.1 Data Sources and Variable Construction

A.1.1 Electricity and Energy Expenditure Data

I use the Residential Energy Consumption Survey (RECS) data from RECS (2023) to analyze household-level annual electricity and total energy expenditure and consumption variables. The RECS is a nationally representative survey of US households' energy consumption and expenditures, housing unit characteristics, and demographic information. I use 2020 RECS microdata, because it is the first year the survey started including state identifiers. Including state identifiers in the analysis is essential because electricity prices vary across states due to differences in electricity generation costs, taxes, and other factors. Due to the lack of state identifiers in earlier RECS data, I cannot control for state fixed effects in the regression analysis. The data includes the annual total electricity expenditure, in dollars, and electricity consumption, in BTUs, of US households in 2020, together with information on whether there is on-site electricity generation using rooftop solar panels, the type of heating fuel used, and other housing unit and household characteristics. I calculate the unit electricity price paid by households in 2020 as the ratio of the annual electricity expenditure to the annual electricity consumption for each household.

To describe the effect of having on-site solar generation on unit electricity price paid by households and annual electricity expenditure, I estimate the following regression specification.

$$\ln(Y_i) = \alpha + \beta \text{Solar}_i + \gamma \mathbf{X}_i + \delta_s + \epsilon_i,$$

where Y_i is the outcome variable of interest (annual electricity expenditure, annual total energy expenditure, electricity unit cost, or total energy unit cost) for household i ; Solar_i is an indicator variable for whether household i generates solar power on-site; $\mathbf{X}_i$ is a vector of household and housing unit characteristics (household size, income category, age of head, etc.); δ_s is state controls; and ϵ_i is the error term. I estimate four specifications for each outcome variable, all with state controls, but with and without additional household and housing unit controls using the survey weights provided in the RECS data.

I present the results of Weighted Least Squares (WLS) estimates of this regression with and without additional controls in Table 12. The results of the regression show that the effect of having solar generation on the unit electricity price, presented in columns (1) and

Table 11: Regression results for the effect of solar generation on annual electricity and total energy expenditure

Variables	Electricitiy		Total Energy	
	(1)	(2)	(3)	(4)
Solar	-721.956 (15.938)	-709.430 (15.881)	-688.5617 (24.286)	-609.9241 (27.094)
State Controls	Yes	Yes	Yes	Yes
Additional Controls	Yes	No	Yes	No
Observations	15,044	15,044	15,044	15,044

Notes: Standard errors are in parentheses. Fixed effects indicate the inclusion of state fixed effects and additional controls refer to inclusion of factors such as household and housing unit characteristics.

(2), is negative and statistically significant. Specifically, having solar generation reduces the unit electricity price by \$0.021 per BTU, which is a substantial reduction given that the average unit electricity price in the entire sample is around \$0.041 per BTU. Thus, having solar generation reduces the unit electricity price by almost 50%. The direction of the effect, presented in columns (3) and (4), is similar when the dependent variable is the unit energy price, but the magnitude of the effect is smaller. The average unit energy price in the entire sample is around \$0.026 per BTU, and having solar generation reduces the unit energy price by \$0.009 per BTU, which is around 36% of the average unit energy price.

In order to understand the total nominal magnitude of the effect of having on-site solar generation on household's annual electricity expenditure, I run an alternative regression where the dependent variable is the annual electricity expenditure of households and annual electricity consumption is an additional control. The results of this regression are presented in columns (1) and (2) Table 11. The results show that having solar generation reduces the annual electricity expenditure of households by more than \$700 annually, which is a substantial reduction given that the average annual electricity expenditure in the entire sample is around \$1,500. The effect of having solar generation on the annual total energy expenditure of households is also negative and statistically significant, but the magnitude of the effect is smaller. The results of the regression, presented in columns (3) and (4), show that having solar generation reduces the annual total energy expenditure of households by around \$600 annually and the average annual total energy expenditure in the entire sample is around \$2,171. Thus, having solar generation reduces the annual total energy expenditure of households by around 30%.

Table 12: Regression results for the effect of solar generation on electricity and total energy unit costs

Variables	Electricitiy		Total Energy	
	(1)	(2)	(3)	(4)
Solar	-0.0214 (0.000)	-0.0214 (0.000)	-0.0093 (0.000)	-0.0091 (0.000)
State Controls	Yes	Yes	Yes	Yes
Additional Controls	Yes	No	Yes	No
Observations	15,044	15,044	15,044	15,044

Notes: Standard errors are in parentheses. Fixed effects indicate the inclusion of state fixed effects and additional controls refer to inclusion of factors such as household and housing unit characteristics.

A.1.2 Policy Shocks Instrument

Data Scope and Unit of Observation

I use the North Carolina Clean Energy Technology Center’s DSIRE ([2025](#)) database to assemble a monthly policy shock panel. The panel is built at two geographic resolutions used in the empirical analysis: state and county levels. I use the state-month panel for specifications with state fixed effects, and county-month panel for specifications with county fixed effects. I retain all program records with residential applicability (statewide, county, municipal, and utility programs). Programs that apply to the entire state are assigned to all counties in the state when constructing the county-month panel. I normalize county identifiers to five-digit FIPS codes.

Dates and Activity Windows

Each program has an activation window defined by its start and end dates, defined at monthly frequency. For programs with missing dates, I follow multiple imputation strategies. First, I search program descriptions for date information. If an active program’s end date is missing, I set its end month to December 2025. If a start date is missing, I use a the date program was first listed in the DSIRE database as a proxy. For inactive programs with missing end dates, I parse the program description for text such as ”expired in YYYY” to infer the end month when possible; otherwise, I drop the record. I assume that a program is active in all months between its start and end dates, inclusive.

Policy Shock, $Z_{j,t}$

The policy shock is a discrete timing shock that flags the onset of a new residential solar

program in location j in month t :

$$Z_{j,t} = \mathbb{1}\{\exists \text{ program } m \text{ with start month } = t, \text{ and serving location } = j\}.$$

If multiple programs begin in the same location and time, $Z_{j,t}$ remains 1, i.e., it is not a count of new programs. For IV, I use a lagged version of this variable, $Z_{j,t-12}$, with $L = 12$ months lag, to allow for a one-year adoption response window.

Policy Generosity, $g_{j,t}$

I construct a monthly policy generosity index $g_{j,t}$ to control measures contemporaneous subsidy intensity in \$ per watt in location j and month t . It aggregates all active programs' incentive values mapped to a common unit of \$ per watt, then sums across concurrent programs:

$$g_{j,t} \equiv \sum_{m \in \mathcal{M}_{j,t}} \text{Generosity per Watt}_{m,t}, \quad \mathcal{M}_{j,t} = \{\text{programs active in } j \text{ at } t\}.$$

Mapping Program Parameters to Generosity per Watt

Each program's parameterization is converted to per-watt generosity using observed NREL (2023) average state-year price and size benchmarks. I compute the average net cost per watt in state s , year y , denote by $\bar{P}_{s,y}$, as total price divided by total watts in a state-year. Let $\bar{S}_{s,y}$ be the average system size (in watts) in same state-year. If a state-year average is unavailable, then I use the national-year average from the same NREL (2023) sample.

For a program m applicable to state s and month t in year y , I use the following mapping rules:

- **Rebate in (\$/W):** If the program offers a fixed rebate amount, use the amount as is.
- **Percentage of cost (%):** Multiply the percentage by the average net cost per watt $\bar{C}_{s,y}$.
- **Flat amount (\$):** Divide the flat amount by the average system size $\bar{S}_{s,y}$.
- **Production based and other non-capital incentives:** Excluded from $g_{j,t}$.

If a program lists multiple parameter rows, I sum the implied generosity per watt across rows in a month. Negative or nonsensical generosity values are set to zero.

Appendix B Theoretical Appendix

B.1 Initial and Terminal Recursive Equilibria

Prior to the availability of the clean energy technology for adoption, households face no adoption decision and are in state $s = 0$. The state space is:

$$\tilde{\mathcal{Z}} = \mathcal{A} \times \mathcal{L}, \quad B(\tilde{\mathcal{Z}}) = B(\mathcal{A}) \times P(\mathcal{L}),$$

and the distribution of households is given by $\Phi \in \tilde{\mathcal{M}}$, where $\tilde{\mathcal{M}}$ is the set of Borel probability measures on $(\tilde{\mathcal{Z}}, B(\tilde{\mathcal{Z}}))$.

Definition 2 *Given the labor income tax rate τ^ℓ , exogenous dirty energy price $\bar{q}$, and ambient air pollution function $X(\Phi)$, a recursive competitive equilibrium consists of: a value function $V : \tilde{\mathcal{Z}} \times \tilde{\mathcal{M}} \rightarrow \mathbb{R}$, household policy functions $c, a' : \tilde{\mathcal{Z}} \times \tilde{\mathcal{M}} \rightarrow \mathbb{R}$, aggregate factor demands $K, L : \tilde{\mathcal{M}} \rightarrow \mathbb{R}$, factor price functions $r, w : \tilde{\mathcal{M}} \rightarrow \mathbb{R}$, a transfer function $T : \tilde{\mathcal{M}} \rightarrow \mathbb{R}$, a pollution function $X : \tilde{\mathcal{M}} \rightarrow \mathbb{R}$, and law of motion $H : \tilde{\mathcal{M}} \rightarrow \tilde{\mathcal{M}}$ such that:*

1. **Household optimization.** *Given the pricing functions $r(\Phi), w(\Phi)$, transfer function $T(\Phi)$, V solves the following Bellman equation:*

$$V(a, \ell; \Phi) = \max_{c \geq 0} U(c, X(\Phi)) + \beta \mathbb{E} \{ V[w(\Phi)(1 - \tau^\ell)\ell + (1 + r(\Phi))a + T(\Phi) - c - \bar{q}e(c), \ell'; \Phi'] | \ell \},$$

subject to $\Phi' = H(\Phi)$,

(17)

and c is the associated consumption policy function, $a' = w(\Phi)(1 - \tau^\ell)\ell + (1 + r(\Phi))a + T(\Phi) - c - \bar{q}e(c)$ is the savings policy function, and $\mathbb{E}$ is the conditional expectation operator.

2. **Factor prices.** *Factor prices $r(\Phi)$ and $w(\Phi)$ satisfy the firm's first-order conditions:*

$$r(\Phi) = F_K(K(\Phi), L(\Phi)),$$

$$w(\Phi) = F_L(K(\Phi), L(\Phi)).$$

3. **Government budget.** *Given the factor prices $r(\Phi)$ and $w(\Phi)$ and the tax rate τ^ℓ , the government runs a balanced budget every period such that $T(\Phi)$ satisfies:*

$$T(\Phi) = w(\Phi)\tau^\ell \int_{\tilde{\mathcal{Z}}} \ell d\Phi.$$

4. **Pollution.** The ambient air pollution $X(\Phi)$ is a function of the aggregate energy consumption and satisfies:

$$X(\Phi) = \Omega \left(\int_{\tilde{\mathcal{Z}}} e(c(a, \ell; \Phi)) d\Phi \right).$$

5. **Market clearing.** For all $\Phi \in \tilde{\mathcal{M}}$,

$$\begin{aligned} L(\Phi) &= \int_{\tilde{\mathcal{Z}}} \ell d\Phi, \\ K(H(\Phi)) &= \int_{\tilde{\mathcal{Z}}} a'(a, \ell; \Phi) d\Phi, \\ \int_{\tilde{\mathcal{Z}}} [c(a, \ell; \Phi) + a'(a, \ell; \Phi) + \bar{q}e(c(a, \ell; \Phi))] d\Phi &= F(K(\Phi), L(\Phi)) + (1 - \delta)K(\Phi) - K(H(\Phi)). \end{aligned}$$

6. Aggregate law of motion H is generated by π and a' , explicitly stated in Appendix B.2.

Next, I characterize the recursive competitive equilibrium in the initial steady state as follows:

Definition 3 Given the labor income tax rate τ^ℓ the stationary recursive equilibrium is a value function V , household policy functions c, a' , aggregate production factors K, L , prices $r, w, \bar{q}$, government transfer T , pollution function X , and a measure Φ , with $\Phi \in \tilde{\mathcal{M}}$ invariant under H , such that the household optimization, factor prices, government budget, pollution, and market clearing conditions above hold, and Φ satisfies:

$$\Phi = H(\Phi).$$

Once adoption is available and irreversible, the terminal steady state has all households in state $s = 1$. Then $\mathcal{Z} = \mathcal{A} \times \mathcal{L} \times \{1\}$. The recursive equilibrium is defined analogously, with $\bar{q}$ replaced by $\underline{q}$ and $X(\Phi) = 0$ for all Φ .

B.2 Explicit Statement of the Aggregate Law of Motion in the Steady States

First, define the Markov transition function: $Q_\Phi : \tilde{\mathcal{Z}} \times B(\tilde{\mathcal{Z}}) \rightarrow [0, 1]$ by:

$$Q_\Phi((a, \ell), (\mathcal{A}, \mathcal{L})) = \sum_{\ell' \in \mathcal{L}} \begin{cases} \Pi(\ell' | \ell) & \text{if } a'(a, \ell; \Phi) \in \mathcal{A}, \\ 0 & \text{otherwise,} \end{cases}$$

for all $(a, \ell) \in \tilde{\mathcal{Z}}$ and $(\mathcal{A}, \mathcal{L}) \in B(\tilde{\mathcal{Z}})$. Thus $Q_\Phi((a, \ell), (\mathcal{A}, \mathcal{L}))$ is the probability that an agent with current assets a and productivity ℓ ends up with assets $a' \in \mathcal{A}$ and productivity $\ell' \in \mathcal{L}$ tomorrow. Then, the aggregate law of motion for the initial steady state distribution is given by:

$$\begin{aligned}\Phi'(\mathcal{A}, \mathcal{L}) &= H(\Phi)(\mathcal{A}, \mathcal{L}) = \int_{\tilde{\mathcal{Z}}} Q_\Phi((a, \ell), (\mathcal{A}, \mathcal{L})) \\ &= \int Q_\Phi((a, \ell), (\mathcal{A}, \mathcal{L})) \Phi(da \times d\ell).\end{aligned}$$

B.3 Explicit Statement of the Aggregate Law of Motion During the Transition Path

Define the Markov transition functions $Q_t : \mathcal{Z} \rightarrow [0, 1]$ induced by the transition probabilities π and optimal policies $a_{t+1}(a, \ell, s)$ and $S_t(a, \ell, s)$ as:

$$Q_t((a, \ell, s), (\mathcal{A}, \mathcal{L}, \{0, 1\})) = \sum_{\ell' \in \mathcal{L}} \begin{cases} \pi(\ell' | \ell) & \text{if } a_{t+1}(a, \ell, s) \in \mathcal{A}, \\ 0 & \text{otherwise,} \end{cases}$$

for all $(a, \ell, s) \in \mathcal{Z}$ and $(\mathcal{A}, \mathcal{L}, \{0, 1\}) \in B(\mathcal{Z})$. Then, for all $(\mathcal{A}, \mathcal{L}, \{0, 1\}) \in B(\mathcal{Z})$, the aggregate law of motion for the transition distribution is given by:

$$\Phi_{t+1}(\mathcal{A}, \mathcal{L}, \{0, 1\}) = [\Gamma_t(\Phi_t)](\mathcal{A}, \mathcal{L}, \{0, 1\}) = \int Q_t((a, \ell, s), (\mathcal{A}, \mathcal{L}, \{0, 1\})) d\Phi_t.$$

Appendix C Quantitative Appendix

C.1 Calibration Details

As a sensitivity check, I adjust the baseline uniform labor income tax rate assumption to match the 2000 US federal income tax brackets and rates presented in Table 13, sourced from IRS (2000). I compute the uniform labor income tax rate τ^ℓ such that the average labor income tax liability under the 2000 tax brackets equals the average labor income tax liability under the uniform tax rate in the model's initial steady state. Using the 2000 tax brackets, I find that the uniform labor income tax rate that matches the average tax liability is $\tau^\ell = 0.1953$.

Table 13: Income Tax Brackets and Rates

Income Bracket (\$2000)	Marginal Tax Rate (%)
0 - 26,250	15%
26,250 - 63,550	28%
63,550 - 132,600	31%
132,600 - 288,350	36%
288,350 - above	39.6%

C.2 Computation Details

Steady States

Two steady states are computed:

1. **Initial steady state:** Clean energy technology is unavailable ($s = 0$), exogenous dirty energy price is $\bar{q}$, and the entire proceeds from the exogenous and flat labor income tax rate τ^ℓ is distributed lump-sum to households equally. After initializing parameters, guess K and T , and solve the Bellman equation (17) using Value Function Iteration with linear interpolation and a uniform asset grid. I use Golden Section Search for optimization and Howard's step for speed improvements. I compute the invariant distribution of (a, ℓ) by iterating over the density function on a finer uniform asset grid until convergence. Finally, I compute the K and T implied by the invariant distribution and compare to the initial guesses, updating using a dampening parameter $a = 0.95$ until convergence within a tolerance threshold equal to 10^{-6} .
2. **Terminal steady state:** The entire population has adopted the clean technology ($s = 1$), the exogenous clean energy price is $\underline{q}$, and the entire proceeds from the exogenous and flat labor income tax rate τ^ℓ is distributed lump-sum to households equally. The same algorithm is used to find $(K_\infty, T_\infty, \Phi_\infty)$.

Transition Path Algorithm

Given a finite horizon $T = 190$, the model is solved under perfect foresight using an outer-loop fixed-point algorithm:

1. *Initialization:* Import the value and policy functions for both the initial and terminal steady state calculations. Also import the invariant distribution at the initial steady state.

2. *Guessing*: Guess sequences $\{K_t, T_t, Z_t\}_{t=1}^T$, where $t = 1$ is the period when the transition starts and $t = T$ is sufficiently far in the future so that I can assume that the economy is sufficiently close to the new steady state. Also calculate the factor prices $\{r_t, w_t\}_{t=1}^T$ implied by evaluating the firm's first-order conditions at the guessed capital path.
3. *Backward induction*: at $t = T + 1$, the economy is in the terminal steady state, so $K_{T+1} = K_\infty$, $T_{T+1} = T_\infty$, and $V_{T+1} = V_\infty$ and I can use the terminal value function from the previous step in the right-hand side of the period- T Bellman equation, given by equations (9) and (10). From there, I solve for the household's value and policy functions with backward induction for $t = 1, \dots, T$.
4. *Forward simulation*: Using the policy function from the previous step, I simulate the economy forward, starting from the imported density function at the initial steady state. Using the simulated distribution of households, I compute K_t , T_t , and Z_t implied by the distribution at each $t = 1, \dots, T$.
5. *Convergence*: Compare the implied $\{K_t, T_t, Z_t\}_{t=1}^T$ to the initial guesses and modify the guesses using a dampening parameter in $(0, 1)$ until convergence of the guess and updated paths within a tolerance threshold equal to 10^{-6} .

C.3 Welfare Change Calculations

For each household state (a, ℓ, s) , I compute three welfare change measures between the baseline and counterfactual paths. **Equivalent Variation (EV)**

I define the consumption equivalent variation (EV) as the percentage change in consumption that, if applied to the baseline consumption path, would yield the same lifetime utility as under the counterfactual path. Thus, a positive EV indicates that the household is better off under the counterfactual scenario.

Formally, let $c_t(a, \ell, s)^\tau$ be the consumption policy function at time $t = 1, \dots, T$, under a scenario with adoption subsidy τ for a household with state (a, ℓ, s) at the initial steady state $t = 0$. Then, the EV from subsidizing the clean energy technology adoption at rate $\tau = 0.3$ relative to not subsidizing is computed as $\lambda^{EV}(a, \ell, s)$ that solves:

$$\mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U((1 + \lambda^{EV}(a, \ell, s))c_t(a, \ell, s)^{\tau=0}) \right] = \mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U(c_t(a, \ell, s)^{\tau=0.3}) \right],$$

Under the CRRA utility specification, the EV can be computed in closed form as:

$$\lambda^{EV}(a, \ell, s) = \begin{cases} \exp \{ (1 - \beta)[V_1(a, \ell, s)^{\tau=0.3} - V_1(a, \ell, s)^{\tau=0}] \} - 1, & \text{if } \sigma = 1, \\ \left[\frac{V_1(a, \ell, s)^{\tau=0.3}}{V_1(a, \ell, s)^{\tau=0}} \right]^{\frac{1}{1-\sigma}} - 1, & \text{if } \sigma \neq 1, \end{cases}$$

where $V_1(a, \ell, s)^{\tau=0}$ is the value function at $t = 1$ under the scenario with no adoption subsidy and $V_1(a, \ell, s)^{\tau=0.3}$ is the value function at $t = 1$ under the scenario with a 30% adoption subsidy for a household with state (a, ℓ, s) at $t = 1$.

In order to compute the welfare change due to the introduction of the clean energy technology, I compute the CEV as:

$$\mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U((1 + \lambda^{EV}(a, \ell, s))c_t(a, \ell, s)^{\text{initial ss}}) \right] = \mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U(c_t(a, \ell, s)^{\text{transition}}) \right],$$

where $c_t(a, \ell, s)^{\text{initial ss}}$ is the consumption path under no technology transition (the technology remains available and the economy remains in the initial steady state) and $c_t(a, \ell, s)^{\text{transition}}$ is the consumption path under the technology transition, without any subsidies. If $\lambda^{EV}(a, \ell, s) > 0$, then the household is better off with the introduction of the clean energy technology and supports the transition.

Compensating Variation (CV)

I define the compensating variation (CV) as the percentage change in consumption that would yield the same lifetime utility as baseline consumption path, but under the counterfactual path. Thus, a positive CV indicates that the household is worse off under the counterfactual scenario.

Formally, let $c_t(a, \ell, s)^\tau$ be the consumption policy function at time $t = 1, \dots, T$, under a scenario with adoption subsidy τ for a household with state (a, ℓ, s) at the initial steady state $t = 0$. Then, the CV from subsidizing the clean energy technology adoption at rate $\tau = 0.3$ relative to not subsidizing is computed as $\lambda^{CV}(a, \ell, s)$ that solves:

$$\mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U(c_t(a, \ell, s)^{\tau=0}) \right] = \mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U((1 + \lambda^{CV}(a, \ell, s))c_t(a, \ell, s)^{\tau=0.3}) \right],$$

Under the CRRA utility specification, the CV can be computed in closed form as:

$$\lambda^{CV}(a, \ell, s) = \begin{cases} \exp \{ (1 - \beta)[V_1(a, \ell, s)^{\tau=0} - V_1(a, \ell, s)^{\tau=0.3}] \} - 1, & \text{if } \sigma = 1, \\ \left[\frac{V_1(a, \ell, s)^{\tau=0}}{V_1(a, \ell, s)^{\tau=0.3}} \right]^{\frac{1}{1-\sigma}} - 1, & \text{if } \sigma \neq 1, \end{cases}$$

where $V_1(a, \ell, s)^{\tau=0}$ is the value function at $t = 1$ under the scenario with no adoption subsidy and $V_1(a, \ell, s)^{\tau=0.3}$ is the value function at $t = 1$ under the scenario with a 30% adoption subsidy for a household with state (a, ℓ, s) at $t = 1$. Note that $\lambda^{CV}(a, \ell, s)$ and $\lambda^{EV}(a, \ell, s)$ are related by:

$$\lambda^{CV}(a, \ell, s) = -\frac{\lambda^{EV}(a, \ell, s)}{1 + \lambda^{EV}(a, \ell, s)}.$$

Net Present Value (NPV) Change

I define the net present value (NPV) change as the absolute change in the present discounted value of consumption between the baseline and counterfactual paths. Thus, a positive NPV change indicates that the household is better off under the counterfactual scenario.

Formally, let $c_t(a, \ell, s)^\tau$ be the consumption policy function at time $t = 1, \dots, T$, under a scenario with adoption subsidy τ for a household with state (a, ℓ, s) at the initial steady state $t = 0$. Then, the NPV change from subsidizing the clean energy technology adoption at rate $\tau = 0.3$ relative to not subsidizing is computed as:

$$\text{NPV}^{\text{change}}(a, \ell, s) = \mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U(c_t(a, \ell, s)^{\tau=0.3}) \right] - \mathbb{E}_1 \left[\sum_{t=1}^T \beta^t U(c_t(a, \ell, s)^{\tau=0}) \right].$$

These three welfare change measures could provide different rankings of households' welfare changes, as they measure different concepts of welfare change. The definition of share of households that are better off under the counterfactual scenario depends on the welfare measure used. Unless otherwise specified, I use the EV measure in the main text.