Equity in Transition: Aggregate and Heterogeneous Effects of Clean Energy Technology Subsidies

Yağmur Menzilcioğlu*

November 1, 2025†

Click here for the latest version

Abstract

I study the aggregate and distributional effects of clean energy subsidies on US residential rooftop solar panel adoption. Using installation-level data on residential solar system installations, I provide new evidence on 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, aggregate welfare, and the distribution of gains across households. Uniform refundable subsidies financed by a flat labor income tax raise aggregate welfare and accelerate adoption, while 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 local pollution externalities reveals that clean energy subsidies can enhance both efficiency and equity, contrary to the view that they are inherently regressive.

Keywords: Clean energy transition, Fiscal policy design, Heterogeneous agent model, Residential solar installations, Learning spillovers, Unequal 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, Joakim Weill, Harun Alp, Cristina Fuentes-Albero, Dan Cao, Adrien Bilal, Sadhika Bagga, Maarten De Ridder, Simone Lenzu, Cesar Santos, Pascual Restrepo, Stephen Terry, Vincenzo Quadrini, Juan Herreño, Devika Chirimar, and Diego Mayorga Cordova for their valuable feedback and suggestions. I received support for this research through a Dissertation Fellowship at the Federal Reserve Board of Governors. The views expressed in this paper are solely my own and do not reflect the views of the Federal Reserve Board or the Federal Reserve System. All errors are my own.

1 Introduction

Renewable ("clean") energy technology subsidies are among the most widely used environmental policies in the United States (US), yet their aggregate efficiency and welfare implications remain uncertain. This paper asks: How do such subsidies affect technology adoption and welfare across the joint distribution of income and wealth in the US? Focusing on residential rooftop solar panel adoption, I quantify both the direct fiscal incidence and the indirect welfare consequences that arise through general equilibrium effects, cost declines, and pollution externalities. I find that uniform refundable subsidies financed by a flat income tax raise aggregate welfare and accelerate adoption, whereas progressive financing or nonrefundable subsidy designs slow diffusion and reduce welfare. When pollution damages are incorporated, the same subsidy becomes universally welfare-improving and progressive, as cleaner air yields the largest gains for more exposed households.

These results challenge the perception that clean energy technology subsidies are inherently regressive. In the US, the top income quintile receives over half of all federal tax credits for residential energy efficiency improvements, according the Internal Revenue Service's (IRS) Statistics of Income (SOI) data. Although these programs are financed primarily by progressive income taxes, they function as transfers from the general taxpayer base to households that are already affluent enough to adopt costly clean technologies. In this sense, the benefit incidence of subsidies appears regressive even if the tax incidence of financing is not.

Yet this static view overlooks two key externalities. First, wealthier households consume more energy and thus are responsible for a larger share of residential emissions. Thus, subsidizing their adoption of clean energy technologies can yield larger emissions reductions. Because poorer households are more exposed to local air pollution, these environmental improvements disproportionately benefit them. Second, early adopters create learning-by-doing spillovers that reduce future installation costs, making adoption more affordable for later – often less wealthy – households. In dynamic general equilibrium, policies that initially appear to favor the rich may therefore yield greater long-run gains for the poor.

To quantify these mechanisms, I develop and calibrate a heterogeneous agent dynamic general equilibrium model with incomplete markets, irreversible technology adoption, endogenous cost reductions from learning-by-doing, and unequal exposure to pollution damages. Using detailed microdata on household income, wealth, and expenditures, installation-level solar costs, and the timing of federal and state policy changes, I parameterize the model and evaluate alternative subsidy designs and financing schemes—comparing refundable, nonrefundable, and income-capped credits under flat versus progressive tax financing. The analysis captures both the aggregate efficiency and heterogeneous welfare implications of policies that subsidize clean technology adoption.

Broadly, prior research falls into two categories: (i) empirical studies on the aggregate and distributional effects of environmental and clean energy policies, and (ii) quantitative macroeconomic analyses of aggregate and heterogeneous effects of climate change policies in general equilibrium. My paper connects these strands by linking empirical evidence on technology diffusion and pollution exposure to a structural dynamic general equilibrium framework that captures both aggregate and heterogeneous welfare effects of clean energy technology subsidies.

Borenstein and Davis (2024) extensively document adoption patterns and the distribution of clean energy tax credits across income groups in the US, showing that participating and subsidy receipt rise sharply with income. Their findings highlight the apparent regressivity of residential clean energy incentives, but do not quantify broader welfare, such as health co-benefits from reducing local air pollution, or general equilibrium effects. Vona (2023) provides a comprehensive overview of the empirical evidence on the multiple effects of climate policies on well-being, emphasizing that clean-energy and energy-efficiency subsidies often reinforce the unequal distribution of benefits when households face borrowing constraints and high up-front adoption costs. Similarly, Levinson (2019) shows in a static framework that taxing energy use would be both more cost effective and less regressive than subsidizing energy efficient appliances or taxing inefficient appliances. These studies underscore the policy relevance of distributional incidence but stop short of modeling the dynamic mechanisms – such as learning-by-doing and endogenous cost declines – that determine how regressivity evolves over time. My analysis builds on this evidence by explicitly modeling these dynamics within a general equilibrium framework.

A growing set of empirical papers examines these mechanisms more directly. The paper by Gao, Rai, and Nemet (2022) provides one of the few empirical analyses of learning-by-doing effects in US residential solar installations, finding that economies of scale reduce both hardware and non-hardware costs. Their results suggest that early adopters generate localized cost spillovers that could make subsequent adoption more affordable for later adopters. On the environmental side, Banzhaf, Ma, and Timmins (2019) review the extensive literature documenting the strong association between ambient air pollution, poverty, and race – the so-called environmental justice gap. Because poorer households tend to live in areas with higher local air pollution, they are likely to gain disproportionately from policies that reduce emissions from residential energy use. These empirical findings motivate the two externalities at the core of my model – learning-by-doing and unequal pollution exposure – and provide the basis for the parameters I use to calibrate the dynamic effects of clean-energy subsidies.

A separate literature develops quantitative macroeconomic models to analyze the aggregate and distributional effects of climate policies in a general equilibrium framework. Most of these papers focus on carbon pricing rather than subsidizing clean energy technologies. Känzig (2023) finds that a restrictive carbon policy shock raises energy prices, reduces emissions, spurs clean innovations, but also lowers economic activity and disproportionately burdens poorer households. Benmir and Roman (2022) examine a carbon pricing path that achieves net-zero emissions in the US by 2050 and show that it induces large redistributions of income and wealth from poorer to richer households. Fried, Novan, and Peterman (2024) analyze how alternative uses of carbon tax revenues affect welfare across and within generations, showing that optimal revenue recycling can mitigate regressivity by reducing distortionary taxes on capital and increasing the progressivity of labor taxation. In earlier work, Fried, Novan, and Peterman (2018) highlight the importance of not only long-run outcomes, but also the transitional welfare effects of how carbon tax revenues are recycled. Using detailed household expenditure and emissions data, Belfiori, Carroll, and Hur (2024) similarly document that low-income households have higher emissions per dollar of spending, making a uniform carbon tax regressive unless revenues are redistributed progressively. Together, these papers identify an efficiency-equity trade-off similar to the one I study, but they do so for carbon pricing. My paper extends this analysis to subsidy-based instruments – the main policy tool used in the US – and examines how dynamic cost declines and pollution heterogeneity reshape that trade-off.

While economists generally view carbon pricing as the most efficient instrument for reducing emissions, in practice – particularly in the US – climate and environmental policies have relied far more on subsidies for clean energy technologies. As Borenstein and Davis (2024) emphasize, relatively little is known about the economic efficiency or the distributional consequences of such subsidy-based approaches. My contribution is to fill this gap by evaluating clean-energy subsidies within a dynamic, heterogeneous agent general equilibrium model that embeds both learning-by-doing and pollution externalities. This approach allows me to quantify how the fiscal design of subsidies shapes aggregate welfare and household-level outcomes simultaneously.

More recent work examines heterogeneity in the adoption of clean energy technologies directly. Kuhn and Schlattmann (2024) develop a quantitative life-cycle model with unequal adoption rates of carbon-neutral goods by income, highlighting the reduction-redistribution trade-off inherent in different policy mixes. Lanteri and Rampini (2025) study clean technology investment by heterogeneous firms facing financial frictions in a dynamic general equilibrium model. They find that constrained firms optimally invest in older, dirtier technologies, generating a positive relationship between firm size and energy efficiency. Their framework provides a natural laboratory for analyzing the distributional effects of environmental policy across firms, although they that exercise for future work. My paper complements these papers by emphasizing the dual role of learning and pollution externalities in shaping adoption dynamics and welfare and by focusing on households rather than firms.

In summary, this paper bridges the micro-level empirical literature on adoption and pollution with the macroeconomic literature on environmental policy in general equilibrium. It contributes by unifying these approaches in a single quantitative framework that links micro evidence on learning and pollution exposure to macro-level welfare outcomes, providing new insights into how clean-energy subsidies affect both efficiency and the distribution of welfare gains.

My analysis yields three main contributions. First, I provide new empirical estimates of localized learning-by-doing in US residential solar panel installations. Using installation-level data merged with state and utility policy shocks – changes in subsidy generosity, eligibility, or program timing that were plausibly exogenous to local installation trends – I find that each doubling of cumulative installed capacity reduces system costs by about 7%, with stronger effects when adoption is policy-driven rather than market-driven. Second, I develop a quantitative heterogeneous agent model that jointly captures private adoption incentives, dynamic cost spillovers, heterogeneous pollution exposure, and general equilibrium feedbacks. Third, I use the model to evaluate the aggregate and distributional welfare consequences of alternative subsidy and financing arrangements.

The results show that uniform refundable subsidies financed by flat labor income taxes raise aggregate welfare and accelerate adoption, whereas progressive financing slows down adoption and reduces welfare gains for liquidity-constrained households by depressing short-run wages and transfers. Nonrefundable tax credits, which mirror the structure of the US federal residential solar credit, further exclude low-income households but do not slow down diffusion. Income-capped subsidies, while intended to improve the distribution of gains, slow down adoption, weaken learning spillovers, and generate aggregate welfare losses that disproportionately affect middle-wealth households. When pollution damages are included, the nonrefundable uniform subsidy becomes universally welfare-improving and strongly progressive, as cleaner air disproportionately benefits households with higher exposure to local pollution. 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. Although the analysis focuses on residential solar adoption, the framework is general and can be applied to study the diffusion of other clean or productivity-enhancing technologies—such as electric vehicles, energy storage, or digital infrastructure—where adoption frictions, learning spillovers, and heterogeneous financing constraints shape both efficiency and equity.

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 present the quantitative results on the distributional and welfare effects of a benchmark uniform subsidy under uniform financing scheme. In Section 6, I conduct policy experiments to evaluate the effectiveness of different policy mixes in achieving the outlined policy objectives and majority support. Finally in Section 7, I conclude and discuss the implications of the results for policy design and implementation.

2 Data and Empirical Motivation

Understanding how energy use, clean-technology adoption, and pollution exposure vary across income groups in the US is crucial for assessing the equity implications of clean energy policies. Table 1 summarizes these distributions for 2015 – the earliest year with complete data – using multiple cross-sectional data sources.

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.[1] 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.

Next, I combine the US Energy Information Agency's (2023) 2015 Residential Energy Consumption Survey (RECS) microdata with income categories aligned to the same quintiles.[2] Weighting by household survey weights, I calculate each group's share of total residential energy consumption and report it in the second row of Table 1. The top quintile accounts for roughly one quarter of total household energy use.

I then compute adoption and subsidy patterns for residential rooftop solar panel deployment across income quintiles. Using the RECS 2015 indicator for on-site solar power generation, I calculate the share of total adopters by income quintile. Adoption rises sharply

[1] For 2015, these limits are $17,929, $35,583, $62,600, $108,429, and the lower limit for top 5% is $146,778.

[2] 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.

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 adoptors	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.

with income, as shown in the third row of Table 1. The top income quintile represents nearly half of all rooftop solar adopters.

To measure the distribution of federal Residential Clean Energy Credits (RCECs), I use the US Internal Revenue Service (2023) 2015 SOI data, aligning income categories as closely as possible with the quintiles above. I calculate each income group's share of total RCEC value by dividing the total amount of RCECs claimed by each group by the total amount of RCECs claimed by all income groups and report it in the fourth row of Table 1. The top quintile receives almost half of all RCEC value. Because these credits are non-refundable – limited by tax liability – they primarily benefit high-income households.

Together, these facts show that higher-income households consume more energy, adopt clean technologies earlier, and capture most federal subsidies. Yet, they also account for a larger share of emissions. Since lower-income households experience greater exposure to local air pollution, subsidizing cleaner technologies for high-income households could still generate progressive environmental benefits.

To document this channel, I merge county-level mortality damages from the Air Pollution Emission Experiments and Policy Analysis (AP4) model, detailed in Dennin et al. (2024), with county-level median income data from the US Census Bureau's (2022) 2017 ACS 5-Year Estimates. Mortality damages per capita decline monotonically with income – from \$4,811 for the bottom quintile to \$2,354 for the top quintile – as shown in the last row of Table 1, confirming that pollution damages are disproportionately borne by poorer households. These stylized facts motivate the model's two central externalities: learning-by-doing and unequal pollution exposure.

2.1 Benefits of Residential Rooftop Solar Panel System Deployment

Deploying solar panels for on-site power generation provides both private and social benefits. Private benefits include reduced electricity bills, increased property values, and

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.

reduced exposure to electricity price volatility. Social benefits include reduced emissions of greenhouse gases and local air pollutants, reduced strain on the electricity grid, and increased energy security.

2.1.1 Private Benefits

Using the 2020 RECS, I regress household electricity expenditure on a solar indicator with state fixed effects and standard controls. Households with on-site solar spend roughly $700 less per year on grid electricity, which is about half the sample average of $1,400 (Table A.1). This difference is purely accounting, not causal, but illustrates magnitude.

For calibration, the key fact is that energy savings are meaningful relative to household budgets. Using the 2015 Consumer Expenditure Survey (CES) by the Bureau of Labor Statistics (2024), I compute the share of residential energy in total expenditure by income quintile. As shown in Table 2, the bottom income quintile devotes 8.5% of total expenditure to energy, compared with only 3.3% for the top quintile. These differences anchor heterogeneous marginal utilities of consumption in the model.

2.1.2 Social Benefits

Social gains stem mainly from avoided local air pollution. In 2023, the residential sector accounted for 15% of end-use energy consumption in the US, according to US Energy Information Administration (2024). Much of this energy is generated from fossil fuels, which emit local air pollutants harmful to human health, such as particulate matter (PM).[3] 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.

[3]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, with diameters that are 2.5 micrometers or smaller, called $PM_{2.5}$, pose the greatest health risks, because they can penetrate deep into the lungs and the bloodstream.

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

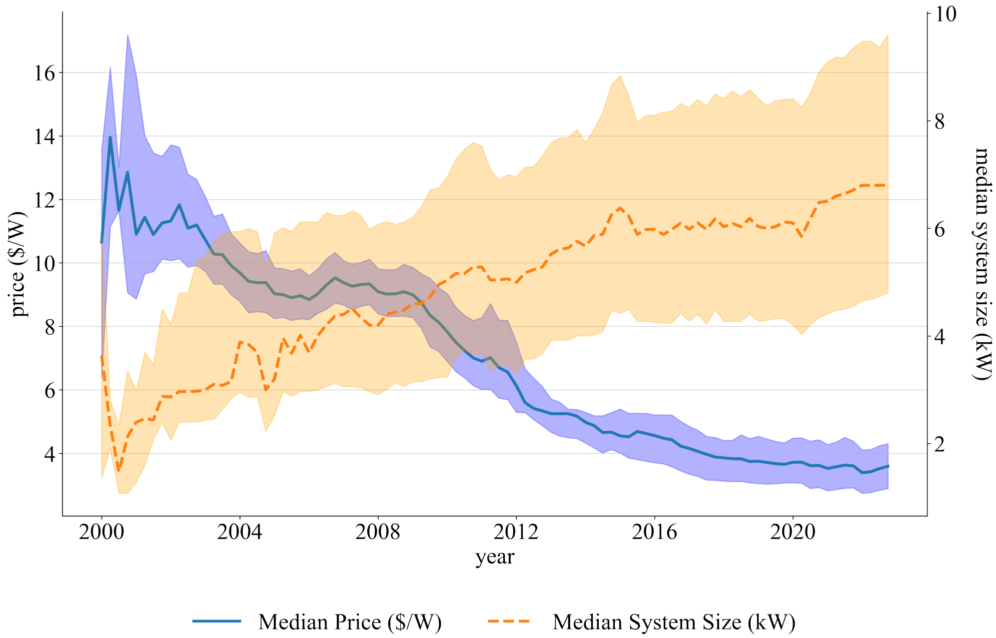

Note: The shaded area represents the 25th and 75th percentiles of the distribution of prices of residential solar panel installations.

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.

Using the National Renewable Energy Laboratory's (NREL) (2023) 2022 Tracking the Sun report data, Figure 1 shows the median installation prices per watt (W) and the median system size of residential solar panels installations in kilowatts (kW) in the US from 2000 to 2022. The figure shows that the median installation price of residential solar panels in the US has declined by almost 65% from 2000 to 2022, 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 2: Ratio of median gross and median net system prices of residential solar PV systems to median annual household income

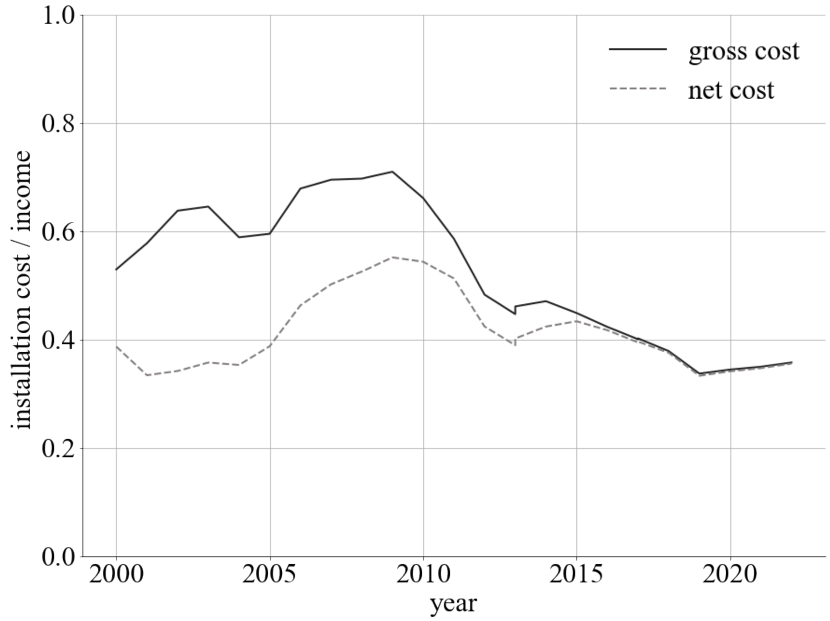

installation price per watt and the median system size. The second measure is the median net total cost, which accounts for the state- and utility-level incentives and rebates deducted from the gross total cost for residential solar installations. Importantly, the net total cost measure does not account 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.

2.3 Did Rooftop Solar Panel Installations Experience Learning Effects?

The idea behind the learning effect is straightforward: as cumulative experience with a technology increases, its costs tend to fall. For technologies with positive externalities — like residential rooftop solar — this creates a dynamic spillover: subsidizing early adopters can lower costs for future adopters. In practice, such learning can come from improved installation techniques, better coordination with permitting and interconnection, streamlined soft costs, and supply-chain efficiencies.

A growing empirical literature documents these learning effects in residential solar. Nemet et al. (2016) show that experienced installers consistently quote lower prices than novice installers, holding system characteristics constant. The cheapest decile of systems is disproportionately installed by firms with extensive prior experience, suggesting accumulated know-how translates into lower prices. O'Shaughnessy (2018) finds that in more concentrated local markets, average installation costs are lower, consistent with high-volume firms moving down their cost curves; however, if markets become too concentrated, reduced competition can offset these gains. Nemet et al. (2020) document significant within-county knowledge spillovers across firms between 2008 and 2014: local cumulative experience lowers installation costs, especially for firms above a size threshold. They also find smaller, but still present, spillovers within firms across counties. Bollinger and Gillingham (2023) estimate that each doubling of installer experience in California reduces soft costs by about \$0.12/W, implying nontrivial but localized learning-by-doing, with relatively weak spillovers across firms. By contrast, Gao, Rai, and Nemet (2022) argue that traditional learning-by-doing is only part of the story once one accounts for "learning-by-searching" (innovation and R&D) and "learning-by-interacting" (supplier networks), suggesting that measured "learning" may bundle several mechanisms.

I test for learning-by-doing in recent US residential rooftop solar installation data by estimating how installation prices respond to cumulative past installations. The baseline learning model assumes that installation costs decline with cumulative experience according to a power law, as formulated by Arrow (1962):

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

where p_t is the net installation price per W at time t (after rebates and incentives), p_0 is the initial price, $\mathcal{I}_t$ is cumulative installed capacity before t (in number of systems or total kW), ξ is the learning-by-doing elasticity, and λ captures exogenous secular cost declines over time (global PV cost improvements, supply chain maturation, etc.) unrelated to local experience. The implied learning rate is $1 - 2^{-\xi}$, the percent cost reduction from a doubling of cumulative installed capacity. Mukoyama (2006) provides a summary of other functional forms used in the learning literature.

Because learning may operate at multiple margins, I allow for both state-level and firm-level experience. Let $s(i)$ denote the state of installation i, and $f(i)$ the installer firm. I estimate the following specification using installation-level data from NREL's (2023) Tracking the Sun data, expanded to include a measure of local incentive generosity $g_{s(i),t(i)}$ generated from state- and utility-level rebate programs data from North Carolina Clean Energy Technology Center's (2025) Database of State Incentives for Renewables & Efficiency (DSIRE):

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

where p_i is the net price per watt of installation i, $\alpha_{j(i)}$ are fixed effects (state, county, or firm, depending on the column), $\mathcal{I}^{\text{state}}s(i),,t-12$ is cumulative installed residential capacity in state $s(i)$, lagged 12 months, $\mathcal{I}^{\text{firm}}f(i),,t-12$ is cumulative capacity installed by firm $f(i)$ lagged 12 months, ξ^{state} and ξ^{firm} are the associated elasticities, $t(i)$ is the installation month, λ is the common time-decay parameter, $g_{s(i),,t(i)}$ is contemporaneous incentive generosity in state $s(i)$, X_i includes installation-level controls (system size, hardware, financing, etc.), and ε_i is the error term. I describe the incentive generosity variable construction in Appendix A.1.

A simple OLS estimate of equation (2) may be biased. Areas (or firms) with lower costs may attract more installations, mechanically generating a negative relationship between price and cumulative experience even without causal learning. Local demand shocks, installer entry/exit, or policy changes could also jointly move both prices and cumulative adoption. To address these endogeneity concerns, I estimate an instrumental variables (IV) version of the model.

The IV strategy uses policy timing shocks to instrument for local cumulative installed capacity. I construct, for each location j (state or county) and month t, a binary "policy shock" variable equal to 1 if a new residential solar incentive (rebate, grant, tax credit, net metering provision) begins in j in month t, and 0 otherwise. These policy onsets are taken from the DSIRE (2025). I lag these shocks by 12 months so that they predict the stock of cumulative installations at t without directly moving prices in t. Appendix A.1.2 details construction of both the policy shock series and the incentive generosity measure.

Intuitively, a state that rolled out a new incentive 12 months ago should have accumulated more installs by now, even if current prices are high, and that boost in cumulative installed capacity may generate learning-driven cost reductions today. The key advantage is that the timing of policy introductions is plausibly exogenous to unobserved local cost shocks in later periods.

The two-stage least squares (2SLS) design is:

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

where $Z_{s(i),t-12}$ is the lagged policy shock indicator, $\mu_{s(i),t-12}$ are fixed effects, $g_{s(i),t-12}$ controls for the generosity of incentives, $W_{s(i),t-12}$ includes additional covariates, and $u_{s(i),t-12}$ is the error term. The first stage links cumulative installed capacity to lagged policy shocks, conditional on fixed effects, current incentive generosity, and time trends.

The second stage replaces $\mathcal{I}^{\text{state}}$ in equation (2) with its fitted values:

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

where $\widehat{\mathcal{I}^{\text{state}}}$ are the predicted cumulative installations from equation (3). In both stages I include contemporaneous incentive generosity $g_{j,t}$ so that any direct price effect of current subsidies is partialled out. Identification then comes from the timing of past policy introductions, not from the current level of subsidies.

The exclusion restriction is that, conditional on current incentive generosity, fixed effects, and common time trends, lagged policy onsets affect current installation prices only through their impact on cumulative installed capacity (i.e., learning-by-doing), and not through any direct price subsidy in period t. This is most credible when policies are primarily adoption incentives rather than direct per-watt price buydowns at the time of installation. I therefore interpret the IV estimates as the causal effect of cumulative installed capacity on prices under policy-driven expansion, rather than under purely organic market growth.

Table 3 reports the main results. Columns (1), (3), and (5) present OLS estimates with different fixed effects and columns (2), (4), and (6) present the corresponding IV estimates. State fixed effects are used in columns (1)–(4), and county fixed effects in (5)–(6). The first-stage F-statistics in the IV columns (2), (4), and (6) exceed 10, indicating strong instruments.

The estimates provide mixed but broadly supportive evidence of learning-by-doing. Location-level cumulative installations (state or county) generally exhibit positive, statistically significant elasticities: higher cumulative capacity is associated with lower prices. The magnitudes of these elasticities imply learning rates between roughly 1 and 8% cost reduction per doubling of cumulative capacity, depending on the specification.

A key pattern is that IV estimates of state-level learning elasticities are substantially larger than their OLS counterparts. In columns (2) and (4), the IV coefficient on lagged cumulative state capacity is roughly four times the OLS coefficient in columns (1) and (3). At the county level, the IV estimate in column (6) is nearly six times the OLS estimate in column

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
log Cumulative Installs (Firm, 12m lag)	−0.0073	−0.0073	−0.0097	−0.0097	−0.0073	−0.0073
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
log Cumulative Installs (State, 12m lag)	0.0226	0.1027	0.0244	0.0996		
	(0.0022)	(0.0054)	(0.0022)	(0.0055)		
log Cumulative Installs (County, 12m lag)					0.0145	0.1249
					(0.0014)	(0.0055)
t	0.0029	0.0029	0.0030	0.0030	0.0031	0.0031
	(0.0001)	(0.0000)	(0.0001)	(0.0000)	(0.0000)	(0.0000)
Policy Generosity $\times 10^{-4}$	−0.0034	−0.0034	−0.0034	−0.0034	−0.0091	−0.0091
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0005)	(0.0005)
Has DC Optimizer	0.0303	0.0303			0.0347	0.0347
	(0.0022)	(0.0022)			(0.0022)	(0.0022)
Ground Mounted	0.0229	0.0229			0.0308	0.0308
	(0.0053)	(0.0069)			(0.0054)	(0.0069)
Has Microinverter	0.0098	0.0098			0.0109	0.0109
	(0.0021)	(0.0024)			(0.0022)	(0.0024)
Inverter Loading Ratio	0.0418	0.0418			0.0425	0.0425
	(0.0038)	(0.0040)			(0.0038)	(0.0040)
log Size	−0.1040	−0.1040			−0.0897	−0.0897
	(0.0016)	(0.0019)			(0.0016)	(0.0020)
Has Tracking Bin	−0.3999	−0.3999			−0.4023	−0.4023
	(0.0089)	(0.0298)			(0.0089)	(0.0298)
First stage						
Policy Shock (State, 12m lag)		−0.7520		−0.7518		
		(0.0070)		(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
R^2	0.060	0.013	0.053	0.013	0.058	0.008
Location FE	state	state	state	state	county	county

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 $\log(\mathcal{I}^{\text{firm}}_{f(i),t-12})$, $\log(\mathcal{I}^{\text{state}}_{s(i),t-12})$, and $\log(\mathcal{I}^{\text{county}}_{c(i),t-12})$. "Policy Shock (State/County, 12m lag)" equals 1 if a new residential PV incentive began in that entity 12 months earlier.

(5). This suggests that OLS attenuates the true learning effect, likely due to simultaneity and measurement error: places with already-low prices attract adoption even absent new learning, biasing OLS downward. By contrast, policy-driven adoption (the IV source of variation) appears to generate stronger subsequent cost declines.

My preferred estimate is the IV specification with state fixed effects in column (2). It addresses endogeneity using policy timing shocks, absorbs persistent state-level heterogeneity, and retains meaningful cross-time variation in cumulative installations. The implied state-level learning elasticity of 0.1027 corresponds to a learning rate of about 7% per doubling of cumulative capacity. This value disciplines the model parameter ξ, and the estimated common time-decay term λ (around -0.003 per month) implies secular cost declines of roughly 4% per year that are not driven by local adoption. As expected, higher contemporaneous incentive generosity $g_{s(i),t(i)}$ is associated with lower prices, though the magnitudes are modest. Other controls behave as expected: larger systems are cheaper per watt, and higher-end technologies (e.g. optimizers, microinverters) are more expensive.

In summary, the regression evidence supports three conclusions. First, there is economically meaningful learning-by-doing at the state level: cumulative local adoption lowers prices for future adopters. Second, this effect is stronger when cumulative installations are driven by policy shocks rather than by purely endogenous market expansion. Third, secular cost declines unrelated to local adoption remain important. These findings align with the broader literature and provide quantitative discipline for the model: they pin down how much today's adopters lower tomorrow's prices for everyone else.

Taken together, these empirical patterns provide the foundation for the model developed in the next section. The descriptive evidence shows that (i) clean-technology adoption and subsidy benefits are highly skewed toward high-income households, (ii) poorer households face higher exposure to pollution damages, and (iii) installation costs decline with cumulative experience, consistent with learning-by-doing. These observations highlight two externalities – technological learning and unequal pollution exposure – that shape the equity and efficiency of energy transition policies but are not captured in static incidence analyses. To quantify their joint implications for adoption dynamics and welfare, I now develop a heterogeneous agent dynamic general equilibrium model with incomplete markets, irreversible clean technology investment, endogenous cost declines, and pollution damages.

3 Model

Motivated by these empirical observations, I develop a heterogeneous agent dynamic stochastic general equilibrium (DSGE) model with incomplete markets in the Bewley-Huggett-Aiyagari tradition, building on Bewley (1977), Huggett (1993), and Aiyagari (1994), augmented to include costly clean energy technology adoption and environmental externalities. Time is discrete and the horizon is infinite. There is a continuum of infinitely-lived households. Households supply labor, earn wage income, accumulate assets in physical capital, and rent capital to firms. Labor income is stochastic due to an idiosyncratic productivity shock. Households self-insure by saving 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 good consumption. This implies (as in Table 2) that low-consumption households spend a higher share of their budget on energy. Households can meet their energy needs using one of two technologies: an old fossil (dirty) technology, and a new renewable (clean) technology. Energy from the two technologies is a perfect substitute. Dirty energy use generates ambient air pollution, which reduces utility and creates a negative externality. Before the clean technology becomes available, all households use dirty energy and the economy is in the initial steady state.

Once the clean technology becomes available, households may switch by making a one-time investment (e.g., installing rooftop solar panels). After adoption, the household uses clean energy permanently. Clean energy has a lower per-unit cost than dirty energy, but adoption requires a significant, one-time, and irreversible upfront expenditure. That expenditure cannot be recovered. The effective adoption cost declines over time as more households adopt, due to learning-by-doing spillovers. In the long run, all households adopt the clean technology and the economy converges to a new (terminal) steady state.

Firms combine capital and labor to produce a final good used for both consumption and investment. The externality from dirty energy use is not internalized by firms or consumers, and instead reduces effective output through pollution damages. This welfare loss is larger in utility terms for poorer households, given diminishing marginal utility of consumption. Both firms and households take prices as given. The government taxes labor income and uses the revenue to subsidize the clean technology's adoption cost, rebating any surplus as a lump-sum transfer.

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:

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

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:

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

where $B(\mathcal{A})$ is 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 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 good consumption, strictly decreasing and convex in pollution, 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 $e_t = e(c_t)$ maps 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 adoption choice S_t is the main deviation from a standard Aiyagari model. S_t is chosen at the start of the period, is irreversible, and permanently upgrades the household to the clean technology. The law of motion is:

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

Irreversibility is a natural first approximation to lumpy household-level adoption (e.g., rooftop solar). Extensions could allow depreciation or replacement.

3.2 Producers

A unit mass of competitive firms produces 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. Firms take factor prices (r_t, w_t) as given and maximize static profits each period.

3.3 Government

The government taxes labor income at rate τ_t^ℓ, part of the revenue to subsidize clean-technology adoption at rate τ_t, and rebates the remainder uniformly as T_t to each household in each period t. Thus, the government budget constraint is:

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

3.4 Ambient Air Pollution

Ambient air pollution, X_t, is determined by the flow of energy in period t with the following mapping:

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

where $\Omega(\cdot)$ is increasing. Pollution is thus a within-period externality: only energy consumed via the dirty technology generates X_t, and the disutility is contemporaneous.

3.5 Learning-by-Doing Spillover

The one-time adoption cost p_t declines with cumulative adoption due to learning-by-doing spillover. As in subsection 2.3, I assume a power-law learning function. Cumulative adoption before period t, denoted by Z_t, is given by:

$$Z_t = \int_{\mathcal{Z}} s_t d\Phi_t, \tag{7}$$

and include an exogenous time decay component as in equation (1). The adoption cost function is:

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

where $p_0 > 0$ is the initial adoption cost, $\xi > 0$ is the learning-by-doing elasticity that captures the rate of cost reduction with each doubling of cumulative adoption, and $\lambda > 0$ captures the secular cost declines unrelated to local spillovers, such as global supply chain improvements and technological change. As more households adopt the clean technology, p_t falls for future adopters. This spillover makes adoption socially beneficial beyond the private gain and means that subsidies do more than redistribute cash: they accelerate cost declines for later adopters.

3.6 Market Clearing

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

$$K_t = \int_{\mathcal{Z}} a_t d\Phi_t,$$
$$L_t = \int_{\mathcal{Z}} \ell_t d\Phi_t$$

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

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

where δ is the depreciation rate of aggregate capital stock.

3.7 Formulation

The model admits a recursive formulation. A household's decision problem depends on its individual state $z_t = (a_t, \ell_t, s_t)$ and on the aggregate distribution Φ_t. Φ_t is the only aggregate state: X_t is a contemporaneous flow determined by current dirty energy use (and therefore by Φ_t), so it does not enter as an independent state. Formal definitions of the initial and terminal stationary competitive equilibria are provided in Appendix B.1.

3.7.1 Transitional Dynamics

The main object of interest is the transition from the initial steady state with $s_t = 0$ for all households (no clean technology) to the terminal steady state with $s_t = 1$ for all households (full adoption). During the transition, households choose the irreversible adoption decision $S_t \in 0, 1$ such that

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

The Bellman equation for a household that has not yet adopted (state $(a_t, \ell_t, 0)$ at the start of period t) is:

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

where $\Gamma_t : \mathcal{M} \to \mathcal{M}$ 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.

After adoption, the household is in state $(a_t, \ell_t, 1)$ and solves:

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

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 all constraints.*

2. ***Factor prices.***

$$r_t = F_K(K_t, L_t),$$
$$w_t = F_L(K_t, L_t).$$

3. ***Government budget constraint.***

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

4. ***Ambient air pollution.***

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

5. ***Cumulative adoption.***

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

6. ***Adoption cost.***

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

7. ***Market clearing.***

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

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

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 key innovations of the model are: (i) heterogeneous utility damages from pollution, (ii) an irreversible binary household adoption decision with subsidized upfront cost, (iii) a learning-by-doing spillover that lowers that cost as cumulative adoption rises. These features allow me to evaluate the welfare and distributional consequences of clean-energy subsidies along more than a simple transfer margin: subsidies both reallocate resources and accelerate cost declines for future adopters. The model also embeds heterogeneous pollution damages, allowing welfare comparisons across the joint income and wealth distribution.

The next section describes the calibration used in the quantitative analysis. Standard parameters follow the literature; parameters governing adoption costs, spillovers, and pollution are disciplined using microdata and the reduced-form estimates above.

4 Quantitative Analysis

Having laid out the structure of the model, I now turn to its quantitative implementation. The goal is to evaluate the distributional and welfare effects of clean energy subsidies by calibrating the model to match key features of the US economy and residential energy sector. 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.

4.1 Functional Forms

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:

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

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 demand is affine in goods consumption:

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

where $\eta_0 > 0$ captures baseline energy needs (e.g., grid connection) and η_1 the marginal energy intensity of 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.

Ambient air pollution is linear in aggregate dirty energy use:

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

where γ is the pollution intensity of dirty energy use.

4.2 Calibration

Model parameters are divided into three groups: (i) standard macroeconomic parameters; (ii) parameters matched to data moments; and, (iii) parameters estimated empirically. Each model period corresponds to one year. The initial steady state represents the US economy in 2000, before large-scale rooftop solar adoption. Table 4 summarizes all baseline parameters.

4.2.1 Baseline Economy (No Pollution Disutility)

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.

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

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	$0.001 \to 0.15$ (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)

(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. The total factor productivity of goods production, A, is normalized to 1. 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 (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 $a' \geq \underline{a}$. 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).[4] Under my baseline parameterization, 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 set $\underline{a} = -0.5$, which corresponds to a debt-to-income ratio of approximately 31% in the initial steady state, consistent with the 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 matches empirical moments from household energy use, residential solar installation data, and macroeconomic aggregates. Unit 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 A.2. 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 to 15% by the terminal steady state, matching the observed residential solar adoption growth in the RECS (2023).

The flat 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 are rates.[5] 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.[6] 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 C.1: This progressive tax

[4] In 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.

[5] See https://www.irs.gov/pub/irs-prior/i1040tt--2000.pdf.

[6] See https://www.cbo.gov/sites/default/files/108th-congress-2003-2004/reports/08-29-2003AverageTaxRates.pdf.

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 2025, as summarized in the Lane (2025).

The initial installation cost p_0 is set so that the model reproduces the largest ratio of median system price to median household income in the observed sample, which is approximately 70%. I obtain the median gross 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 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. This calibration anchors the affordability of adoption in observed 2000 conditions.

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 IV 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:

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

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 elasticitiy 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

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

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

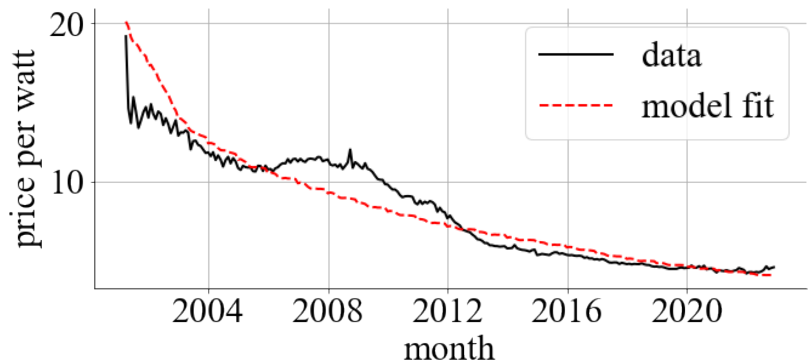

and the (unit-free) cost index

$$\mathrm{idx}_t \;=\; \left(\frac{\mathrm{eff}_t}{\mathrm{eff}_0}\right)^{-\xi} \exp\bigl(-\lambda\,\tilde{t}\bigr), \qquad \tilde{t} \equiv \min\{t,\,99\}. \tag{12}$$

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 $\mathrm{idx}_t \leftarrow \max\{\mathrm{idx}_t, 10^{-12}\}$ to maintain positivity.

Given a baseline level p_0, the adoption cost path is

$$\mathrm{p}_t \;=\; \mathrm{p}_0 \times \mathrm{idx}_t, \tag{13}$$

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

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

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 $\iota : \mathcal{Z} \to \mathbb{R}_+$ and quantile cutoffs $\{\kappa_j\}_{j=0}^{5}$ such that

$$\Phi\big(\{z :\ \iota(z) \le \kappa_j\}\big) = \frac{j}{5}, \qquad j = 0, 1, \ldots, 5,$$

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

$$Q_j \ :=\ \{\, z \in \mathcal{Z} :\ \kappa_{j-1} < \iota(z) \le \kappa_j \,\}, \qquad j = 1, \ldots, 5.$$

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

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

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:

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

where $\omega_j = 1$ by default (equal weighting); when available, I set $\omega_j = 1/\widehat{\sigma}^2_{Q_j}$ 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

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

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 $\eta_1 = 0.74$.

4.2.2 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 $PM_{2.5}$ (µg/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 $PM_{2.5}$ in year 2000 using the baseline model's implied dirty-energy integral with pollution preferences shut down. The baseline $PM_{2.5}$, denoted by X_0, is obtained from the EPA (2024a) and is 13.52 µg/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 $PM_{2.5}$ standard, equal to 9 µg/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 $PM_{2.5}$ 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:

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

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 µg/m^3 reduction in $PM_{2.5}$ at the baseline pollution level. Since baseline $\bar{c}$ equals c_0^{med}, this simplifies to:

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

I calculate the MWTP per 1 µg/m^3 reduction in $PM_{2.5}$ at the baseline pollution level using: (i) a long-term all cause mortality $PM_{2.5}$ concentration-response (C-R) of 7% per 10 µg/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:

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

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:

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

where $\widehat{\psi}_g^D$ is the share of total mortality damages borne by income tercile g from Dennin et al. (2024), and $\widehat{\psi}_g^C$ is the share of total consumption expenditure by income tercile g from the 2014 CES. The empirical pollution burden ratio is approximately $\widehat{\Upsilon} = 5.83$, 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:

$$\omega^{\star} = \arg\min_{\omega \geq 0} \left(\Upsilon(\omega) - \widehat{\Upsilon} \right)^2.$$

The implied ω is $\omega = 2.84$.

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.

4.3 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.

Let c_t^{base} and c_t^{cf} denote the consumption paths under the baseline and counterfactual scenarios, respectively, for $t = 1, \dots, T$. Formally, EV and CV, denoted by λ^{EV} and λ^{CV}, and lifetime utility change, denoted by ΔV, are defined as follows:

$$\begin{aligned}
\text{EV:}\quad & \mathbb{E}_t\left[\sum_{t=1}^{T}\beta^t U\big((1+\lambda^{EV})\,c_t^{\text{base}}\big)\right] = \mathbb{E}_t\left[\sum_{t=1}^{T}\beta^t U\big(c_t^{\text{cf}}\big)\right],\\
\text{CV:}\quad & \mathbb{E}_t\left[\sum_{t=1}^{T}\beta^t U\big(c_t^{\text{base}}\big)\right] = \mathbb{E}_t\left[\sum_{t=1}^{T}\beta^t U\big((1+\lambda^{CV})\,c_t^{\text{cf}}\big)\right],\\
\Delta\text{V:}\quad & \Delta V \equiv \mathbb{E}_t\left[\sum_{t=1}^{T}\beta^t U\big(c_t^{\text{cf}}\big)\right] - \mathbb{E}_t\left[\sum_{t=1}^{T}\beta^t U\big(c_t^{\text{base}}\big)\right].
\end{aligned}$$

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. In my analysis, I only report EV as the main welfare metric, as it provides a clear interpretation in terms of consumption changes. The computation of the EV metric is described in detail in Appendix C.3.

In sum, the calibration aligns the model with observed macro aggregates, household energy data, and empirically estimated cost dynamics. The computational procedure then traces out equilibrium transitions consistent with individual optimization and general equilibrium feedbacks. This foundation enables the next section's analysis of how subsidy design and financing shape adoption, inequality, and welfare in the clean energy transition.

5 Quantitative Results

This section presents the quantitative results of the model. I begin by evaluating how well the initial stationary equilibrium reproduces salient features of the joint income-wealth distribution observed in the data, validating the model as a credible tool for policy analysis. I then introduce a uniform adoption subsidy for clean energy technologies and trace its effects on adoption incentives, prices, and welfare across heterogeneous households. These baseline results provide a benchmark for understanding the equity and efficiency consequences of adoption subsidies before introducing pollution externalities.

5.1 Model Fit

Before turning to the welfare effects of subsidies, I assess the model's fit to the US wealth distribution in 2000 using PSID data. Table 5 compares the model's stationary equilibrium to the data across net-worth quintiles for income, expenditure, and wealth shares, as well as the ratio of energy expenditure to total expenditure.

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.

The model reproduces the joint distribution of income, expenditure, and wealth reasonably well. It slightly overpredicts the income and expenditure shares of the lower quintiles and underpredicts those of the top quintile, but the discrepancies are modest. As in many heterogeneous agent models, the wealth concentration of the top quintile is somewhat underrepresented; this could be improved with heterogeneity in discount factors or returns. The model matches energy expenditure shares across the net-worth distribution closely, indicating that it captures observed consumption patterns well.

5.2 Baseline Results

To evaluate the welfare implications of the clean-energy transition, I compute welfare changes for each household along the transition path. All results in this section abstract from pollution disutility to isolate the macroeconomic and adoption channels. Appendix C.3 provides definitions of the welfare metrics used.

Aggregate Dynamics

Figure 4 plots aggregate dynamics under a uniform labor income tax with learning-by-doing (LBD) in adoption costs, comparing scenarios with (blue) and without (orange) adoption subsidies. The panels show: (a) cumulative adopters, (b) the technology-cost-to-median-income ratio, (c) the capital-labor ratio, and (d) lump-sum transfers.

The capital-labor ratio at the terminal steady state is 12% lower than in the initial steady state capital-labor ratio, driven by reduced energy use due to efficiency gains. This implies a 4% fall in the real wage and a 24% rise in the real interest rate. The 50% decline in energy prices increases consumption by 4.9% for the poorest and 1.6% for the richest households. Although aggregate capital and output fall, aggregate consumption rises because energy becomes much cheaper.

The transition dynamics are non-monotonic and differ sharply with subsidies. Without subsidies, the capital-labor ratio initially rises as households save to finance adoption. This raises wages and lowers interest rates, benefiting labor-reliant households and hurting savers. As adoption spreads and energy costs fall, the capital-labor ratio eventually declines, reversing those short-run price effects.

The aggregate path of adoption has an S-shape, consistent with diffusion of innovations theory discussed by Rogers (2003). Adoption starts slowly, as only high-income, high-wealth households can afford the upfront costs. As adoption costs fall through learning-by-doing, more households find it optimal to adopt, accelerating the process. Eventually, adoption saturates as most households that can benefit have already adopted, and the stock of new adopters tapers off.

Subsidies alter this pattern by lowering adoption costs and reducing the need to save beforehand. Early adoption and higher savings among adopters raise the capital-labor ratio more strongly and for longer, about 15 years. Higher wages from this expansion benefit all households, while the lower interest rate hurts the asset-rich but helps borrowers. Although higher wages increase tax revenues, financing the subsidies reduces net transfers substantially, hurting especially the poor, who depend more on transfers. Overall, the transition exhibits rich, non-monotonic dynamics in wages, interest rates, and transfers, all of which shape heterogeneous welfare outcomes.

Aggregate Welfare Effects

First, I evaluate the welfare effects of the technology transition in the absence of any policy change. This counterfactual compares an economy that remains permanently in its initial steady state to one that undergoes the clean-technology diffusion driven solely by LBD and exogenous cost decline, without subsidies. Figure 5 shows the equivalent-variation (EV) welfare impacts of this technology transition across the joint income-wealth distribution. The welfare effects of the technology transition are monotonically increasing in wealth: wealthier households gain more, while the less wealthy households gain less, and even some experience small welfare losses. The distributional pattern arises because without subsidies, aggregate savings decline sharply during the medium term of the transition and recover only slowly. This depresses wages and raises interest rates, which hurts labor-reliant poor households and benefits capital-reliant rich households. Moreover, the fall in wages also reduces government tax revenues, leading to lower transfers that disproportionately affect the poor.

Figure 6 shows the EV welfare impacts of subsidizing clean energy technology adoption across the joint income-wealth distribution. On average, most groups gain, but welfare gains rise with income and wealth: only the poorest households experience losses. Middle-income

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

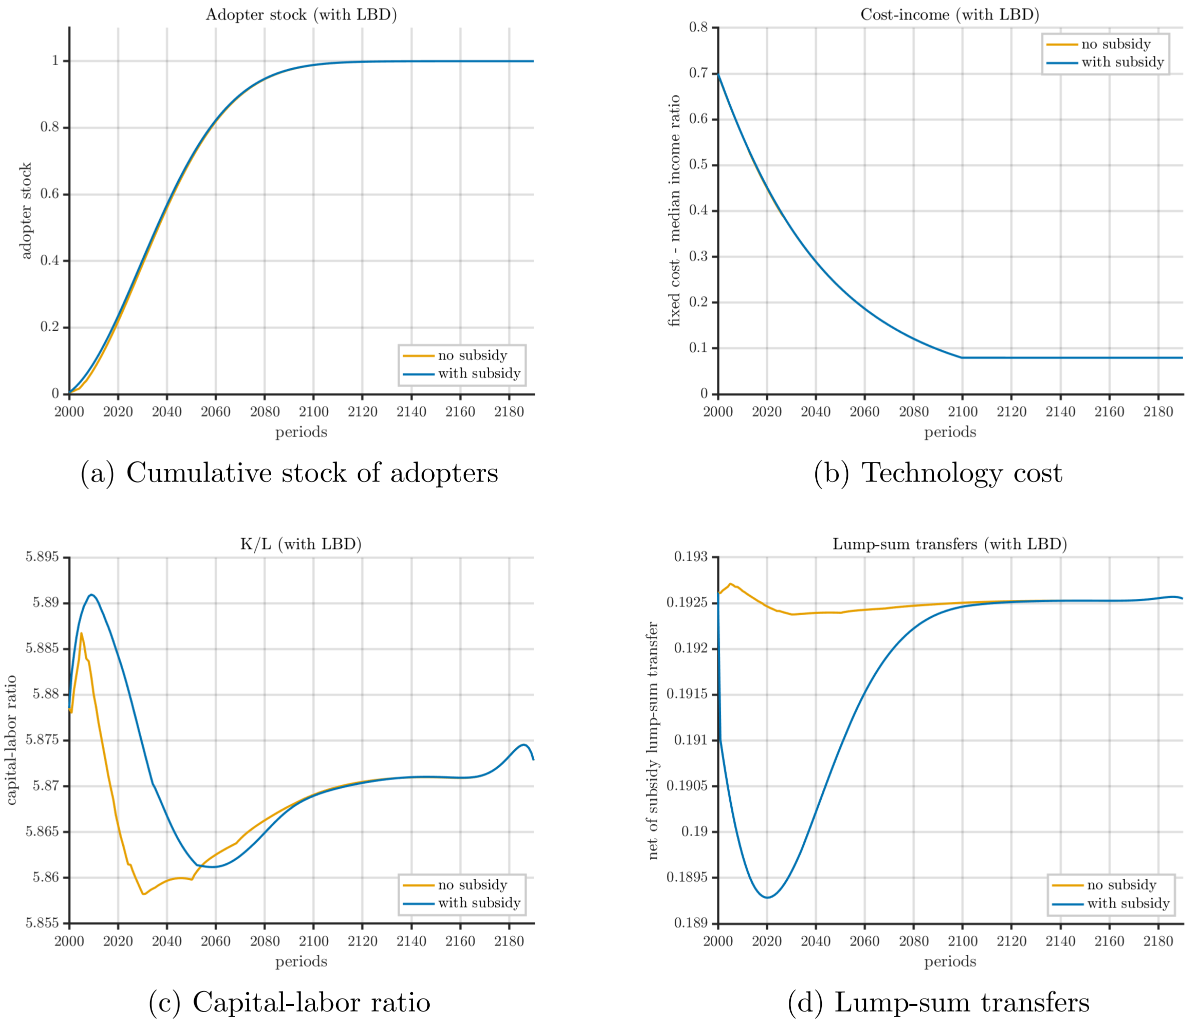

(a) Cumulative stock of adopters

(b) Technology cost

(c) Capital-labor ratio

(d) Lump-sum transfers

groups see the largest average gains.

While Figures 5–6 illustrate the distributional patterns of welfare changes, it is also useful to compare their aggregate magnitudes. In the absence of any policy, the clean-technology transition alone raises aggregate consumption-equivalent welfare by 0.3 percent relative to an economy that remains in the initial steady state. Introducing the adoption subsidy on top of this transition yields an additional 0.05 percent gain. Thus, the subsidy accounts for roughly one-sixth of the total welfare improvement associated with the clean-energy transition. Although smaller in aggregate magnitude, the policy plays a pivotal role in broadening participation in the transition: it mitigates welfare losses among lower-wealth households and shifts part of the gains from high-wealth to middle-income groups.

Table 6 presents a detailed decomposition of the welfare changes induced by adoption

Figure 5: Welfare impact of transitioning to clean energy technology across the income-wealth distribution, measured by equivalent variation (EV) as a percentage of initial consumption.

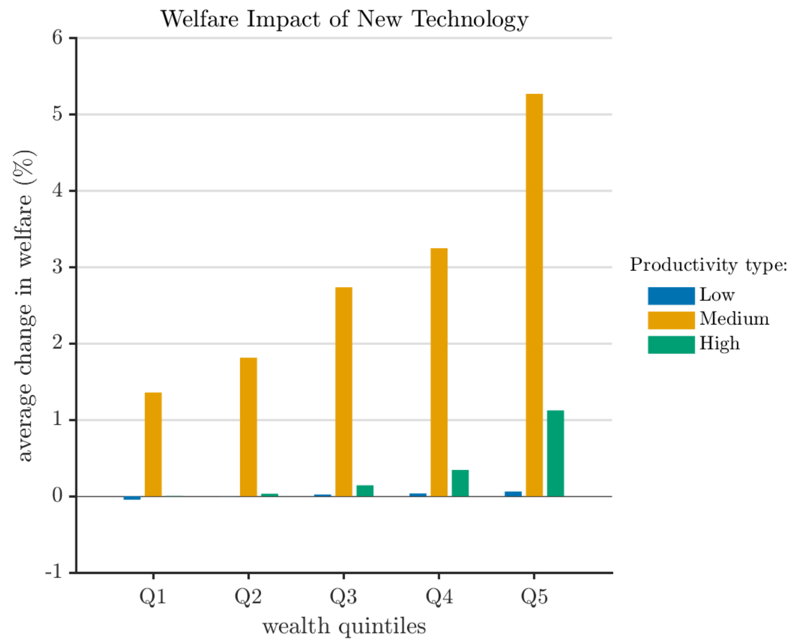

Figure 6: Welfare impact of subsidizing solar panel adoption cost across the income-wealth distribution, measured by equivalent variation (EV) as a percentage of initial consumption.

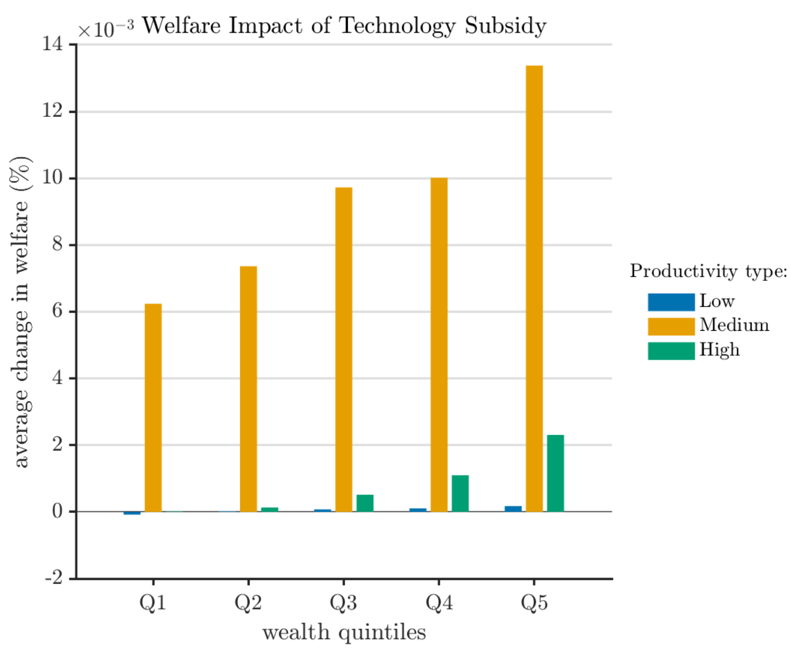

subsidies, breaking down the contributions from direct subsidies, LBD-induced cost reductions, price effects, and fiscal transfer effects. Before I discuss the welfare decomposition, I first highlight the total welfare effects and their distributional patterns across the wealth terciles.

Panel A of Table 6 summarizes the EV metric for each welfare changing component of the subsidy. I provide the definitions of these welfare metrics in Appendix C.3. The last row of Panel A of Table 6 summarizes the aggregate welfare change across all households. On average, I find that the subsidy leads to a small welfare gain across all households, with 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 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: (i) the direct subsidy effect, which captures the immediate benefit to adopters from the subsidy; (ii) the LBD-induced cost change effect, which reflects how the subsidy accelerates adoption and thereby reduces future technology costs for all households; (iii) the price effect, which accounts for changes in equilibrium prices (wages, interest rates, energy prices) induced by the subsidy; and (iv) 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 cost effects are positive for essentially all households. The subsidy lowers the upfront cost of adoption directly, while accelerated adoption generates spillovers that reduce future costs. On average, these two channels raise welfare by EV gains of 0.24% and 0.004%, respectively. The latter is small in magnitude, consistent with the fact

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

Panel A. Aggregate Welfare Decomposition

Component	Avg. EV (%)	Winner (%)
Direct subsidy	+0.24	100.0
LBD-induced cost change	+0.00	99.5
Price effect	−0.02	39.3
Transfer effect	−0.17	0.0
Total	+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 shares are [0.360, 0.323, 0.317].

that LBD affects future adopters more than current ones. Panel B shows that these gains are broadly shared across the asset distribution: all terciles weakly benefit from both the direct subsidy effect and the LBD cost effect.[7]

The price effect introduces heterogeneity. On average, it reduces welfare by an EV change of −0.02% and strictly benefits only 39.3% of households. Panel B shows that this channel is sharply distributional: the winner share is 92.9% in the bottom asset tercile, 18.2% in the middle tercile, and 0.0% in the top tercile.

This pattern reflects how factor prices move during the transition. As shown in Figure 4, subsidies accelerate adoption and raise the capital-labor ratio more sharply than in the no-subsidy counterfactual. In the short run this pushes up wages and lowers the interest rate. Low-asset households rely primarily on labor income and therefore benefit from higher wages. High-asset households rely more on capital income and are hurt by lower returns. Thus, the price effect is progressive in incidence: it favors the asset-poor and penalizes the asset-rich.

The transfer effect is negative for everyone. Because the subsidy is financed out of the same tax base, it reduces lump-sum transfers even though rising wages increase tax revenue. On average, this channel lowers welfare by an EV change of −0.17%, and no households gain from it (winner share 0%). The distributional bite of this channel is most severe for the

[7] For instance, the top asset tercile has a 98.3% winner share from the LBD channel. Winner share is defined as the fraction of households with strictly positive welfare gains; the remaining 1.7 are indifferent, not worse off.

asset-poor, who rely more on transfers for consumption. This is visible in Panel B: despite benefiting from the price effect, low-asset households still face weaker overall support once reduced transfers are taken into account.

Aggregating all four channels, the total welfare effect of the subsidy is positive on average. The direct subsidy effect is the dominant contributor to this average gain and offsets the strong negative effect coming from reduced transfers. Learning spillovers (the LBD cost effect) are present but quantitatively modest at baseline. At the same time, the fiscal incidence of the policy is not neutral: low-wealth households bear a relatively larger share of the financing burden through reduced transfers, while high-wealth households are more exposed to the fall in interest rates.

In summary, the subsidy delivers broad gains through lower adoption costs and, to a lesser extent, through learning spillovers. But those gains are partially offset by two general equilibrium forces: a transfer channel that hurts everyone (especially the asset-poor who depend on transfers) and a price channel that hurts the asset-rich (through lower capital returns). The net result is that the subsidy enjoys majority support, but the burden of financing it is regressive in the sense that low-wealth households pay more, relative to their resources, via foregone transfers.

Importantly, these results come from a simplified baseline environment: subsidies are refundable, labor income taxation is flat, and pollution externalities are turned off. In the next section, I relax these assumptions—introducing progressive taxation, nonrefundability, and pollution damages—and examine how each changes both aggregate welfare and the distribution of winners and losers.

6 Sensitivity Analysis

The baseline results evaluated the welfare effects of subsidizing clean energy adoption under a simplified policy: a uniform labor income tax financing a permanent, refundable adoption subsidy. Real-world policies are more complex. The US tax system is progressive: the federal residential solar investment tax credit has historically been nonrefundable and temporary and environmental policy is motivated in part by pollution damages, which the baseline abstracted from.

This section extends the analysis along two dimensions. First, Section 6.1 studies alternative financing and subsidy designs that mirror US policy practice. Second, Section 6.2 incorporates pollution damages into utility, allowing me to quantify gains from improved air quality and to revisit the distributional incidence of subsidies.

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 differ 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. Importantly, these three experiments (progressive, nonrefundable, and income-capped) are all solved in the same progressive-tax environment, so their initial and terminal steady states are comparable to one another, but not to the benchmark with a uniform tax. In this progressive financing environment, the policy-free energy transition (no subsidy) also raises aggregate welfare by about 0.3%, so the marginal subsidy effects reported below can be read against a transition of a similar order of magnitude as in the benchmark.

Experiment 1: Progressive Labor Income Tax Financing

I first repeat the baseline subsidy experiment, but replace the uniform labor income tax with a progressive labor income tax schedule. Table 7 presents the detailed welfare decomposition for this experiment, analogous to Table 6 for the benchmark. Note that both the initial and terminal steady states differ under progressive taxation, even before introducing subsidies: effective tax rates shift savings, adoption timing, and thus the baseline against which subsidies are evaluated.

Under progressive financing, subsidizing adoption still raises aggregate welfare, with an average EV gain of 0.05%, the same as under uniform financing. However, two changes emerge. First, the aggregate winner share falls from 93.9% under uniform financing to 83.3%. Second, the decline in support is concentrated among low-wealth households: within-tercile winner shares fall most in the bottom tercile.

To understand this, I compare adoption dynamics in Figure 7. Panels (a)–(b) show that aggregate adoption without subsidies is slower under progressive financing. Higher-income (and typically higher-wealth) households face higher effective tax rates, have less liquidity, and adopt more slowly. Lower-income households face lower effective tax rates, but they are not the early adopters in any case. As a result, early aggregate adoption is delayed.

When subsidies are introduced, adoption accelerates under both financing schemes, and the aggregate adoption paths converge. But the composition differs. Under progressive financing,

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

Panel A. Aggregate Welfare Decomposition				
Component		Avg. EV (%)		Winner (%)
Direct subsidy		+0.23		100.0
LBD-induced cost change		+0.00		99.8
Price effect		−0.02		32.3
Transfer effect		−0.17		0.0
Total		+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 shares are [0.335, 0.343, 0.322].

subsidies disproportionately speed up adoption among middle- and high-wealth households, who were previously slowed by higher tax burdens. Figure 7, panels (c)–(d), shows that adoption across wealth terciles becomes more similar once subsidies are present.

These adoption patterns map into different general equilibrium effects. Figure 8 compares aggregate savings (through the capital-labor ratio) and transfers with and without subsidies under progressive financing.

Relative to uniform financing (Figure 4), progressive financing produces a smaller short-run increase in the capital-labor ratio once subsidies are introduced (panel (a)). The intuition is that under progressive taxation, early adopters face higher marginal tax rates; after adopting, they save less than they would under uniform financing. The muted rise in the capital-labor ratio means wages rise by less in the short run. Since low-wealth households rely more on labor income, they benefit less. This weaker wage response is the first reason support falls among the bottom tercile.

At the same time, net lump-sum transfers fall more under progressive financing (panel (b)). Even though wages rise somewhat and generate tax revenue, financing the subsidy consumes a larger share of that revenue, so the net-of-subsidy transfer drops more than under uniform financing. Because low-wealth households depend more on lump-sum transfers, they are hit harder by this reduction.

Comparing Tables 6 and 7, and Figures 4 and 8, the mechanism is clear: progressive

Figure 7: Aggregate and across asset tercile adopter stock, benchmark and experiment 1.

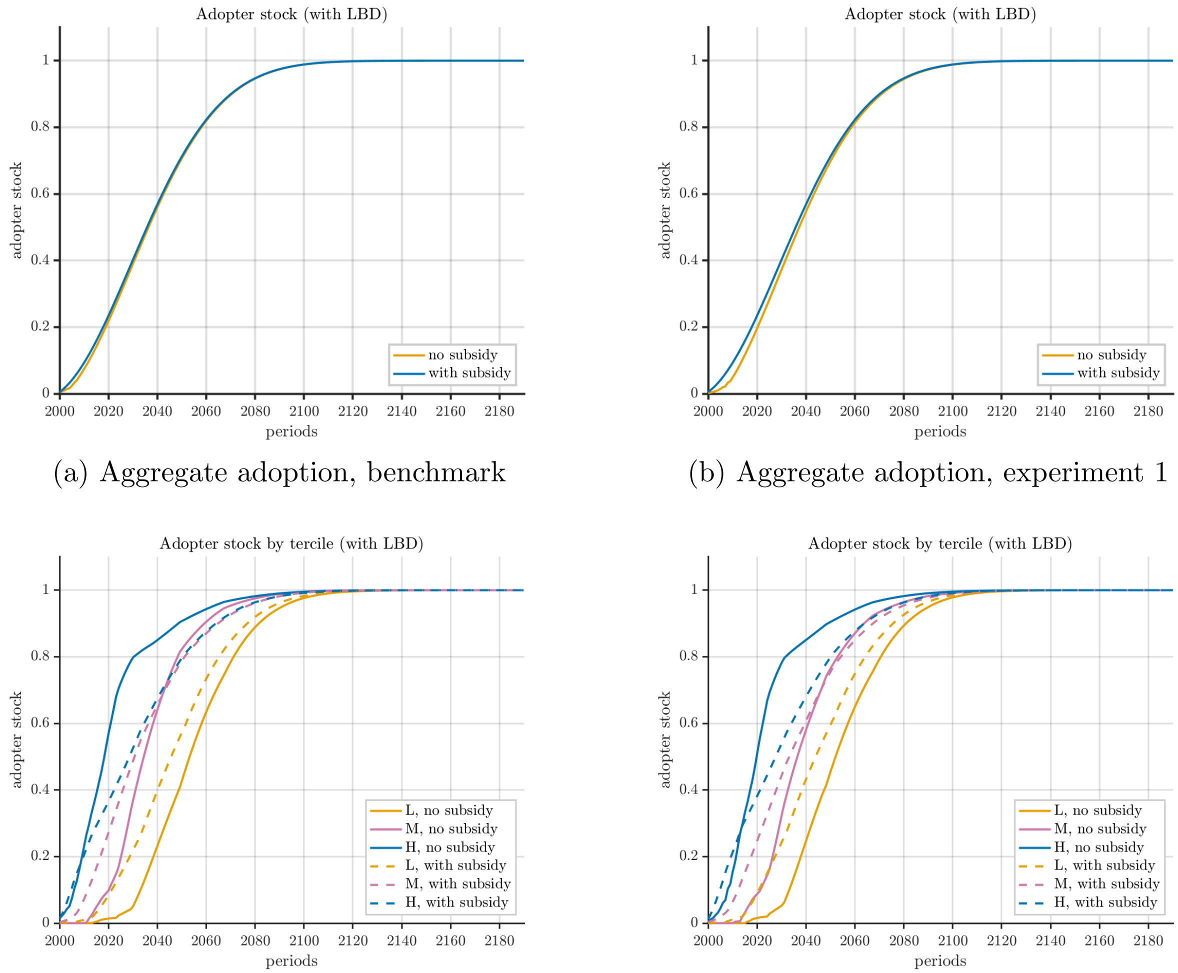

(a) Aggregate adoption, benchmark

(b) Aggregate adoption, experiment 1

(c) Asset-tercile adoption, benchmark

(d) Asset-tercile adoption, experiment 1

financing weakens the wage boost and deepens the transfer cut. Both disproportionately harm low-wealth households, reducing their winner share. The LBD-induced cost channel remains small in either case.

Experiment 2: Nonrefundable Subsidy under Progressive Financing

I next consider a nonrefundable subsidy under progressive financing, mirroring the historical US federal solar ITC. Under a nonrefundable design, households can only claim the subsidy against tax liability: low-income households with little or no tax liability receive no benefit. This design aims to limit fiscal cost (and thus transfer reductions) but may exclude exactly the liquidity-constrained households that policy is often meant to help.

Table 8 shows that aggregate welfare still improves: the average EV gain remains 0.05%, similar to experiment 1. Direct subsidy effects and LBD-induced cost effects remain positive

Figure 8: Aggregate dynamics under progressive financing policy with LBD, with and without adoption subsidies, experiment 1.

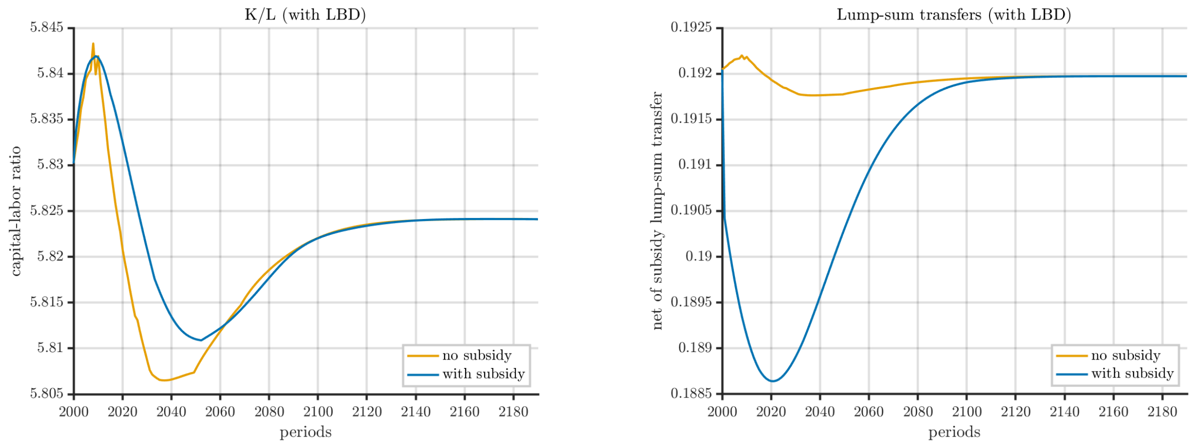

(a) Capital-labor ratio

(b) Lump-sum transfers

for essentially all households, although the magnitude of the direct subsidy channel falls slightly because some low-income households are no longer eligible. The LBD effect remains small.

The average price effect is still negative, but weaker than in experiment 1, and the share of households who benefit from the price channel falls from 32.3% to 23%. The transfer effect becomes slightly less negative on average: making the subsidy nonrefundable does reduce the fiscal pressure somewhat. But transfer effects are still strictly nonpositive for all households.

The distributional pattern shifts. The aggregate winner share falls further, from 83.3% in experiment 1 to 78.1%. Almost all of this decline comes from the bottom asset tercile: their within-tercile winner share drops from 55.8% to 44.4%. In other words, limiting refundability mainly withdraws support from the poor.

Figure 9 compares adoption under refundable (experiment 1) and nonrefundable (experiment 2) subsidies. Panels (a)–(b) show that aggregate adoption paths are nearly identical, so refundability does not change overall diffusion much. Panels (c)–(d) show why distributional incidence changes: with a nonrefundable subsidy, early adoption slows among high-wealth households and accelerates slightly among low- and middle-wealth households. Eligibility becomes state-contingent: high-wealth households delay if they are temporarily in a low-income state (and therefore can't claim the full credit), while some liquidity-constrained households adopt earlier if they happen to draw a temporarily high-income state and thus qualify. Because income follows a persistent AR(1) process, these timing frictions matter in the short- to medium-run.

These shifts in adoption timing feed into savings, prices, and transfers (Figure 10). With a

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

Panel A. Aggregate Welfare Decomposition

Component	Avg. EV (%)	Winner (%)
Direct subsidy	+0.21	100.0
LBD-induced cost change	+0.00	100.0
Price effect	−0.00	23.0
Transfer effect	−0.16	0.0
Total	+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 shares are [0.335, 0.343, 0.322].

nonrefundable subsidy, aggregate savings barely rise in the short run relative to the no-subsidy case (panel (a)). That weaker savings response implies a smaller short-run increase in the capital-labor ratio and hence a smaller wage boost. Lower wage gains reduce the labor-income benefit to low-wealth households, which helps explain their weaker support.

In the medium run, the capital-labor ratio falls more under the nonrefundable subsidy than under the refundable subsidy. As more middle- and high-wealth households adopt, they finance adoption out of savings, reducing asset accumulation. The resulting decline in capital lowers wages and raises interest rates. That shift hurts low-wealth households (who rely on wages) and benefits high-wealth households (who rely on capital income).

Panel (b) shows that net-of-subsidy lump-sum transfers fall less under nonrefundable subsidies than under refundable subsidies, because fiscal costs are lower. But that improved transfer channel is not enough to offset the weaker wage gains for the poor. Overall, nonrefundability slightly reduces aggregate welfare gains but substantially reduces support among low-wealth households.

In short: making the subsidy nonrefundable modestly alleviates fiscal pressure, but further erodes support among low-wealth households by excluding many of them from eligibility and by muting short-run wage gains. Incorporating permanent productivity shocks could change this result, and is left for future work.

Figure 9: Aggregate and across asset tercile adopter stock, experiment 1 and experiment 2.

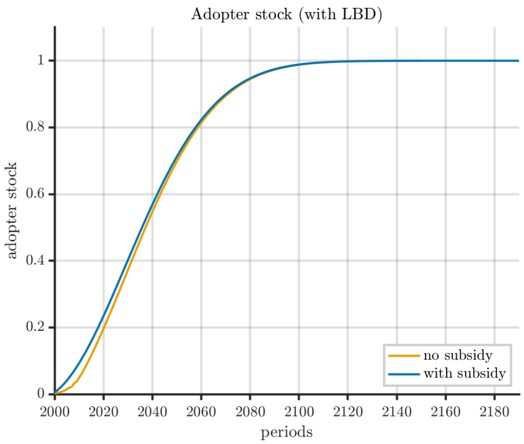

(a) Aggregate adoption, experiment 1

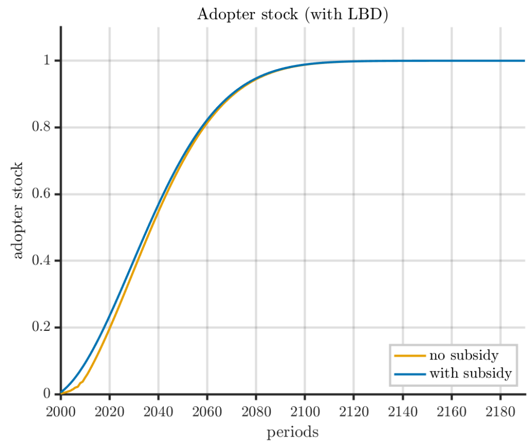

(b) Aggregate adoption, experiment 2

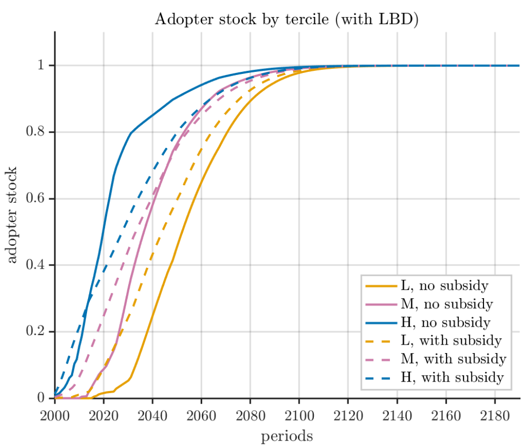

(c) Asset-tercile adoption, experiment 1

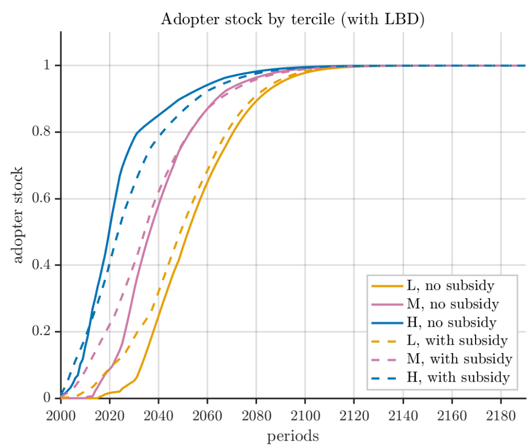

(d) Asset-tercile adoption, experiment 2

Experiment 3: Income-Capped Subsidy under Progressive Financing

Finally, I analyze income-capped subsidies that restrict eligibility to middle- and lower-income households. I consider refundable subsidies 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.

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

Figure 10: Aggregate dynamics under progressive financing policy with LBD, with and without nonrefundable adoption subsidies, experiment 2.

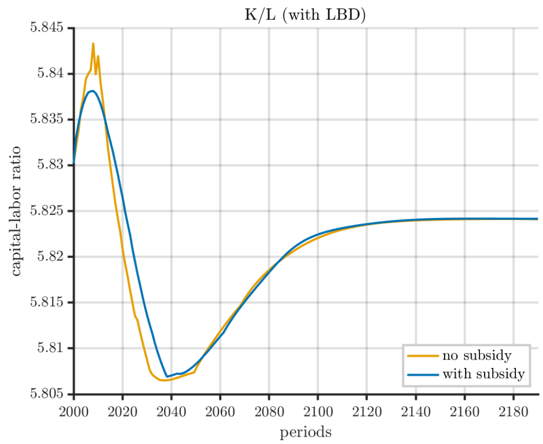

(a) Capital-labor ratio

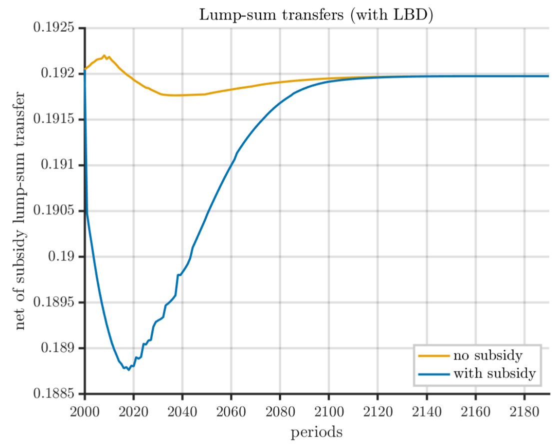

(b) Lump-sum transfers

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 LBD 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 reports the results. This is the only policy among the three variants that reduces aggregate welfare: the average EV change is -0.09%. Only 2.2% of households strictly benefit. The direct subsidy and LBD channels remain positive, but the direct subsidy channel shrinks sharply because all households at or above the cap (here set to median income) are ineligible. The average price effect turns positive – general equilibrium prices help on net – but the transfer effect remains strictly negative, and still large.

Figure 11 shows why. Panels (a)–(b) reveal that aggregate adoption slows sharply under income-capped subsidies. In fact, panel (b) shows that with income caps, subsidies fail to raise aggregate adoption above the no-subsidy path. Panels (c)–(d) show that the cap prevents the "early adopter push" among middle- and high-wealth households that existed under uniform access. Low-wealth households adopt slightly faster in the very short run, but by the medium run their adoption pace falls back to the no-subsidy level, while high-wealth households eventually speed up.

Figure 12 shows the resulting macro dynamics. In the short run, aggregate savings and the capital-labor ratio barely move relative to the no-subsidy path (panel (a)), because the

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

Panel A. Aggregate Welfare Decomposition

Component	Avg. EV (%)	Winner (%)
Direct subsidy	+0.04	100.0
LBD-induced cost change	+0.00	100.0
Price effect	+0.02	77.0
Transfer effect	−0.15	0.0
Total	−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].

composition of adopters does not shift much. In the medium run, however, aggregate savings fall sharply. Two forces drive this. First, because the cap slows aggregate adoption, the policy fails to accelerate LBD. Costs stay high, so per-household subsidy payouts remain large whenever the subsidy is actually claimed. Financing those payouts causes net-of-subsidy lump-sum transfers to fall even more than under the uncapped subsidy (panel (b)), despite the tighter eligibility. Second, the larger drop in transfers depresses savings, which lowers the capital-labor ratio and wages in the medium run. Lower wages cut labor income for low- and middle-wealth households and slow their adoption even further. At the same time, the lower capital-labor ratio raises interest rates, benefiting high-wealth households, who rely on capital income and subsequently speed up adoption.

This explains the sign flip in the average price effect: higher interest rates now benefit the asset-rich enough to outweigh wage losses in the aggregate, even though wage-dependent households are worse off. But the transfer effect remains strongly negative and is borne by everyone.

Figure 13 shows the welfare impact across the income-wealth distribution. Almost all groups experience welfare losses, with the largest losses in the middle of the distribution. Middle-income, middle-wealth households typically do not qualify for the subsidy but still face adverse general equilibrium price and transfer effects. Low-wealth households lose less

Figure 11: Aggregate and across asset tercile adopter stock, experiment 1 and experiment 3.

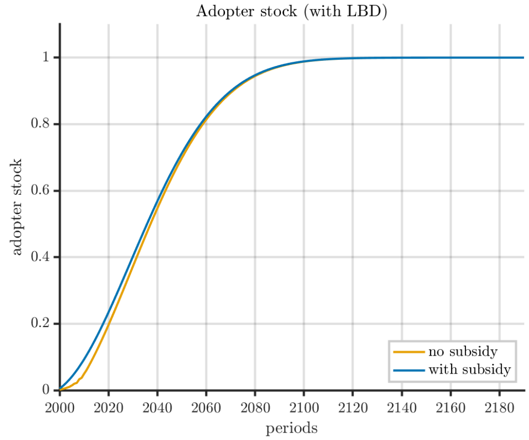

(a) Experiment 1

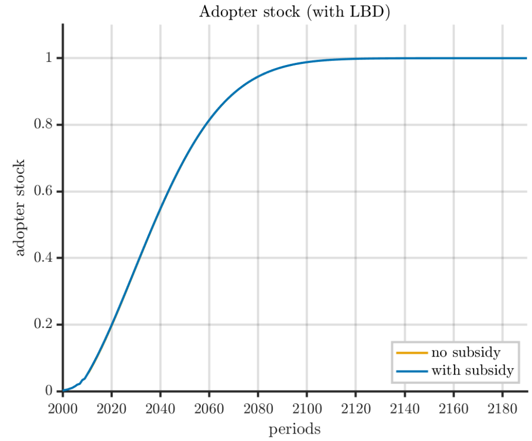

(b) Experiment 3

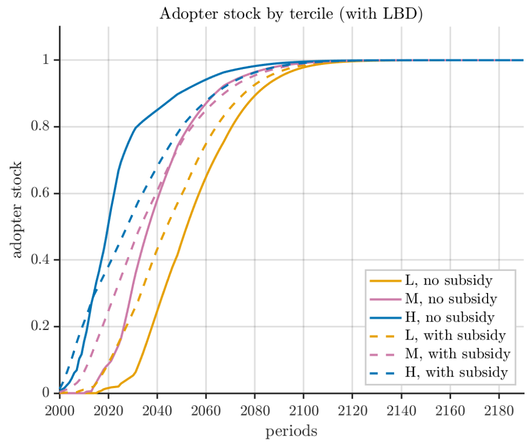

(c) Asset-tercile adoption, experiment 1

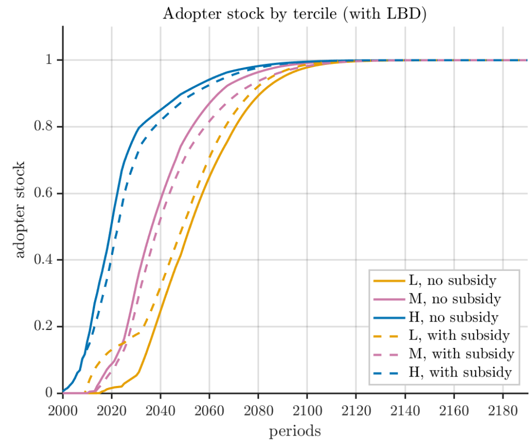

(d) Asset-tercile adoption, experiment 3

because they are more likely to qualify directly. High-wealth households also lose less because higher interest rates partially offset the transfer losses. The policy is therefore both inefficient (negative aggregate EV) and distributionally perverse (large welfare losses concentrated in the middle).

6.2 Extension: Activating Pollution Damages and Inequality

The previous experiments abstracted from pollution damages, even though environmental damages – and their unequal incidence – are core justifications for clean energy policy. I now activate the pollution term in utility (equation (11)), which penalizes high ambient pollution more for lower-consumption households, consistent with evidence on environmental inequality Banzhaf, Ma, and Timmins (2019) and Sergi et al. (2020).

Figure 12: Aggregate dynamics under progressive financing policy with LBD, with and without income-capped adoption subsidies, experiment 3.

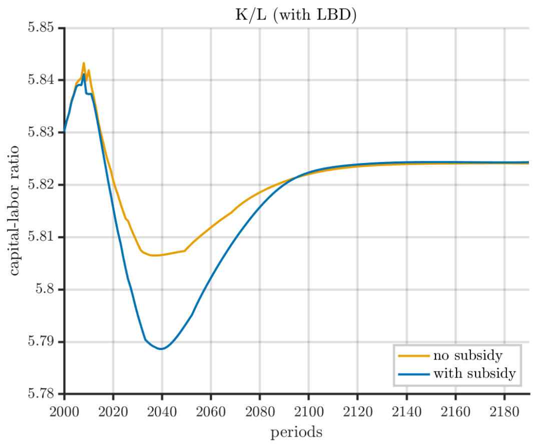

(a) Capital-labor ratio

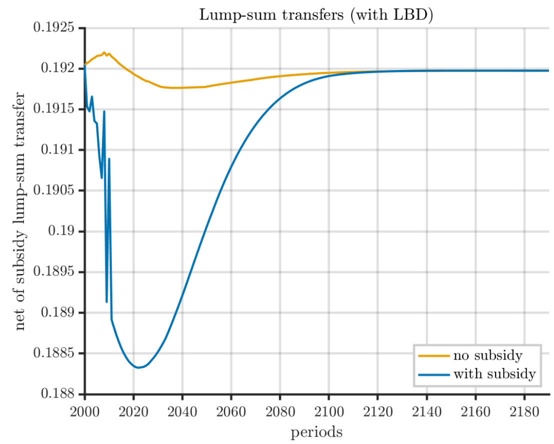

(b) Lump-sum transfers

Figure 13: Welfare impact of subsidizing solar panel adoption cost across the income-wealth distribution, experiment 3

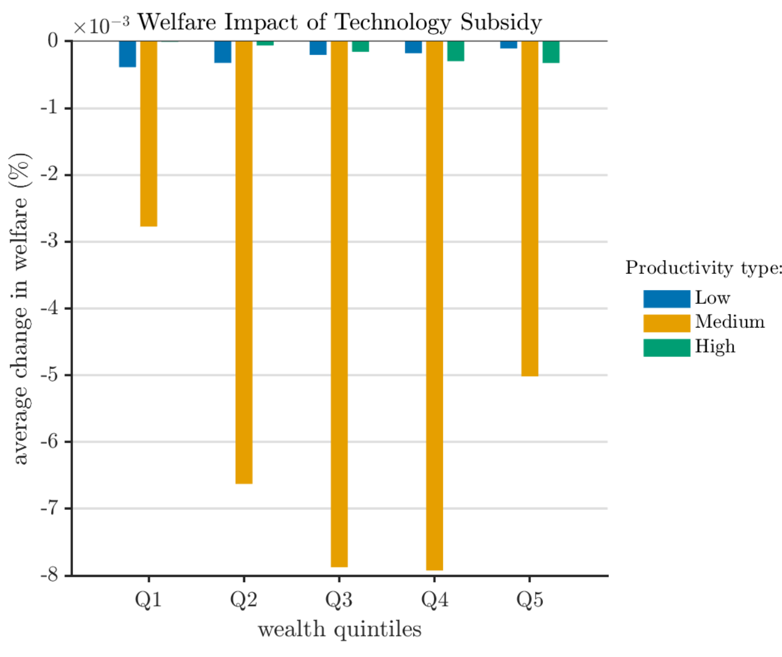

I introduce pollution damages into experiment 2 (nonrefundable subsidy under progressive financing) and recompute welfare. The aggregate EV change from transitioning to clean energy without subsidies is now 12%, reflecting the large benefits from reduced pollution. Because pollution now enters utility directly, this 12% transition gain is not quantitatively comparable to the 0.3% transition gains in the tax experiments: it reflects the environmental

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. EV (%)	Winner (%)
Direct subsidy	+0.20	100.0
LBD-induced cost change	+0.00	100.0
Price effect	−0.01	56.7
Pollution effect	+2.15	100.0
Transfer effect	−0.32	0.0
Total	+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 shares are [0.355, 0.312, 0.333].

externality rather than just technology diffusion. I then evaluate the marginal welfare effects of subsidizing adoption in this setting. Table 10 summarizes the results. Accounting for pollution damages dramatically raises the aggregate welfare gain from subsidizing adoption: the average EV increases from 0.05% in experiment 2 to 2.02%, and the winner share rises from 78.1% to 100%. Every household is strictly better off once pollution damages are internalized.

Figure 14 compares aggregate adoption with and without pollution damages (no-subsidy and subsidy cases). When pollution damages are active, baseline (no-subsidy) adoption is slower. The reason is precautionary: pollution enters utility in a way that increases effective risk aversion and prudence for $X > \bar{X}$, especially at low consumption levels.

Differentiating equation (11) gives

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

For $X > \bar{X}$, both $|u_{cc}|$ and u_{ccc} increase. Thus risk aversion $A(c; X) = -u_{cc}/u_c$ and prudence $P(c; X) = -u_{ccc}/u_{cc}$ rise with pollution. Pollution therefore increases the marginal value of

Figure 14: Aggregate adopter stock, experiment 2 and extension with pollution damages.

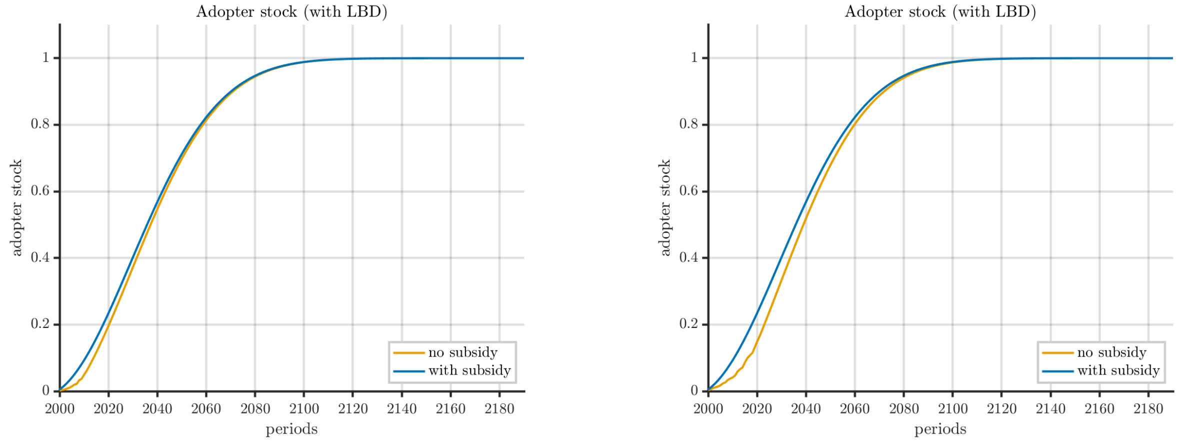

(a) Experiment 2

(b) Experiment 2 with pollution damages

Figure 15: Aggregate dynamics under progressive financing policy with LBD, with and without nonrefundable adoption subsidies, for pollution extension of experiment 2.

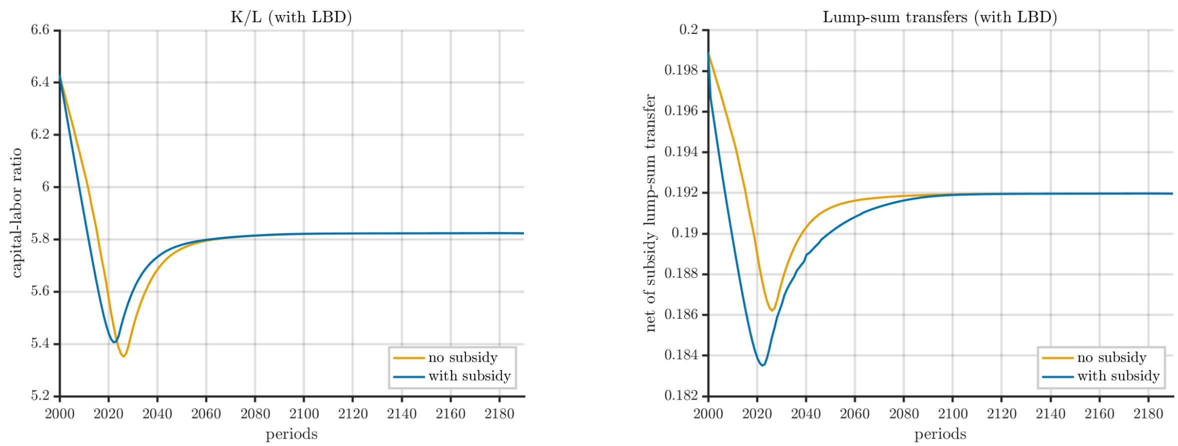

(a) Capital-labor ratio

(b) Lump-sum transfers

liquid wealth. Households save more for self-insurance, and they become more reluctant to give up liquidity for an irreversible investment like solar.

Formally, let $V^N(a, \ell, 0)$ and $V^A(a, \ell, 0)$ be the value functions from not adopting and adopting. Adoption requires paying p_t up front, reducing liquid assets from a to $a - p_t$. The adoption surplus is:

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

A second-order expansion around a gives:

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

thus, adopting entails an additional second-order welfare loss from giving up liquid wealth, $\frac{1}{2}|V^A_{aa}(a,\ell,0)|p_t^2$. Because $V^A_{aa} < 0$, this term measures the curvature-induced penalty from converting liquid assets into an illiquid investment. The curvature $|V^A_{aa}|$ 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 $a^\star(\ell;p_t)$ denote the adoption threshold such that $\Delta(a^\star,\ell;p_t) = 0$. Because the adoption surplus falls with greater curvature, $\partial a^\star(\ell;p_t)/\partial|V^A_{aa}| > 0$, pollution raises the wealth level required for adoption and slows the extensive-margin response in the early transition, as seen in Figure 14. This delay feeds back through LBD, 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 15, panel (a)), before recovering as pollution stabilizes and the precautionary motive dominates. This pattern contrasts with the benchmark without pollution damages (Figure 10, 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 14 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 LBD—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

Figure 16: Welfare impact of subsidizing solar panel adoption cost across the income-wealth distribution, extension with pollution

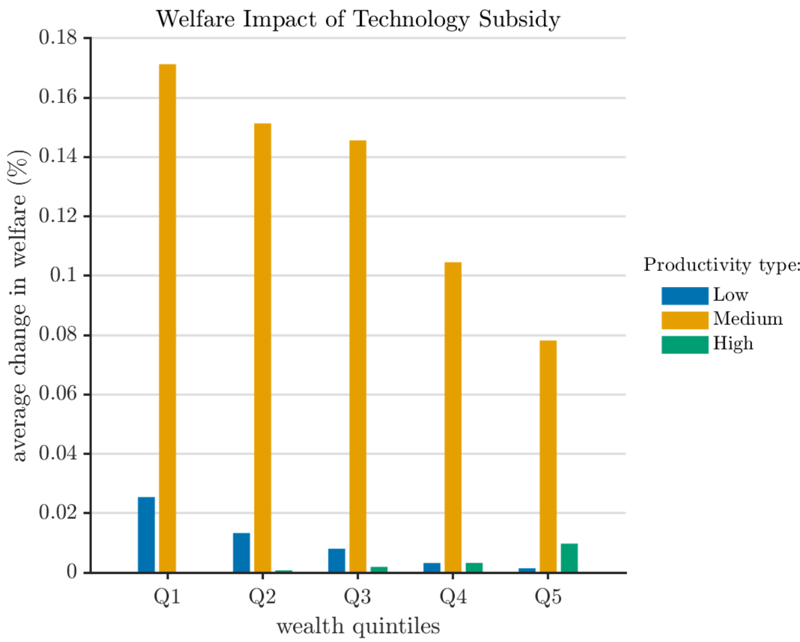

indirectly mitigating the precautionary motive that suppresses early saving. As a result, households adopt more rapidly, pollution declines sooner, and the welfare impact of the subsidy expands well beyond its pure LBD effects. Table 10 quantifies this amplification: the equivalent variation rises from 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 15, panel (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 LBD 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 16 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.

In summary, once pollution damages are accounted for, subsidies internalize two externalities at once: LBD and environmental harm. The result is (i) much larger aggregate welfare gains, (ii) universal support, and (iii) a strongly progressive distribution of benefits. This stands in contrast to the baseline without pollution, where gains were modest, support was incomplete, and financing burdens could be regressive.

6.3 Robustness

Table 11 presents a broad set of robustness checks examining how welfare gains and political support vary with key model parameters and policy design features. The qualitative conclusions are broadly robust but not universal. While many specifications deliver positive aggregate welfare gains and broad-based winners, several experiments produce the opposite outcome. For example, the income-capped subsidy generates a significant aggregate welfare loss (-0.15%) with no households strictly better off and losses concentrated in the middle of the distribution. More generally, policy designs that materially restrict subsidy access of early adopters or parameter combinations that substantially slow diffusion can yield aggregate welfare losses and highly skewed incidence.

The results can be grouped into four categories: technology parameters, subsidy design, pollution-induced utility damages, and financing design. Importantly, only changes to technology parameters or subsidy design are quantitatively comparable to the baseline results, because these reparameterizations and subsidy designs only change the transition dynamics while preserving the steady states. Changes to financing design or pollution-induced utility damages alter the steady state itself, making quantitative comparisons less meaningful. I therefore focus on qualitative patterns for those cases.

Technology Parameters

Varying the speed of diffusion of clean technologies (θ_t) produces the strongest quantitative differences. Faster diffusion substantially magnifies welfare gains by accelerating cost declines, adoption rates, and learning spillovers, raising aggregate welfare to about 3 percent. Conversely, slower diffusion sharply reduces welfare and leaves only a small share of households better off, highlighting that sluggish diffusion can offset the benefits of subsidy policies. Adjusting the LBD elasticity (ξ) or exogenous cost decay (λ) yields smaller effects: faster learning marginally increases welfare, whereas slower learning or slower exogenous cost decline slightly reduce it. These patterns suggest that dynamic spillovers matter primarily through their interaction with diffusion rather than the exact curvature of the learning function.

Subsidy Design

Reducing the baseline nonrefundable subsidy rate from 30% to 15% increases aggregate welfare by 25 percentage points and maintains universal winners. This finding underscores that even moderate subsidies can effectively catalyze adoption and generate broad-based benefits while limiting fiscal costs. However, replicating experiment 3's income-capped subsidy when accounting for pollution preferences maintains negative aggregate welfare and concentrated losses, indicating that restricting subsidy access of earlier adopters can undermine policy effectiveness regardless of other model features. Making the subsidy refundable maintains positive welfare gains across the distribution, but reduces aggregate welfare by 10 percentage points compared to the baseline nonrefundable design, reflecting higher fiscal costs that dampen net benefits when households face disutility from pollution. Finally, making the refundable subsidy income capped again yields negative aggregate welfare and concentrated losses, reinforcing that limiting subsidy access of early adopters can negate policy benefits.

Heterogeneous Pollution Preferences

Adjusting pollution-related parameters also leaves the main result that subsidies are universally welfare-enhancing intact. Importantly, these exercises are not quantitatively comparable to the baseline, as they change the steady-state equilibra. Thus, I focus on qualitative patterns. Equalizing pollution exposure across households ($\omega = 0$) raises aggregate welfare by EV of about 0.78% and maintains universal winners. However, interestingly, shutting down pollution heterogeneity increases welfare gains of lower-wealth households relative to other groups', as an equal change in utility from pollution reduction represents a larger fraction of their baseline utility. The model's conclusion that subsidies are universally beneficial remains robust to smaller pollution inequality parameter (ω half its baseline value) or lower pollution disutility (ν halved).

Table 11: Robustness of welfare gains and political support across parameterizations

Panel A. Consumption-equivalent welfare gains (EV, %)				
Scenario	**Bottom**	**Middle**	**Top**	**Aggregate**
Baseline (prog. tax + nonref. subs. + pollution)	1.30	2.85	1.83	2.02
Technology Parameters (quantitatively comparable)				
$\theta_t \times 1.25$ (fast diffusion)	1.91	4.28	2.73	3.01
$\theta_t \times 0.5$ (slow diffusion)	−0.07	−0.32	−0.17	−0.19
$\xi \times 1.5$ (faster learning)	1.30	2.85	1.84	2.02
$\xi \times 0.5$ (slower learning)	1.30	2.86	1.84	2.03
$\lambda \times 0.5$ (slower exog. decay)	1.18	1.85	1.42	1.50
Subsidy Design (quantitatively comparable)				
$\tau \times 0.5$ (less subsidy)	1.42	3.25	2.06	2.27
Income-capped subsidy	−0.11	−0.19	−0.15	−0.15
Refundable subsidy	1.23	2.71	1.74	1.92
Refundable, income capped subsidy	−0.11	−0.20	−0.16	−0.16
Pollution Preferences (steady states change; qualitative comparison only)				
$\omega = 0$ (uniform pollution exposure)	0.84	0.73	0.78	0.78
$\omega \times 0.5$ (less pollution exposure inequality)	1.23	1.95	1.55	1.59
$\nu \times 0.5$ (less pollution disutility)	0.63	1.40	0.86	0.98
Financing Design (steady states change; qualitative comparison only)				
Flat tax financing	0.79	1.64	1.10	1.18
Less progressive tax	0.82	1.69	1.12	1.23
More progressive tax	1.65	3.63	2.29	2.58

Panel B. Share of households with EV > 0 (%)				
Scenario	**Bottom**	**Middle**	**Top**	**Aggregate**
Baseline (prog. tax + nonref. subs. + pollution)	100.00	100.00	100.00	100.00
Panel A. Consumption-equivalent welfare gains (EV, %)				
$\theta_t \times 1.25$ (fast diffusion)	100.00	100.00	100.00	100.00
$\theta_t \times 0.5$ (slow diffusion)	0.15	7.83	18.17	8.54
$\xi \times 1.5$ (faster learning)	100.00	100.00	100.00	100.00
$\xi \times 0.5$ (slower learning)	100.00	100.00	100.00	100.00
$\lambda \times 0.5$ (slower exog. decay)	100.00	100.00	100.00	100.00
Technology Parameters (quantitatively comparable)				
$\tau \times 0.5$ (less subsidy)	100.00	100.00	100.00	100.00
Income-capped subsidy	0.00	0.00	0.00	0.00
Refundable subsidy	100.00	100.00	100.00	100.00
Refundable, income capped subsidy	0.00	0.00	0.00	0.00
Pollution Preferences (steady states change; qualitative comparison only)				
$\omega = 0$ (uniform pollution exposure)	100.00	100.00	100.00	100.00
$\omega \times 0.5$ (less pollution exposure inequality)	100.00	100.00	100.00	100.00
$\nu \times 0.5$ (less pollution disutility)	100.00	100.00	100.00	100.00
Financing Design (steady states change; qualitative comparison only)				
Flat tax financing	100.00	100.00	100.00	100.00
Less progressive tax	100.00	100.00	100.00	100.00
More progressive tax	100.00	100.00	100.00	100.00

Notes: "Bottom/Middle/Top" report average consumption-equivalent welfare gain (EV) within each initial wealth tercile. "Aggregate" reports the average EV across all households in that scenario. Rows are grouped into (i) technology parameters and (ii) subsidy design, which keep the same initial steady state as the baseline and can be compared quantitatively, and (iii) pollution preferences and (iv) financing design, which require recomputing the steady state. For groups (iii) and (iv), EV levels are internally valid but not quantitatively comparable to the baseline or to groups (i)–(ii). Panel B reports the fraction of households with strictly positive EV.

Financing Design

The model's qualitative conclusions also hold under alternative financing designs. Similar to adjustments to pollution preference parameters, these exercises change the steady-state equilibria, so I focus on qualitative patterns. Switching from the baseline progressive labor income tax financing to uniform, less progressive, or more progressive financing maintains positive aggregate welfare gains and broad-based winners.

Overall, the robustness exercises show that the paper's main conclusions hold under a broad but not universal set of assumptions. Subsidies generally enhance aggregate welfare and yield widespread benefits when technology diffusion is sufficiently fast and early (wealthy) adopters have subsidy access. However, several experiments demonstrate that these conclusions can weaken or even reverse when diffusion is sluggish and subsidy eligibility of early adopters is restricted. In particular, income-capped subsidies, and parameterizations that restrict diffusion, can generate negative aggregate welfare changes. Thus, the robustness analysis highlights the boundaries of the main results: the welfare and equity benefits of clean energy subsidies are not automatic, but depend on policy designs that sustain diffusion and maintain access for early adopters.

7 Conclusion

This paper examines the equity and efficiency of clean energy subsidies through the case of US residential rooftop solar. By 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 alternative subsidy and financing designs shape adoption, prices, and welfare across the income-wealth distribution.

The analysis shows that static incidence measures overstate the regressivity of residential solar subsidies. Once learning-by-doing spillovers are accounted for, subsidies accelerate cost declines, expand adoption, and generate broad welfare gains – even when direct fiscal transfers appear to favor higher-wealth households. When unequal pollution exposure and its health damages are incorporated, the gains become both universal and progressive: every household benefits, and the poorest benefit most from cleaner air and faster cost declines.

The results also show that policy design matters. Progressive tax financing, intended to improve fairness, can unintentionally dampen early adoption, weaken short-run wage gains, and erode support among low-wealth households. Nonrefundability further excludes liquidity-constrained households. In contrast, broadly available refundable subsidies speed diffusion and raise aggregate welfare. The most equitable path therefore combines: (i) broad early subsidies that accelerate learning and reduce future costs, and (ii) targeted support that relaxes liquidity constraints and internalizes pollution damages.

More broadly, the paper offers a quantitative framework for evaluating environmental policy in environments with heterogeneous households, incomplete markets, and dynamic externalities. The results for US residential solar suggest that when policy is designed to internalize both learning and pollution damages, an accelerated clean energy transition can also be an equitable one.

While the analysis focuses on clean energy adoption, the underlying framework is broadly applicable to other settings in which technology diffusion interacts with inequality and public finance. The same structure could be used to study the diffusion of electric vehicles, home energy efficiency retrofits, or carbon capture systems, as well as non-environmental innovations such as broadband expansion, digital payments, or health technologies. In all these contexts, irreversible investment decisions, learning spillovers, and heterogeneous financing constraints generate similar trade-offs between efficiency, equity, and fiscal cost. Extending the framework to these domains would deepen our understanding of how technology policy can jointly promote innovation and inclusion.

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.

Arrow, Kenneth J. (1962). "The Economic Implications of Learning by Doing". *The Review of Economic Studies* 29.3, pp. 155–173.

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, Elisa, Daniel Carroll, and Sewon Hur (2024). *Unequal Climate Policy in an Unequal World*. Working Paper 427. Federal Reserve Bank of Dallas.

Benmir, Ghassane and Josselin 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, Bryan and Kenneth 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, Stephie, 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.

Fried, Stephie, Kevin Novan, and William B. Peterman (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.

Mukoyama, Toshihiko (2006). "Rosenberg's "learning by using" and technology diffusion". *Journal of Economic Behavior & Organization* 61.1, pp. 123–144.

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.

Rogers, Everett M. (2003). *Diffusion of Innovations*. 5th. Free Press.

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.

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 PM2.5 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.

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

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 A.2. The results of the regression show that the effect of having solar generation on the unit electricity price, presented in columns (1) and

Table A.1: Regression results for the effect of solar generation on annual electricity and total energy expenditure

Variables	Electricitiy (1)	(2)	Total Energy (3)	(4)
Solar	-721.956	-709.430	-688.5617	-609.9241
	(15.938)	(15.881)	(24.286)	(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 A.1. 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 A.2: Regression results for the effect of solar generation on electricity and total energy unit costs

Variables	Electricitiy (1)	(2)	Total Energy (3)	(4)
Solar	-0.0214	-0.0214	-0.0093	-0.0091
	(0.000)	(0.000)	(0.000)	(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 YYY" 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 $V : \tilde{\mathcal{Z}} \times \tilde{\mathcal{M}} \to \mathbb{R}$, household policy functions $c, a' : \tilde{\mathcal{Z}} \times \tilde{\mathcal{M}} \to \mathbb{R}$, aggregate factor demands $K, L : \tilde{\mathcal{M}} \to \mathbb{R}$, factor price functions $r, w : \tilde{\mathcal{M}} \to \mathbb{R}$, a transfer function $T : \tilde{\mathcal{M}} \to \mathbb{R}$, a pollution function $X : \tilde{\mathcal{M}} \to \mathbb{R}$, and law of motion $H : \tilde{\mathcal{M}} \to \tilde{\mathcal{M}}$ such that:*

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

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

   and c is the associated consumption policy function, $a' = w(\Phi)(1-\tau^\ell(\Phi))\ell + (1+r(\Phi))a + T(\Phi) - c - \bar{q}e(c)$ 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 $\Phi \in \tilde{\mathcal{M}}$,*

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

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: $Q_\Phi : \tilde{\mathcal{Z}} \times B(\tilde{(\mathcal{Z})}) \to [0,1]$ 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 $Q_t : \mathcal{Z} \to [0, 1]$ 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 C.1, 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 C.1: 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 (B.1) 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 functon 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 $t = 1, \ldots, T$.

4. *Forward simulation:* Using the policy function from the previous step, I simulate the economy forward, starting from the invariant distribution at the initial steady state. At each period $t = 1, \ldots, T$, I evolve the joint distribution of households over assets, income, and adoption status, given the transition matrix of idiosyncratic income shocks Π and the endogenous policy functions for asset accumulation and technology adoption.

   (a) **Adoption desire.** For each household currently in state $(a, y, s = 0)$ (a non-adopter), I evaluate its discrete adoption choice based on the period-t value functions, yielding an indicator $I_t(a, y) = 1$ if adoption is preferred and 0 otherwise. The total mass of households that *wish* to adopt is then

   $$M_t^{\text{want}} = \sum_{a,y} f_t(a, y, 0)\, I_t(a, y),$$

   where $f_t(a, y, 0)$ denotes the joint distribution of assets and income among non-adopters.

   (b) **Calvo adoption constraint.** To capture gradual diffusion due to market frictions and supply constraints, I impose a Calvo-type restriction: only a fixed fraction $\theta_t \in (0, 1)$ of non-adopters can realize their adoption decision in any given period. If the desired adoption mass exceeds the allowed fraction, the realized adoption flow is scaled down proportionally:

   $$M_t^{\text{real}} = \min\left(M_t^{\text{want}},\ \theta_t \sum_{a,y} f_t(a, y, 0) \right).$$

The probability that a household wanting to adopt actually does so is

$$\rho_t = \begin{cases} 1, & \text{if } M_t^{\text{want}} \le \theta_t \sum_{a,y} f_t(a,y,0), \\ \dfrac{\theta_t \sum_{a,y} f_t(a,y,0)}{M_t^{\text{want}}}, & \text{otherwise.} \end{cases}$$

Hence, only a random fraction ρ_t of the willing adopters transition from $s = 0$ to $s = 1$ (adopt), while the rest remain non-adopters. The resulting realized adoption flow $\rho_t f_t(a, y, 0) I_t(a, y)$ enters the law of motion for the distribution.

(c) **Aggregate updates.** The updated distribution $f_{t+1}(a', y', s')$ is then computed by integrating the policy functions and transition probabilities over assets and income. From this distribution, I compute the implied aggregates $\{K_t, T_t, Z_t\}$ for each t, which will be used to update the outer-loop guesses in the next iteration.

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^{-3}.

When the model incorporates the pollution disutility in household preferences, I follow the same algorithm with the addition of updating the ambient air pollution level sequence $\{X_t\}_{t=1}^{T}$ at each period t in the outer loop based on the aggregate energy consumption implied by the distribution at period t.

C.3 Welfare Change Calculations

For each household state (a, ℓ, s), I compute the consumption equivalent variation (EV) welfare change measure between the baseline and counterfactual paths.

Equivalent Variation (EV)

I define 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. I consider two counterfactual scenarios: (i) the technology transition without any subsidies relative to no technology transition, and (ii) subsidizing the clean energy technology adoption at rate $\tau = 0.3$ relative to not subsidizing. The first EV calculation captures the welfare change from introducing the clean energy technology, while the second EV calculation captures the welfare change from subsidizing the clean energy technology adoption at rate $\tau = 0.3$ relative to not subsidizing.

In order to compute the welfare change due to the introduction of the clean energy technology, I compute the EV for a household in initial steady state (a, ℓ) as $\lambda^{EV,tech}(a, \ell)$ that solves:

$$\mathbb{E}_1\left[\sum_{t=1}^{T}\beta^t U((1+\lambda^{EV,tech}(a,\ell))c_t(a,\ell)^{\text{initial ss}})\right] = \mathbb{E}_1\left[\sum_{t=1}^{T}\beta^t U(c_t(a,\ell,0)^{\text{transition}})\right],$$

where $c_t(a, \ell)^{\text{initial ss}}$ is the consumption path under no technology transition (the technology remains available and the economy remains in the initial steady state) and $c_t(a, \ell, 0)^{\text{transition}}$ is the consumption path under the technology transition, without any subsidies. If $\lambda^{EV,tech}(a, \ell) > 0$, then the household is better off with the introduction of the clean energy technology and supports the transition. Under the CRRA utility specification, the EV of the technology transition can be computed in closed form as:

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

where $V_1(a, \ell, 0)^{\text{transition}}$ is the value function at $t = 1$ under the scenario with no adoption subsidy during the technology transition and $V(a, \ell)^{\text{initial ss}}$ is the value function at the initial steady state when the clean technology is unavailable.

Formally, let $c_t(a, \ell, s)^{\tau}$ be the consumption policy function at time $t = 1, \ldots, T$, 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 for a household in initial steady state (a, ℓ) is computed as $\lambda^{EV,subs}(a, \ell)$ that solves:

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

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

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

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