Utility-scale solar entry and county-level business and population dynamics in the Western United States

by

Lauren Phillips

B.S., University of Tennessee, Knoxville, 2025

A THESIS

submitted in partial fulfillment of the requirements for the degree

MASTER OF SCIENCE

Department of Agricultural Economics
College of Agriculture

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2026

Approved by:

Major Professor
Juhee Lee

Copyright

© Lauren Phillips 2026.

Abstract

This study examines how county entry into utility-scale solar generation is associated with business, employment, and demographic adjustment across seven western U.S. states from 2010 to 2022. Solar entry is defined as the first audited year in which qualifying county-level solar capacity reaches 1 MW. Historical generator information is used to distinguish counties with pre-panel exposure from new adopters. The main analysis includes 112 adopting counties across 12 entry cohorts and 108 never-treated counties before outcome-specific exclusions. An entry cohort consists of counties that first meet the audited solar-entry threshold in the same calendar year.

The empirical design uses cohort-specific inverse-probability-weighted group-time difference-in-differences comparisons, with baseline population, income per capita, unemployment, and state used to improve comparability between adopting and never-treated counties. The primary estimates average supported cohort-event contrasts over event years 0–5, weighted by treated cohort size. The estimated private-establishment contrast is 1.59 log points relative to the weighted never-treated trajectory ($p = 0.035$), with a similar positive result for total covered establishments. Private and total covered employment estimates are positive but do not reach the 10% significance level. Population and cumulative net migration estimates are also positive and reach the 10% significance level but not the 5% level; these demographic results are more sensitive to alternative estimators and geographic restrictions.

The private-establishment contrast remains positive across complementary estimators and leave-one-state-out analyses, although the statistical evidence weakens under higher-capacity entry definitions. A separate descriptive analysis of 80 adopters observed over a common five-year post-entry window shows larger mean changes among metropolitan counties for private establishments, private employment, population, and cumulative net migration. However, the

distributions overlap substantially, and these comparisons do not identify metropolitan-specific treatment effects.

Overall, the evidence provides more consistent support for differences in business presence following initial utility-scale solar entry than for broad workplace employment or demographic adjustment. Residual covariate imbalance, concentrated weights in some early cohorts, and unfavorable matched-business pre-treatment diagnostics limit strong causal interpretation. The findings underscore the importance of distinguishing business presence, workplace employment, and residential adjustment when evaluating the local economic consequences of utility-scale solar development.

Table of Contents

Table of Contents v
List of Figures viii
List of Tables ix
Acknowledgements x
Dedication xi
Chapter 1. Introduction 1
Chapter 2. Background and conceptual framework 7
2.1 Solar investment and the local retention of economic activity 7
2.2 Establishments and employment as distinct private-sector margins 8
2.3 Demographic adjustment and competing local channels 10
2.4 Conceptual ordering and scope of inference 11
Chapter 3. Literature review 12
3.1 Utility-scale solar and local economic outcomes 12
3.2 Utility-scale solar and demographic adjustment 14
3.3 Economic and demographic adjustment in a common place-based framework 16
3.4 Synthesis, research gap, and contribution of this study 18
Chapter 4. Data and study area 20
4.1 Geographic setting and observational unit 20
4.2 Business and labor-market measures 22
4.3 Population and migration 23
4.4 Baseline covariates and descriptive metropolitan classification 24
Chapter 5. Treatment and outcome construction 26
5.1 Audited county entry 26
5.2 Historical evidence and the estimation sample 27
5.3 Stocks and workplace employment 29
5.4 Cumulative net migration 30
Chapter 6. Empirical methodology and identification 31
6.1 Target parameters and conditional comparison 31
6.2 Baseline balancing, propensity scores, and ATT weights 32

6.3 Cohort-time estimation, event-time aggregation, and interpretation 34
6.4 Bootstrap inference and dependence 37
6.5 Doubly robust adjustment 39
6.6 Risk-set matching 40
6.7 Descriptive heterogeneity by metropolitan status 42
6.8 Identification assumptions and diagnostics 43
Chapter 7. Results 46
7.1 Baseline comparability and achieved balance 46
7.2 Economic and business outcomes 48
7.3 Demographic outcomes 51
7.4 Cross-domain interpretation 53
7.5 Descriptive heterogeneity by metropolitan status 53
Chapter 8. Robustness and sensitivity 57
8.1 Alternative counterfactual estimators 57
8.1.1 Doubly robust results 57
8.1.2 Risk-set matched results 58
8.2 Transformation, covariate, timing, and sample checks 61
8.3 Higher-capacity entry concepts 65
8.4 Geographic sensitivity 66
Chapter 9. Discussion 68
9.1 What the establishment result contributes 68
9.2 Why employment and resident responses require caution 70
9.3 How the supporting evidence changes the interpretation 72
9.4 Descriptive heterogeneity by metropolitan status 73
9.5 Implications for local development and future evaluation 74
Chapter 10. Limitations and scope of inference 76
Chapter 11. Conclusion 79
References 80
Appendix A. Supporting diagnostics 85
A.1 Weight concentration and covariate balance 85
A.2 Joint pre-treatment tests and event-time support 86

A.3 Complete estimator comparisons 87
A.4 Matching quality and propensity-score overlap 88

List of Figures

Figure 1. Study area and audited county solar-entry classification, 2010–2022. 21
Figure 2. Baseline balance before and after cohort-specific IPW. 47
Figure 3. Main IPW event-study estimates for business and employment outcomes. 50
Figure 4. Main IPW event-study estimates for population and cumulative net migration. 52
Figure 5. County outcomes after utility-scale solar entry. 54
Figure 6. Metropolitan and nonmetropolitan distributions of post-entry outcomes. 55
Figure 7. Complementary estimator comparisons for the event-years 0–5 average. 60
Figure A1. Cohort-specific raw propensity-score distributions. 90

List of Tables

Table 1. Audited county treatment classification, 2010–2022 28
Table 2. Main IPW business and employment estimates, event years 0–5 49
Table 3. Main IPW demographic estimates, event years 0–5 51
Table 4. Private-establishment IPW sensitivity estimates 63
Table 5. Population and cumulative-migration IPW sensitivity estimates 64
Table 6. Leave-one-state-out IPW sensitivity 67
Table A1. Cohort-specific effective control sample size and weight concentration 85
Table A2. Baseline covariate balance under IPW and matching 86
Table A3. Joint tests of pre-treatment event coefficients 86
Table A4. Treated county and cohort support by event year 87
Table A5. Complete main and complementary estimator comparisons 88
Table A6. Cohort-level risk-set matching diagnostics 89

Acknowledgements

I would like to thank my major professor and advisor, Dr. Juhee Lee, for her continued support through this journey. I want to thank her for always showing support and lifting me up with kind, positive words when the completion of this degree has seemed challenging, and for her guidance in helping me navigate my next steps.

I would like to thank my committee members, Dr. Jason Bergtold and Dr. Micah Cameron-Harp, for their support in this endevour, providing useful feedback and thoughtful suggestions.

I would like to thank one of my fellow graduate students, Hasif Md Asraf Mahmud, for his continuous coding help and providing insightful tips and tricks throughout the duration of this thesis, to help me become the best version of myself as a researcher.

Lastly, I would like to thank my family and my friends for their support throughout this journey and helping me navigate the challenges of this program. Kansas State and the Ag Econ Department has given me a sense of community and I will forever be grateful for my time in this program and the life-long friendships and connections I have made. All of my friends and family have been my main source of strength and motivation and I am deeply grateful.

Dedication

I am dedicating this thesis to my grandfather who has always taught me that I can do anything I put my mind to because when one door closes, another one opens.

Chapter 1. Introduction

Utility-scale solar generation has expanded rapidly across the United States, supported by declining technology costs, larger projects, and sustained investment (U.S. Department of Energy, 2021; Bolinger et al., 2023). This expansion has implications beyond the electricity sector. For counties that host solar facilities, it raises a more localized question: how does the arrival of utility-scale solar development relate to changes in business activity, employment, and the resident population? The answer matters because economic benefits attributed to renewable investment at the national or regional level need not accrue to the communities where generating facilities are located.

Solar facilities are tied to particular sites, but the economic activity associated with them need not remain in the host county. Equipment may be manufactured elsewhere, engineering and construction services may be supplied by outside firms, workers may commute rather than relocate, and project revenues may accrue to nonresident owners. At the same time, development can increase demand for site preparation, transportation, lodging, contracting, and other local services. Whether this activity translates into sustained changes in the host economy depends partly on how much spending and employment is retained locally. Project-supported job estimates and regional economic multipliers therefore cannot be treated as direct measures of changes in host-county outcomes (Jenniches, 2018). Installed capacity alone likewise reveals little about the form or magnitude of local adjustment.

This distinction is especially relevant when examining establishments and employment. The number of operating establishments measures the presence of covered business locations, whereas payroll employment measures jobs at those workplaces. Establishment counts may rise without a proportionate increase in aggregate employment if additional locations are small.

Conversely, existing establishments may hire more workers without opening new locations. These are separate margins of local economic adjustment, not interchangeable measures of economic development. An increase in establishments does not necessarily imply new firm formation; it may reflect additional locations of existing firms or fewer closures. Examining establishment and employment outcomes together nevertheless makes it possible to assess whether business presence and workplace activity follow similar paths after solar entry.

Residential outcomes provide another perspective. Solar-related workers may commute from neighboring counties, and temporary construction demand may not induce permanent relocation. Population changes also reflect births, deaths, housing conditions, amenities, and economic forces unrelated to solar development. Net migration is a complementary indicator of residential adjustment, but a net flow cannot distinguish changes in arrivals from changes in departures. Regional growth models recognize that households and firms respond jointly to local economic conditions (Carlino and Mills, 1987), while internal-migration research emphasizes determinants beyond contemporaneous employment opportunities (Molloy et al., 2011). Changes in workplace employment should therefore not be assumed to produce corresponding changes in population or migration, or to precede them in a fixed causal sequence.

Existing place-based research suggests that renewable-energy investment can influence host communities, although the observed outcomes differ across technologies and settings. Brown et al. (2012) report modest gains in county employment and personal income following wind development, while Brunner and Schwegman (2022a) find broader local economic gains without comparable population growth. These findings provide a useful benchmark, but their implications may not extend to solar development. Wind and solar projects differ in construction activity, operating labor needs, land use, and supply-chain relationships. Evidence from Spanish

municipalities shows that renewable-energy investments can generate employment without all associated jobs being located in the host municipality (Fabra et al., 2024). A study of Spanish provinces likewise finds that wind and solar employment responses depend on project size, investment period, and regional economic linkages (Barrutiabengoa Ortubai and Ortiz Durán, 2026). Both the technology and the geography of measured outcomes matter when assessing local economic change.

Recent demographic evidence offers a particularly relevant comparison. Using U.S. county data and staggered treatment timing, Cary and Ortez (2026) find population declines following wind development but no statistically discernible population response to utility-scale solar facilities. Their analysis also considers neighboring-county spillovers and rural–urban differences. The solar finding is relevant here, although differences in treatment definitions, samples, outcome measures, and comparison designs limit direct comparisons. This literature therefore underscores why local business presence, workplace employment, population, and migration should be examined as distinct outcomes rather than inferred from one another.

Two questions remain comparatively underexamined. The first concerns the stock of operating establishments in counties entering utility-scale solar generation, an outcome distinct from project-level job creation and county payroll employment. The second concerns whether business presence, workplace employment, population, and net migration follow similar patterns when measured around the same solar-entry event. Existing solar studies provide important evidence on specific outcomes, while much of the multi-outcome U.S. county literature focuses on wind. Examining the four margins within a common county-level framework allows their timing and relative strength to be compared without presuming that business activity causes employment growth or that employment growth subsequently causes migration.

Such a comparison faces substantial identification challenges. Solar facilities are not located randomly: counties with favorable solar resources, transmission access, interconnection opportunities, available land, or broader development advantages may be more likely to attract investment and may also experience different economic and demographic trends. In addition, the start of generation need not mark the beginning of economic exposure. Permitting, contracting, and construction may precede the recorded year of entry, creating potential anticipation effects. Historical installations can also be missed when contemporaneous capacity data begin after a county's initial exposure. Because entry occurs in different years, conventional two-way fixed-effects difference-in-differences estimators can be difficult to interpret when effects differ by cohort or time since entry (Goodman-Bacon, 2021; Sun and Abraham, 2021). Reliable historical classification, explicit comparisons, and careful assessment of identifying assumptions are therefore essential.

This study examines 252 counties and county equivalents in Arizona, California, Colorado, Nevada, New Mexico, Oregon, and Utah over 2010–2022. Solar entry is defined as the first audited year in which aggregate qualifying county solar capacity reaches 1 MW. Historical operating information is used to identify counties that crossed this threshold before the panel began, rather than classifying them as new adopters or untreated comparisons. Excluding 29 pre-treated counties and three 2010 entrants without a pre-entry baseline leaves 112 adopting counties across 12 entry cohorts (2011–2022) and 108 never-treated counties before outcome-specific exclusions. The treatment is an absorbing indicator of initial county entry rather than a continuous measure of solar deployment. Accordingly, the estimand captures adjustment following initial entry, not the marginal effect of an additional megawatt (MW) of installed capacity or megawatt-hour (MWh) of electricity generation.

The primary analysis uses cohort-specific, inverse-probability-weighted group-time comparisons, following the framework of Callaway and Sant'Anna (2021) and the weighting logic of Abadie (2005). Each entry cohort is compared with counties that remain untreated through 2022, with weights based on pre-entry population, income per capita, unemployment, and state. Event-study estimates trace differences before and after entry, and the principal summary averages supported event years 0–5. Diagnostics for balance, overlap, effective sample size, and pre-entry trends assess the credibility of the comparison. Doubly robust and matched specifications, alternative capacity thresholds and timing rules, geographic restrictions, and other sensitivity analyses provide additional evidence. These procedures improve transparency but do not make solar siting exogenous or rule out residual confounding.

The outcome structure distinguishes four principal dimensions of county adjustment. Private establishments are the primary business outcome, and private employment measures payroll jobs at private workplaces. Total covered establishments and employment provide broader corroborating outcomes. Population measures the stock of county residents, while cumulative net migration is expressed as net migrants per 1,000 baseline residents. Each outcome is analyzed separately over event time relative to the same audited solar-entry year, allowing their trajectories to be compared without assuming causal relationships among the outcomes. A complementary fixed-window comparison describes observed post-entry changes among metropolitan and nonmetropolitan adopters. Because that comparison lacks group-specific counterfactual trajectories, it is interpreted as descriptive variation, not a causal estimate of treatment-effect heterogeneity.

The study contributes to the literature by examining business presence alongside employment and residential adjustment specifically around the onset of utility-scale solar

generation. Distinguishing the establishment stock from payroll employment provides information that conventional jobs-per-megawatt measures cannot capture, while the common treatment definition makes cross-outcome comparisons more coherent. The historically audited entry measure also reduces the risk of treating previously exposed counties as new adopters, and the cohort-specific design avoids using already-treated counties as controls. Balance, overlap, changing event-time support, and sensitivity estimates are reported alongside the main results to show where the comparisons are better supported and where they remain uncertain.

The results reveal a differentiated pattern. Averaged over supported event years 0–5, the private-establishment estimate is 1.59 log points—approximately 1.6%—above the weighted never-treated trajectory, with a similar estimate for total covered establishments. Estimates for private and total covered employment are positive but statistically imprecise and do not reach the 10% significance level. Population and cumulative net migration also have positive estimates in the main specification, but they reach the 10% rather than the 5% level and are more sensitive to alternative estimators and geographic exclusions. The establishment association remains positive under complementary estimators and leave-one-state-out comparisons, although it becomes smaller and less precise under higher-capacity entry definitions. Residual covariate imbalance and concentrated control weights in some early cohorts limit a strong causal interpretation. The evidence is therefore more consistent with a detectable difference in business presence following initial solar entry than with a broad, precisely estimated expansion in payroll employment or residential adjustment.

The remaining sections review the conceptual framework and literature, describe the study area and treatment audit, present the data and methods, and discuss the results, sensitivity analyses, policy implications, and limitations.

Chapter 2. Background and conceptual framework

2.1 Solar investment and the local retention of economic activity

Utility-scale solar is place-based in its physical location, but its economic footprint depends on where project spending, labor, and ownership are located. Development and construction can create demand for land-related services, site preparation, transportation, electrical installation, lodging, and other forms of contracting. Some of these activities may be supplied by firms and workers in the host county, while others may be provided by specialized contractors operating across county or state boundaries. The composition of activity also changes once a facility begins operating, when construction-related expenditures decline and a smaller set of operation and maintenance activities remains. Labor requirements vary across renewable technologies and project stages (Cameron and van der Zwaan, 2015), and green employment differs across occupations and locations (Curtis and Marinescu, 2023). County entry into utility-scale solar generation therefore marks exposure to a local development event, not a direct count of where every solar-related job or expenditure occurs.

The local response depends in part on how much of this activity is retained within the host economy. Project-related spending retained within the county can increase demand for other local goods and services, creating a possible local-multiplier channel (Moretti, 2010). Renewable-energy projects may also affect local public finances through property taxes, payments in lieu of taxes, or other arrangements that support public spending or tax relief. Evidence from wind development shows that such fiscal responses can form an important part of local adjustment, although the institutional treatment of solar projects varies across states and counties (Brunner and Schwegman, 2022b). Land use creates another channel. Utility-scale solar occupies physical space, and its local consequences depend on the prior use and economic value of the site, land ownership,

and proximity to residents and businesses. Ong et al. (2013) distinguish the land directly occupied by solar infrastructure from the broader area associated with a facility, underscoring that equal amounts of generating capacity can involve different local land-use conditions.

These channels can reinforce or offset one another. Local procurement, project-related income, and fiscal revenues may support county activity, while reliance on nonlocal suppliers, commuting workers, tax concessions, or displacement of alternative land uses may limit the amount retained locally. Project timing adds further variation. Because treatment is dated by the first audited capacity crossing, the recorded entry year need not coincide with the beginning of planning, contracting, or construction. Construction can produce a short-lived increase in demand, whereas operation generally supports a different and less labor-intensive set of recurring activities. Slattery et al. (2011) illustrate the importance of distinguishing construction from operation and local from statewide effects in the context of wind development. Their estimates do not determine the expected magnitude of a county-level solar response, but the distinction is useful for interpreting annual data, which may smooth brief construction peaks and more readily capture changes that persist across years. The net effect of solar entry on county business and labor outcomes is therefore theoretically ambiguous and likely to vary across places. The empirical analysis estimates reduced-form changes in county outcomes after entry rather than assigning those changes to any single mechanism.

2.2 Establishments and employment as distinct private-sector margins

Private establishments and private employment capture related but distinct dimensions of local economic activity. An establishment is an operating business location, whereas employment measures payroll jobs at those locations. Examining both outcomes within the private ownership domain keeps the comparison conceptually consistent: a private establishment count is paired with

employment measured over the same broad ownership sector. Total covered establishments and employment provide corresponding measures with broader ownership coverage and are used to assess whether the observed pattern depends on the private-sector restriction.

There is no reason to expect establishments and employment to move proportionately. The establishment stock reflects the net result of openings, closures, relocations, and organizational changes in active business locations. Employment can change through staffing decisions at both existing and newly observed establishments. An increase in establishments may therefore be accompanied by little aggregate employment growth if the additional locations are small or if job gains are offset by losses elsewhere. Conversely, employment may expand without an increase in establishment counts when incumbent businesses add workers. Annual measurement can also smooth short construction-related changes. Estimates of the labor required by renewable technologies help motivate the employment outcome, but they do not determine the magnitude or direction of county-wide payroll adjustment (Wei et al., 2010).

Broader research on place-based investment also suggests that a major project can affect incumbent firms through demand, labor-market, and agglomeration linkages (Greenstone et al., 2010). In the solar setting, those linkages could influence the survival, location, or staffing of businesses that are not themselves solar firms. Similar movements in private and total covered outcomes would indicate that the pattern is not confined to one ownership definition, whereas divergence would direct attention to differences in the populations covered by the measures. In either case, establishments and employment should be interpreted as separate margins: one captures the local presence of operating business locations, and the other captures the scale of payroll activity at covered workplaces.

2.3 Demographic adjustment and competing local channels

Residential adjustment is shaped by a broader set of forces than workplace activity alone. New or expanding employment opportunities may attract residents, but relocation also depends on expected earnings, housing availability, public services, amenities, and moving costs. Solar development may alter some of these conditions through land-use, fiscal, or amenity channels, and the direction of those effects need not be uniform across locations. Evidence on property values near utility-scale solar facilities, for example, shows that effects vary across property types and local settings (Hu et al., 2025). More generally, research on internal migration emphasizes that residential location decisions respond to multiple economic and non-economic factors and need not adjust immediately to a change in local labor demand (Molloy et al., 2011).

Population and cumulative net migration therefore provide complementary measures of demographic adjustment. Population is a stock that reflects natural increase as well as migration. Cumulative net migration more directly summarizes net residential movement relative to the pre-entry population, but it does not distinguish additional arrivals from fewer departures. Neither measure identifies who moves or whether a move is directly connected to solar-sector employment. The demographic outcomes are consequently interpreted as changes in county residence relative to the comparison trajectory, not as counts of workers relocating for a solar project.

The timing of demographic adjustment may differ from that of business or employment responses. Commuting can accommodate changes in workplace demand without a change in residence, and households may wait to relocate until new opportunities appear durable. Causality can also run in the other direction: population growth can expand local demand and support additional business activity. The simultaneous population-employment framework of Carlino and Mills (1987) cautions against imposing a fixed sequence between labor-market and residential

growth. For this reason, the analysis evaluates business, employment, population, and migration under a common solar-entry definition while estimating each outcome separately. The objective is to compare their timing and relative strength, not to infer one outcome mechanically from another.

2.4 Conceptual ordering and scope of inference

A useful organizing sequence begins with solar entry, followed by possible changes in local business activity, payroll employment, and, potentially, residential adjustment. This sequence is plausible but neither deterministic nor expected to hold uniformly. Some channels may bypass intermediate stages. Land-related payments or fiscal changes could influence household resources and residential decisions without creating additional establishments, while nonresident contractors could increase workplace activity without changing the resident population. Differences in the timing, magnitude, or precision of outcome-specific estimates can therefore be economically consistent rather than contradictory.

The empirical design treats these outcomes as separate reduced-form responses. The demographic specifications do not control contemporaneously for establishments or employment, and the analysis does not decompose responses into direct and indirect effects. Identifying mediation would require additional data and stronger assumptions about mediator-outcome confounding (Imai et al., 2010). The ordering is therefore interpretive, not a structural sequence from solar entry to business growth, employment, and migration.

Chapter 3. Literature review

3.1 Utility-scale solar and local economic outcomes

Research on the local economic consequences of renewable-energy investment has increasingly moved from engineering employment factors and input-output projections toward ex post comparisons of observed outcomes in host communities. Much of the U.S. county-level evidence has developed around wind power. Brown et al. (2012), studying wind development from 2000 to 2008, estimate modest increases in county employment and personal income associated with additional installed capacity. Brunner and Schwegman (2022a) find increases in GDP per capita, personal income, median household income, and housing values following commercial wind installations, but comparatively limited evidence of an increase in overall employment. These studies establish that renewable investment can coincide with economic changes along some county-level margins without similar responses along others. Their estimates, however, describe wind development under treatment definitions and local production relationships that cannot be applied directly to the initial entry of utility-scale solar generation.

Solar-specific evidence is more limited but particularly relevant to the question of how much project-related activity remains in the host community. Fabra et al. (2024) study renewable-energy investments in more than 3,900 Spanish municipalities during 2017–2021. They report stronger employment and unemployment responses to solar projects than to wind projects, particularly around construction, while finding that new jobs often accrue outside the municipality hosting the facility. They also document increases in local public spending and per capita income, although the broader economic gains are modest, especially for smaller and medium-sized plants. This combination of findings distinguishes employment supported by investment from

employment recorded within the host jurisdiction. The municipal setting and project-based exposure differ from the county-level first-entry comparison considered here.

Barrutiabengoa Ortubai and Ortiz Durán (2026) extend the Spanish evidence by combining installation records with provincial employment data over 2005–2024 and estimating dynamic responses to additions of renewable capacity. Their reported peak employment responses are approximately 10 jobs per MW for solar and 12 for wind. These are dynamic estimates of employment associated with capacity additions, not fixed numbers of permanent operating jobs. The authors also find differences by project phase, installation size, worker education, and economic connections across provinces, with employment per MW generally higher for smaller and medium-sized projects and during the earlier deployment period. The evidence is consistent with the importance of timing and local labor-market structure, while also showing why capacity-based employment responses should not be equated with changes in a county's aggregate payroll employment or establishment stock.

Evidence from the United States that separates solar and wind employment remains comparatively scarce. Fang (2022), in a master's thesis using Quarterly Census of Employment and Wages (QCEW) data and Energy Information Administration (EIA) plant records, distinguishes construction and operating periods and explores instrumental-variable specifications based on wind and solar resources. The thesis reports some evidence of larger employment responses for wind than for solar and identifies data and identification limitations, including the potential value of establishment-level employment records for further analysis. Because this is a master's thesis with acknowledged data and identification limitations, its findings are best treated as suggestive evidence rather than a peer-reviewed benchmark.

Other solar studies examine outcomes that may be locally important without directly measuring employment or the number of businesses. Hao and Michaud (2024), for example, analyze housing values around 70 utility-scale solar facilities built in the U.S. Midwest between 2009 and 2022. They find modest positive average property-value responses that differ by project size. Such findings broaden the set of possible local consequences of solar development, but housing prices, county payroll employment, and the stock of operating establishments are distinct empirical outcomes. The presence of a locally observed change in one cannot establish a change in the others.

The literature reviewed above points to two distinctions central to the present study. First, economic activity associated with construction or a renewable-energy supply chain need not translate into a similarly persistent change in employment or business presence within the host county. Second, solar-related outcomes depend on technology, project scale, stage, and the geographic unit over which they are measured. The studies reviewed here offer substantial evidence on jobs, income, public spending, and property values, but more limited direct evidence on how the stock of operating business establishments changes around a county's initial entry into utility-scale solar generation.

3.2 Utility-scale solar and demographic adjustment

Compared with employment and income, the residential consequences of utility-scale solar development have received less direct empirical attention. Workplace activity and residential change need not coincide: workers may commute, construction demand may be temporary, and household location choices depend on housing and other local conditions. These considerations motivate analysis of population and migration as outcomes in their own right, rather than as changes inferred from employment estimates.

Cary and Ortez (2026) provide an especially relevant U.S. comparison. Using county data for 2010–2023 and a staggered difference-in-differences design, they investigate population responses to utility-scale wind and solar facilities. Their preferred analyses associate wind development with population declines of approximately 2%–3% in host counties and report similarly sized spillovers in adjacent counties. By contrast, they do not detect a statistically discernible population response to solar development. They also examine rural–urban differences and report broadly similar wind-related responses across those settings. The solar finding should not be interpreted as proof that solar has no demographic effect; rather, the estimated response is not statistically distinguishable from zero in their design. Their treatment definitions, geographic scope, comparison groups, and population outcome also differ from the cumulative net-migration measure examined in this thesis.

County-level wind research offers a related perspective on the relationship between local economic gains and residential adjustment. Brunner and Schwegman (2022a) report improvements in several measures of local economic activity and income but find that wind installations do not reverse population decline. Shoeib et al. (2022) likewise find positive associations of additional wind capacity with per capita income, farm income, employment, and poverty reduction in their broader county analysis. Among the rural counties in their comparison, however, farm income is the clearest statistically detectable benefit, with less evidence of widespread economic or demographic improvement. Together, these results caution against treating an improvement in income or workplace activity as sufficient evidence of an accompanying population gain.

European studies reinforce the importance of location and the existing economic structure. García-Riazuelo et al. (2025) apply synthetic-control methods to five renewable-intensive rural European regions and examine population, employment, and GDP per capita. Their findings vary

across regions: renewable deployment does not consistently stem depopulation, and favorable employment and population responses within the regions are more evident in relatively affluent and densely populated subregions. These results concern a mix of wind and photovoltaic capacity in European regions, not a common solar-only effect in U.S. counties. San-Martín González and Soler-Vaya (2024), studying small rural municipalities in Spain's Valencia region, analyze a wider set of correlates of population change, including public services, economic diversification, and renewable energy. Their correlational design further illustrates that demographic trends reflect multiple local factors rather than a single renewable-energy intervention.

This literature provides two relevant lessons. Renewable-energy investment can be associated with changes in local economic conditions without a corresponding improvement in population trends, and demographic responses may vary substantially by technology and place. Direct evidence jointly examining resident population and net migration alongside operating establishments and workplace employment after the same solar-entry event nevertheless remains relatively limited in the studies reviewed here.

3.3 Economic and demographic adjustment in a common place-based framework

Considering economic and demographic outcomes together provides information that separate analyses of employment, income, or population may miss. The underlying relationships are not necessarily sequential: investment can change business demand without generating lasting workplace employment or residential growth, while population change can alter local markets without originating in a new renewable-energy project. Evidence that evaluates several outcomes in the same locations is therefore especially useful for assessing whether local responses move together or diverge.

Multi-outcome U.S. wind studies provide a useful benchmark for this approach. Brunner and Schwegman (2022a) and Shoeib et al. (2022) document uneven responses across economic and demographic outcomes, with patterns also varying by local context. Their findings illustrate the importance of examining different dimensions of local adjustment separately rather than treating economic development as a single outcome.

Shoeib et al. (2021) provide a complementary view using descriptive data, interviews, and documents from eleven rural U.S. counties hosting very large wind developments. They report that wind-related tax revenues can support community services without noticeable increases in service demand or the cost of living, while permanent employment opportunities remain limited. This evidence highlights the possibility that public finance, labor-market conditions, and residential demand respond differently to renewable investment. Its descriptive and qualitative components do not, however, identify a common causal treatment effect across the outcomes examined.

The European evidence makes a similar point across different spatial scales. García-Riazuelo et al. (2025) estimate long-run changes in GDP per capita, employment, and population, finding substantial variation among regions and within them. Benefits observed in one locality cannot be presumed to extend to nearby or economically different communities. Moreover, their combined wind and photovoltaic exposure, long-run regional comparisons, and economic context differ from U.S. county entry into utility-scale solar generation.

These studies help position the contribution of a common solar-entry design. Fabra et al. (2024) examine local employment, public spending, and per capita income in Spanish municipalities; Barrutiabengoa Ortubai and Ortiz Durán (2026) estimate dynamic employment responses to renewable-capacity additions across Spanish provinces; and Cary and Ortez (2026) examine U.S. population responses and spatial spillovers. Comparing results across these separate

studies cannot establish whether establishments, payroll employment, population, and migration respond similarly to a single entry event because their geographic units, samples, treatments, and estimands differ. Examining the four outcomes under a consistent county-level treatment definition permits a more coherent assessment of their direction, timing, and statistical precision while leaving the causal relationships among the outcomes unestimated.

3.4 Synthesis, research gap, and contribution of this study

The reviewed evidence establishes that renewable-energy investment may be followed by measurable changes in host-community economic conditions, but those changes are neither uniform across technologies nor confined to one outcome. Wind studies have contributed substantially to U.S. county-level evidence on income, employment, business-related activity, and population. Solar studies have provided increasingly detailed evidence on job creation, the geography of employment, public spending, income, and property values. Yet project-supported jobs, employment at covered workplaces, and the number of operating establishments describe different dimensions of local activity. Likewise, business and employment changes do not imply a predetermined response in resident population or migration.

Two comparatively underexamined questions motivate this study. The first is how the stock of operating establishments changes in counties entering utility-scale solar generation, rather than how many jobs are supported by solar construction or by the renewable-energy supply chain. The second is whether the measured paths of establishments, payroll employment, population, and net migration appear similar or different when they are evaluated around the same initial solar-entry event. These are relative gaps in the literature reviewed here, not a claim that previous renewable-energy research has never considered economic and demographic outcomes together.

This study addresses those questions using a historically audited county entry measure for utility-scale solar in seven western U.S. states. It distinguishes operating private establishments from private payroll employment and county resident population from cumulative net migration, with total covered establishments and employment providing broader corroborating measures. Because solar capacity accumulates and installation histories can precede the observed panel, separating already-exposed counties from new entrants is important for interpreting the entry event. The common first-entry definition and cohort-specific comparisons provide a consistent basis for describing outcome-specific trajectories without assuming that establishments cause employment growth or that payroll changes cause residential movement.

The contribution is therefore both substantive and empirical: it brings business presence into the analysis of solar entry alongside labor-market and demographic outcomes and examines their respective trajectories under a common, historically informed comparison design. Evidence on covariate balance, effective comparison support, pre-entry trends, and sensitivity analyses helps clarify the credibility and limits of these comparisons. The following sections describe the data, historical treatment audit, and estimation methods used to address these questions.

Chapter 4. Data and study area

4.1 Geographic setting and observational unit

The study covers all 252 counties and county equivalents in Arizona, California, Colorado, Nevada, New Mexico, Oregon, and Utah from 2010 through 2022. The seven-state region includes counties with solar-generation capacity established before the observation period, counties entering utility-scale solar generation during the panel, and counties that do not cross the study's capacity threshold by 2022. It also encompasses substantial variation in population size and economic structure, which is relevant both to potential local adjustment and to the construction of comparable untreated trajectories. Bolinger et al. (2023) document the broader national expansion and geographic diversification of utility-scale solar, although their plant-level capacity criterion differs from the county-level entry definition used here. The seven states constitute the defined study area rather than a representative sample of all U.S. counties.

Figure 1 presents the study area and the geographic distribution of audited solar-entry classifications. It distinguishes counties entering in 2011–2022 from never-treated counties, counties exposed before 2010, and counties entering in 2010 without an observed pre-entry baseline. The latter two groups are excluded from the main treatment comparison, as described in Section 5.

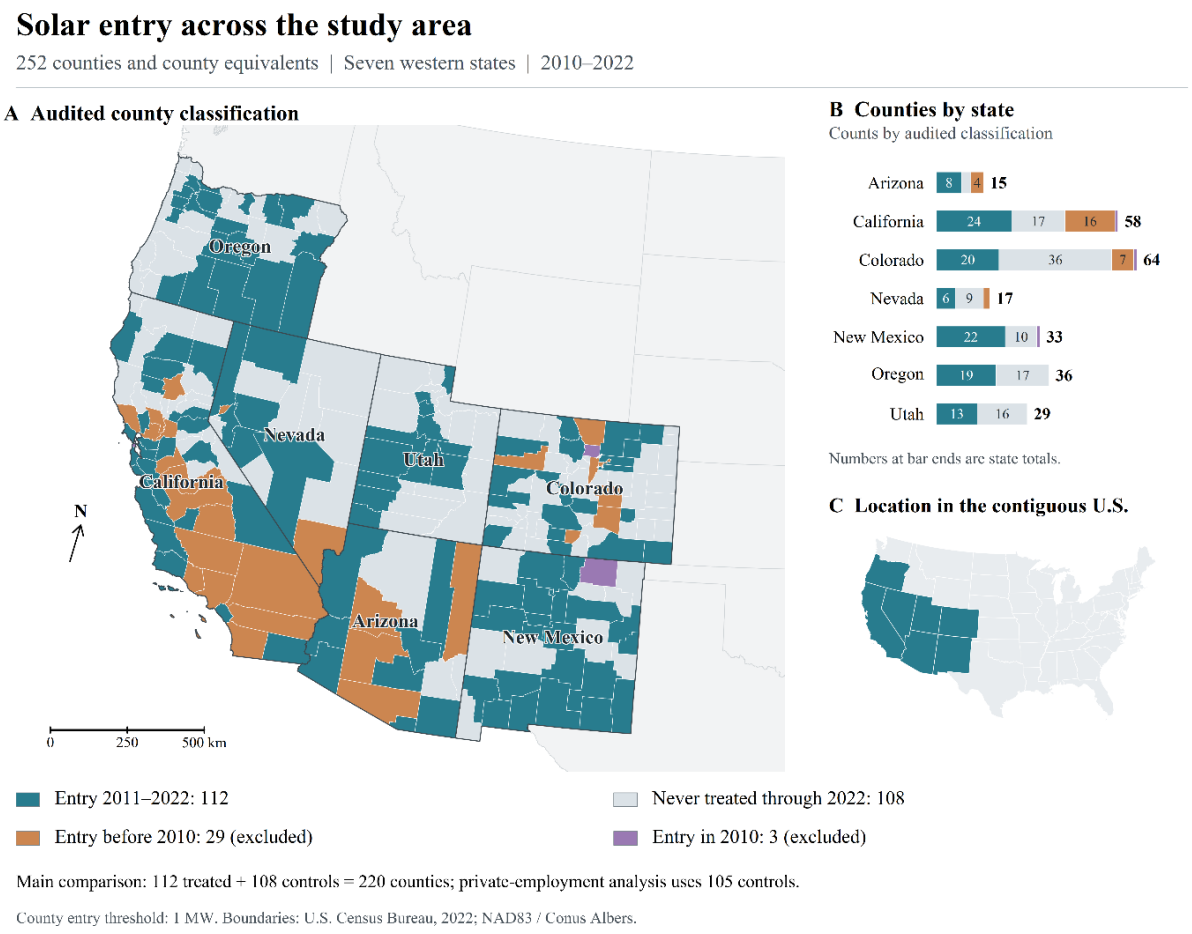


Figure 1. Study area and audited county solar-entry classification, 2010–2022.

The observational unit is the county-year. The full panel contains 3,276 county-year observations (252 counties over 13 years), before exclusions associated with treatment timing and outcome-specific data availability. Five-digit county Federal Information Processing Standards (FIPS) codes link annual observations to the audited solar-entry classifications. Treatment is defined at the county level, so additional generators within an already-treated county do not constitute separate treatment events. This geographic aggregation aligns solar exposure with the available economic and demographic outcomes, but it cannot identify variation in exposure within counties or fully capture commuting relationships and economic spillovers across county boundaries.

4.2 Business and labor-market measures

The business and labor-market outcomes are obtained from annual county-level records in the U.S. Bureau of Labor Statistics (BLS) Quarterly Census of Employment and Wages (QCEW). The data extract uses county geography, annual records identified by quarter code A, the QCEW total-industry aggregation code 10, and the private and total-covered ownership categories (codes 5 and 0, respectively). Four outcomes are constructed: private establishments, private employment, total covered establishments, and total covered employment. An establishment represents an operating business location rather than necessarily a distinct firm, while QCEW employment counts covered payroll jobs at their places of work rather than employed residents (U.S. Bureau of Labor Statistics, 2026b). Annual-average establishment counts are calculated from quarterly establishment counts, whereas annual-average employment reflects monthly employment levels over the year.

Private establishments are the primary business outcome, and private employment provides the corresponding labor-market measure within the same ownership domain. Total covered establishments and employment extend the comparison to a broader set of covered employers, including government establishments and employment. These measures do not capture all forms of local economic activity. In particular, QCEW coverage excludes unincorporated self-employment, unpaid family workers, and certain agricultural and domestic workers. Nor do county-wide establishment and employment totals distinguish solar-related firms or jobs from activity in other industries.

The administrative nature of QCEW data also affects interpretation. Changes in reported county totals can reflect actual business and employment adjustments, but they may also arise from changes in industry or geographic classification, employer reporting arrangements, and the

reorganization or relocation of establishments. BLS applies editing, validation, and, where necessary, imputation procedures before publication (U.S. Bureau of Labor Statistics, 2026a). The present analysis does not apply additional employment imputation in the R estimation workflow; this should not be interpreted as a claim that the published QCEW data contain no source-level imputation. Consistent annual industry and ownership definitions improve comparability across the panel, but the aggregate measures cannot identify the specific business-level processes underlying every observed change.

4.3 Population and migration

Resident population and annual net migration are obtained from the U.S. Census Bureau Population Estimates Program. The analysis combines the CO-EST2019-alldata series for 2010–2019 with the CO-EST2024-alldata series for 2020–2022, corresponding to the county population-estimate releases used in constructing the panel (U.S. Census Bureau, 2020, 2025). Population is the primary demographic outcome. Annual net migration is accumulated relative to the relevant last pre-entry year, and cumulative net migration is expressed as net migrants per 1,000 baseline residents, as defined in Section 5.4. Both measures are residence-based rather than workplace-based.

Because the demographic panel combines two Census vintages, the series is not treated as a single homogeneous release. Revisions to earlier estimates and the transition to the 2020-based population estimates can introduce discontinuities around 2020. The pre-COVID sensitivity analysis truncates the observation window before 2020 and therefore avoids both the pandemic-period observations and the splice between the two Census vintages. Differences in that sensitivity analysis, however, cannot be attributed separately to pandemic-related changes and vintage

comparability. No interpolation is applied to either demographic outcome in the estimation workflow, and population change is not used as a substitute for the reported net-migration measure. The Population Estimates Program updates a population base using births, deaths, domestic migration, and international migration. Successive vintages can revise estimates for earlier years, and consistency adjustments at subnational levels can produce differences between the sum of demographic components and total population change (U.S. Census Bureau, 2024). For this reason, the analysis uses the Census-reported net-migration component rather than deriving migration as a residual from population growth. This distinction is important because population change also reflects natural increase and statistical adjustments. The combination of vintages nevertheless remains a measurement limitation and should be kept in mind when interpreting demographic changes around 2020.

4.4 Baseline covariates and descriptive metropolitan classification

The main balancing covariates are population, income per capita, unemployment, and state. County personal-income statistics are obtained from the U.S. Bureau of Economic Analysis (BEA), and county unemployment rates are constructed from BLS Local Area Unemployment Statistics workbooks (U.S. Bureau of Economic Analysis, 2026; U.S. Bureau of Labor Statistics, 2025). For each cohort-specific comparison, the time-varying covariates are measured in the last pre-entry year ($g - 1$). Population and income per capita enter in logarithms, unemployment enters as a rate, and state is represented by indicator variables.

The expanded-covariate sensitivity additionally includes log total covered establishments and the overall CDC/ATSDR Social Vulnerability Index (SVI). The SVI series is constructed from the 2010, 2014, 2016, 2018, 2020, and 2022 releases using the overall ranking field, with linear interpolation between available releases and interpolation of the missing 2018 observation for Rio

Arriba County. Because SVI is a relative vulnerability ranking and its underlying definitions and rankings are not strictly comparable across releases, the interpolated series should not be interpreted as a fully comparable cardinal measure over time (Agency for Toxic Substances and Disease Registry [ATSDR], 2026). SVI is neither an outcome nor a covariate in the main specification. Moreover, although the value used in the sensitivity model is indexed to $(g-1)$, interpolation can depend partly on a release dated in or after the treatment-entry year. For this reason, the expanded specification is reported as an exploratory covariate sensitivity rather than as a fully predetermined balancing specification.

Metropolitan status is used only in the descriptive heterogeneity analysis. Counties are classified using the 2013 Rural-Urban Continuum Codes (RUCC), with codes 1–3 classified as metropolitan and codes 4–9 as nonmetropolitan (U.S. Department of Agriculture, Economic Research Service, 2013). The classification is held fixed throughout the descriptive analysis. The fixed-window sample contains 80 audited adopters from the 2011–2018 entry cohorts that are observed in $(g-1)$ and in every year from (g) through $(g+4)$. Of these counties, 39 are metropolitan and 41 are nonmetropolitan, and each county receives equal weight. This analysis characterizes differences between the two groups without identifying separate causal treatment effects by metropolitan status.

Chapter 5. Treatment and outcome construction

5.1 Audited county entry

The EIA utility-scale reporting convention applies at the generating-plant level and concerns plants with at least 1 MW of combined nameplate capacity (U.S. Energy Information Administration, 2026-a, 2026-b). The treatment used in this study is constructed at a different geographic level. Qualifying capacity from generator records identified by the EIA prime-mover codes PV and CP is aggregated to the county-year level to identify the initial entry of utility-scale solar generation. The EIA reporting convention therefore does not itself define county treatment. The 1 MW threshold is selected to provide a consistent and observable measure of initial county exposure to solar generation, using the conventional scale of EIA utility-scale reporting. It serves as an operational entry criterion rather than an economically determined threshold at which local effects are expected to begin. Because qualifying generator capacity is aggregated at the county level, crossing 1 MW does not necessarily indicate the construction of a single facility of that size or distinguish electricity consumed on-site from electricity delivered elsewhere. The analysis therefore focuses on initial county entry rather than the ownership, end use, or incremental effects of individual solar installations. Alternative capacity thresholds of 5, 10, 25, and 50 MW are examined in Section 8.3 to assess sensitivity to the definition of entry.

Let i index counties, t calendar years, and C_{it} the county's qualifying solar capacity in year t. The audited first-entry year and absorbing treatment indicator are defined as

$$G_i = \min\{t: C_{it} \geq 1 \text{ MW}\}, \qquad D_{it} = \mathbf{1}\{t \geq G_i\}. \tag{1}$$

The observed annual panel begins in 2010, so the first threshold crossing visible in that panel need not represent a county's true first exposure to qualifying solar capacity. The treatment-

timing audit therefore incorporates historical generator information when the panel alone is insufficient to establish first entry. A county that does not cross the threshold through 2022 is assigned $G_i = \infty$. Once a county enters treatment, it remains treated in the main design; additional generators within the county do not constitute new treatment events. The main design therefore evaluates adjustment following initial county entry into utility-scale solar generation rather than the marginal effect of an additional MW of capacity or differences across individual power plants.

5.2 Historical evidence and the estimation sample

The treatment-timing audit uses annual Form EIA-860 solar workbooks, generator operating-year information, and the EIA prime-mover codes PV and CP. When location information is inconsistent across annual records, the construction procedure resolves the inconsistency using the latest mappable location associated with the relevant plant-code and generator-ID combination. Historical generator information is then aggregated to the county level to identify qualifying solar exposure before 2010. These historical records are used to establish county treatment timing rather than to define multiple plant-level treatments.

The audited county-entry classification, rather than the raw annual capacity panel alone, defines treatment in the main analysis. The raw panel-based timing measure is retained only for the explicitly labeled panel-timing sensitivity analysis. This distinction is important because counties with qualifying solar exposure before the observation window should not be classified as untreated merely because their earlier entry is not visible in the 2010–2022 panel.

Relative to the panel-only timing measure, the historical audit reclassifies 29 counties as pre-treated based on evidence of threshold crossings before 2010. The remaining counties comprise 115 within-panel adopters and 108 never-treated counties, and none of the 252 counties has an ambiguous pre-2010 threshold classification based on the available generator records (Table

1). Among the 115 within-panel adopters, three enter in 2010 and therefore lack the required 2009 pre-entry baseline. Excluding the 29 pre-treated counties and these three 2010 entrants leaves 112 treated counties in 12 estimable entry cohorts from 2011 through 2022, together with 108 never-treated controls before outcome-specific exclusions. The cohort sizes for 2011–2022 are 17, 7, 8, 10, 9, 9, 11, 9, 12, 6, 9, and 5, respectively. Previously treated counties are not used as untreated controls in comparisons for later entry cohorts.

Table 1. Audited county treatment classification, 2010–2022

State	Never treated	Pre-treated	Panel adopters
Arizona	3	4	8
California	17	16	25
Colorado	36	7	21
Nevada	9	2	6
New Mexico	10	0	23
Oregon	17	0	19
Utah	16	0	13
All counties	**108**	**29**	**115**

Notes: Pre-treated counties entered before 2010. Panel adopters include three counties entering in 2010 that lack a 2009 baseline. The main estimation sample therefore contains 112 adopters entering during 2011–2022 and 108 never-treated controls before outcome-specific exclusions. Private-employment models use 105 never-treated controls because of the zero-valued private-employment observations described in Section 5.3.

Four county-year capacity observations in the annual panel are missing: Ventura County, California, in 2010 and 2011, and Tooele County, Utah, in 2021 and 2022. These observations are not recoded as zero. Treatment in the main analysis follows the audited county-entry classification, while the complete-capacity sensitivity excludes counties with any missing capacity observations.

The replication materials begin with the final audited datasets and verify the consistency of the treatment assignments used in estimation. They do not independently reconstruct the historical generator records from the original EIA workbooks.

5.3 Stocks and workplace employment

For a strictly positive county-year outcome Q_{it} —private establishments, private employment, total covered establishments, total covered employment, or resident population—the cohort-specific transformed outcome is

$$Z_{it}^{Q}(g) = 100\{\log Q_{it} - \log Q_{i,g-1}\}. \quad (2)$$

Equation (2) gives the county's log change from the calendar year before cohort g enters treatment ($g-1$). Never-treated controls use that same calendar baseline in each cohort-specific comparison. The transformed outcome at event time −1 is zero by construction, not an estimated pre-entry effect. Event year 0 is not used as the reference period because it marks the first year of audited solar entry and may already reflect economic responses to solar development, which would otherwise be absorbed into the baseline comparison. The last pre-entry year is used as the common reference because it provides a directly observed, cohort-specific baseline while retaining early adopting counties in the estimation sample. Using an average of several pre-entry years would require additional historical observations unavailable for the earliest cohorts and would change the outcome transformation. However, reliance on a single baseline year does not eliminate concerns about temporary pre-entry shocks or anticipatory responses before recorded solar entry. Estimates are in log points: one log point corresponds approximately to 1% at the magnitudes observed here. More exactly, a contrast of β log points corresponds to $100[exp(\beta/100) - 1]$ percent when translated from logs. Exponentiating a log contrast gives the corresponding relative geometric-change contrast, not a percentage change in the arithmetic mean county count. For private establishments and population, the exact-percentage sensitivity instead uses $100\left(Q_{it}/Q_{i,g-1} - 1\right)$.

The final analysis panel contains 18 zero-valued private-employment observations, for which the logarithm is undefined: Hinsdale County, Colorado, in 2017–2021 (five); Esmeralda

County, Nevada, in 2010–2021 (12); and Wheeler County, Oregon, in 2022 (one). All three are never-treated counties, and their full county-year paths are excluded from private-employment weighting and regression models, leaving 105 never-treated controls and the same 112 treated counties. Zero values are not replaced with one, and no additional private-employment values are imputed. None of the three counties is selected in the common matched design, so its 81 distinct controls have valid private-employment paths. Other outcomes retain counties with valid data.

5.4 Cumulative net migration

Let M_{it} denote annual net migration, P_{it} resident population, and $A_{it} = \sum_{s=2010}^{t} M_{is}$ the running sum of annual net migration beginning in 2010. For an entry cohort g, the transformed migration outcome is

$$Z_{it}^{M}(g) = 1000 \frac{A_{it} - A_{i,g-1}}{P_{i,g-1}}. \tag{3}$$

At the baseline year $(g-1)$, Equation (3) equals zero by construction. Never-treated controls are evaluated relative to the same cohort-specific calendar baseline, $g-1$. For $t \geq g$, its numerator accumulates net migration from g through t. For $t < g-1$, the numerator equals the negative of net migration accumulated from $t+1$ through $g-1$. The denominator is fixed at county population in the cohort-specific baseline year. The unit is cumulative net migrants per 1,000 baseline residents, not an annual migration rate.

The headline average combines cumulative migration measures evaluated at different post-entry horizons; it is not the cumulative change exactly five years after entry for every county. Because pre-entry values are signed accumulations relative to the baseline, a lead coefficient does not isolate migration in a single earlier year.

Chapter 6. Empirical methodology and identification

6.1 Target parameters and conditional comparison

Let $Q_{it}(g)$ and $Q_{it}(\infty)$ denote the potential outcome for county i in year t under first county-level solar entry in year g and under no entry through the observation period, respectively. Applying the outcome-specific transformations in Equations (2) and (3) to these potential-outcome paths yields $Z_{it}(g; g)$ and $Z_{it}(\infty; g)$. In both expressions, the second argument identifies the entry cohort whose last pre-entry calendar year, $g - 1$, serves as the common baseline. Under no anticipation, entry in year g does not affect the outcome at $g - 1$.

For post-entry years $t \geq g$, the cohort-time target is

$$ATT(g, t) = E[Z_{it}(g; g) - Z_{it}(\infty; g) \mid G_i = g], \qquad t \geq g. \tag{4}$$

Equation (4) gives the average difference for counties entering in cohort g between their transformed outcome under entry and the outcome they would have experienced without entry. The latter is unobserved. Analogous pre-entry contrasts serve as placebo comparisons and are not interpreted as treatment effects. In the estimation equations below, $Z_{it}(g)$ denotes the observed outcome transformed using cohort g's baseline.

The conditional parallel-trends assumption requires that, absent solar entry, cohort-g adopters and never-treated counties would have experienced the same expected changes relative to $g - 1$, conditional on the baseline covariates. This motivates reweighting eligible never-treated counties by baseline characteristics rather than presuming that their unadjusted mean is an adequate counterfactual. Reweighting does not itself establish conditional parallel trends; the identifying assumptions and relevant diagnostics are discussed in Section 6.8.

6.2 Baseline balancing, propensity scores, and ATT weights

For each entry cohort g, the baseline covariate vector $X_{i,g-1}$ comprises log resident population, log income per capita, the unemployment-rate level, and six state indicators, with Arizona as the reference category. Population reflects county scale and market size, which may be related to solar siting opportunities and subsequent economic growth. Income characterizes purchasing power and local economic conditions, while unemployment captures labor-market slack that may be associated with investment opportunities and later employment adjustment. State indicators accommodate broad differences in policy and institutional settings. These variables provide substantive grounds for baseline adjustment, but they do not exhaust the determinants of solar adoption.

Covariates are chosen using a substantive pre-treatment rationale rather than an automated search for variables that improve outcome fit. Guidance on confounder selection emphasizes the role of causal knowledge in adjustment decisions (VanderWeele, 2019). Baseline population, income, unemployment, and state may be related to both solar entry and subsequent local adjustment, but their inclusion does not identify every underlying siting mechanism. No contemporaneous or post-entry business, employment, population, or migration measure is added to the balancing specification.

Within each cohort-specific risk set, continuous covariates are standardized using the mean and sample standard deviation calculated across that cohort's adopters and eligible never-treated counties. All time-varying covariates are measured in the last pre-entry year, $g-1$, rather than contemporaneously with post-entry outcomes. Under the maintained no-anticipation assumption, pre-entry outcome levels can be used in balancing or matching because they precede treatment, whereas post-entry values may already have responded to solar development and should not be

conditioned on when estimating its overall association. The risk-set matching analysis uses additional pre-entry levels and slopes, as described in Section 6.6. The interpolated SVI, included only in the expanded-covariate sensitivity, has the separate timing and comparability limitations noted in Section 4.4.

The cohort-specific propensity score is

$$p_g(X_i) = \Pr\big(G_i = g \mid G_i \in \{g, \infty\}, X_{i,g-1}\big). \tag{5}$$

It is estimated within the risk set consisting only of the current entry cohort and eligible never-treated counties. A logistic model with an intercept and the baseline covariates is fitted by Newton–Raphson. A ridge penalty of 10^{-4} is imposed on the slope coefficients, but not the intercept, to stabilize separated state–cohort comparisons. The penalty is fixed and is not selected according to outcome results. This is a finite-sample stabilized propensity-weighting implementation; estimating a score does not guarantee adequate balance or overlap (Rosenbaum and Rubin, 1983; Abadie, 2005).

For each eligible never-treated control county j, fitted probabilities are bounded and converted into capped odds:

$$\bar{p}_{jg} = \min\{0.98, \max\big(0.02, \hat{p}_{jg}\big)\}, \qquad q_{jg} = \min\{\bar{p}_{jg}/\big(1 - \bar{p}_{jg}\big), 20\}. \tag{6}$$

The 0.02 lower bound and 0.98 upper bound limit extreme fitted scores, and the second restriction caps the odds at 20. Neither operation removes observations or creates genuine common support when comparable controls are lacking. Let $\mathcal{C}_Q$ denote the eligible never-treated pool for outcome Q. Let n_g be the number of counties in cohort g. The normalized ATT weights are

$$w_{jg} = n_g \frac{q_{jg}}{\sum_{h \in \mathcal{C}_Q} q_{hg}}, \qquad \sum_{j \in \mathcal{C}_Q} w_{jg} = n_g. \tag{7}$$

The control weights sum to the treated cohort size. Thus, the weighted control mean uses the normalized weights w_{jg}/n_g and is directly comparable to the equally weighted treated mean. For private employment, the propensity model and weights are re-estimated using the separate pool of 105 never-treated counties; this is not implemented by merely dropping three counties after calculating weights with the 108-control pool. Appendix Table A1 reports cohort-level effective sample size (ESS) and maximum control weights. Appendix Figure A1 displays fitted-score distributions before bounding so that extreme fitted probabilities remain visible in the overlap assessment (Austin and Stuart, 2015).

6.3 Cohort-time estimation, event-time aggregation, and interpretation

The primary cohort-time inverse-probability-weighted (IPW) estimator compares the transformed outcome for counties in entry cohort g with the corresponding weighted change among eligible never-treated counties:

$$\widehat{ATT}^{IPW}(g,t) = \frac{1}{n_g} \sum_{i:G_i=g} Z_{it}\,(g) - \frac{1}{n_g} \sum_{j \in \mathcal{C}_Q} w_{jg}\, Z_{jt}(g). \tag{8}$$

Equation (8) subtracts the normalized, weighted mean change of eligible never-treated controls from the equally weighted mean change among counties entering in year g. The outcome transformation uses the common cohort-specific baseline $g-1$ defined in Sections 5.3–5.4, and control weights are constructed as in Equation (7). Other treated cohorts never enter the comparison pool. The cohort-time organization follows the staggered group-time framework of Callaway and Sant'Anna (2021), while Equations (5)–(8) describe the particular stabilized weighting implementation used here. The resulting estimate is a weighted comparison of changes;

its interpretation as a causal cohort-time ATT depends on the identifying assumptions discussed in Section 6.8.

Never-treated counties are selected as the preferred comparison group because they remain below the audited entry threshold throughout the observation period, providing a consistent pool of untreated counties across cohort-time comparisons. This choice also avoids including future adopters as controls when project planning or construction may influence outcomes before their recorded entry year. Although not-yet-treated counties could provide additional comparisons under appropriate identifying assumptions, their inclusion would require time-varying control eligibility and change the comparison design. The use of never-treated counties nevertheless does not eliminate potential differences between adopting and comparison counties, making covariate balance, overlap, and the conditional parallel-trends assumption important for interpretation.

The main analysis uses IPW as a weighting-based benchmark that retains all estimable entry cohorts and adjusts the never-treated comparison trajectory using observed pre-entry characteristics. This approach permits direct assessment of achieved covariate balance and weight concentration without specifying a separate outcome-regression model. Doubly robust adjustment and risk-set matching provide complementary comparisons, although matching requires a more restricted set of entry cohorts. The preference for IPW does not imply that weighting eliminates residual imbalance, limited overlap, or unobserved determinants of solar siting.

Event time is $k = t - g$. Let G_k denote the set of estimable entry cohorts with the required observations in calendar year $g + k$. The corresponding event-time coefficient is

$$\hat{\theta}_k = \frac{\sum_{g \in \mathcal{G}_k} n_g \widehat{ATT}(g,g+k)}{\sum_{g \in \mathcal{G}_k} n_g}. \tag{9}$$

Equation (9) weights each supported cohort in proportion to its number of treated counties. The event-study figures display event years $k = -5, \ldots, +8$, with the last pre-entry year ($k = -1$) normalized to zero by construction. Cohort support changes across event years: later post-entry lags exclude recent adopters, while more distant pre-entry leads exclude early adopters. Appendix Table A4 reports the treated-county and cohort counts underlying the event-time paths. Changes in the plotted coefficients should therefore not be interpreted as changes experienced by an unchanged set of counties.

The headline parameter averages supported cohort-time estimates from entry through event year +5:

$$\hat{\theta}_{0:5} = \frac{\sum_g \sum_{k=0}^{5} \mathbf{1}\{g+k \le T\} n_g \widehat{ATT}(g,g+k)}{\sum_g \sum_{k=0}^{5} \mathbf{1}\{g+k \le T\} n_g}, \quad T = 2022. \tag{10}$$

In Equation (10), T = 2022, and the average includes only observed cohort-event cells with $g + k \le T$ for event years $k = 0, \ldots, 5$. Cohort-event cells receive weights proportional to treated cohort size, so the headline is a treated-county-event-weighted average rather than the equally weighted mean of the six plotted post-entry coefficients. A recent entry cohort contributes fewer observable post-entry years. The complete-horizon sensitivity retains only cohorts observed at every event year from 0 through 5; the pre-COVID specification instead uses $T = 2019$. These restrictions alter the cohorts and post-entry horizons represented by the estimate, not merely its standard error. The event-years 0–5 window is selected to summarize adjustment during the initial and subsequent years following solar entry while retaining the broader set of adopting cohorts. Requiring complete follow-up through event year +5 would restrict the analysis to 71 treated counties entering during 2011–2017, compared with 112 adopters in the main specification. The

complete-horizon sensitivity analysis therefore assesses whether the results remain similar when attention is restricted to cohorts observed throughout the same post-entry window. Because construction activity may precede audited solar entry, while operation and maintenance continue afterward, the event-years 0–5 average does not separately identify responses associated with these project stages. Accordingly, the estimate summarizes early post-entry adjustment rather than establishing whether any observed economic response is temporary or persistent.

The interpretation operates at three levels. Within a county, the outcome transformation measures change relative to the last pre-entry year. Within a cohort and calendar year, IPW compares that change with the weighted change among eligible never-treated controls. Across supported cohorts and post-entry years, the headline parameter averages these cohort-time contrasts, weighting each observed cohort-event cell by the number of treated counties in its cohort. A positive estimate means that the treated trajectory rises more, or falls less, than the weighted comparison trajectory; it does not imply that every treated county experiences positive absolute growth. This distinction between comparison units and aggregation is consistent with the interpretive emphasis in Callaway and Sant'Anna (2021) and Roth et al. (2023).

6.4 Bootstrap inference and dependence

Main inverse-probability-weighted (IPW) inference uses 499 stratified county-bootstrap draws. In each draw, treated counties are sampled with replacement separately within their original entry cohorts, preserving cohort sizes. Eligible never-treated counties are sampled once with replacement to form a common control pool reused across the cohort comparisons in that draw. The resampling retains each county's full annual observation path. Baseline covariate standardization, propensity-score estimation, score bounds, odds caps, and weight normalization

are recalculated within each bootstrap draw. Private employment uses a separate set of bootstrap draws based on its eligible pool of 105 never-treated controls.

For a reported parameter $\hat{\psi}$ and B successful bootstrap estimates, the standard error is

$$\widehat{SE}(\hat{\psi}) = \left[\frac{1}{B-1}\sum_{b=1}^{B}\left(\hat{\psi}^{*(b)} - \overline{\psi}^{*}\right)^2\right]^{1/2}. \quad (11)$$

Equation (11) gives the sample standard deviation across successful bootstrap estimates. The 95% confidence intervals use the normal reference, $\hat{\psi} \pm z_{0.975}\widehat{SE}$, and two-sided p-values use $2\Phi\left(-\left|\hat{\psi}/\widehat{SE}\right|\right)$. The intervals are not bootstrap percentile intervals, and the reported event-study confidence intervals are pointwise rather than simultaneous bands.

Throughout the manuscript, three asterisks denote $p < 0.01$, two denote $p < 0.05$, and one denotes $p < 0.10$. Each estimate receives the strongest applicable mark, based on its unrounded p-value. A single star therefore indicates significance at the 10% level but not the 5% level; no star indicates $p \geq 0.10$ when a test is defined. Because the confidence intervals are reported at 95%, an estimate carrying a single star can have a displayed interval that includes zero. Significance statements identify the relevant level. Statistical thresholds do not establish substantive importance; magnitudes, uncertainty intervals, and design diagnostics remain central to interpretation (Wasserstein and Lazar, 2016).

The calculations retain all finite successful bootstrap draws and report their counts; all 499 requested draws succeeded in the reported run. The base random seed is 7,352,026. Resampling complete county time paths preserves within-county temporal dependence in the bootstrap procedure, an important concern for difference-in-differences inference (Bertrand et al., 2004). Reusing the same sampled never-treated pool across entry cohorts also preserves the dependence arising from common controls in the cohort-specific comparisons. No separate formal

autocorrelation test is reported; temporal dependence within counties is addressed through the county-path resampling procedure. However, this bootstrap does not explicitly model spatial dependence across different counties. The reported standard errors remain conditional on the sampling and comparison assumptions and do not account for arbitrary spatially correlated disturbances, all possible common local shocks, or measurement error.

6.5 Doubly robust adjustment

For each entry cohort g and calendar year t, an ordinary least-squares (OLS) outcome regression $\hat{m}_{g,t}(X)$ is fitted among eligible never-treated counties to predict the cohort-normalized outcome using the same baseline covariates as the main propensity-score model in Section 6.2. Define the residual $\hat{R}_{it}(g) = Z_{it}(g) - \hat{m}_{g,t}(X_i)$. The doubly robust weighted residual contrast is

$$\widehat{ATT}^{DR}(g,t) = \frac{1}{n_g}\sum_{i:G_i=g} \hat{R}_{it}(g) - \frac{1}{n_g}\sum_{j\in\mathcal{C}_Q} w_{jg}\,\hat{R}_{jt}(g). \quad (12)$$

Equation (12) subtracts the fitted control-outcome surface from the observed normalized outcomes before comparing the treated-cohort mean with the weighted never-treated mean. The control-outcome regression includes an intercept, standardized baseline covariates, and state indicators. Aliased all-zero columns are handled by QR decomposition without changing fitted control predictions. The residual contrast uses the normalized, capped control weights defined in Equation (7); cohort-event aggregation follows Equations (9) and (10).

The 499-draw county bootstrap refits the propensity-score and outcome-regression models in each replication, as described in Section 6.4. The procedure combines propensity weighting and regression adjustment in the spirit of Sant'Anna and Zhao (2020), but retains the score bounds and capped odds used in the main IPW specification. Consequently, the usual double-robustness claim of consistency when either nuisance model is correctly specified cannot be asserted without qualification: persistent score truncation can invalidate the propensity-correct branch of that result,

and limited overlap can require extrapolation by the outcome model. These estimates provide an alternative adjusted comparison rather than eliminating concerns about incomplete comparability.

6.6 Risk-set matching

Risk-set matching provides an alternative comparison design to the main inverse-probability-weighted (IPW) analysis. It restricts treated counties to the 2013–2022 entry cohorts so that three pre-entry calendar years ($g-3, g-2, g-1$) are available. Each cohort is compared only with never-treated counties. Six matching features are measured before entry: log population, log income per capita, the unemployment-rate level, and log private establishments at $g-1$, together with ordinary least-squares slopes of log population and log private establishments over $g-3$ through $g-1$. The two slopes summarize pre-entry trajectories rather than post-entry responses. All six features are standardized using the combined cohort-treated and eligible never-treated risk set.

Within each state, the distance between treated county *i* and candidate never-treated control *j* is defined as

$$d_{ijg} = \left[\sum_{\ell=1}^{6}\left(V_{i\ell g}^{*} - V_{j\ell g}^{*}\right)^{2}\right]^{1/2}, \tag{13}$$

Here, V^{*} denotes the six standardized matching characteristics. Each treated county is assigned its three nearest never-treated controls within the same state, with no distance caliper. Controls can be reused across treated counties; exact ties in distance are resolved by county FIPS code. Let c_{jg} denote the number of times control *j* appears among the $3n_g$ links for cohort g. The cohort-time matched estimate is

$$\widehat{ATT}^{M}(g,t) = \frac{1}{n_g}\sum_{i:G_i=g} Z_{it}\,(g) - \frac{1}{3n_g}\sum_j c_{jg}\, Z_{jt}(g). \quad (14)$$

Equation (14) compares the equally weighted transformed outcomes of adopters in cohort g with the link-weighted average of their matched control outcomes. When a control is selected for more than one treated county, its contribution is weighted by its number of selections. The common matching design contains 88 treated counties in ten entry cohorts, 81 distinct never-treated controls, and 264 treated–control links (three per treated county). The same matching assignments are used across outcomes. This design changes both the comparison construction and the set of treated cohorts relative to the main IPW analysis. The IPW restriction to 2013–2022 cohorts in Tables 4–5 helps distinguish changes attributable to cohort coverage from differences between comparison methods, but does not make the estimators equivalent.

Matched inference uses 399 bootstrap draws, resampling treated counties within each entry cohort while retaining each sampled county's three original control links. Matches are not recomputed in these draws. Inference is therefore conditional on the realized matching design and does not account for uncertainty from selecting the nearest neighbors. It is not presented as a generally valid unconditional bootstrap procedure for nearest-neighbor matching. Abadie and Imbens (2008) establish that the ordinary bootstrap can be invalid for fixed-neighbor matching estimators, reinforcing the need for this qualification. Moreover, shared controls can induce dependence across treated-county contrasts that the treated-only resampling does not fully capture as independent control-sampling uncertainty.

The estimates from these alternative designs are compared in Section 8.1 and Appendix Table A5. Appendix Table A6 reports control reuse, effective sample size, and matching-distance summaries. A common matching sample does not imply uniformly close matches, so these diagnostics remain important when assessing the matched comparisons.

6.7 Descriptive heterogeneity by metropolitan status

The metropolitan-status analysis describes variation in post-entry outcomes among audited adopting counties rather than estimating metropolitan-specific solar-entry treatment effects. Following the fixed 2013 Rural-Urban Continuum Codes classification described in Section 4.4, codes 1–3 identify metropolitan counties and codes 4–9 identify nonmetropolitan counties. The sample consists of 80 adopters in the 2011–2018 entry cohorts, each observed in the last pre-entry year ($g_i - 1$) and throughout the common post-entry window (g_i through $g_i + 4$). Of these counties, 39 are metropolitan and 41 are nonmetropolitan. Holding the follow-up window fixed avoids comparing group averages based on different post-entry horizons. These are adopter-only descriptive comparisons: no never-treated counterfactual, propensity weighting, or covariate adjustment is used.

For private establishments, private employment, and resident population, let Y_{it} denote the relevant outcome for county i in year t, and let g_i denote county i's audited entry year. The county-level descriptive measure is the average of the five annual log-point changes relative to the last pre-entry year:

$$\overline{\Delta Y}_i = \frac{1}{5}\sum_{k=0}^{4} 100\left[\ln\left(Y_{i,g_i+k}\right) - \ln\left(Y_{i,g_i-1}\right)\right],$$

Cumulative net migration is measured differently. Annual net migration is summed from entry through the fourth post-entry year and scaled by population in the last pre-entry year:

$$C_i = 1000 \times \frac{\sum_{k=0}^{4} M_{i,g_i+k}}{P_{i,g_i-1}},$$

Figures 5 and 6 present the resulting county-level measures and equally weighted distributions for metropolitan and nonmetropolitan adopters. Two-sided Welch tests compare the

group means while allowing unequal variances. Because four outcomes are examined and the p-values are not adjusted for multiple comparisons, these tests are interpreted as exploratory descriptive evidence. The analysis documents descriptive differences among adopting counties and does not identify differential causal effects of solar entry by metropolitan status.

6.8 Identification assumptions and diagnostics

A causal interpretation of the cohort-time contrasts requires conditional parallel trends: for each entry cohort, absent solar entry, adopting counties and comparable never-treated counties would have experienced the same expected changes in the transformed outcomes, conditional on the baseline covariates. Cohort-specific reweighting improves comparability on observed characteristics but cannot eliminate differences arising from unmeasured local shocks, anticipated growth, transmission and interconnection opportunities, or other determinants of solar siting and economic adjustment. State indicators capture broad differences across the seven states but not every time-varying state or county shock. A staggered difference-in-differences estimator does not, by itself, establish this identifying assumption.

No anticipation further requires that the relevant potential outcomes and pre-entry balancing information not respond to solar development before the audited entry year. Project planning, permitting, construction, and expectations about future investment can precede recorded generation. Pre-entry event-study contrasts offer diagnostics for some departures from this condition, but annual observations and baseline normalization may leave anticipatory responses undetected. Spillovers pose a separate concern: commuting, contracting, relocation, and shared infrastructure may connect treated counties with never-treated comparison counties. Thus, a county without its own qualifying solar entry need not be economically unaffected by development

elsewhere. Related renewable-energy research highlights the relevance of geographic spillovers (Cary and Ortez, 2026).

Identification also depends on sufficient overlap between each adopting cohort and its eligible never-treated comparison pool. Bounding estimated propensity scores and capping control odds stabilize the calculations, but they do not create comparable controls where support is weak. Persistent truncation changes the effective weighting rule; even a correctly specified uncapped propensity model does not automatically confer the conventional ATT identification guarantee after such truncation. The estimates are therefore interpreted as stabilized weighted contrasts intended to approximate cohort-specific ATT parameters under adequate comparability, with their credibility assessed alongside balance, weight concentration, and pre-entry diagnostics. Finally, the absorbing first-entry indicator pools counties with different project characteristics and subsequent capacity paths. Threshold and persistence checks examine this heterogeneity without fully resolving it. These identification concerns motivate the diagnostic framework used to assess balance, overlap, weight concentration, and pre-treatment comparability.

For each baseline covariate, the standardized mean difference (SMD) is the treated mean minus the unweighted or reweighted control mean, divided by the square root of the average of the unweighted treated and control variances. The same unweighted denominator is retained before and after weighting within each cohort, so changes in the SMD reflect changes in mean balance rather than changes in the standardizing scale. Figure 2 displays the simple average across cohorts of the absolute SMDs; Appendix Table A2 also reports the maximum absolute post-weighting difference. Both summaries matter because a low average can conceal a poorly balanced cohort. SMDs are descriptive balance diagnostics, not statistical tests of whether the groups are equivalent (Austin, 2009).

Weight concentration is summarized by the cohort-specific effective sample size (ESS) of the eligible, normalized never-treated control weights:

$$ESS_g = \frac{\left(\sum_{j \in \mathcal{C}_Q} w_{jg}\right)^2}{\sum_{j \in \mathcal{C}_Q} w_{jg}^2}. \tag{15}$$

The ESS is the number of equally weighted controls that would yield the same degree of weight concentration. Consequently, a large nominal comparison pool can provide limited effective support when a small number of counties receive most of the weight. Appendix Table A1 reports cohort-specific ESS alongside maximum weights and the share of fitted propensity scores at the imposed bounds. These diagnostics should be considered jointly with achieved covariate balance; weighting stability alone does not establish overlap or identification.

To assess pre-entry contrasts jointly, let the vector of estimated lead coefficients contain event years $-5, -4, -3,$ and -2, and let its estimated covariance matrix be obtained from the bootstrap. The reported joint statistic is

$$W = \hat{\theta}_{pre}{}' \hat{\Sigma}_{pre}^{+} \hat{\theta}_{pre}, \qquad W \sim \chi_r^2 \text{ (under the null)}, \tag{16}$$

In Equation (16), the superscript + denotes the Moore–Penrose generalized inverse of the covariance matrix and r denotes its rank. The chi-square distribution with r degrees of freedom provides the reference distribution under the null; the reported tests use reference rank four. Event year −1, normalized to zero by construction, is excluded. Appendix Table A3 presents the joint lead tests for all six outcomes and all three estimators.

Failure to reject zero pre-entry contrasts does not validate conditional parallel trends, particularly when leads are imprecisely estimated (Roth, 2022). A rejection raises concern about the comparison but does not, by itself, identify the source of the deviation. Specifications are neither selected nor discarded based on the joint test. Formal sensitivity bounds of the kind

proposed by Rambachan and Roth (2023) are not implemented; nonrejection is therefore not treated as evidence that the post-entry contrasts are robust to unobserved differences in untreated trends.

Chapter 7. Results

7.1 Baseline comparability and achieved balance

Before entry, adopting counties are substantially larger than never-treated counties. Averaging the cohort-specific pre-entry baseline means, adopters have approximately 6,224 private establishments and 205,989 residents, compared with 1,091 establishments and 34,600 residents in the unweighted comparison pool. The corresponding IPW-weighted control means are 2,548 establishments and 91,480 residents. These quantities summarize cohort-specific baseline comparisons rather than statistics from a single common baseline year. Weighting brings the comparison group closer to adopters on these dimensions without eliminating the differences in county scale.

Figure 2 reports improved but incomplete balance on the baseline covariates. The mean absolute standardized difference across cohorts falls from 0.982 to 0.218 for log population, from 0.307 to 0.122 for log income per capita, and from 0.300 to 0.177 for unemployment. These averages do not, however, describe the worst-balanced cohort. The largest residual cohort-specific standardized differences are 0.584 for log population, 0.448 for unemployment, and 0.738 for the New Mexico indicator (Appendix Table A2). Thus, the IPW design yields a more comparable weighted control trajectory, not a fully balanced counterfactual.

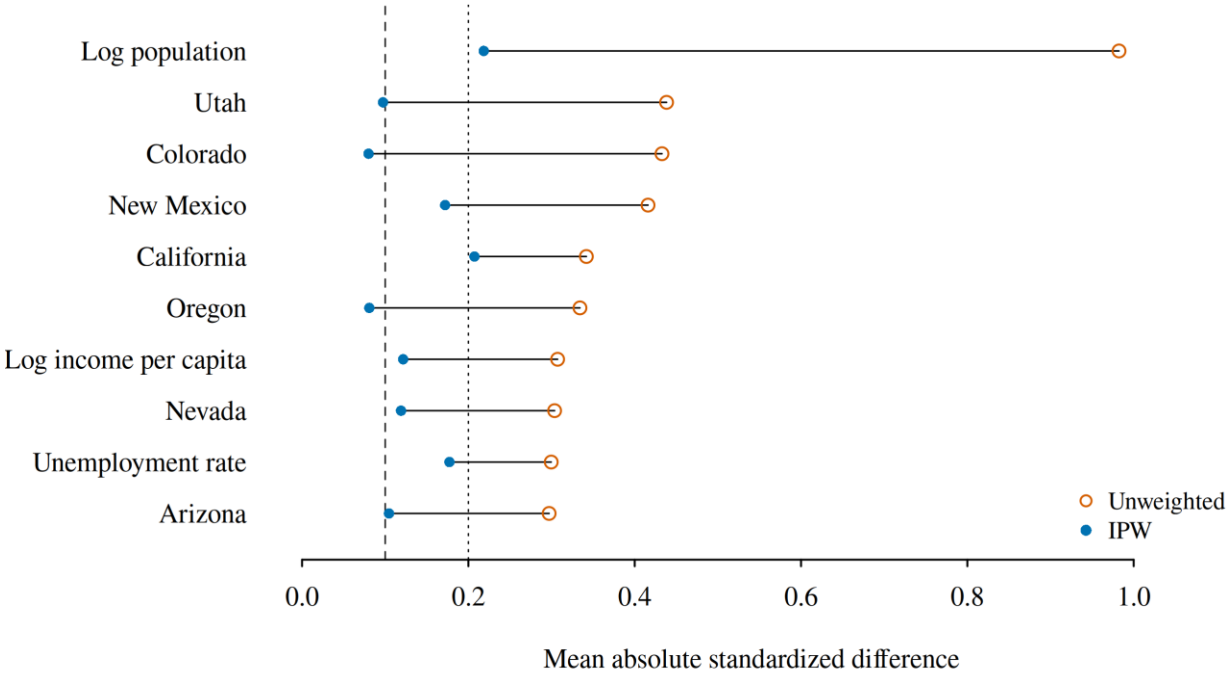


Figure 2. Baseline balance before and after cohort-specific IPW.

Notes: Points give the simple mean across 12 cohorts of absolute standardized mean differences; orange open circles are unweighted and blue filled circles are IPW. Dashed guides indicate 0.10 and 0.20, not formal acceptance thresholds. Covariates are measured at $g-1$. Sample: 112 adopters and 108 never-treated controls, 2010–2022. These are point diagnostics without confidence intervals; Appendix Table A2 reports means and maximum residual differences.

Weight concentration imposes a further qualification. Although the nominal never-treated control pool contains 108 counties, the ESS is only 1.85 for the 2011 cohort and 2.38 for the 2013 cohort. Their largest normalized control weights are 12.41 and 5.16, respectively, and the corresponding proportions of risk-set scores at the imposed bounds are 68% and 75%. By contrast, the 2019 cohort has an ESS of 66.92. Appendix Table A1 reports the cohort-by-cohort diagnostics, including those based on the separate private-employment control pool. Concentrated weights are particularly relevant to the headline estimates because earlier cohorts contribute more observed post-entry cohort-event cells to the event-years 0–5 average. Some early-cohort comparisons therefore depend heavily on a small number of controls.

7.2 Economic and business outcomes

Table 2 reports the main IPW results averaged over supported event years 0–5. For private establishments, the primary business outcome, the estimated difference relative to the weighted never-treated trajectory is 1.59 log points (SE 0.75; 95% CI 0.11 to 3.07; p = 0.035). This positive contrast reaches the 5% significance level and corresponds to approximately a 1.6% difference on the multiplicative scale. It concerns the county stock of operating private establishments; the outcome cannot distinguish establishment openings from closures, reclassification, or other changes in that stock, nor does it identify solar suppliers or newly formed independent firms.

The corresponding private-employment estimate is also positive, at 1.39 log points, but is less precise (SE 0.92; 95% CI −0.41 to 3.19; p = 0.131). Its confidence interval includes both a slight decline and a positive difference exceeding three log points. The result neither establishes a county-wide private-employment gain nor demonstrates that employment did not respond. The difference in statistical significance between the establishment and employment estimates does not itself establish that the two coefficients differ statistically.

The total-coverage outcomes show a similar distinction. Total covered establishments have an estimated contrast of 1.64 log points (SE 0.69; 95% CI 0.29 to 3.00; p = 0.018), comparable in direction and magnitude to private establishments. Total covered employment, by contrast, is estimated at 0.85 log points (SE 0.69; 95% CI −0.49 to 2.20; p = 0.214). Both establishment outcomes reach the 5% level, whereas neither employment estimate reaches the 10% level.

Table 2. Main IPW business and employment estimates, event years 0–5

Outcome	Estimate	SE	95% CI	p-value	T/C/G
Private establishments	1.59**	0.75	[0.11, 3.07]	0.035	112/108/12
Private employment	1.39	0.92	[-0.41, 3.19]	0.131	112/105/12
Total covered establishments	1.64**	0.69	[0.29, 3.00]	0.018	112/108/12
Total covered employment	0.85	0.69	[-0.49, 2.20]	0.214	112/108/12

Notes: All outcome estimates and SE are in log points (100 × log change). Main audited 1 MW entry; 2011–2022 cohorts and 2010–2022 observations. Private employment excludes the full county-year paths of three never-treated controls with zero-valued private-employment observations. 499 county-bootstrap draws. Estimates average supported cohort-event cells for event years 0–5, weighted by treated cohort size; baseline is event year −1. T/C/G denotes treated counties/eligible controls/cohorts. Confidence intervals (CI) use normal critical values and bootstrap standard errors (SE); p-values are two-sided and unadjusted. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$, using unrounded p-values. One-star estimates can have 95% intervals including zero.

Figure 3 places these average contrasts in event time. Private-establishment estimates are relatively small around the audited entry year and become larger at subsequent lags. At event year +5, the estimated contrast is 4.14 log points, but it is based on 71 treated counties in seven cohorts, compared with 112 counties in 12 cohorts at event year 0. Support falls further to 42 counties in four cohorts at +8 (Appendix Table A4). The plotted increase across lags therefore need not represent a single unchanged set of counties experiencing progressively larger effects. Changing cohort composition, together with wider uncertainty at later event years, limits the interpretation of the post-entry profile.

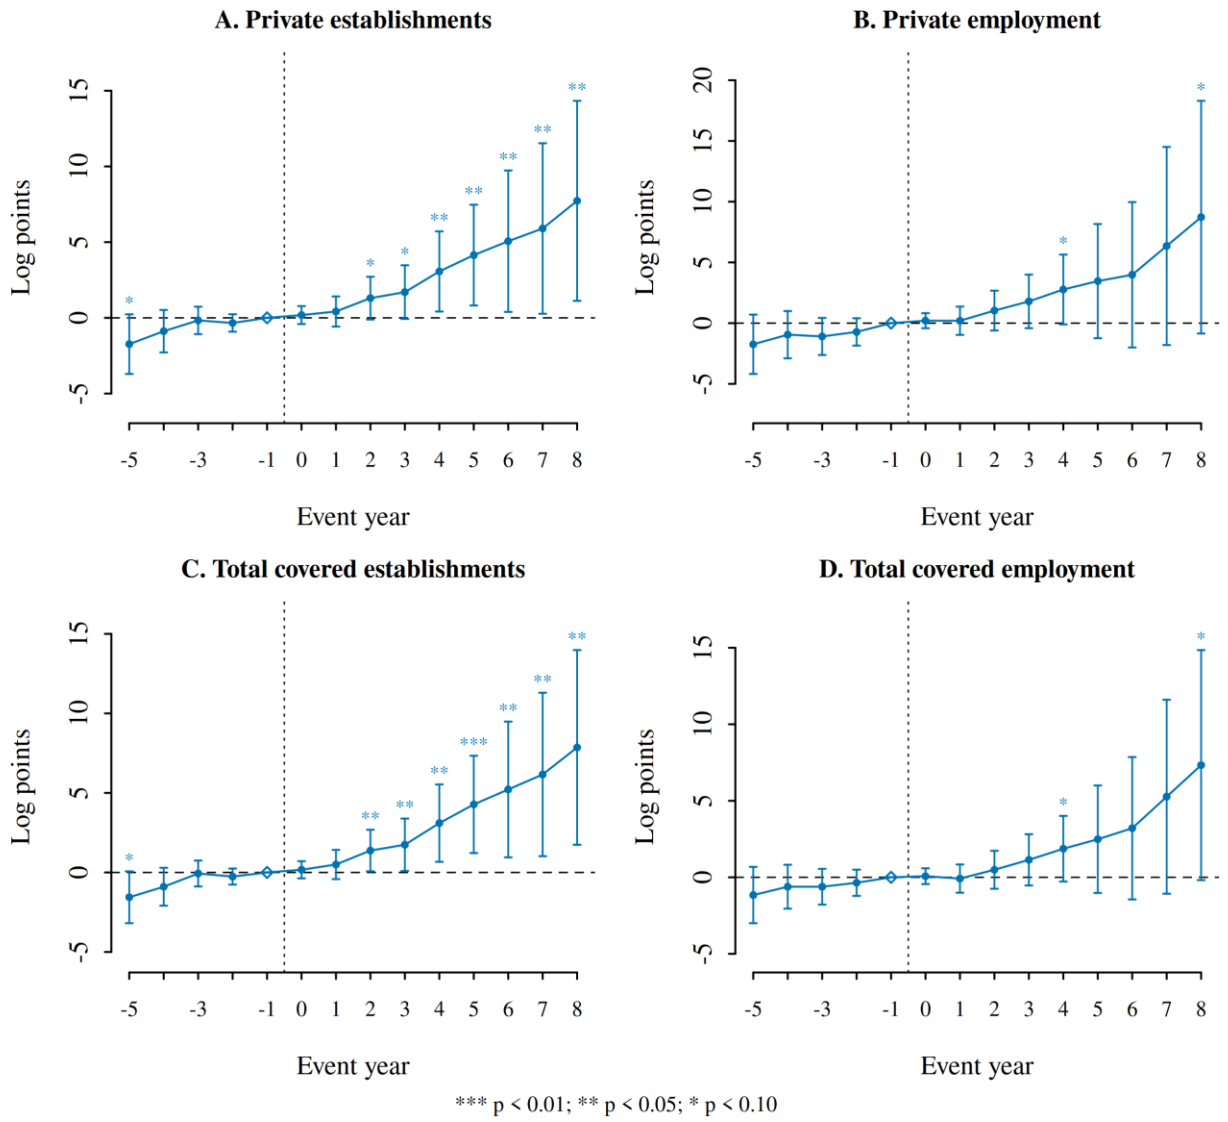


Figure 3. Main IPW event-study estimates for business and employment outcomes.

Notes: All panels are in log points and use audited county 1 MW entry. Points are cohort-size-weighted ATT estimates; bars are pointwise 95% normal-reference intervals from 499 county-bootstrap draws. The open diamond marks event year −1, normalized to zero; the vertical guide separates pre-entry and entry years. Sample: 112 adopters in 12 cohorts, 108 controls except 105 for private employment, calendar years 2010–2022. Support changes across event years (Appendix Table A4); intervals are not simultaneous bands. *** p < 0.01; ** p < 0.05; * p < 0.10, using unrounded p-values. One-star estimates can have 95% intervals including zero.

Joint IPW lead tests do not reject zero pre-treatment coefficients at the 10% level for private establishments (p = 0.190), private employment (p = 0.381), or total covered employment (p = 0.619). The total covered-establishment leads reject at 10% but not 5% (p = 0.053), qualifying that positive post-entry result. Appendix Table A3 reports the tests using four-degree-of-freedom

reference distributions. Nonrejection does not establish parallel trends or exclude anticipation, particularly given incomplete balance and low early-cohort ESS. Although the alternative estimators in Section 8.1 retain positive establishment contrasts, matched business-outcome leads raise additional comparability concerns.

7.3 Demographic outcomes

Table 3 reports a population contrast of 0.79 log points (SE 0.45; 95% CI −0.10 to 1.68; p = 0.081) in the main IPW analysis. This positive estimate reaches the 10% significance level but not the 5% level, and its 95% confidence interval includes zero. It therefore provides tentative evidence of a modest population difference relative to the weighted never-treated trajectory, rather than a precisely established residential response to solar entry.

The cumulative net-migration estimate is 9.38 net migrants per 1,000 baseline residents (SE 5.35; 95% CI −1.12 to 19.87; p = 0.080). Averaged across the supported cohort-event contrasts for event years 0–5, this point estimate represents approximately nine more cumulative net migrants per 1,000 baseline residents relative to the weighted comparison trajectory. It is neither an annual migration rate nor the cumulative difference measured only at event year +5, and it does not identify whether migration was related to solar employment. Like population, migration reaches the 10% but not the 5% level; its 95% interval includes zero and extends to nearly 20 net migrants per 1,000 residents. The two estimates point in the same direction but are related demographic measures, not independent evidence of a causal effect.

Table 3. Main IPW demographic estimates, event years 0–5

Outcome	Estimate	SE	95% CI	p-value	T/C/G
Population	0.79*	0.45	[-0.10, 1.68]	0.081	112/108/12
Cumulative net migration	9.38*	5.35	[-1.12, 19.87]	0.080	112/108/12

Notes: Population is in log points; cumulative net migration is per 1,000 baseline residents, not an annual rate. Main audited 1 MW entry; 2011–2022 cohorts and 2010–2022 observations. 499 county-bootstrap draws. Estimates average supported cohort-event cells for event years 0–5, weighted by treated cohort size; baseline is event year −1.

T/C/G denotes treated counties/eligible controls/cohorts. Confidence intervals (CI) use normal critical values and bootstrap standard errors (SE); p-values are two-sided and unadjusted. *** p < 0.01; ** p < 0.05; * p < 0.10, using unrounded p-values. One-star estimates can have 95% intervals including zero.

Figure 4 reports the corresponding event-study profiles. Population differences are larger at some later event years, but those individual estimates do not supersede the main average over supported event years 0–5. As with the business outcomes, the composition of contributing cohorts changes at longer lags (Appendix Table A4), limiting a simple interpretation of the later coefficients as a continuous trajectory for the same counties. Joint IPW pre-treatment lead tests do not reject zero leads for population (p = 0.789) or cumulative net migration (p = 0.431; Appendix Table A3). Nonrejection, however, does not establish parallel trends or exclude anticipatory and other unobserved differences.

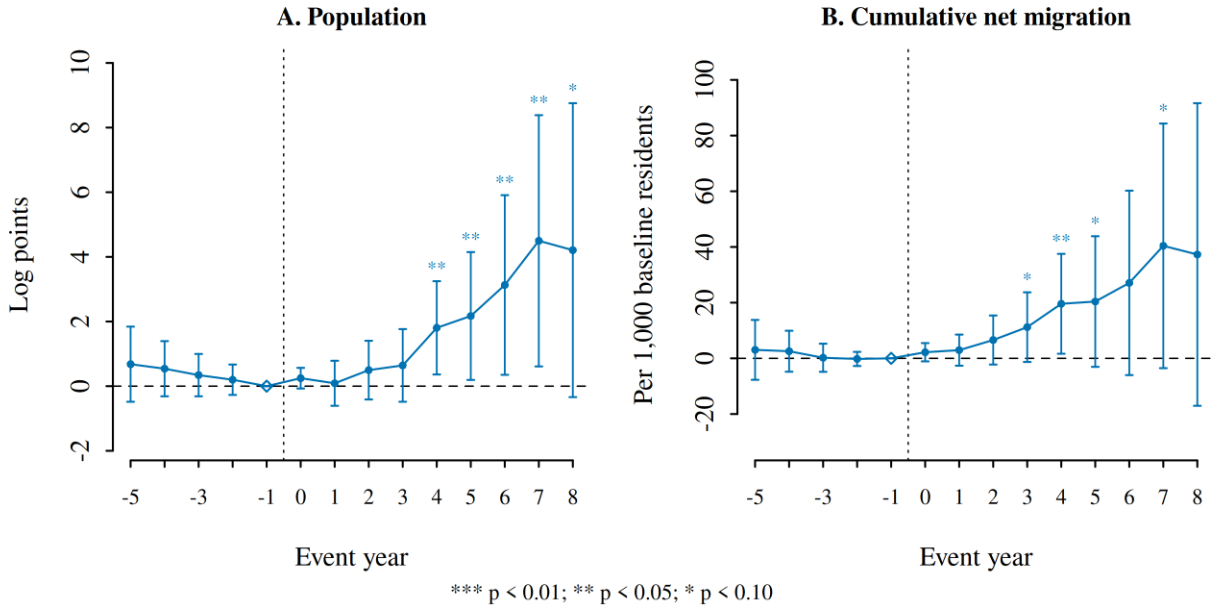


Figure 4. Main IPW event-study estimates for population and cumulative net migration.

Notes: Population is in log points; migration is net migrants per 1,000 baseline residents accumulated relative to $g-1$. Points are cohort-size-weighted estimates; bars are pointwise 95% normal-reference intervals from 499 county-bootstrap draws. The open diamond is the normalized −1 baseline. Sample: 112 adopters, 108 never-treated controls, 12 cohorts, 2010–2022. Support is reported in Appendix Table A4. Different vertical scales reflect different units; migration is not an annual rate. *** p < 0.01; ** p < 0.05; * p < 0.10, using unrounded p-values. One-star estimates can have 95% intervals including zero.

The sensitivity evidence in Section 8 places additional limits on the demographic interpretation. Doubly robust outcome adjustment attenuates the population estimate and leaves neither demographic outcome significant at the 10% level. The leave-one-state-out analysis shows that omitting New Mexico substantially reduces both the population and migration estimates. Matching yields a migration confidence interval narrowly above zero, but uses fewer treated cohorts and inference conditional on the selected matches. Together, these checks leave the positive demographic direction suggestive rather than uniformly robust across specifications.

7.4 Cross-domain interpretation

Across outcomes, the main IPW results provide clearer evidence of establishment differences than of employment or residential adjustment. Private and total covered establishment estimates reach the 5% level; neither employment measure reaches 10%. Population and cumulative net migration reach 10% but not 5%, with 95% confidence intervals including zero. This ordering reflects differences in statistical evidence, not formal tests of coefficient differences.

These outcome-specific patterns do not establish a causal sequence from changes in establishment counts to employment, population, or migration; mediation is not estimated, and contributing cohorts change across event years. Demographic evidence is also sensitive to some adjustments and state exclusions. These weighted contrasts remain subject to the identification, balance, and weight-concentration limitations discussed in Sections 6.8 and 7.1.

7.5 Descriptive heterogeneity by metropolitan status

The fixed-window sample contains 80 audited adopters from the 2011–2018 entry cohorts (39 metropolitan and 41 nonmetropolitan counties; Section 6.7). Figure 5 maps each county's mean log-point changes in private establishments, private employment, and population from entry through $g + 4$ relative to $g - 1$, alongside cumulative net migration from g through $g + 4$, scaled

per 1,000 residents at $g-1$. Positive and negative changes occur in both metropolitan-status groups, and their magnitudes vary substantially within each group. The map therefore shows county-level heterogeneity rather than a uniform metropolitan–nonmetropolitan divide.

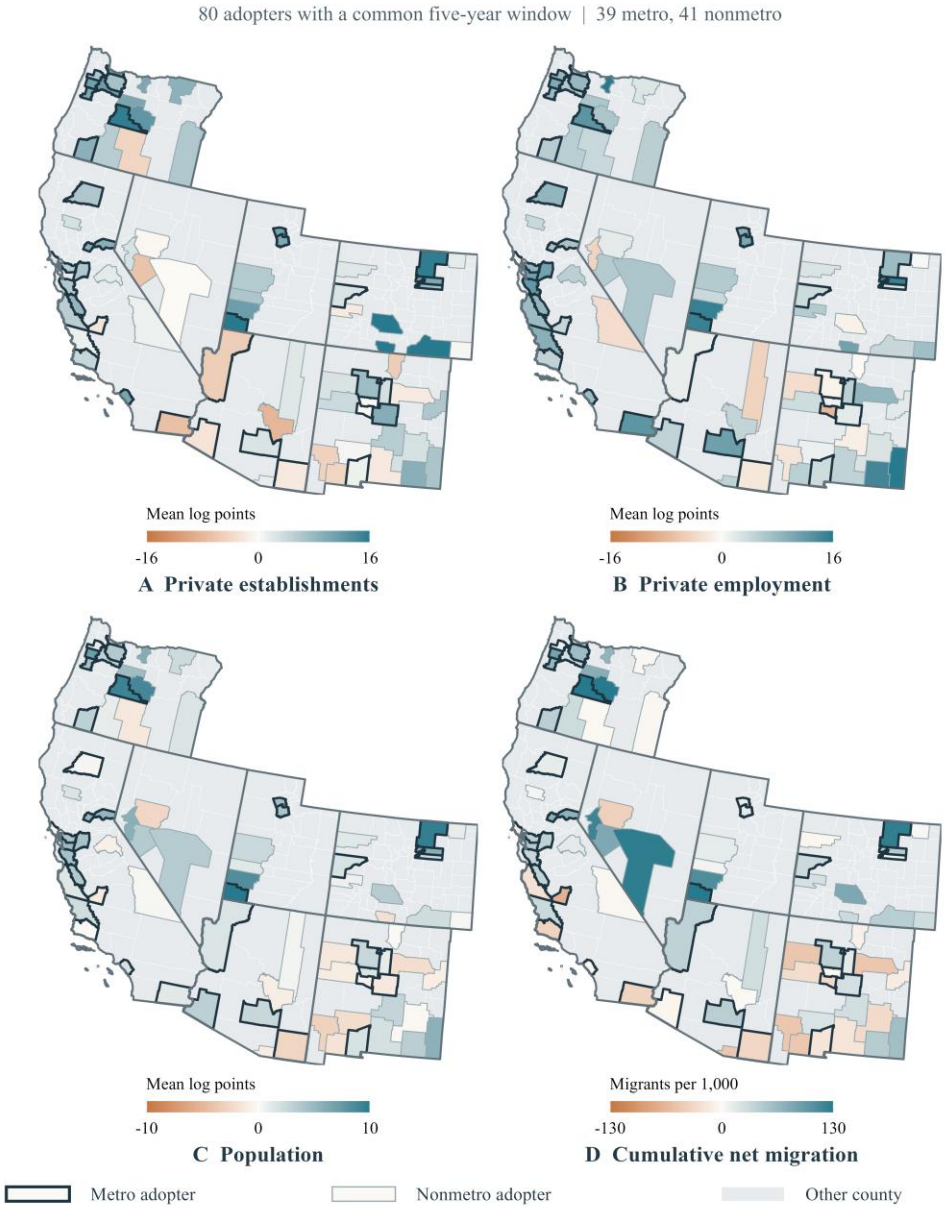


Figure 5. County outcomes after utility-scale solar entry.

Notes: Colored counties are the 80 adopters in the common fixed-window sample defined in Section 6.7. Dark, heavier outlines denote metropolitan adopters; lighter, thinner outlines denote nonmetropolitan adopters; gray counties are outside the sample. Panels report mean log-point changes from g through $g+4$ relative to $g-1$ for private establishments, private employment, and population, and cumulative net migrants per 1,000 baseline residents for migration. Scales are symmetric around zero; endpoint colors cap the display, not the underlying values. State boundaries are emphasized. The map uses the 2022 U.S. Census cartographic county file (U.S. Census Bureau, 2022).

Figure 6 compares the equally weighted county distributions. Metropolitan adopters have higher sample means for all four outcomes: 5.79 versus 3.86 log points for private establishments; 6.56 versus 4.10 for private employment; 2.91 versus 1.28 for population; and 26.49 versus 18.98 net migrants per 1,000 baseline residents for cumulative migration. Nevertheless, the interquartile ranges overlap in every panel, indicating considerable within-group dispersion. The two-sided Welch tests yield unadjusted p-values of 0.207, 0.039, 0.013, and 0.496, respectively. Only the private-employment and population comparisons have unadjusted p-values below 0.05; no multiple-comparison adjustment is applied across the four outcomes.

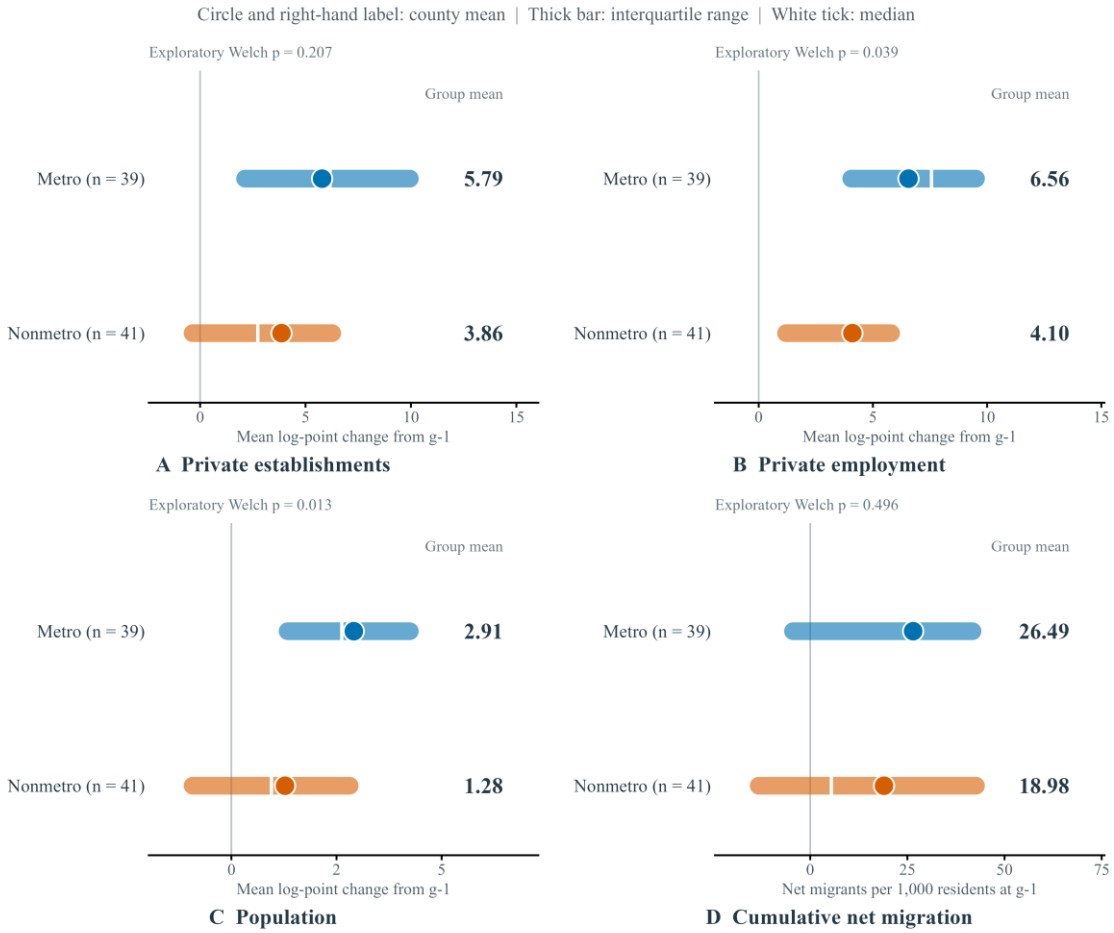


Figure 6. Metropolitan and nonmetropolitan distributions of post-entry outcomes.

Notes: Circles and right-hand labels report equally weighted county means; thick horizontal bars span the 25th–75th percentiles; white ticks mark medians; and the vertical line marks zero. P-values are from two-sided Welch tests allowing unequal group variances and are unadjusted across the four outcomes. These are exploratory adopter-only comparisons, not causal estimates of metropolitan-specific effects of solar entry.

These descriptive comparisons supplement the main weighted estimates. Metropolitan adopters show higher average post-entry changes in this restricted sample, alongside substantial within-group variation. Because the analysis includes only adopters, without a never-treated counterfactual or covariate adjustment, these contrasts do not identify causal effects of solar entry or differences in treatment effects by metropolitan status.

Chapter 8. Robustness and sensitivity

8.1 Alternative counterfactual estimators

Doubly robust (DR) adjustment and risk-set matching examine whether the main inverse-probability-weighted (IPW) pattern is sensitive to how the never-treated comparison trajectory is constructed. They serve as complementary checks rather than replacements for the main specification. Figure 7 compares all six outcomes across the three estimators using outcome-specific units and pointwise 95% normal-reference intervals, and Appendix Table A5 reports the estimates, standard errors, intervals, p-values, and sample information. Significance markers follow the same two-sided, unadjusted convention used elsewhere and do not rank the estimators by causal credibility. Because matching changes both the treated-cohort sample and the inference conditions, differences across estimators cannot be attributed solely to the estimation method.

8.1.1 Doubly robust results

The DR estimator described in Section 6.5 augments propensity weighting with a regression for untreated outcome changes among eligible never-treated counties. It uses 112 adopters in 12 cohorts, 499 bootstrap draws, and outcome-specific control pools. Because the propensity-score bounds and odds cap are retained, the estimator does not guarantee the usual double-robust result under weak overlap; instead, it probes sensitivity to an additional model of the untreated trajectory.

For private establishments, the DR estimate is 1.69 log points (SE 0.84; 95% CI 0.04 to 3.34; p = 0.044), while private employment is 1.15 log points (SE 1.24; 95% CI −1.28 to 3.57; p = 0.354). Total covered establishments show a contrast of 1.49 log points (SE 0.73; 95% CI 0.06 to 2.91; p = 0.041), compared with 0.21 log points for total covered employment (SE 0.80; 95% CI −1.36 to 1.79; p = 0.791). Thus, both establishment estimates retain 5%-level evidence of

positive differences, whereas neither employment estimate reaches the 10% level. As in the main results, this difference in statistical significance is not a test that the establishment and employment coefficients differ.

The demographic contrasts are less stable. The DR population estimate is 0.31 log points (SE 0.46; 95% CI −0.59 to 1.20; p = 0.503), and cumulative net migration is 6.24 net migrants per 1,000 baseline residents (SE 4.29; 95% CI −2.18 to 14.66; p = 0.146). Neither reaches the 10% level, unlike the corresponding main IPW estimates. All six DR joint pre-treatment lead tests also fail to reject at the 10% level (Appendix Table A3), although nonrejection does not establish parallel trends. Overall, the establishment pattern persists under this additional adjustment, while the demographic evidence becomes weaker and the underlying overlap concern remains.

8.1.2 Risk-set matched results

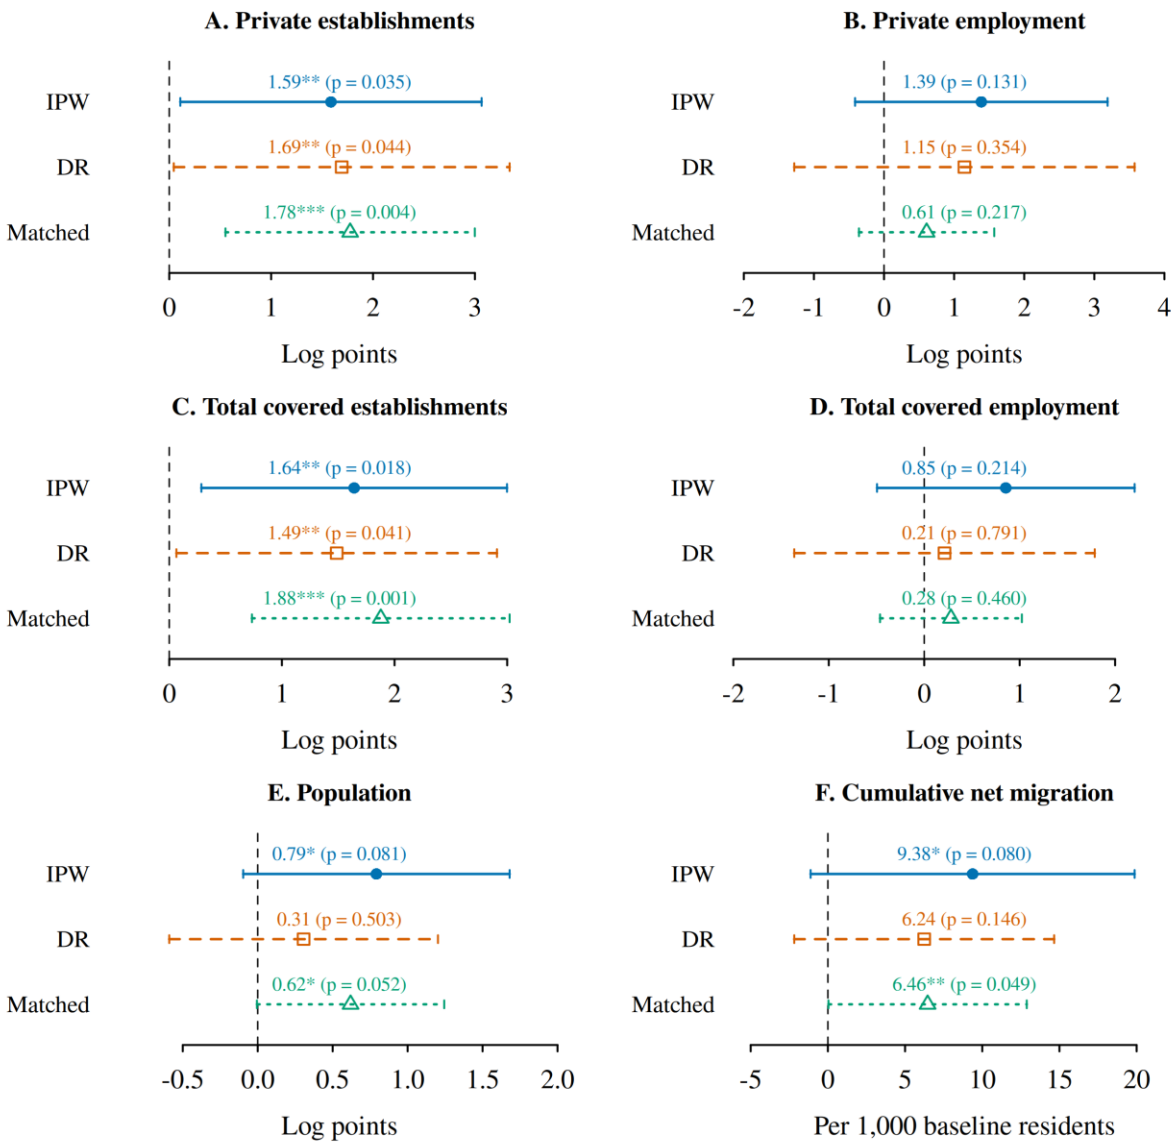
Risk-set matching, detailed in Section 6.6, pairs each adopter with three never-treated counties in the same state using six pre-entry characteristics, including economic levels and prior population and establishment trends. Requiring three pre-entry years yields 88 adopters in ten cohorts (2013–2022), 264 links, and 81 distinct controls. Appendix Table A6 reports cohort sizes, control reuse, effective sample sizes, and matching distances. Inference uses 399 bootstrap draws that resample treated counties while retaining fixed matches, so it is conditional on the realized design and does not capture neighbor-selection uncertainty or fully account for dependence from reused controls.

The matched private-establishment estimate is 1.78 log points (SE 0.62; 95% CI 0.55 to 3.00; p = 0.004), whereas private employment is 0.61 log points (SE 0.49; 95% CI −0.36 to 1.57; p = 0.217). Total covered establishments are estimated at 1.88 log points (SE 0.58; 95% CI 0.73 to 3.02; p = 0.001), and total covered employment at 0.28 log points (SE 0.38; 95% CI −0.46 to

1.02; p = 0.460). Both establishment estimates reach the 1% level, while neither employment estimate reaches the 10% level. These positive establishment contrasts are not, however, decisive corroboration of the causal interpretation.

Matched balance remains incomplete: the mean absolute standardized differences after matching are 0.347 for log population and 0.294 for log private establishments (Appendix Table A2). More importantly, joint pre-treatment lead tests reject zero coefficients for each matched business outcome: private establishments at 10% (p = 0.059), private employment at 1% (p = 0.0095), total covered establishments at 5% (p = 0.043), and total covered employment at 5% (p = 0.014; Appendix Table A3). These unfavorable diagnostics, alongside conditional fixed-match inference, prevent the smaller matched p-values from being interpreted as stronger identification.

For population, the matched estimate is 0.62 log points (SE 0.32; 95% CI −0.01 to 1.24; p = 0.052), reaching 10% but not 5%. Cumulative net migration is 6.46 per 1,000 baseline residents (SE 3.28; 95% CI 0.04 to 12.88; p = 0.049), with the interval only narrowly above zero. Demographic lead tests do not reject at 10% (Appendix Table A3), but this does not remove concerns about the restricted treated-cohort sample, conditional inference, or unobserved pre-entry differences. The restricted-cohort IPW checks in Tables 4–5 assess sensitivity to cohort coverage without making the counterfactual designs equivalent. Overall, establishment estimates remain consistently positive across alternative estimators, whereas demographic contrasts are not uniformly robust and the identification concerns in Section 6.8 remain.



*** p < 0.01; ** p < 0.05; * p < 0.10 (two-sided, unadjusted)

Figure 7. Complementary estimator comparisons for the event-years 0–5 average.

Notes: Blue filled circles denote the main IPW estimator; orange open squares, doubly robust DiD (DR); and green open triangles, risk-set matched DiD. Labels above the intervals show estimates, significance stars, and p-values. Bars are pointwise 95% normal-reference intervals. *** p < 0.01; ** p < 0.05; * p < 0.10, using unrounded, two-sided, unadjusted p-values; no marker indicates p ≥ 0.10. One-star estimates can have 95% intervals containing zero. Panels A–E use log points; panel F uses cumulative net migrants per 1,000 baseline residents. Baseline: event year −1; averages use supported event years 0–5. IPW/DR: 112 adopters, 12 cohorts, 108 controls (105 for private employment), 499 draws. Matching: 88 adopters, ten cohorts, 81 distinct controls, 399 draws conditional on 264 links. The matched design changes the treated population. Appendix Table A5 reports numerical estimates; Appendix Tables A2–A3 report balance and joint lead tests.

8.2 Transformation, covariate, timing, and sample checks

Tables 4–5 summarize the retained sensitivity analyses for private establishments, population, and cumulative net migration, including the higher-capacity entry checks discussed in Section 8.3. The main rows use 499 county-bootstrap draws; each alternative specification uses 249 draws, retaining the cohort-stratified common-control bootstrap procedure and refitting the propensity model. These checks assess whether the main estimates depend on the outcome transformation, balancing specification, treatment timing, capacity classification, data completeness, or available cohort and calendar-year coverage. Because some specifications change treatment definitions or the counties contributing to the estimates, differences across them are interpreted as sensitivity to these analytical choices rather than changes in a common treatment effect.

The exact-percentage transformation assesses dependence on the log-point scale. The expanded-covariate specification adds baseline log total covered establishments and an interpolated Social Vulnerability Index (SVI) measure; the latter is not strictly predetermined for every cohort and must be interpreted cautiously. The pre-COVID restriction limits observation years to 2019 and entry cohorts to 2011–2019, leaving 92 treated counties in nine cohorts. It changes both follow-up support and exposure to pandemic-period and Census-series changes. Panel-only timing instead defines entry by the first observed 1 MW crossing, without the historical audit corrections, yielding 119 estimable adopters. It serves as a sensitivity check on the timing audit rather than as a preferred alternative treatment definition.

The persistence check requires no observed post-entry capacity value below 1 MW; missing capacity by itself does not violate this rule, and 111 treated counties remain. A separate complete-capacity restriction excludes counties with any missing capacity record, yielding 112

treated counties and 107 controls. Restricting entry cohorts to 2013–2022 removes the earliest estimable cohorts (88 treated counties); the 2013–2020 window additionally removes the two latest cohorts (74). The complete 0–5 horizon check retains 71 treated counties in the seven 2011–2017 entry cohorts, each observed through event year +5. Counties excluded by treatment-cohort restrictions are not reassigned to the never-treated comparison pool.

Private-establishment estimates remain positive across the transformation, covariate, timing, persistence, data-completeness, and cohort-support checks (Table 4). The exact-percentage specification yields 1.75 percentage points, with a 95% confidence interval excluding zero. Expanded covariates, panel-only timing, persistence, complete capacity records, the two later-cohort windows, and complete 0–5 follow-up also yield intervals excluding zero. The pre-COVID estimate is 1.48 log points ($p = 0.084$); its 95% interval includes zero, so this check supports the positive direction only at the 10% level. The larger point estimates in later-cohort samples do not resolve the severe early-cohort weight concentration in the main analysis; those checks describe different treated populations.

Table 4. Private-establishment IPW sensitivity estimates

Specification	Estimate	SE	95% CI	p-value	T/C/G
Main audited 1 MW	1.59**	0.75	[0.11, 3.07]	0.035	112/108/12
Exact percentage change	1.75**	0.82	[0.15, 3.34]	0.032	112/108/12
Expanded covariates	1.64**	0.77	[0.14, 3.15]	0.033	112/108/12
Pre-COVID through 2019	1.48*	0.86	[-0.20, 3.17]	0.084	92/108/9
Panel-only timing	1.60**	0.72	[0.18, 3.02]	0.027	119/108/12
5 MW entry	0.78	0.60	[-0.40, 1.95]	0.196	104/138/12
10 MW entry	0.57	0.60	[-0.60, 1.74]	0.341	89/158/12
25 MW entry	0.84	0.73	[-0.59, 2.27]	0.251	70/178/12
50 MW entry	0.97	0.65	[-0.31, 2.25]	0.139	59/192/12
Persistent entry	1.52**	0.73	[0.08, 2.96]	0.038	111/108/12
Complete capacity records	1.64**	0.73	[0.21, 3.07]	0.025	112/107/12
Cohorts 2013–2022	1.89**	0.82	[0.27, 3.50]	0.022	88/108/10
Cohorts 2013–2020	2.03**	0.90	[0.27, 3.80]	0.024	74/108/8
Complete 0–5 horizon	1.98**	0.85	[0.30, 3.65]	0.021	71/108/7

Notes: Unit: log points, except the exact-percentage row, which is percentage points. Main row: 499 draws; other rows: 249 draws. Pre-COVID ends in 2019; complete horizon retains 2011–2017 cohorts. Higher thresholds change treatment and control definitions; their controls need not be solar-free. The expanded model includes interpolated SVI and is not strictly predetermined in every cohort. Estimates average supported cohort-event cells for event years 0–5, weighted by treated cohort size; baseline is event year −1. T/C/G denotes treated counties/eligible controls/cohorts. Confidence intervals (CI) use normal critical values and bootstrap standard errors (SE); p-values are two-sided and unadjusted. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$, using unrounded p-values. One-star estimates can have 95% intervals including zero.

Demographic estimates also remain positive across the displayed specifications, but their inferential support is less consistent (Table 5). Exact-percentage population change does not reach the 10% level ($p = 0.105$). Population contrasts reach the 5% level with expanded covariates, pre-COVID observations, panel-only timing, and a complete 0–5 horizon, and the 10% level with complete capacity records; persistence and both 2013-onward cohort restrictions do not reach 10%. For cumulative net migration, expanded covariates reach the 5% level, whereas pre-COVID, panel-only timing, persistence, complete capacity records, and complete 0–5 follow-up reach only the 10% level. Neither later-cohort restriction reaches 10%. The pre-COVID migration result is especially close to the conventional cutoff: its unrounded p-value is 0.05003, and its 95% interval

has a slightly negative lower bound (−0.00154) despite the displayed −0.00. Accordingly, it is marked with one star rather than two. Overall, these checks provide less consistent statistical evidence for the demographic outcomes than for the private-establishment outcome.

Table 5. Population and cumulative-migration IPW sensitivity estimates

Specification	Population	p-value P	Migration	p-value M	T/C/G
Main audited 1 MW	0.79* (0.45) [-0.10, 1.68]	0.081	9.38* (5.35) [-1.12, 19.87]	0.080	112/108/12
Exact percentage change	0.82 (0.50) [-0.17, 1.81]	0.105	Not applicable	—	112/108/12
Expanded covariates	0.87** (0.40) [0.08, 1.65]	0.030	10.24** (4.90) [0.62, 19.85]	0.037	112/108/12
Pre-COVID through 2019	1.03** (0.46) [0.12, 1.93]	0.026	11.92* (6.08) [−0.00, 23.83]	0.050	92/108/9
Panel-only timing	0.83** (0.40) [0.04, 1.62]	0.040	9.64* (5.28) [-0.71, 20.00]	0.068	119/108/12
5 MW entry	1.10** (0.44) [0.24, 1.96]	0.012	8.98* (4.72) [-0.27, 18.22]	0.057	104/138/12
10 MW entry	0.92** (0.40) [0.14, 1.70]	0.021	8.32** (3.87) [0.74, 15.91]	0.031	89/158/12
25 MW entry	1.02** (0.50) [0.05, 2.00]	0.039	8.05* (4.43) [-0.63, 16.74]	0.069	70/178/12
50 MW entry	1.19*** (0.39) [0.42, 1.96]	0.003	4.33 (3.88) [-3.27, 11.93]	0.264	59/192/12
Persistent entry	0.74 (0.47) [-0.19, 1.68]	0.117	8.96* (5.27) [-1.38, 19.29]	0.089	111/108/12
Complete capacity records	0.84* (0.47) [-0.08, 1.76]	0.072	9.95* (5.13) [-0.10, 20.01]	0.052	112/107/12
Cohorts 2013–2022	0.48 (0.45) [-0.40, 1.37]	0.285	6.10 (4.51) [-2.74, 14.94]	0.176	88/108/10
Cohorts 2013–2020	0.57 (0.56) [-0.53, 1.66]	0.308	7.10 (5.17) [-3.04, 17.24]	0.170	74/108/8
Complete 0–5 horizon	1.07** (0.50) [0.08, 2.06]	0.034	12.42* (6.66) [-0.63, 25.48]	0.062	71/108/7

Notes: Cells show estimate (SE), then 95% CI. Population is in log points, except the exact-percentage row (percentage points); migration is per 1,000 baseline residents. P and M identify the separate population and migration p-values. Exact-percentage migration is not defined. The pre-COVID migration lower limit is −0.00154 before rounding (p = 0.05003), so its interval includes zero. Main row: 499 draws; others: 249. Sample restrictions are described in Section 8.2, and higher-capacity comparison groups in Section 8.3; interpolated SVI requires caution. Estimates average supported cohort-event cells for event years 0–5, weighted by treated cohort size; baseline is event year −1. T/C/G denotes treated counties/eligible controls/cohorts. Confidence intervals (CI) use normal critical values and bootstrap standard errors (SE); p-values are two-sided and unadjusted. *** p < 0.01; ** p < 0.05; * p < 0.10, using unrounded p-values. One-star estimates can have 95% intervals including zero.

8.3 Higher-capacity entry concepts

Four alternative entry definitions use the first crossing of 5, 10, 25, or 50 MW of qualifying county solar capacity. At each threshold, the first crossing observed in the 2010–2022 panel determines entry timing unless the audit's conservative historical capacity lower bound establishes that the threshold was crossed before 2010; those counties are excluded as pre-treated. The historical evidence need not establish an exact pre-panel entry year for counties so excluded. Conversely, when the lower bound does not establish an earlier crossing, treatment classification relies on the observed panel. This does not rule out an unobserved pre-2010 crossing.

Changing the capacity threshold changes both the timing and composition of the treated and comparison groups. A control county under the 25 MW definition, for example, may have solar capacity below 25 MW and need not be solar-free. The number of treated counties falls from 104 at 5 MW to 59 at 50 MW, while the eligible control pool increases from 138 to 192 (Tables 4–5). These estimates therefore concern different entry concepts and county populations, rather than interchangeable definitions of the same treatment. Their differences do not identify a dose–response relationship or the marginal effect of additional installed capacity.

Private-establishment estimates remain positive at all four higher thresholds, ranging from 0.57 to 0.97 log points, but none reaches the 10% significance level, and every 95% confidence interval includes zero (Table 4). These results provide weaker statistical evidence for positive establishment contrasts under higher-capacity entry definitions than under the main audited 1 MW definition. They do not, however, establish that establishment responses are absent at those thresholds.

Population estimates reach the 5% level at 5, 10, and 25 MW and the 1% level at 50 MW. Cumulative net migration reaches the 5% level at 10 MW and the 10% level at 5 and 25 MW, but

does not reach the 10% level at 50 MW (Table 5). The contrast with the establishment results does not establish that greater installed capacity generates demographic adjustment without business adjustment. Statistical precision, treatment timing, treated-county composition, and exposure within the comparison group differ across thresholds, precluding a causal interpretation of the contrast between these outcome patterns.

8.4 Geographic sensitivity

The leave-one-state-out analyses exclude both treated and never-treated counties from each of the seven states in turn, refit the propensity model, and reconstruct the supported cohort–event comparisons. Table 6 reports private establishments, population, and cumulative net migration in their respective units. Each estimate pertains to the counties remaining in the regional sample rather than to a treatment effect in the omitted state.

The private-establishment point estimate remains positive following every state omission, ranging from 1.19 log points without Colorado to 2.08 without California. All seven estimates reach the 10% significance level. The estimates excluding Arizona, California, New Mexico, or Utah also reach the 5% level, whereas those excluding Colorado, Nevada, or Oregon reach only the 10% level and have 95% confidence intervals that include zero. Thus, no single state exclusion reverses the positive establishment contrast, although the strength of the statistical evidence varies across the remaining regional samples.

The demographic estimates are more sensitive to geographic composition. Population ranges from 0.19 log points without New Mexico to 1.09 without California. It reaches the 5% level after excluding Arizona, California, or Colorado and the 10% level after excluding Utah, but does not reach 10% after excluding Nevada, New Mexico, or Oregon. Cumulative net migration falls from 9.38 per 1,000 baseline residents in the full sample to 0.65 without New Mexico, while

the estimate without California is 12.21. Migration reaches the 5% level after excluding California or Colorado and the 10% level after excluding Arizona, Oregon, or Utah; it does not reach 10% after excluding Nevada or New Mexico. The attenuation following the New Mexico exclusion is particularly important when interpreting the positive demographic averages in the main analysis.

Table 6. Leave-one-state-out IPW sensitivity

Omitted state	Private establishments	Population	Migration	T/C/G
None	1.59** (0.75) [0.11, 3.07]	0.79* (0.45) [-0.10, 1.68]	9.38* (5.35) [-1.12, 19.87]	112/108/12
Arizona	1.82** (0.86) [0.13, 3.50]	1.03** (0.44) [0.17, 1.90]	9.15* (5.17) [-0.99, 19.29]	104/105/12
California	2.08** (0.88) [0.35, 3.81]	1.09** (0.51) [0.09, 2.09]	12.21** (5.80) [0.84, 23.57]	88/91/12
Colorado	1.19* (0.71) [-0.19, 2.58]	1.03** (0.46) [0.13, 1.93]	10.91** (5.44) [0.26, 21.57]	92/72/12
Nevada	1.45* (0.76) [-0.05, 2.94]	0.60 (0.45) [-0.29, 1.49]	7.93 (5.33) [-2.51, 18.37]	106/99/12
New Mexico	1.38** (0.69) [0.02, 2.73]	0.19 (0.42) [-0.64, 1.02]	0.65 (3.84) [-6.88, 8.17]	90/98/12
Oregon	1.51* (0.84) [-0.14, 3.17]	0.71 (0.56) [-0.39, 1.81]	11.88* (6.28) [-0.42, 24.18]	93/91/12
Utah	1.75** (0.83) [0.13, 3.37]	0.93* (0.49) [-0.04, 1.90]	11.10* (6.14) [-0.94, 23.14]	99/92/12

Notes: Cells show estimate (SE), then 95% CI. Private establishments and population are in log points; migration is per 1,000 baseline residents. Each omission excludes both treated and control counties and refits weights. The "None" row reports the main 499-draw estimate; state-omission rows use 249 draws. Estimates average supported cohort-event cells for event years 0–5, weighted by treated cohort size; baseline is event year −1. T/C/G denotes treated counties/eligible controls/cohorts. Confidence intervals (CI) use normal critical values and bootstrap standard errors (SE); p-values are two-sided and unadjusted. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$, using unrounded p-values. One-star estimates can have 95% intervals including zero.

Excluding a state changes the composition of both the treated and comparison groups and can alter the overlap and concentration of the weights. These exercises therefore assess the dependence of the regional contrasts on particular state samples; they do not identify spillovers or state-specific mechanisms. In conjunction with the alternative estimators, the geographic checks provide more consistent evidence for a positive establishment contrast than for demographic outcomes, without resolving the identification concerns discussed in Section 6.8.

Chapter 9. Discussion

9.1 What the establishment result contributes

The clearest main finding is a positive private-establishment contrast following audited county entry into utility-scale solar generation. In the main inverse-probability-weighted analysis, private and total covered establishment estimates reach the 5% significance level, whereas neither employment outcome reaches the 10% level. Differences in statistical significance do not, however, establish that the outcome coefficients differ. The contribution is to examine business presence, workplace employment, population, and migration under the same audited entry definition, rather than to treat project employment as a sufficient measure of local adjustment. The establishment outcome measures operating business locations, not necessarily new firms or solar suppliers, and the regional average does not imply a gain in every adopting county.

This distinction complements employment-oriented studies of renewable energy. Engineering and employment-factor approaches estimate the labor requirements associated with a technology or its deployment (Cameron and van der Zwaan, 2015). The county-level analysis instead compares the observed stock of business establishments with the trajectory of reweighted never-treated counties. These measures answer different questions and need not move together. An increase in establishment counts could coexist with a limited change in aggregate payroll employment, but the data do not reveal whether this reflects smaller workplaces, changes among existing firms, or shifts in industry composition.

The event-study profile is consistent with business adjustment extending beyond the audited entry year, but it does not identify the underlying channel. Local-multiplier reasoning suggests one possible demand mechanism through expenditures on contracting and related services (Moretti, 2010). An alternative is that counties attracting solar investment were already positioned

for business growth, or that anticipated business expansion influenced solar siting decisions, in ways not fully captured by observed baseline characteristics. Neither explanation is established by the estimates. Furthermore, later event-time coefficients draw on fewer and different cohorts, while some early cohorts rely on highly concentrated control weights. The rising profile therefore cannot be read as an identified cumulative effect for an unchanged group of counties.

The headline estimate is a modest average difference in relative establishment trajectories over supported event years 0–5. The practical importance of this relative difference depends on the scale of each county's existing business sector, as similar proportional changes may correspond to substantially different absolute changes in business presence across counties. It is neither a permanent annual growth premium nor a count of establishments attributable to an individual project. The positive exact-percentage sensitivity estimate indicates that the direction is not an artifact of the logarithmic transformation. Nevertheless, the weaker evidence under the pre-COVID restriction and higher-capacity entry definitions limits how broadly the main 1 MW first-entry result can be generalized. These qualifications matter when interpreting the estimate as evidence about local economic development rather than a project-specific impact calculation.

The aggregate stock measure also cannot distinguish openings from closures, relocations, or reporting changes, or identify how incumbent firms respond. Research on productivity responses among existing plants (Greenstone et al., 2010) and on broader adjustments in energy-producing local economies (Allcott and Keniston, 2018) provides relevant conceptual comparisons, not mechanisms identified here. Establishment-level transitions and industry-specific evidence would be needed to determine which forms of business adjustment account for the observed county-wide contrast.

9.2 Why employment and resident responses require caution

The positive but statistically imprecise private-employment estimate does not contradict the establishment finding. County workplace employment combines many industries and may respond differently from the number of operating business locations. Moreover, jobs located in a host county need not be held by its residents when workers commute, and construction activity may be temporary. Fabra et al. (2024) similarly distinguish project-associated employment from benefits retained within host communities. The present county-level data cannot identify who fills the observed jobs or how much project spending remains local; separating workplace and residential outcomes avoids presuming that these margins coincide.

Evidence of resident adjustment is more tentative. The population estimate of 0.79 log points corresponds to approximately a 0.8% difference in relative geometric population change, while the cumulative net-migration estimate represents 9.38 additional net migrants per 1,000 baseline residents relative to the weighted comparison trajectory, averaged across supported event years 0–5. Although these estimates indicate potentially meaningful demographic differences, both reach only the 10% significance level and have 95% confidence intervals that include zero, preventing a precise or geographically uniform interpretation of residential adjustment. Population is attenuated under doubly robust adjustment, and both demographic estimates weaken substantially when New Mexico is omitted. Cary and Ortez (2026), who examine population responses to wind and solar development using a different design, provide context for the possibility of technology-specific and spatially heterogeneous outcomes rather than independent confirmation of these estimates.

Population and migration also capture different, though related, demographic processes. Population changes reflect natural increase as well as migration, while the cumulative net-

migration outcome is defined relative to the pre-entry population and compared across supported post-entry horizons. Its main coefficient is not an annual migration rate, a count of solar workers moving to the county, or solely an event-year +5 difference. The confidence intervals remain wide, encompassing values consistent with little or no demographic change as well as substantively positive differences. These limits are especially important because the two demographic outcomes are not independent corroborating measures.

Housing availability, land use, amenities, and moving costs could also shape residential responses. Hu et al. (2025) document variation in solar-related property responses across settings, while the migration literature emphasizes mobility as a separate margin of local adjustment (Molloy et al., 2011). Yet this analysis does not observe housing costs, landowner income, movers' characteristics, or other features needed to distinguish those channels. These factors are plausible explanations for differences between workplace and resident outcomes, not mechanisms tested in the present study.

Finally, event-time patterns do not establish a sequence from establishments to employment and then to population or migration. Later event years draw disproportionately on earlier adopters because fewer cohorts have long follow-up periods (Appendix Table A4). The complete-horizon sensitivity check follows a fixed subset of adopting cohorts through event year +5, whereas the main average includes later entrants with shorter observed follow-up. Neither specification identifies mediation or justifies reading the outcome profiles as a causal sequence. The distinction between business, workplace, and residential responses remains informative without implying a verified pathway among them.

9.3 How the supporting evidence changes the interpretation

The alternative estimators show how sensitive the main establishment pattern is to the construction of the counterfactual. Doubly robust adjustment retains positive establishment contrasts for the main treated-cohort population, while within-state risk-set matching produces positive contrasts for a smaller set of adopters (Appendix Table A5). The agreement in direction is informative, but it is not independent confirmation of a causal effect. The doubly robust estimator retains the propensity-score bounds and weight cap, and matching changes cohort coverage, reuses selected controls, and reports inference conditional on fixed matches (Appendix Table A6). These differences matter when comparing the apparent strength of the estimates.

The diagnostics qualify that agreement. Weighting improves average baseline balance, yet some cohorts have large residual differences and very low effective control sample sizes (Appendix Tables A1–A2). The main IPW lead tests do not reject the null of zero pre-treatment coefficients for private establishments or population, although the total covered-establishment test rejects at the 10% level. All four matched business-outcome lead tests reject at one of the reported significance levels (Appendix Table A3). These results weaken the case for treating the matched establishment estimates as stronger causal evidence merely because their p-values are smaller. Achieved balance is a diagnostic rather than proof of comparability (Austin, 2009), and nonrejection of pre-treatment tests cannot establish the parallel-trends assumption (Roth, 2022).

The demographic conclusions are less stable across estimators. Doubly robust adjustment attenuates the population contrast and leaves neither demographic estimate significant at the 10% level. Matching retains a 10%-level population contrast and a migration estimate whose 95% confidence interval lies only narrowly above zero. Those results concern a smaller treated-cohort sample and are subject to the matching diagnostics and conditional inference just described. They

do not resolve the uncertainty surrounding the positive demographic direction in the main IPW results.

The remaining checks place further limits on generalization. Higher-capacity thresholds change treatment timing and the composition of both groups; their weaker establishment evidence cannot be read as a dose-response estimate or as a direct falsification of the 1 MW entry contrast. Leave-one-state-out results preserve a positive establishment estimate across all omissions, whereas population and migration are particularly sensitive to removing New Mexico (Table 6). Taken together, these analyses support the relative consistency of the establishment direction within the audited first-entry design, while leaving the identifying assumptions and the broader demographic interpretation unsettled.

9.4 Descriptive heterogeneity by metropolitan status

The metropolitan comparison describes variation among adopters rather than differential effects of solar entry. In the fixed-window sample of 80 counties, metropolitan adopters have larger mean post-entry changes for private establishments, private employment, population, and cumulative net migration. Yet Figure 5 shows positive and negative county outcomes in both groups, and the interquartile ranges in Figure 6 overlap for all four measures. Only the employment and population Welch comparisons have unadjusted p-values below 0.05; the four comparisons are exploratory and have not been adjusted for multiple testing.

The differences cannot be attributed to metropolitan status or interpreted as metropolitan-specific treatment effects. The analysis includes no never-treated counterfactual and does not adjust for pre-entry differences between the two adopter groups. Metropolitan classification may coincide with market size, sectoral structure, commuting networks, infrastructure, housing conditions, and other determinants of county trajectories, none of which is isolated here.

Estimating causal heterogeneity would require credible group-specific counterfactuals, sufficient support within each group, and a formal comparison of the resulting treatment-effect estimates. The present evidence instead documents substantial variation in observed post-entry outcomes within as well as between metropolitan categories.

9.5 Implications for local development and future evaluation

Local evaluation should distinguish establishment stocks, workplace employment, and residential outcomes rather than use installed solar capacity or project-job estimates as sufficient measures of economic benefit. The positive establishment contrast concerns business presence after county entry; it does not identify which businesses or households benefit or quantify fiscal returns, resident earnings, distributional gains, or decarbonization benefits.

Observed local changes need not represent aggregate gains: activity may shift across jurisdictions, while commuting, supply chains, and spillovers complicate distinctions between host and comparison counties (Kline and Moretti, 2014). Research on wind-related fiscal responses illustrates another dimension of local development that establishment counts cannot measure (Brunner and Schwegman, 2022b). Neither fiscal nor interregional effects are estimated here.

These estimates are better suited to framing local evaluation questions than to forecasting benefits for a particular installation. Project scale, construction duration, local contracting, land ownership, commuting, and supplier locations may shape where expenditures and jobs accrue. The distinction between construction and operation in renewable-energy studies (Slattery et al., 2011) and research on host-community employment benefits (Fabra et al., 2024) underscore the need for project-specific information. These channels are not separately estimated by the county first-entry design. Sensitivity to state composition and the descriptive metropolitan differences also caution against applying a single regional average to every prospective host county.

Future research could strengthen the design with longer verified pre-entry histories, better measures of announcements, construction and operation, and information on neighboring-county exposure. Linked establishment records could distinguish openings, closures, and incumbent-firm adjustment. Industry-level employment, commuting, housing, and origin–destination migration data could test channels that remain unobserved here. A causal analysis of geographic heterogeneity would require credible treated and untreated comparisons within each group, not differences among adopters alone. These extensions reflect the limits of the present evidence.

Chapter 10. Limitations and scope of inference

The scope of inference remains bounded by the maintained conditional parallel-trends assumption and several features of the comparison design. Most importantly, weighting improves observed comparability but does not eliminate residual imbalance or limited effective control support. Some early cohorts depend heavily on a small number of never-treated counties, and ridge stabilization, propensity-score bounds, and odds caps cannot create suitable counterfactuals where overlap is weak. The normal-reference bootstrap reflects the implemented sampling procedure but cannot remove the practical fragility associated with highly concentrated weights. Neither the doubly robust nor the matched estimator resolves this limitation automatically. The reported contrasts should therefore be interpreted in light of the balance, overlap, and pre-treatment diagnostics rather than as consequences of the estimator alone.

Treatment timing also has a narrower interpretation than the full development process surrounding a solar facility. The audited treatment identifies the first year in which qualifying county capacity reaches the entry threshold; it does not date project announcement, permitting, construction onset, or other anticipatory activity. Such activity may affect outcomes before recorded generation begins. The historical audit substantially improves classification of pre-panel exposure, but the replication materials do not independently reconstruct historical records from the original EIA workbooks. In addition, the higher-capacity checks rely on conservative historical lower bounds when complete pre-panel threshold histories are unavailable. County-level entry further abstracts from project size, proximity to facilities, the number of projects within a county, and subsequent capacity expansion.

Spatial and temporal aggregation impose additional limits. Economic activity need not stop at county boundaries: workers can commute, contractors can serve multiple counties, households

can relocate across nearby jurisdictions, and infrastructure can connect treated and comparison areas. The leave-one-state-out analyses are therefore sensitivity checks rather than spillover estimators. Event-study support also declines at longer leads and lags, so later coefficients are based on different sets of cohorts rather than a fixed group followed over time. Annual data may smooth short-lived construction responses. These features limit the interpretation of dynamic patterns even when the plotted coefficients are informative about timing.

The outcome measures themselves capture broad county-level margins rather than project-specific mechanisms. QCEW employment represents covered workplace employment and omits some forms of self-employment and noncovered activity; it is not a measure of employment among county residents. Establishment stocks do not distinguish openings, closures, relocations, organizational changes, or firms directly connected to solar development. Population reflects both migration and natural increase, while net migration does not distinguish inflows from outflows. The demographic series also combines Census vintages around 2020, creating a comparability concern that motivates the pre-COVID sensitivity analysis. The interpolated SVI measure used only in the expanded-covariate sensitivity can also incorporate information from a release dated in or after treatment for some cohorts and should not be viewed as a uniformly predetermined control.

Statistical inference requires similar caution. Reported event-study intervals are pointwise rather than simultaneous, and p-values are not adjusted for the multiple outcomes, event years, and sensitivity analyses considered in the thesis. The stated outcome hierarchy helps organize interpretation but does not constitute a multiple-testing correction or formal preregistration. Likewise, an establishment estimate that reaches the 5% level and an employment estimate that does not reach the 10% level do not by themselves establish that the two coefficients differ statistically. Findings that reach only the 10% level warrant particular caution in a multi-outcome

setting. Matched inference is conditional on the realized matching links and does not fully account for rematching uncertainty or dependence arising from reused controls. The analysis also does not identify causal mediation among business, employment, and demographic outcomes.

Finally, the metropolitan-status analysis has a deliberately narrower purpose. It compares post-entry outcomes among 80 adopters observed over a common $g-1$ to $g+4$ window and reports unadjusted Welch comparisons across four outcomes. It does not construct never-treated counterfactuals for metropolitan and nonmetropolitan counties, balance the two adopter groups on pre-entry characteristics, or estimate group-specific treatment effects. Figures 5 and 6 therefore document descriptive geographic heterogeneity in observed post-entry paths rather than causal heterogeneity in the response to solar entry.

Taken together, these limitations do not negate the information contained in the estimated contrasts, but they define the scope of the conclusions that can be drawn from them. The evidence is most informative about relative county-level trajectories following audited solar entry under the stated comparison design. Stronger claims about project-specific effects, mechanisms, spatial spillovers, or heterogeneous causal responses would require additional data and research designs tailored to those questions.

Chapter 11. Conclusion

Audited county entry into utility-scale solar generation is followed by a modest positive difference in private establishments relative to a propensity-weighted never-treated trajectory in the seven-state western sample. Private and total covered establishment estimates are statistically significant at the 5% level. Population and cumulative net migration estimates are positive but reach only the 10% significance level. The fixed-window descriptive analysis shows higher mean post-entry changes among metropolitan adopters on all four principal outcomes, but overlapping distributions and the absence of a never-treated counterfactual preclude interpreting these differences as causal heterogeneity.

Alternative estimators retain positive establishment contrasts, although rejections in matched business-outcome pre-treatment lead tests raise additional comparability concerns. Higher-capacity entry definitions provide weaker statistical evidence for establishment differences, while demographic estimates are sensitive to geographic composition. Concentrated early-cohort weights and residual imbalance further limit causal interpretation. Overall, the evidence is most consistent with a difference in business presence following initial solar entry, but it does not establish a causal sequence linking establishments, employment, population, and migration. For local evaluation, these findings highlight the importance of distinguishing business presence, workplace employment, and resident adjustment, documenting treatment timing and diagnostic limitations, and accounting for geographic variation rather than relying on a single measure of jobs or installed capacity.

References

Abadie, A. (2005). Semiparametric difference-in-differences estimators. The Review of Economic Studies, 72(1), 1–19. https://doi.org/10.1111/0034-6527.00321

Abadie, A., & Imbens, G. W. (2008). On the failure of the bootstrap for matching estimators. Econometrica, 76(6), 1537–1557. https://doi.org/10.3982/ECTA6474

Agency for Toxic Substances and Disease Registry [ATSDR]. (2026, May 11). SVI data & documentation download. https://www.atsdr.cdc.gov/place-health/php/svi/svi-data-documentation-download.html

Allcott, H., & Keniston, D. (2018). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. The Review of Economic Studies, 85(2), 695–731. https://doi.org/10.1093/restud/rdx042

Austin, P. C. (2009). Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. Statistics in Medicine, 28(25), 3083–3107. https://doi.org/10.1002/sim.3697

Austin, P. C., & Stuart, E. A. (2015). Moving towards best practice when using inverse probability of treatment weighting (IPTW) using the propensity score to estimate causal treatment effects in observational studies. Statistics in Medicine, 34(28), 3661–3679. https://doi.org/10.1002/sim.6607

Barrutiabengoa Ortubai, J. M., & Ortiz Durán, R. (2026). Measuring the impact of installed renewable capacity on employment. Energy Policy, 211, 115106. https://doi.org/10.1016/j.enpol.2026.115106

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? The Quarterly Journal of Economics, 119(1), 249–275. https://doi.org/10.1162/003355304772839588

Bolinger, M., Seel, J., Mulvaney Kemp, J., Warner, C., Katta, A., & Robson, D. (2023). Utility-scale solar, 2023 edition: Empirical trends in deployment, technology, cost, performance, PPA pricing, and value in the United States. Lawrence Berkeley National Laboratory. https://emp.lbl.gov/utility-scale-solar

Brown, J. P., Pender, J., Wiser, R., Lantz, E., & Hoen, B. (2012). Ex post analysis of economic impacts from wind power development in U.S. counties. Energy Economics, 34(6), 1743–1754. https://doi.org/10.1016/j.eneco.2012.07.010

Brunner, E. J., & Schwegman, D. J. (2022a). Commercial wind energy installations and local economic development: Evidence from U.S. counties. Energy Policy, 165, 112993. https://doi.org/10.1016/j.enpol.2022.112993

Brunner, E. J., & Schwegman, D. J. (2022b). Windfall revenues from windfarms: How do county governments respond to increases in the local tax base induced by wind energy installations? Public Budgeting & Finance, 42(3), 93–113. https://doi.org/10.1111/pbaf.12318

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

Cameron, L., & van der Zwaan, B. (2015). Employment factors for wind and solar energy technologies: A literature review. Renewable and Sustainable Energy Reviews, 45, 160–172. https://doi.org/10.1016/j.rser.2015.01.001

Carlino, G. A., & Mills, E. S. (1987). The determinants of county growth. Journal of Regional Science, 27(1), 39–54. https://doi.org/10.1111/j.1467-9787.1987.tb01143.x

Cary, M., & Ortez, M. (2026). Gone with the wind: Renewable energy projects and population change. Energy Economics, 158, 109303. https://doi.org/10.1016/j.eneco.2026.109303

Curtis, E. M., & Marinescu, I. (2023). Green energy jobs in the United States: What are they, and where are they? Environmental and Energy Policy and the Economy, 4, 202–237. https://doi.org/10.1086/722677

Fabra, N., Gutiérrez, E., Lacuesta, A., & Ramos, R. (2024). Do renewable energy investments create local jobs? Journal of Public Economics, 239, 105212. https://doi.org/10.1016/j.jpubeco.2024.105212

Fang, C. (2022). The employment impact of wind and solar energy in the US [Master's thesis, Tufts University]. Tufts Digital Library.

García-Riazuelo, Á., Duarte, R., & Sarasa, C. (2025). The long-term socioeconomic impacts of renewable energy deployment: Lessons from case studies in European rural regions. Journal of Regional Science, 65(4), 1094–1117. https://doi.org/10.1111/jors.12775

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014

Greenstone, M., Hornbeck, R., & Moretti, E. (2010). Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. Journal of Political Economy, 118(3), 536–598. https://doi.org/10.1086/653714

Hao, S., & Michaud, G. (2024). Assessing property value impacts near utility-scale solar in the Midwestern United States. Solar Compass, 12, 100090. https://doi.org/10.1016/j.solcom.2024.100090

Hu, C., Chen, Z., Liu, P., Zhang, W., He, X., & Bosch, D. (2025). Impact of large-scale solar on property values in the United States: Diverse effects and causal mechanisms. Proceedings of the National Academy of Sciences, 122(24), e2418414122. https://doi.org/10.1073/pnas.2418414122

Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. Psychological Methods, 15(4), 309–334. https://doi.org/10.1037/a0020761

Jenniches, S. (2018). Assessing the regional economic impacts of renewable energy sources—A literature review. Renewable and Sustainable Energy Reviews, 93, 35–51. https://doi.org/10.1016/j.rser.2018.05.008

Kline, P., & Moretti, E. (2014). People, places, and public policy: Some simple welfare economics of local economic development programs. Annual Review of Economics, 6(1), 629–662. https://doi.org/10.1146/annurev-economics-080213-041024

Molloy, R., Smith, C. L., & Wozniak, A. (2011). Internal migration in the United States. Journal of Economic Perspectives, 25(3), 173–196. https://doi.org/10.1257/jep.25.3.173

Moretti, E. (2010). Local multipliers. American Economic Review, 100(2), 373–377. https://doi.org/10.1257/aer.100.2.373

Ong, S., Campbell, C., Denholm, P., Margolis, R., & Heath, G. (2013). Land-use requirements for solar power plants in the United States (NREL/TP-6A20-56290). National Renewable Energy Laboratory. https://www.osti.gov/servlets/purl/1086349

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. The Review of Economic Studies, 90(5), 2555–2591. https://doi.org/10.1093/restud/rdad018

Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. Biometrika, 70(1), 41–55. https://doi.org/10.1093/biomet/70.1.41

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. American Economic Review: Insights, 4(3), 305–322. https://doi.org/10.1257/aeri.20210236

Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. Journal of Econometrics, 235(2), 2218–2244. https://doi.org/10.1016/j.jeconom.2023.03.008

San-Martín González, E., & Soler-Vaya, F. (2024). Depopulation determinants of small rural municipalities in the Valencia Region (Spain). Journal of Rural Studies, 110, 103369. https://doi.org/10.1016/j.jrurstud.2024.103369

Sant'Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. Journal of Econometrics, 219(1), 101–122. https://doi.org/10.1016/j.jeconom.2020.06.003

Shoeib, E. A. H., Hamin Infield, E., & Renski, H. C. (2021). Measuring the impacts of wind energy projects on U.S. rural counties' community services and cost of living. Energy Policy, 153, 112279. https://doi.org/10.1016/j.enpol.2021.112279

Shoeib, E. A. H., Renski, H. C., & Hamin Infield, E. (2022). Who benefits from renewable electricity? The differential effect of wind power development on rural counties in the United States. Energy Research & Social Science, 85, 102398. https://doi.org/10.1016/j.erss.2021.102398

Slattery, M. C., Lantz, E., & Johnson, B. L. (2011). State and local economic impacts from wind energy projects: Texas case study. Energy Policy, 39(12), 7930–7940. https://doi.org/10.1016/j.enpol.2011.09.047

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

U.S. Bureau of Economic Analysis. (2026, April 22). Personal income by county. https://www.bea.gov/data/income-saving/personal-income-by-county

U.S. Bureau of Labor Statistics. (2025, March 17). Local Area Unemployment Statistics overview. https://www.bls.gov/lau/lauov.htm

U.S. Bureau of Labor Statistics. (2026a, August 31). Quarterly Census of Employment and Wages: Calculation. https://www.bls.gov/opub/hom/cew/calculation.htm

U.S. Bureau of Labor Statistics. (2026b, August 31). Quarterly Census of Employment and Wages: Concepts. Handbook of Methods. https://www.bls.gov/opub/hom/cew/concepts.htm

U.S. Census Bureau. (2020). County and metro area population estimates. Vintage 2019 release materials. https://www2.census.gov/programs-surveys/popest/datasets/2010-2019/counties/totals/

U.S. Census Bureau. (2022). 2022 cartographic boundary files: Counties [Data set]. https://www.census.gov/geographies/mapping-files/2022/geo/carto-boundary-file.html

U.S. Census Bureau. (2024, December). Methodology for the United States population estimates: Vintage 2024. Nation, states, counties, and Puerto Rico—April 1, 2020 to July 1, 2024. https://www2.census.gov/programs-surveys/popest/technical-documentation/methodology/2020-2024/methods-statement-v2024.pdf

U.S. Census Bureau. (2025). Population estimates and components of change: Counties, metropolitan and micropolitan statistical areas, and Puerto Rico municipios. Vintage 2024 release materials. https://www2.census.gov/programs-surveys/popest/datasets/2020-2024/counties/totals/

U.S. Department of Agriculture, Economic Research Service. (2013). 2013 Rural-Urban Continuum Codes [Data set]. https://www.ers.usda.gov/data-products/rural-urban-continuum-codes/

U.S. Department of Energy. (2021). Solar futures study (DOE/GO-102021-5621). https://www.energy.gov/sites/default/files/2021-09/Solar%20Futures%20Study.pdf

U.S. Energy Information Administration. (2026-a). Electricity generation, capacity, and sales in the United States. https://www.eia.gov/energyexplained/electricity/electricity-in-the-us-generation-capacity-and-sales.php

U.S. Energy Information Administration. (2026-b). Form EIA-860 detailed data with previous form data (EIA-860A/860B). https://www.eia.gov/electricity/data/eia860/index.php

VanderWeele, T. J. (2019). Principles of confounder selection. European Journal of Epidemiology, 34(3), 211–219. https://doi.org/10.1007/s10654-019-00494-6

Wasserstein, R. L., & Lazar, N. A. (2016). The ASA statement on p-values: Context, process, and purpose. The American Statistician, 70(2), 129–133. https://doi.org/10.1080/00031305.2016.1154108

Wei, M., Patadia, S., & Kammen, D. M. (2010). Putting renewables and energy efficiency to work: How many jobs can the clean energy industry generate in the US? Energy Policy, 38(2), 919–931. https://doi.org/10.1016/j.enpol.2009.10.044

Appendix A. Supporting diagnostics

A.1 Weight concentration and covariate balance

Table A1 reports the complete cohort weight diagnostics used to assess the main comparisons. The 108-control pool applies to all outcomes except private employment, which uses 105 controls. Normalized weights sum to the treated cohort size, so a maximum weight should be interpreted relative to that cohort's treated count. The bounded-score percentage refers to the combined treated and control risk set, not just to controls.

Table A1. Cohort-specific effective control sample size and weight concentration

Cohort	Treated	ESS 108	Max weight 108	Bounded %	ESS 105	Max weight 105
2011	17	1.85	12.41	68.0	1.84	12.44
2012	7	20.42	0.92	66.1	20.07	0.93
2013	8	2.38	5.16	75.0	2.36	5.18
2014	10	17.75	1.43	52.5	17.53	1.44
2015	9	13.79	1.82	55.6	13.67	1.83
2016	9	37.75	0.74	42.7	37.51	0.73
2017	11	39.65	0.82	30.3	39.54	0.82
2018	9	38.47	0.77	29.9	38.78	0.74
2019	12	66.92	0.59	2.5	67.27	0.58
2020	6	37.78	0.62	45.6	37.50	0.62
2021	9	33.33	0.97	32.5	33.78	0.95
2022	5	28.06	0.45	51.3	27.85	0.45

Notes: Main audited 1 MW IPW/DR propensity design, 2011–2022 cohorts. ESS is a count-equivalent weight-dispersion measure. Max weight is the maximum normalized control weight, with cohort weights summing to the treated count. Columns labeled 108 use the common control pool; 105 use the private-employment pool. Bounded % is the percentage of combined risk-set scores at 0.02 or 0.98 in the 108-control model. Covariates are measured at $g-1$; these are point diagnostics, not estimates with confidence intervals.

Table A2 reports mean absolute imbalance before and after adjustment and the largest post-adjustment absolute difference across cohorts. IPW diagnostics refer to the common 108-control design. Matching diagnostics cover the six matching features and the ten 2013–2022 cohorts; state is matched exactly and therefore is not a continuous matching feature in this panel.

Table A2. Baseline covariate balance under IPW and matching

Design	Covariate	Mean before	Mean after	Maximum after
IPW	Log income per capita	0.307	0.122	0.422
IPW	Log population	0.982	0.218	0.584
IPW	Unemployment rate	0.300	0.177	0.448
IPW	Arizona	0.297	0.105	0.476
IPW	California	0.342	0.207	0.537
IPW	Colorado	0.433	0.080	0.322
IPW	Nevada	0.304	0.119	0.301
IPW	New Mexico	0.416	0.172	0.738
IPW	Oregon	0.334	0.081	0.271
IPW	Utah	0.438	0.097	0.235
Matched	Log population	0.886	0.347	0.764
Matched	Log income per capita	0.300	0.195	0.414
Matched	Unemployment rate	0.301	0.207	0.434
Matched	Log private establishments	0.764	0.294	0.687
Matched	Population pre-trend	0.206	0.189	0.508
Matched	Establishment pre-trend	0.189	0.136	0.347

Notes: Entries are absolute standardized mean differences (unitless). Means give each cohort equal weight; Maximum after is the largest post-adjustment absolute difference. The denominator is the square root of the average unweighted treated and control variances and is held fixed before/after. IPW: 112 treated/108 controls/12 cohorts. Matching: 88 treated/81 distinct controls/10 cohorts, 264 links. Levels use $g-1$; slopes use $g-3$ through $g-1$. These descriptive diagnostics have no confidence intervals.

A.2 Joint pre-treatment tests and event-time support

Table A3 contains all joint lead tests, including unfavorable matched-business diagnostics. The same four lead years are tested for all estimators, although the counties contributing to individual leads differ. No test is used as a gatekeeper for including an outcome or estimator.

Table A3. Joint tests of pre-treatment event coefficients

Outcome	IPW W	IPW p	DR W	DR p	Matched W	Matched p
Private establishments	6.12	0.190	3.66	0.453	9.09	0.059*
Private employment	4.19	0.381	2.40	0.662	13.39	0.0095***
Total covered establishments	9.34	0.053*	5.58	0.232	9.86	0.043**
Total covered employment	2.64	0.619	1.33	0.856	12.44	0.014**
Population	1.71	0.789	2.12	0.713	0.63	0.960
Cumulative net migration	3.82	0.431	2.12	0.714	0.51	0.973

Notes: W is the joint Wald statistic for event years −5 through −2, using the bootstrap covariance and its generalized inverse. All reported covariance ranks and chi-squared reference degrees of freedom are four. Event year −1 is normalized and excluded. IPW and DR use 499 draws and 112 treated counties, with 108 controls (105 for private employment); matching uses 399 draws conditional on fixed matches, 88 treated counties, and 81 controls. P-values are unadjusted. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. Stars denote rejection of zero joint leads, not a favorable diagnostic. Nonrejection is not evidence that parallel trends is true.

Table A4 documents the support for Figures 3–4 and the underlying alternative event paths. Normalization makes event year −1 zero for every outcome. The matching cohort restriction reduces support at entry and later lags relative to IPW and doubly robust DiD. Private-employment exclusions change the control pool, not the treated support shown here.

Table A4. Treated county and cohort support by event year

Event year	IPW DR counties	IPW DR cohorts	Matched counties	Matched cohorts
-5	70	8	70	8
-4	80	9	80	9
-3	88	10	88	10
-2	95	11	88	10
-1	112	12	88	10
0	112	12	88	10
1	107	11	83	9
2	98	10	74	8
3	92	9	68	7
4	80	8	56	6
5	71	7	47	5
6	60	6	36	4
7	51	5	27	3
8	42	4	18	2

Notes: Counts apply to every outcome within an estimator. IPW and DR use cohorts 2011–2022; matching uses 2013–2022. Calendar support is 2010–2022. Event year −1 is the normalized baseline; later lags exclude recent cohorts. Control counts are 108 for common IPW/DR outcomes, 105 for private employment, and 81 distinct matched controls. This is a support table; confidence intervals are not applicable.

A.3 Complete estimator comparisons

Table A5 reports the primary and complementary estimators together for all six outcomes. Separate columns preserve the distinction between the full-cohort IPW/DR comparisons and the

smaller matched design. The table supplies the numerical uncertainty underlying Figure 7, including estimates whose confidence intervals include zero.

Table A5. Complete main and complementary estimator comparisons

Outcome	IPW	DR	Matched
Private establishments	1.59** (0.75) [0.11, 3.07] p = 0.035	1.69** (0.84) [0.04, 3.34] p = 0.044	1.78*** (0.62) [0.55, 3.00] p = 0.004
Private employment	1.39 (0.92) [-0.41, 3.19] p = 0.131	1.15 (1.24) [-1.28, 3.57] p = 0.354	0.61 (0.49) [-0.36, 1.57] p = 0.217
Total covered establishments	1.64** (0.69) [0.29, 3.00] p = 0.018	1.49** (0.73) [0.06, 2.91] p = 0.041	1.88*** (0.58) [0.73, 3.02] p = 0.001
Total covered employment	0.85 (0.69) [-0.49, 2.20] p = 0.214	0.21 (0.80) [-1.36, 1.79] p = 0.791	0.28 (0.38) [-0.46, 1.02] p = 0.460
Population	0.79* (0.45) [-0.10, 1.68] p = 0.081	0.31 (0.46) [-0.59, 1.20] p = 0.503	0.62* (0.32) [-0.01, 1.24] p = 0.052
Cumulative net migration	9.38* (5.35) [-1.12, 19.87] p = 0.080	6.24 (4.29) [-2.18, 14.66] p = 0.146	6.46** (3.28) [0.04, 12.88] p = 0.049

Notes: Cells show estimate (SE), 95% CI, and two-sided unadjusted p-value. Business, employment, and population use log points; cumulative migration uses net migrants per 1,000 baseline residents. All estimates average supported cohort-event cells for event years 0–5, weighted by treated cohort size, relative to event year −1. IPW and DR use 112 treated counties in 12 cohorts and 108 controls (105 for private employment), with 499 bootstrap draws. Matching uses 88 treated counties in ten cohorts and 81 distinct controls, with 399 draws conditional on fixed matches. Intervals use normal critical values. Matching changes the treated population; estimates are not a formal test of differences between estimators. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$, using unrounded p-values. One-star estimates can have 95% intervals including zero.

A.4 Matching quality and propensity-score overlap

Table A6 summarizes the realized matching design within each cohort. Control ESS accounts for repeated use of a control within that cohort; maximum reuse counts its appearance among the three-neighbor links. Distances are Euclidean distances in the six standardized matching features and have no natural substantive unit. They are not an acceptance test, and no caliper is retrospectively introduced based on this table.

Table A6. Cohort-level risk-set matching diagnostics

Cohort	Treated	Controls	Links	Control ESS	Max reuse	Median distance	Max distance
2013	8	14	24	11.08	3	1.64	4.18
2014	10	21	30	16.67	3	1.67	4.65
2015	9	19	27	16.20	3	1.85	3.44
2016	9	23	27	19.70	3	1.33	4.49
2017	11	28	33	25.33	2	1.39	3.53
2018	9	22	27	18.69	3	1.45	2.85
2019	12	31	36	27.00	3	1.82	3.43
2020	6	15	18	13.50	2	1.19	4.70
2021	9	22	27	19.70	2	1.85	2.70
2022	5	13	15	11.84	2	1.47	2.44

Notes: Matching uses three nearest never-treated controls within state, with replacement across treated counties. Controls count distinct selected counties within the cohort; Links counts selections including reuse. Control ESS equals squared total selection counts divided by the sum of squared control selection counts. Max reuse counts the largest number of selections of one control. Distances summarize all links in the six standardized features; they are unitless and no caliper is imposed. Levels use $g-1$ and trends use $g-3$ through $g-1$. Overall: 88 treated counties, 81 distinct controls, ten cohorts, 264 links. These are design diagnostics without confidence intervals.

Figure A1 shows the empirical cumulative distributions of raw fitted adoption probabilities for the main 108-control design. The curves use probabilities before the imposed bounds, so separation is visible rather than compressed into the endpoints. These are descriptive overlap diagnostics, not confidence bands or treatment-effect estimates. The private-employment design has a separate 105-control fit, summarized in Table A1.

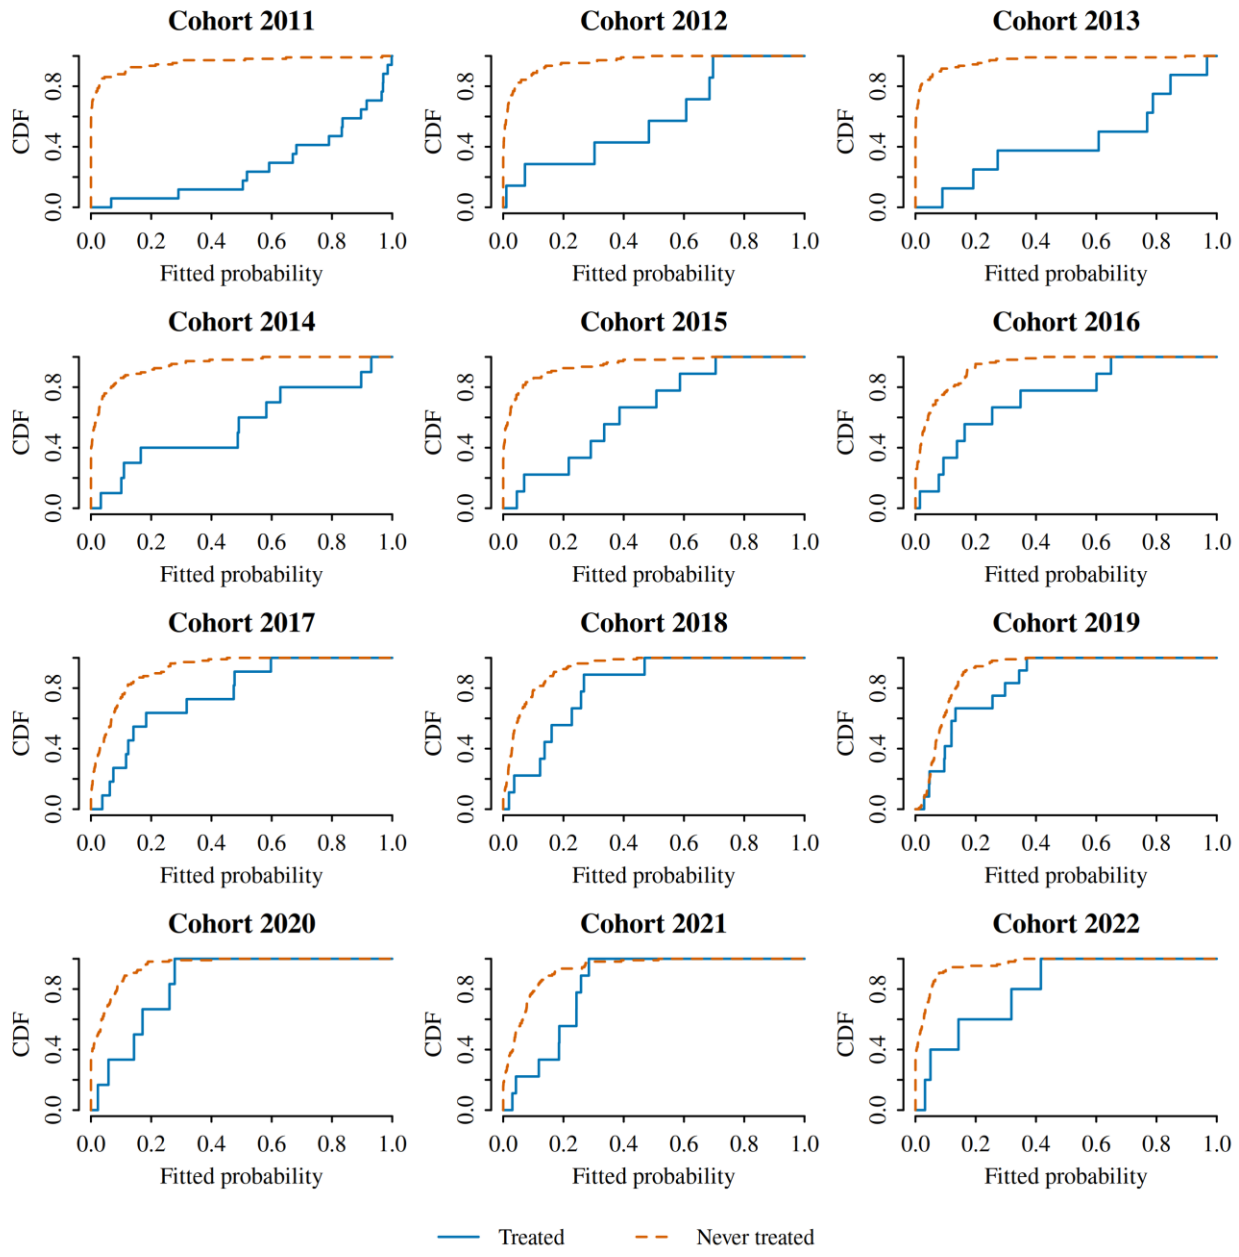


Figure A1. Cohort-specific raw propensity-score distributions.

Notes: Blue solid curves show treated counties; orange dashed curves show the 108 never-treated controls. The vertical axis is the empirical cumulative distribution function (CDF), the fraction of that group with a score at or below the horizontal value. Curves use fitted probabilities before bounds of 0.02 and 0.98 and are unweighted. The main IPW/DR score model uses covariates measured at $g-1$; 112 treated counties span 12 cohorts during 2011–2022. Curves closer together indicate more similar fitted-score distributions, but do not establish covariate balance. The private-employment fit excludes three controls and is summarized in Table A1. These are overlap diagnostics without confidence bands.